

ON FOUNDATIONS OF PROTECTING COMPUTATIONS

zur Erlangung des akademischen Grades eines

Doktors der Naturwissenschaften

von der KIT-Fakultät für Informatik
des Karlsruher Instituts für Technologie (KIT)

genehmigte

Dissertation

von

Thomas Agrikola

geboren am 03. April 1991 in Karlsruhe

Tag der mündlichen Prüfung: 07. Mai 2021

1. Referent: Prof. Dr. Jörn Müller-Quade
2. Referent: Prof. Dr. Dennis Hofheinz

ON FOUNDATIONS OF PROTECTING COMPUTATIONS

Dissertation by Thomas Agrikola in partial fulfillment of the requirements for the degree of Doctor of Natural Sciences (Dr. rer. nat.) accepted by the KIT Department of Informatics of the Karlsruhe Institute of Technology (KIT)

Day of examination: May 7, 2021

First referee: Prof. Dr. Jörn Müller-Quade, KIT

Second referee: Prof. Dr. Dennis Hofheinz, ETH Zürich

Published by KIT, Karlsruhe, Germany

DOI: [10.5445/IR/1000133798](https://doi.org/10.5445/IR/1000133798)

ABSTRACT

Information technology systems have become indispensable to uphold our way of living, our economy and our safety. Failure of these systems can have devastating effects. Consequently, securing these systems against malicious intentions deserves our utmost attention.

Cryptography provides the necessary foundations for that purpose. In particular, it provides a set of building blocks which allow to secure larger information systems. Furthermore, cryptography develops concepts and techniques towards realizing these building blocks. The *protection of computations* is one invaluable concept for cryptography which paves the way towards realizing a multitude of cryptographic tools. In this thesis, we contribute to this concept of protecting computations in several ways.

PROTECTING COMPUTATIONS OF PROBABILISTIC PROGRAMS. An indistinguishability obfuscator (**IO**) compiles (deterministic) code such that it becomes provably unintelligible. This can be viewed as the ultimate way to protect (deterministic) computations. Due to very recent research, such obfuscators enjoy plausible candidate constructions.

In certain settings, however, it is necessary to protect *probabilistic* computations. The only known construction of an obfuscator for probabilistic programs is due to Canetti, Lin, Tessaro, and Vaikuntanathan, TCC, 2015 and requires an indistinguishability obfuscator which satisfies extreme security guarantees. We improve this construction and thereby reduce the requirements on the security of the underlying indistinguishability obfuscator.

Agrikola, Couteau, and Hofheinz, PKC, 2020

PROTECTING COMPUTATIONS IN CRYPTOGRAPHIC GROUPS. To facilitate the analysis of building blocks which are based on cryptographic groups, these groups are often overidealized such that computations in the group are protected from the outside. Using such overidealizations allows to prove building blocks secure which are sometimes beyond the reach of standard model techniques. However, these overidealizations are subject to certain impossibility results. Recently, Fuchsbauer, Kiltz, and Loss, CRYPTO, 2018 introduced the algebraic group model (**AGM**) as a relaxation which is closer to the standard model but in several aspects preserves the power of said overidealizations. However, their model still suffers from implausibilities. We develop a framework which allows to transport several security proofs from the **AGM** into the standard model, thereby evading the above impossibility results, and instantiate this framework using an indistinguishability obfuscator.

Agrikola, Hofheinz, and Kastner, EUROCRYPT, 2020

PROTECTING COMPUTATIONS USING COMPRESSION. Perfect compression algorithms admit the property that the compressed distribution is truly random leaving no room for any further compression. This property is invaluable for several cryptographic applications such as “honey encryption”

or password-authenticated key exchange. However, perfect compression algorithms only exist for a very small number of distributions. We relax the notion of compression and rigorously study the resulting notion which we call “pseudorandom encodings”. As a result, we identify various surprising connections between seemingly unrelated areas of cryptography. Particularly, we derive novel results for adaptively secure multi-party computation which allows for protecting computations in distributed settings. Furthermore, we instantiate the weakest version of pseudorandom encodings which suffices for adaptively secure multi-party computation using an indistinguishability obfuscator.

Agrikola, Couteau, Ishai, Jarecki, and Sahai, TCC, 2020

PUBLICATIONS

- [AH18] Thomas Agrikola and Dennis Hofheinz. “Interactively Secure Groups from Obfuscation.” In: *Public-Key Cryptography - PKC 2018 - 21st IACR International Conference on Practice and Theory of Public-Key Cryptography, Rio de Janeiro, Brazil, March 25-29, 2018, Proceedings, Part II*. (Work partially included in my master’s thesis.) 2018, pp. 341–370. DOI: [10.1007/978-3-319-76581-5_12](https://doi.org/10.1007/978-3-319-76581-5_12). EPRINT: <https://eprint.iacr.org/2018/010>. URL: https://doi.org/10.1007/978-3-319-76581-5_12 (cit. on pp. [9](#), [91](#), [92](#), [110](#), [111](#), [115](#), [116](#), [120](#)).
- [ACH20] Thomas Agrikola, Geoffroy Couteau, and Dennis Hofheinz. “The Usefulness of Sparsifiable Inputs: How to Avoid Subexponential iO.” In: *Public-Key Cryptography - PKC 2020 - 23rd IACR International Conference on Practice and Theory of Public-Key Cryptography, Edinburgh, UK, May 4-7, 2020, Proceedings, Part I*. Ed. by Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas. Vol. 12110. Lecture Notes in Computer Science. Springer, 2020, pp. 187–219. DOI: [10.1007/978-3-030-45374-9_7](https://doi.org/10.1007/978-3-030-45374-9_7). EPRINT: <https://eprint.iacr.org/2018/470>. URL: https://doi.org/10.1007/978-3-030-45374-9_7 (cit. on pp. [3–5](#), [17–19](#), [41](#), [46](#), [71](#), [81](#), [245](#)).
- [ACI+20] Thomas Agrikola, Geoffroy Couteau, Yuval Ishai, Stanislaw Jarecki, and Amit Sahai. “On Pseudorandom Encodings.” In: *Theory of Cryptography - 18th International Conference, TCC 2020, Durham, NC, USA, November 16-19, 2020, Proceedings, Part III*. Ed. by Rafael Pass and Krzysztof Pietrzak. Vol. 12552. Lecture Notes in Computer Science. Springer, 2020, pp. 639–669. DOI: [10.1007/978-3-030-64381-2_23](https://doi.org/10.1007/978-3-030-64381-2_23). URL: https://doi.org/10.1007/978-3-030-64381-2_23 (cit. on pp. [4](#), [5](#), [12](#), [15](#), [17–19](#), [147](#), [149](#), [153](#), [155](#), [245](#)).
- [AHK20] Thomas Agrikola, Dennis Hofheinz, and Julia Kastner. “On Instantiating the Algebraic Group Model from Falsifiable Assumptions.” In: *Advances in Cryptology - EUROCRYPT 2020 - 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10-14, 2020, Proceedings, Part II*. Ed. by Anne Canteaut and Yuval Ishai. Vol. 12106. Lecture Notes in Computer Science. Springer, 2020, pp. 96–126. DOI: [10.1007/978-3-030-45724-2_4](https://doi.org/10.1007/978-3-030-45724-2_4). EPRINT: <https://eprint.iacr.org/2020/070>. URL: https://doi.org/10.1007/978-3-030-45724-2_4 (cit. on pp. [4](#), [5](#), [9](#), [10](#), [17–19](#), [87](#), [89](#), [91](#), [130](#), [245](#)).

*We can only see a short distance ahead,
but we can see plenty there that needs to be done.*

— Alan Turing [Tur50]

ACKNOWLEDGMENTS

First and foremost, I would like to thank my advisor Dennis Hofheinz for his support and for awakening my interest in cryptography in the first place. Your technical intuition for cryptography is enviable. I also would like to thank Jörn Müller-Quade for making it possible to finish my PhD at KIT when Dennis decided to leave for ETH Zürich.

Furthermore, I would like to thank my co-authors Geoffroy Couteau, Yuval Ishai, Stanisław Jarecki, Julia Kastner and Amit Sahai for many inspiring discussions, for your ability to make hard problems seem easy and for your trust. It was an honor and a pleasure to work with you.

I cannot emphasize enough how grateful I am to Geoffroy. You were like a second advisor to me and without your support during a rather frustrating time, I might not be where I am now. Thank you!

Special thanks also goes to Michael Kloof and Sven Maier for thoroughly proofreading this thesis.

Also, I would like to thank all my former and present colleagues for the great working atmosphere and for making my time at the institute so enjoyable. In particular, I would like to thank Akin Ünal for decorating your office beyond comparison; Björn Kaidel for encouraging me to major in cryptography; Bogdan Ursu for the countless hours of trying to best me at table tennis; Brandon Broadnax for a great time at the conference in Brazil and for always decimating my cough drops; Geoffroy Couteau for many fun and entertaining discussions about everything; Jessica Koch for the fun time we had when co-organizing the practice part for the security lecture together; Julia Hesse for Elin; Julia Kastner for sharing your limitless knowledge about squirrels; Lisa Kohl for fun hangman games during exam supervisions; Michael Kloof for your unparalleled love for formalisms; Sven Maier for an unending supply of Doctor Who DVDs; Valerie Fetzer for a great canoeing trip.

Finally, I would like to thank my family for their encouragement and support. Furthermore, I would like to thank my better half Dani for all her support and patience during my PhD. Together, we can overcome anything.

CONTENTS

1	INTRODUCTION	1
1.1	Protecting Computations of Probabilistic Programs	5
1.2	Protecting Computations in Cryptographic Groups	8
1.3	Protecting Computations – Using Compression?	12
1.4	Structure of this Thesis	17
2	PRELIMINARIES	19
2.1	Notations	19
2.2	Hardness Assumptions	20
2.2.1	Efficient Falsifiability	22
2.3	One-Way Functions	22
2.4	Puncturable Pseudorandom Functions	23
2.5	Obfuscation	25
2.5.1	Probabilistic Indistinguishability Obfuscation	27
2.6	Assuming Adversarial Knowledge	30
2.7	Public-Key Encryption and its Variants	33
2.7.1	Fully Homomorphic Encryption	34
2.8	Non-Interactive Zero-Knowledge Proof Systems	35
I	DOUBLY-PROBABILISTIC INDISTINGUISHABILITY OBFUSCATION	
3	INTRODUCTION	41
3.1	Technical Overview	43
4	PRELIMINARIES	47
4.1	Perfect Puncturable PRFs	47
4.2	Extremely Lossy Functions	47
4.2.1	Instantiating Extremely Lossy Functions	48
5	DOUBLY-PROBABILISTIC IO	51
6	CONSTRUCTION	59
6.1	Overview	59
6.2	Constructing Doubly-Probabilistic IO	60
7	LEVELED HOMOMORPHIC ENCRYPTION	73
II	INSTANTIATING THE ALGEBRAIC GROUP MODEL FROM OBFUSCATION	
8	INTRODUCTION	87
8.1	Technical Overview	90
8.2	Related Work	92
9	PRELIMINARIES	93

9.1	Notations and Cryptographic Groups	93
9.2	The Algebraic Group Model	93
9.3	Subset Membership Problem	94
9.4	Dual-mode NIWI Proof System	95
9.5	Re-Randomizable Public-Key Encryption	97
10	STATISTICALLY CORRECT PIO	99
10.1	Statistically Correct pIO	100
10.2	Puncturable Pseudorandom Functions	102
10.3	Construction	104
11	HOW TO SIMULATE EXTRACTION	109
11.1	Group Schemes and Wrappers	110
11.2	An Algebraic Wrapper	111
12	CONSTRUCTION	115
12.1	Main Theorem and Security Analysis	117
13	SIGNED ELGAMAL	129
III ON PSEUDORANDOM ENCODINGS		
14	INTRODUCTION	147
14.1	Flavors of Pseudorandom Encodings	149
14.1.1	Randomized, Computational Pseudorandom Encodings	151
14.2	Implications of our Results	154
14.2.1	New Results for Adaptively Secure Computation . . .	154
14.2.2	Steganography and Covert Multi-Party Computation .	154
14.2.3	Other Results	156
14.3	Negative Results	156
14.3.1	Deterministic, Statistical Pseudorandom Encodings . .	156
14.3.2	Deterministic, Computational Pseudorandom Encodings	157
14.3.3	Randomized, Statistical Pseudorandom Encodings . .	157
14.4	Open questions and subsequent work	158
14.5	Overview of Techniques	159
15	THE PSEUDORANDOM ENCODING HYPOTHESIS	165
15.1	The Pseudorandom Encoding Hypothesis with Setup	169
15.2	Static-to-Adaptive Transformation	171
16	PSEUDORANDOM ENCODINGS AND INVERTIBLE SAMPLING	177
16.1	The Invertible Sampling Hypothesis	177
16.2	The Invertible Sampling Hypothesis with Setup	178
16.3	Equivalence between PREH and ISH	181
16.3.1	Every Inverse Samplable Distribution can be Pseudo- randomly Encoded	181

16.3.2	Every Pseudorandomly Encodable Distribution can be Inverse Sampled	184
17	A TAXONOMY OF PSEUDORANDOM ENCODINGS	189
17.1	Deterministic Encoding Algorithm	189
17.1.1	Information-theoretic Guarantees and Compression	190
17.1.2	Computational Guarantees and Pseudoentropy	195
17.2	Randomized Encoding Algorithm	200
17.2.1	(Generalized) Extractable One-way Functions	200
17.2.2	Information-theoretic Guarantees and Privately Verifiable GEOWFs	201
17.2.3	Computational Guarantees and EOWFs with Common Auxiliary Information	204
17.3	Static Pseudorandom Encodings from IO	212
17.4	Bootstrapping Pseudorandom Encodings with URC	216
18	APPLICATIONS OF PSEUDORANDOM ENCODINGS	221
18.1	Adaptively Secure Multi-Party Computation	221
18.1.1	Adaptive MPC	221
18.1.2	Pseudorandom Encodings Imply Adaptive MPC	224
18.1.3	Adaptive MPC Implies Pseudorandom Encodings	226
18.2	Honey Encryption	229
18.3	Keyless Steganography	231
18.4	Deniable Encryption	233
18.5	Non-Committing Encryption	237
18.6	Super-Polynomial Encoding	238
	Outlook and Further Directions	
19	OUTLOOK AND FURTHER DIRECTIONS	245
	BIBLIOGRAPHY	247

LIST OF FIGURES

Figure 1.1	A high-level overview of the contributions of this thesis.	4
Figure 1.2	Overview of the strategy to reduce the amount of hybrids to a polynomial.	7
Figure 1.3	Description of group element representations in the algebraic wrapper.	11
Figure 1.4	Illustration of our notion of pseudorandom encodings.	14
Figure 1.5	An overview of the different notions of pseudorandom encodings.	15
Figure 2.1	Selective security game for pPRFs	24
Figure 2.2	Security games for circuit samplers.	29
Figure 2.3	One-way and extraction games for EOWFs	31
Figure 2.4	One-way and extraction game for EOWFs with common auxiliary input.	32
Figure 5.1	Overview of input samplers and states.	52
Figure 5.2	Simulatability of encodings games.	53
Figure 5.3	Statistical closeness game.	54
Figure 5.4	Restricted correctness games.	56
Figure 5.5	Indistinguishability game for dpIO	56
Figure 6.1	An ℓ -source dpIO scheme.	62
Figure 6.2	Simulator Sim used in proof of Theorem 6.1	63
Figure 6.3	Modified circuit used in proof of Theorem 6.1	68
Figure 6.4	Modified circuit used in proof of Theorem 6.1	70
Figure 7.1	Input samplers for construction of LHE from dpIO	74
Figure 7.2	LHE scheme from dpIO	75
Figure 7.3	Definition of the circuits C and tC.	76
Figure 7.4	Hybrid used in security proof of LHE scheme.	77
Figure 7.5	Circuit sampler used in the security proof LHE	78
Figure 10.1	The oracles used in the definition of correctness of ℓ -expanding pIO	102
Figure 10.2	The definition of our special extracting puncturable PRF F'	104

Figure 10.3	Construction of ℓ -expanding pIO.	105
Figure 11.1	The re-randomization and k -switching games.	113
Figure 12.1	Definition of the algorithms of the group scheme underlying the algebraic wrapper.	118
Figure 12.2	Definition of the algorithms of the algebraic wrapper.	119
Figure 12.3	Games for Lemma 12.4.	121
Figure 12.4	Modified addition and re-randomization circuits used in proof of Lemma 12.4.	123
Figure 12.5	Modified addition and re-randomization circuits used in proof of Lemma 12.4.	124
Figure 12.6	Modified addition and re-randomization circuits used in proof of Lemma 12.4.	125
Figure 12.7	Games used in proof of Lemma 12.5.	126
Figure 13.1	The Schnorr-signed ElGamal encryption scheme.	130
Figure 13.2	Game used in the security proof of Schnorr-signed ElGamal in Theorem 13.1.	135
Figure 13.3	Game used in the security proof of Schnorr-signed ElGamal in Theorem 13.1.	136
Figure 13.4	Game used in the security proof of Schnorr-signed ElGamal in Theorem 13.1.	137
Figure 13.5	Game used in the security proof of Schnorr-signed ElGamal in Theorem 13.1.	138
Figure 13.6	Game used in the security proof of Schnorr-signed ElGamal in Theorem 13.1.	139
Figure 13.7	Game used in the security proof of Schnorr-signed ElGamal in Theorem 13.1.	140
Figure 13.8	Game used in the security proof of Schnorr-signed ElGamal in Theorem 13.1.	141
Figure 14.1	An overview of the connections of pseudorandom encodings.	153
Figure 15.1	The pseudorandomness games.	166
Figure 15.2	The static games experiments with setup.	170
Figure 15.3	The adaptive pseudorandomness games with setup.	171
Figure 15.4	Adaptive pseudorandom encodings.	172
Figure 15.5	Hybrid games for the proof of adaptive correctness.	173
Figure 15.6	Adversary used for proof of adaptive correctness.	174
Figure 15.7	Hybrid games for the proof of adaptive pseudorandomness.	174

Figure 15.8	Adversary used in proof of adaptive pseudorandomness.	175
Figure 15.9	Adversary used for proof of adaptive pseudorandomness.	176
Figure 16.1	The closeness and invertibility experiments.	178
Figure 16.2	The static closeness and invertibility experiments with setup.	179
Figure 16.3	The adaptive closeness and invertibility experiments with setup.	180
Figure 16.4	Hybrids used in the proof of correctness of Lemma 16.1.	182
Figure 16.5	Hybrids used in the proof of pseudorandomness of Lemma 16.1.	183
Figure 16.6	Hybrids used in the proof of closeness of Lemma 16.2.	184
Figure 16.7	Hybrids used in the proof of invertibility of Lemma 16.2.	186
Figure 16.8	Proof sketch for the equivalence of the adaptive notions of pseudorandom encodings and invertible sampling.	188
Figure 17.1	Hardness and extractability games for GEOWFs.	201
Figure 17.2	Adversary and sampler used in proof of Theorem 17.7.	202
Figure 17.3	Hybrids used in the proof of Theorem 17.7.	202
Figure 17.4	Sampler and adversaries used in proof of Theorem 17.10.	205
Figure 17.5	Hybrids used in the proof of Theorem 17.10.	206
Figure 17.6	Hybrids used in the proof of Theorem 17.10.	207
Figure 17.7	Hybrids used in the proof of Lemma 17.9.	209
Figure 17.8	Perfectly correct pseudorandom encoding scheme from IO with universal setup.	213
Figure 17.9	Hybrids used in the proof of pseudorandomness for Theorem 17.13.	214
Figure 17.10	Perfectly correct pseudorandom encoding scheme from IO without universal setup.	215
Figure 17.11	Pseudorandom encoding scheme with common random string.	216
Figure 17.12	Hybrids used in the proof of correctness of Theorem 17.15.	217
Figure 17.13	Hybrids used in the proof of Lemma 17.10.	218
Figure 17.14	Hybrids used in the proof of pseudorandomness of Theorem 17.15.	218
Figure 18.1	Hybrids used in proof of Theorem 18.2.	225
Figure 18.2	Alternative and inverse sampler used in proof of Theorem 18.3.	227
Figure 18.3	Message recovery game.	230

Figure 18.4	DTE games.	231
Figure 18.5	Definition of the security game for keyless stegosystems.	232
Figure 18.6	A keyless stegosystem.	233
Figure 18.7	Indistinguishability of explanation games (static and adaptive).	234
Figure 18.8	Publicly deniable encryption scheme.	235
Figure 18.9	Hybrids used in the proof of IND-CPA security for Theorem 18.6.	236
Figure 18.10	Static indistinguishability of explanation games.	237
Figure 18.11	Non-committing encryption scheme in the global CRS model based on [CDMW09].	238
Figure 18.12	Simulator for the non-committing encryption scheme.	239
Figure 18.13	Pseudorandom encoding scheme with super-polynomial time encoding.	240

LIST OF TABLES

Table 12.1	Overview of proof steps of Lemma 12.4. Changes to previous games are highlighted in boxes.	121
------------	--	-----

ACRONYMS

\mathcal{BPP}	Bounded probability polynomial time. 21, 35
\mathcal{NC}	Nick’s class. 2
\mathcal{NP}	Non-deterministic polynomial time. 1–3, 21, 25, 35, 61, 205, 211
\mathcal{P}	Deterministic polynomial time. 1, 2, 21, 25
\mathcal{UP}	Unambiguous polynomial-time. 94
\mathcal{P}/poly	Non-uniform polynomial time. 21
AGM	Algebraic group model. iii, 9–12, 87–94, 109, 111, 245
AMPC	Fully adaptively secure multi-party computation. 221, 224, 226
AT	Anonymous transfer. 246

BLS	Boneh-Lynn-Shacham signature scheme. 12 , 89
CRS	Common reference string. xv , 15 , 46 , 63 , 95 , 149 , 150 , 152 , 154 , 156 , 158 , 161–163 , 171 , 172 , 178 , 181 , 187 , 194 , 197 , 199 , 204 , 205 , 207 , 212 , 221–224 , 226 , 227 , 231 , 237–239
DDH	Decisional Diffie-Hellman assumption. 42 , 43 , 45 , 46 , 48 , 49 , 80 , 95 , 129 , 240
DE	Deniable encryption. 233
DIO	Differing-inputs obfuscation. 27
DL	Discrete logarithm assumption. 93 , 129 , 130 , 134
DLIN	Decisional linear assumption. 96
DPIO	Doubly-probabilistic indistinguishability obfuscation. xii , 51–56 , 59–62 , 71 , 73–75 , 81 , 245
DTE	Distribution transforming encoder. xv , 156 , 158 , 230 , 231
EDDH	Exponential decisional Diffie-Hellman assumption. 48 , 49
ELF	Extremely lossy function. 7 , 8 , 42 , 44–49 , 61 , 63 , 66 , 80 , 81 , 238–240
EOWF	Extractable one-way function. xii , 31–33 , 88 , 152–154 , 200 , 204 , 205 , 211 , 212 , 216
FHE	Fully homomorphic encryption. 34 , 42 , 43 , 45 , 46 , 80 , 81 , 97
GEOWF	Generalized extractable one-way function. xiv , 153 , 200–202 , 204 , 211
GGM	Generic group model. 8 , 9 , 87 , 93
GKEA	Generalized knowledge-of-exponent assumption. 109
HE	Honey encryption. 230
IACR	International Association of Cryptologic Research. 4 , 17
IND-CCA1	Indistinguishability under (non-adaptive) chosen-ciphertext attack. 46 , 245
IND-CCA2	Indistinguishability under chosen-ciphertext attack. 89 , 129–131

IND-CPA	Indistinguishability under chosen-plaintext attack. xv , 3 , 34 , 73–76 , 78–80 , 129 , 213 , 233–237
IO	Indistinguishability obfuscation. iii , xiv , 2–8 , 10–12 , 15 , 17 , 18 , 21 , 25–27 , 29 , 33 , 41–46 , 60 , 71 , 73 , 80 , 81 , 89 , 91 , 92 , 97 , 99 , 149 , 152–154 , 158 , 159 , 187 , 211 , 215 , 231 , 245
ISH	Invertible sampling hypothesis. 151 , 152
KDM	Key-dependent message attack. 42 , 45 , 46 , 73
KEA	Knowledge-of-exponent assumption. 30–33
KEM	Key-encapsulation mechanism. 129
KLSTS	Key-less stegosystem. 231
LHE	Leveled homomorphic encryption. xii , 8 , 45 , 73–75 , 80 , 81
LPN	Learning-pairity-with-noise assumption. 2 , 245
LWE	Learning-with-errors assumption. 2 , 3 , 36 , 43 , 153 , 158 , 159 , 187 , 200 , 245
MPC	Multi-party computation. 3 , 16 , 17 , 149 , 151–155 , 171 , 187 , 221 , 222 , 229 , 237 , 241 , 246
NCE	Non-committing encryption. 237
NIWI	Non-interactive witness-indistinguishable proof system. 95 , 96 , 115 , 117 , 212
NIZK	Non-interactive zero-knowledge proof system. 22 , 35 , 36 , 44 , 59–61 , 63 , 96 , 117 , 153 , 199 , 200 , 204 , 205 , 210–212
OT	Oblivious transfer. 3 , 224
OWF	One-way function. 23 , 33 , 97 , 153
PAKE	Password-authenticated key exchange. 150 , 158
PIO	Probabilistic indistinguishability obfuscation. xii , xiii , 5–7 , 27–29 , 42–46 , 60 , 80 , 99–102 , 104 , 105 , 116 , 117 , 127
PKE	Public-key encryption. 33 , 46 , 97 , 153 , 221 , 233 , 234 , 237
PKSTS	Public-key stegosystem. 231
PPAD	Polynomial parity argument (directed). 3
PPRF	Puncturable pseudorandom function. xii , 23 , 24 , 47 , 102–105

PPT	Probabilistic polynomial time. 19 , 21–26 , 28 , 31–36 , 47 , 52 , 53 , 55 , 56 , 62 , 63 , 65 , 66 , 68–71 , 73 , 76–80 , 93–97 , 101 , 102 , 105 , 106 , 112 , 113 , 116 , 120–124 , 126 , 130 , 148 , 151 , 154 , 159–161 , 163 , 165–185 , 187 , 197 , 200 , 201 , 205 , 207 , 208 , 210 , 212 , 214 , 216 , 217 , 219 , 221–229 , 231–234 , 236 , 237 , 239–241
PRE	Pseudorandom encoding. 149
PREH	Pseudorandom encoding hypothesis. 14 , 148 , 168
PRF	Pseudorandom function. xii , 3 , 23 , 24 , 29 , 43 , 45 , 47 , 59–61 , 64 , 69–71 , 94 , 100 , 102–105 , 109 , 212 , 214 , 215
PRG	Pseudorandom generator. 2 , 167 , 192 , 212 , 214
QRA	Quadratic residuosity assumption. 153 , 157 , 163 , 199
ROM	Random oracle model. 10 , 88 , 90 , 91
SKE	Symmetric-key encryption. 33 , 230
SMP	Subset membership problem. 94 , 95
SNARG	Succinct non-interactive arguments. 27
SS-CCA	Chosen-coverttext attack. 231
SSB	Somewhere statistically binding. 3
SXDH	Symmetric external Diffie-Hellman assumption. 96
UC	Universal composability framework. 223 , 224 , 226
UPRE	Universal proxy re-encryption. 46
URC	Uniformly random CRS. 171 , 181 , 216 , 219 , 229
VBB	Virtual black-box obfuscation. 25
VLF	Very lossy function. 238–241

1

INTRODUCTION

WHY CRYPTOGRAPHY? In the current information age, our economy is primarily built on information technology. This applies for our society as well. Access to and distribution of information has become an essential part of our daily lives.

Many critical infrastructures rely on information technology systems. This includes, for instance, control systems for nuclear power plants, military drones, smart grids, online payment systems, internal databases, medical records, personal correspondence and social media. Failure of these systems has severe consequences ranging from damage of social lives and reputations, over damage to economic infrastructures, to endangerment of human lives. In fact, on account of this dependency on information technology, cyberwarfare, cyberterrorism and cybercrime have become as devastating as their physical counterparts.

Consequently, the necessity to ensure security of information systems has become pervasive.

WHAT IS CRYPTOGRAPHY? Cryptography provides building blocks which are capable of creating an *asymmetry* in the computational costs between honest users and malicious entities. These building blocks can then be used to ensure security in larger information systems. One primary focus of cryptography is to provide *rigorous mathematical proofs* of such building blocks' properties and guarantees.

One widely used approach towards achieving such an asymmetry is to resort to computational complexity theory. This approach is referred to as asymptotic security. The philosophy behind asymptotic security is to be able to scale the security of a building block to withstand *arbitrarily powerful* adversaries, while the computational cost for honest users remains affordable. That is, any supposed computational adversarial power, even magnitudes of whole computing centers, can be outmatched by computationally limited users.

For most cryptographic building blocks, the task of breaking them is in \mathcal{NP} , hence non-deterministically solvable in polynomial time (since solution candidates can be verified in polynomial time). Consequently, given our current understanding of computational complexity, asymptotic security cannot be achieved unconditionally but requires computational hardness assumptions. Clearly, $\mathcal{P} \neq \mathcal{NP}$ is a necessary assumption.

However, the assumption $\mathcal{P} \neq \mathcal{NP}$ is not sufficient for cryptography since the definition of \mathcal{NP} only considers hardness in the worst case. In the context of cryptography, not only need there be hard problems, but furthermore sampling hard instances of these problems must be efficient since this task is performed by an honest user. The initial foundational study of hardness-on-the-average is due to Levin [Lev86].

SOURCES OF HARDNESS. One central goal of cryptography is to identify and analyze potential sources of computational hardness. Number theory gives rise to several computational hardness assumptions such as (variants of) the factoring assumption [RSA78] and assumptions in cyclic groups [DH76]. Some time later, codes [BFKL94] and lattices [Ajt96] were discovered to yield new sources of computational hardness.

An alternative road towards secure cryptographic building blocks is to base their security on *generic* hardness assumptions which can be instantiated from several sources of hardness. The most fundamental generic hardness assumption is the existence of one-way functions [DH76], i.e., functions which can efficiently be evaluated but are hard to invert. To justify security based on generic assumptions which do not admit any inherent algebraic structure (such as the aforementioned sources of hardness) is rather appealing from a theoretical perspective. Furthermore, if one particular source of hardness was to fall victim to modern cryptanalysis, there remain several other sources that can equally be used.

In the recent past, another generic hardness assumption has been conceived – the assumption that computations within certain programs are protected. In the following, we shed light on several angles of protecting computations.

CODE OBFUSCATION. Compiling program code such that it becomes *provably* unintelligible can be seen as the ultimate way to protect computations.¹ The concept of code obfuscation was first formalized in [Hadoo; BGI+12]. An obfuscator is a compiler which preserves the functionality of programs but makes the original code unintelligible. In search of the “right” definition of unintelligibility, the notion of indistinguishability obfuscation (IO) emerged as the “best possible” [GR07], achievable and useful [GGH+13] notion. An indistinguishability obfuscator guarantees that the obfuscations of two programs (represented as Boolean circuits) which compute the same functionality are indistinguishable from each other. Note that since IO can exist even if $\mathcal{P} = \mathcal{NP}$, an additional mild assumption is usually necessary to instantiate a cryptographic building block from IO.

Starting from the line of work [Lin16; LV16; LT17; Lin17], there has been substantial progress towards instantiating IO from (falsifiable) standard assumptions [AJL+19; Agr19; JLMS19; BDGM20a; GJLS20; JLS20; GP20; BDGM20b]. Most remarkably, Jain, Lin, and Sahai [JLS20] introduce the first instantiation of IO based solely on well-founded assumptions, namely pairing-friendly cryptographic groups, *LWE*, *LPN* and *PRGs* in \mathcal{NC}_0 . Another remarkable candidate is due to Brakerski, Döttling, Garg, and Malavolta [BDGM20b] which is based only (a slightly non-standard variant of) circularly secure *LWE*.

Another line of work seeks to remove pairing based assumptions from the above works [Agr19; AP20]. Very recently, [WW20], building upon [BDGM20a], instantiate IO from lattice-based assumptions in conjunction with a novel but falsifiable oblivious sampling assumption.

Furthermore, there has been significant progress towards instantiating IO from graded multilinear maps (graded encoding schemes) [GGH15; BMSZ16; GMM+16; CVW18; CHVW19]. These current candidates for graded encoding schemes only provide heuristic security (that is, they are proven secure in idealized models) but withstood intense cryptanalysis [CHL+15; MSZ16; CLLT16; ADGM17; CLLT17; CGH17; CVW18].

¹ This should not be confused with practical code obfuscation used in the realm of software engineering which does not provide any provable guarantees, e.g., [CTL98a; CTL98b].

Another incomparable line of work base **IO** directly on assumptions about tensor products [GJK18] and affine determinant programs [BJ+20].

Protecting computations by making code unintelligible emerged as an extremely powerful tool unveiling undreamt-of possibilities. Following the positive result of [GGH+13], a wide variety of applications has been proposed, e. g., functional encryption [GGH+13], deniable encryption [SW14], fully homomorphic encryption [CLTV15], two-round secure multi-party computation [GGHR14], multi-party key exchange [BZ14], secret sharing for \mathcal{NP} [KNY14], **IO** for Turing machines [KLW15], succinct randomized encodings [BGL+15], constant-round concurrent zero-knowledge [CLP15], reductions from the hardness of **PPAD** to the hardness of cryptographic problems [BPR15; GPS16], trapdoor permutations [BPW16; CL18], non-interactive zero-knowledge proof systems (in the common random string model) and non-interactive witness-indistinguishable proof systems [BP15a]. This includes applications that seemed beyond our reach for a long time and became feasible thanks to **IO**.

Not only does **IO** yield intriguing new applications, but it also triggers the development of new techniques. These techniques prove useful even in settings without **IO**. The core technique which was developed for **IO** is “puncturing” [SW14]. This motivated the study and use of puncturable pseudorandom functions (**PRFs**). Puncturing proved useful in further areas and puncturable **PRFs** are at the heart of recent constructions like homomorphic secret sharing [BGI17], oblivious transfer with small communication [BCG+19a], and pseudorandom correlation generators [BCG+19b]. Further, **IO** motivated the design of “somewhere statistically binding (**SSB**) hashing” [HW15; OPWW15] as an abstraction of useful properties for the analysis of **IO**-based constructions. This notion of **SSB** hashing has inspired recent results like laconic **OT** [CDG+17] and trapdoor hash functions [DGI+19] which both have significant applications for **MPC**. Due to [GGG+14], **IO** can be used to achieve two-round **MPC**. This inspired the notion of interactive garbling schemes [FJNT15] which led to the result that two-round **MPC** can be instantiated from **OT** [GS18; BL18a]. Moreover, **IO** techniques proved useful for identity-based encryption [DG17], verifiable random functions [Bit17; GHKW17] and the Fiat-Shamir paradigm and correlation intractability allowing to emulate obfuscation-like properties from key-dependent message security [CCRR18; HL18; CCH+19; PS19]. Furthermore, **IO** motivated the study of multilinear maps. The candidate construction from [GGH15] can indeed be proven secure in certain settings under the **LWE** assumption paving the way for novel **LWE**-based constructions such as privately constrained **PRFs** [CC17], compute-and-compare obfuscation [WZ17] and separations between circular security and **IND-CPA** security for secret-key encryption [GKW17].

Thereby, the study of **IO** helps to improve our understanding of cryptography.

OVERVIEW OF CONTRIBUTIONS. In the following, we provide a high-level overview of the contributions of this thesis, see Figure 1.1.

We improve concepts regarding the protection of computations in several ways. In [ACH20], we improve constructions in the context of probabilistic obfuscation – a variant of indistinguishability obfuscation for randomized programs. To be specific, previous constructions of probabilistic **IO** required

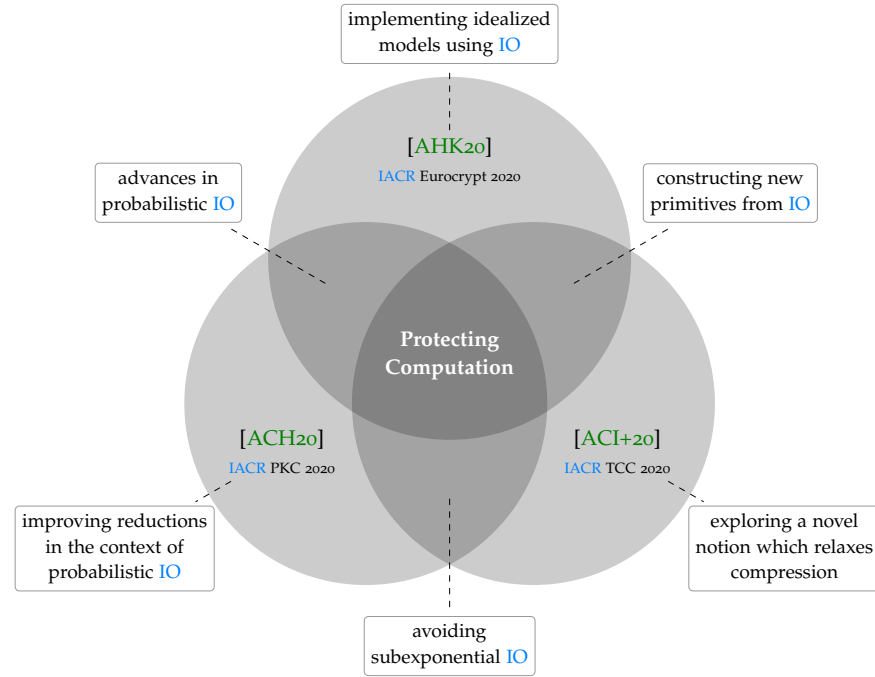


Figure 1.1: A high-level overview of the contributions of this thesis.

² Subexponentially secure IO provides guarantees even against subexponential-time adversaries. Cryptographic hardness assumptions classically only require security against polynomial-time adversaries which is a much weaker requirement.

subexponentially secure² IO , a very strong assumption. We manage to reduce this requirement to standard polynomially secure IO . In [ACI+20], we relax the notion of perfect compression to harness the useful property that compressed samples appear to be uniform randomness for a broader class of distributions. We derive a notion we call “pseudorandom encodings” which we prove to exist for all efficiently samplable distributions assuming IO . Further, our study of pseudorandom encodings reveals interesting and unexpected relations to secure multi-party computation which allows to protect computations in distributed systems. In [AHK20], we use IO to instantiate an idealized model in the context of cryptographic groups which protects – or shields – computations in the group. While these contributions are centered around the protection of computations, this is not their only common ground.

[ACH20; AHK20] both make progress in the field of probabilistic IO . [ACH20] reduces the necessary hardness assumptions and [AHK20] improves the correctness guarantee offered by probabilistic IO . More precisely, previous notions of probabilistic IO information-theoretically leaked whether an output was produced via the original or the obfuscated program while our novel notion of probabilistic IO conceals even this information. Hence, we make progress in the context of protecting randomized computations.

[ACH20] avoids assuming subexponential IO . Surprisingly, [ACI+20] also avoids this assumption in a completely different context: adaptively secure multi-party computation with sublinear communication complexity [CsW19]. Thereby, we make progress towards realizing protection of computations based on weaker and more reliable hardness assumptions.

Furthermore, we make progress in deriving novel building blocks from the assumption that certain computations are protected. In [ACI+20], we construct a novel primitive (our relaxation of compression) from IO . This is also true for [AHK20], where we construct a novel building block from IO

which can be used to transport proofs from an idealized group model to the standard group model.

In the following, we elaborate on these results in more detail.³

1.1 PROTECTING COMPUTATIONS OF PROBABILISTIC PROGRAMS

Presently, almost all conceivable cryptographic building blocks can be instantiated from **IO** (and mild generic assumptions).⁴ However, some building blocks currently require *subexponentially* secure **IO**, i. e., require security even against subexponential-time adversaries with time complexity in $\mathcal{O}(2^{\lambda^\epsilon})$ (for some $0 \leq \epsilon < 1$).⁵ This is a much stronger requirement compared to the standard notion of security which covers polynomial-time adversaries.

There are several promising roads towards achieving **IO**. The approaches based on functional encryption start from subexponential assumptions and directly yield subexponential **IO** “for free”. This includes the recent breakthrough results [BDGM20a; GJLS20; JLS20; GP20; BDGM20b] outlined above. However, this does not apply to the approaches using graded encoding schemes [GGH15; BMSZ16; GMM+16; CVW18; CHVW19] and the incomparable approaches from tensor products [GJK18] or affine determinant programs [BJ+20]. For these approaches, security against polynomial-time adversaries and security against subexponential-time adversaries does make a significant difference. Furthermore, constructing a building block from polynomial **IO** relates the security of that building block more tightly to **IO**. This can give rise to more insightful techniques and help to improve our understanding of that building block. Thus, it is desirable to avoid subexponentially secure **IO**.

RELATED WORK ON AVOIDING SUBEXPONENTIAL IO. Avoiding subexponentially secure **IO** has been studied in [GPS16; GPSZ17; LZ17; KLMR18]. These works follow the approach to replace (subexponentially secure) **IO** with polynomially secure functional encryption. This strong variant of public-key encryption allows to implement several **IO**-based proof techniques without relying on guarantees against subexponential-time adversaries. However, this strategy requires the building block in question to meet certain structural conditions. We seek to provide a more general solution.

PROBABILISTIC OBFUSCATION. A certain class of building blocks requires to protect *probabilistic* computations. The notion of probabilistic indistinguishability obfuscation (**pIO**) was put forth by Canetti, Lin, Tessaro, and Vaikuntanathan [CLTV15] and requires subexponentially secure **IO**. On a high level, **pIO** not only protects the program code, but additionally protects the randomness which is used to evaluate the randomized program from a user of the program. Syntactically, a probabilistic indistinguishability obfuscator compiles a randomized circuit into a deterministic circuit which has “almost” the same functionality. Recall that standard **IO** guarantees indistinguishability of obfuscated programs, provided that these programs behave identically on all inputs. For a useful definition, this condition of identical behavior needs to be relaxed such that **pIO** is applicable for, e. g.,

³ Note that parts of the following are loosely inspired by the introductions of the underlying publications [ACH20; AHK20; ACI+20].

⁴ Except for the notable examples of collision resistant hash functions or fully-homomorphic signatures.

⁵ This definition of subexponential security is a slight simplification which suffices for this overview.

programs which produce identical output distributions. Informally, a probabilistic indistinguishability obfuscator guarantees that the obfuscations of two randomized programs are indistinguishable, provided that these programs induce suitably close output distributions. Currently, the only known construction of **pIO** applies standard **IO** on a de-randomized version of the original program. De-randomized programs derive their randomness deterministically by applying a pseudorandom function⁶ F on their input, that is,

$$\text{piO}(C) := \text{iO}(\overline{C}), \text{ where } \overline{C}(x) := C(x; F(x)).$$

⁶ A pseudorandom function is a function which is indistinguishable from a truly random function.

The reason why polynomially secure **IO** is insufficient for **pIO** is best explained by example. Consider the two randomized programs C_0 and C_1 , where C_b on input of $x \in \{0, 1\}^\lambda$ and (internal) randomness $r \in \{0, 1\}^\lambda$, outputs $x \oplus r \oplus b^\lambda$ (for $b \in \{0, 1\}$). Clearly, these programs both induce the uniform distribution over $\{0, 1\}^\lambda$. Hence, **pIO** should guarantee indistinguishability between $\text{piO}(C_0) = \text{iO}(\overline{C}_0)$ and $\text{piO}(C_1) = \text{iO}(\overline{C}_1)$. However, **IO** does not provide any guarantees for circuits which do not behave fully identically. In fact, there is not even one input x such that the de-randomized circuits \overline{C}_0 and \overline{C}_1 produce the same output.

The only known strategy to deal with this issue is to prove that $\text{iO}(\overline{C}_0)$ and $\text{iO}(\overline{C}_1)$ are indistinguishable for every single input x (applying the security property of **IO** every time). Since the number of inputs is exponential, this proof technique inherently requires (sub)exponential security of **IO**.

We aim at realizing a sufficient protection of randomized computations without having to rely on subexponentially secure **IO**.

DOUBLY-PROBABILISTIC IO. For the purpose of avoiding subexponential **IO**, a natural solution would be to reduce the number of inputs for the probabilistic circuit. However, obfuscation with only polynomially many inputs is trivially solvable by enumerating each input-output-pair and can thus not be useful towards constructing complex cryptographic building blocks.

We observe that in many cases, it suffices to consider probabilistic circuits expecting inputs from a small number of efficiently samplable distributions. To exemplify, consider the two distributions induced by encryption of a single bit. In many applications of **pIO**, the obfuscated circuit expects only inputs which originate from such distributions. Towards reducing the number of meaningful inputs which need to be considered in the security proof, we restrict the correctness guarantee provided by **pIO** to only hold for “well-generated” inputs. To exclude all “non-well-generated” inputs from the security proof, we certify each input in the domain of an input sampler with a non-interactive zero-knowledge proof and define the probabilistic circuit to reject all uncertified inputs.⁷ However, while the number of distributions is small, the support of each such distribution is still exponential.

⁷ This approach only certifies membership, no distributional property.

The restricted correctness property lays the foundation to an orthogonal way to reduce the amount of well-generated inputs. Namely, suppose we could narrow the randomness space of the input samplers to a polynomial amount, the amount of well-generated inputs necessarily shrinks to a polynomial. However, shrinking the randomness space from an exponential size to a polynomial size is far from trivial. Consider the example of bit-encryption distributions from above. For these distributions, it is essential to use

good randomness to ensure secrecy of the encrypted bit. If the randomness space for encryption is only polynomial, exhaustive search can recover the encrypted bit.

Contrary to the above intuition, it is indeed possible to sparsify the randomness space to a polynomial size using a special tool called *extremely lossy functions* introduced by Zhandry [Zha16]. Extremely lossy functions (ELFs) are functions which can be set up in two computationally indistinguishable modes – an injective mode and an extremely lossy mode. In injective mode, the range of the ELF is exponential whereas in extremely lossy mode, the range of the ELF merely comprises a polynomial quantity. Such a notion requires a slightly non-black-box use of adversary. More precisely, such an indistinguishability cannot hold against arbitrary polynomial-time adversaries since finding collisions requires only polynomial resources. Instead, the size of the range in extremely lossy mode needs to be sufficiently large to fool the particular adversary in question.

We use extremely lossy functions to pre-process the randomness for the input sampler, see also Figure 1.2 for an overview. In injective mode, the original distribution is not affected. During the proof, we can unnoticeably (for the particular adversary) switch the ELF into extremely lossy mode, thereby sparsifying the randomness space of the input sampler and hence the amount of well-generated inputs.




Figure 1.2: Overview of the strategy to reduce the amount of hybrids to a polynomial. The randomized program P receives inputs from some input distribution S^{in} and derives its randomness using a pseudorandom function F on this input. The input distribution S^{in} receives an input x from a small input space and a random tape r (from an exponentially large randomness space). Pre-processing r with the extremely lossy function G allows to sparsify the amount of possible inputs to P to a polynomial amount.

Using the strategy above, we are able to use standard polynomially secure IO by “pushing” the subexponential loss to a much more standard and better-understood assumption on cyclic groups which is necessary to instantiate the extremely lossy function. We provide a general framework called “doubly-probabilistic IO” which can be used to avoid subexponentially secure IO in many pIO-based applications without making any assumptions on the structure of the obfuscated programs. We only need to require that inputs for the obfuscated circuit originate from a small number of efficiently samplable distributions.

APPLICATIONS. We are able to apply doubly-probabilistic **IO**, for instance, in the context of leveled homomorphic encryption (**LHE**). **LHE** allows to evaluate fixed-depth circuits on the contents of ciphertexts.

The **LHE** construction due to Canetti et al. [**CLTV15**] requires to obfuscate a probabilistic program which expects two ciphertexts as input, decrypts them, evaluates say the logical NAND on the plaintexts, and re-encrypts the result. Clearly, this program expects two inputs which both come from one of two efficiently samplable distributions: the ciphertext distribution of the plaintext 0 or the ciphertext distribution of the plaintext 1. Hence, by using doubly-probabilistic **IO**, we are able to push the subexponential assumptions for **LHE** away from **IO** to **ELFs** which can be based on well-established and well-analyzed assumptions.

Thus, our notion of doubly-probabilistic **IO** provides a general framework for reducing the necessary assumptions for protecting probabilistic computations. We refer the reader to Part **I** of this thesis for more details.

SUBSEQUENT WORK. Our notion of doubly-probabilistic **IO** has already found further application. Döttling and Nishimaki [**DN18**] are able to avoid subexponentially secure **IO** in the following context. [**DN18**] consider the task of converting ciphertexts produced with *any* public-key encryption scheme into re-randomized ciphertexts under *any possibly different* public-key encryption scheme. They call this notion universal proxy re-encryption and instantiate it using probabilistic **IO**. Furthermore, they observe that using our doubly-probabilistic **IO**, they are able to avoid subexponentially secure **IO**.

1.2 PROTECTING COMPUTATIONS IN CRYPTOGRAPHIC GROUPS

Since the seminal work of Diffie and Hellman [**DH76**], certain cyclic groups have emerged as a major source of computational hardness giving rise to a great variety of cryptographic building blocks such as public-key encryption and digital signatures [**ElG85**], non-committing encryption [**Bea97**], identity-based encryption [**BF01**], non-interactive zero-knowledge proof systems [**GOS12**], and many more.

Protecting computations is a central concept in the context of cryptographic groups. Building blocks or novel hardness assumptions which are based on cryptographic groups are commonly analyzed in an idealized group model. This idealized model assumes that the computations in the group are “protected” in the sense that an adversary has no control over the group itself, but is granted access to a restricted interface providing only generic group operations. This model is called the generic group model (**GGM**) and was formalized by Shoup [**Sho97**] and Maurer [**Mau05**]. An adversary who only uses this generic interface is called a generic adversary. Given a security proof in the **GGM**, no generic adversary can succeed in breaking the proven security properties. In other words, in order to break a generically proven security property in some concrete group, every adversary is required to exploit the concrete encodings of group elements in that group in a nontrivial way. Hence, the **GGM** provides a sanity check for the security of building blocks or novel assumptions. However, there exist building blocks which

can be proven secure in the **GGM** but are insecure when instantiated with *any* concrete group [Deno2]. Hence, generic groups cannot exist. In my master’s thesis which led to the publication [AH18], we demonstrate that fundamental proof strategies developed for the **GGM** can be transported into the standard model at the expense of uniqueness of group element encodings. However, due to this lack of unique group element encodings and the inability to extract a unique bitstring from group element encodings, the applicability of the resulting group is very limited. Without the ability to extract a bitstring which is unique for every encoding of the same group element, it is impossible to extract, e. g., a unique encryption key from a group element. This result is not part of this thesis, but [AHK20], which is part of this thesis, constitutes follow-up work.

ALGEBRAIC GROUPS. The algebraic group model [FKL18] is a relaxation of the generic group model conceived to avoid the above mentioned impossibility results while preserving the power and generality of the **GGM**. While the **GGM** only considers the class of generic adversaries, the **AGM** considers the broader class of so-called algebraic adversaries. In contrast to generic adversaries, algebraic adversaries can make arbitrary use of group element encodings but must explain each output group element in terms of their input group elements. More precisely, if an algebraic adversary outputs a group element h given the group elements g_1, \dots, g_ℓ , the adversary must additionally provide a list of exponents z_1, \dots, z_ℓ , such that $h = \prod_{i=1}^{\ell} g_i^{z_i}$. This formalization of algebraic adversaries is equivalent to the more standard definition that for every adversary who outputs a group element, there exists an efficient extractor which is able to derive the exponents z_1, \dots, z_ℓ as above. Every generic adversary is algebraic, since a generic adversary needs to use a generic interface of oracles to obtain new group elements.

Hence, while the **AGM** does not protect computations to the same extent as the **GGM**, it still offers a meaningful protection of computations in the group. Namely, even though an adversary can derive group elements in an arbitrary way disregarding the actual group operation, ultimately, he has to be able to explain his doing in terms of the group operation. Thus, the **AGM** protects the computations in the group by nullifying any advantage in deriving group elements in any other way than using the group operation directly.

INHERENT OBSTACLES. The **AGM** is equivalent to a very strong form of the knowledge-of-exponent assumption [Dam92; WS07; KP19]. Knowledge-type assumptions basically state that if an adversary can produce an output, this adversary *must know* how he computed this output. To reason about adversarial knowledge is a common concept in cryptography. Adversarial knowledge of information is modeled by the ability to algorithmically extract this information from the adversary’s code using an efficient extractor. Consequently, while escaping the impossibility results of the **GGM**, the **AGM** nevertheless suffers from several inherent obstacles.

Cryptographic hardness assumptions can be classified into falsifiable and non-falsifiable ones based on the complexity of falsifying them. An assumption is falsifiable if an observation can show it to be false. More precisely, a cryptographic assumption is falsifiable if the assumption can be formalized as an interaction between an *efficient* challenger and an adversary such that

the challenger can efficiently provide the adversary with a challenge and is able to efficiently verify the adversary's solution attempt. This classification was put forth by Naor [Na003]. Lacking the property of efficient falsifiability renders assessing the plausibility of assumptions extremely difficult. Hence, it is (arguably) undesirable to use knowledge-type assumptions (such as the [AGM](#)) in cryptography.

Moreover, the [AGM](#) suffers from a further obstacle. Certain knowledge-type assumptions such as the [AGM](#) which allow an additional unstructured input conflict with [IO](#) [BCPR16]. Recall that every adversary in the [AGM](#) must explain his output elements as group operations in terms of his input elements. Suppose an adversary receives an obfuscated program as additional input. Suppose further that this obfuscated program, on input of a group element g , computes $g^{F(g)}$, where F is a pseudorandom function. An adversary who receives as input a group element g , evaluates this obfuscated program on g and outputs the result, is certainly unable to explain his output as algebraic operations in terms of his input group element, and can hence not be algebraic as per [FKL18].

Therefore, it is impossible to instantiate the [AGM](#) from [IO](#). Instead, we aim at an instantiation of a relaxation of the [AGM](#) which preserves the generality of the [AGM](#) and can be based solely on falsifiable assumptions.⁸ Hence, we tackle the following intriguing question:

⁸ The assumption that [IO](#) exists is falsifiable in the sense of Naor.

To what extent can knowledge-type properties be simulated while keeping all assumptions purely falsifiable?

IMPLEMENTING IDEALIZED MODELS. The use of idealized models which protect computations by shielding it from adversaries is a popular paradigm in cryptography. Aside from cryptographic groups, building blocks which use cryptographic hash functions are often analyzed in an idealized model – the random oracle model ([ROM](#)). This model was introduced by Bellare and Rogaway in [BR93] but was implicitly already used in [FS87; GGM84b; GGM86]. In the random oracle model, cryptographic hash functions are modeled as oracles, similar in spirit to the generic group model due to Shoup [Sho97]. Building blocks which come with a proof of security in the [ROM](#) are often considerably more efficient compared to building blocks which are provably secure in the standard model.

Replacing idealized models with concrete standard-model implementations is an intriguing problem in cryptography. Hofheinz and Kiltz [HK12] follow this aim by introducing the notion of programmable hash functions. Programmable hash functions allow to transport certain proofs from the random oracle model into the standard model. Following [HK12], a line of work [FHPS13; HSW13; HSW14] transported proofs from the random oracle model to the standard model using multilinear maps or [IO](#). Our results follow this endeavor in the context of cryptographic groups.

INSTANTIATING THE [AGM](#) FROM [IO](#). In Part II of this thesis, which is based on the publication [AHK20], we develop a suitable relaxation of the [AGM](#) and provide an instantiation based on probabilistic [IO](#) and further mild and falsifiable assumptions.

[AGM](#)-based proofs heavily rely on the ability to extract a decomposition of the adversarially produced group elements. However, since we strive to

avoid knowledge-type assumptions, extracting this information from the adversary’s code is not a viable option. Hence, an alternative extraction mechanism needs to be conceived. We seek to extract a decomposition of group elements *from the group elements themselves*. This type of extraction must exclusively be available during a security proof and must be beyond reach for any user of the group – this includes honest users as well as malicious entities. Hence, we define a special trapdoor which is associated to the public parameters of the group to be necessary to extract the decomposition. This notion of extraction entirely avoids knowledge-type assumptions.

ALGEBRAIC WRAPPER. Albrecht, Farshim, Hofheinz, Larraia, and Paterson [AFH+16] lay the foundations for the construction of groups with non-unique encodings from IO. Groups with non-unique encodings allow each group element to carry auxiliary information inside its representation. We utilize the ability of the [AFH+16] framework to equip any given base group with non-unique encodings. The resulting group virtually “wraps” the group elements of the base group while preserving the original group structure and hardness guarantees. We exploit the thus achieved non-uniqueness of encodings to let each group element carry an *encrypted* decomposition of itself relative to some a priori fixed set of group elements. We refer to this fixed set of group elements as “basis”.

Hence, any valid group element is bound to contain a decomposition of itself. See Figure 1.3 for an overview of wrapped group element representations. In particular, this must be true for any *adversarially produced* group element. Consequently, if the corresponding decryption key is known, this decomposition can be extracted from any such group element representation. Intuitively, this forces any adversary to be *algebraic* in a similar sense as in the AGM. Thus, we denote this notion as “algebraic wrapper”.




Figure 1.3: Description of group element representations in the algebraic wrapper. Each group element h from the base group is accompanied by a ciphertext containing a decomposition of h with respect to the basis $\{g_1, \dots, g_\ell\}$. To ensure consistency between h and the encrypted decomposition, a non-interactive zero-knowledge proof π accompanies the representation forcing the encrypted decomposition to match the group element h .

SIMULATING PROOFS FROM THE AGM. For the purpose of proving some building block secure in the AGM, we need to reduce the hardness of some conjectured-to-be-hard problem in the group (e.g., finding discrete logarithms⁹) to the task of breaking that building block. More precisely, an AGM-based proof transforms any algebraic adversary on the building block into a solver for discrete logarithms in the group. In order to simulate this strategy in the standard model, it is essential to be able to deal with arbitrary (not necessarily algebraic) adversaries. In other words, we need to force an arbitrary adversary to provide a decomposition of his output group elements. The ability to extract the decomposition can be ensured by the algebraic wrapper as follows. The reduction (i.e., the discrete logarithm

⁹ The discrete logarithm $\text{dlog}_g(h)$ of h to the basis g is an integer x such that $g^x = h$.

solver) receives a discrete logarithm challenge “ $d\log_g(h) = \perp$ ” and sets up parameters for the algebraic wrapper with respect to the given base group. Hence, the reduction is able to perform extraction since the trapdoor (i. e., the decryption key) can be generated alongside the wrapper parameters. The reduction then uses the wrapper group to simulate the security game for the adversary on the building block – suitably embedding the discrete logarithm challenge such that a successful breach of the building block’s security allows to extract information about the discrete logarithm challenge.

Generally, using the set $\{g, h\}$ as basis is a good choice. This allows the reduction to generate the group element wrapping h (from the discrete logarithm challenge) since producing group elements requires to know their decomposition with respect to the basis.

LIMITATIONS. Group element representations can only hide the decomposition they carry computationally, i. e., against computationally bounded adversaries. information-theoretically, this decomposition cannot be hidden.¹⁰ However, some **AGM** proofs require that group element representations information-theoretically look the same, independently of how they were derived (or, in our, case which information they carry). These proofs cannot be simulated using the algebraic wrapper. This includes the tight security proof of the Boneh-Lynn-Shacham (**BLS**) [**BLS04**] signature scheme from [**FKL18**].

¹⁰ Since no public-key encryption scheme can be secure information-theoretically.

APPLICATIONS. Using the algebraic wrapper, we are able to transport several **AGM**-based proofs into the standard (group) model. For instance, the algebraic wrapper provides all necessary extraction properties for the tight security proof of Schnorr-signed ElGamal from [**FPS20**].

Thus, the algebraic wrapper provides protection of computations in the wrapped group in a relaxed manner which still suffices for several applications and can be instantiated from **IO**. We refer the reader to Part **II** of this thesis for more details.

1.3 PROTECTING COMPUTATIONS – USING COMPRESSION?

Lastly, we consider the classical notion of compression. Aside from the field of information theory [**Sha48**; **Huf52**; **GM59**; **Flo64**; **Sch72**; **Ris76**; **Pas77**; **ZL77**; **ZL78**; **RJ79**], starting from the seminal work of Goldberg and Sipser [**GS85**], compression has received considerable attention in the field of computational complexity theory [**Wee04**; **TVZ05**; **HLR07**].¹¹ Informally, given some distribution X , compression aims to efficiently encode samples from X as short strings while at the same time being able to efficiently recover these samples.

Prior to our work [**ACI+20**], it was entirely unclear how compression relates to the protection of computation. In this thesis and [**ACI+20**], we rigorously examine the notion and the properties of compression and vary this notion to escape classical impossibility results.¹² Subsequently, we shed light on an entirely unexpected connection between compression and the protection of computations in a distributed setting.

¹¹ Complexity theory considers the case of compressing a single sample, whereas information theory classically considers compression of multiple independent samples.

¹² If one-way functions exist, then there exist incompressible distributions (observation due to Levin).

PROPERTIES OF COMPRESSION. Compression in its “perfect” form guarantees that compressed outputs are completely (information-theoretically) random leaving virtually no room for any further compression without losing information. From this angle, compression can be viewed as unifying two widely studied problems in computational complexity theory: *randomness condensers* [RR99; TV00; TUZ01; DRV12] and *resource-bounded Kolmogorov complexity* [Sol64; Kol68; Cha69; LV90; LV19]. Randomness condensers are functions that efficiently map an input distribution to an output distribution with higher entropy rate, i. e., with higher entropy per length ratio.¹³ More precisely, a randomness condenser guarantees that its output entropy rate is higher than its input entropy rate. A randomness condenser can be viewed as an efficient compression algorithm – without a corresponding efficient decompression algorithm. On the other hand, the resource-bounded Kolmogorov complexity of a string is the smallest description length of an efficient program that outputs this string. For instance, the string “12121212121212” can be produced by a short program whereas the string “32576143764251” has no obvious simpler description than the string itself. Hence, the resource-bounded Kolmogorov complexity provides an intuitive measure of the amount of “randomness” in a particular string. Such a shortest program description can be viewed as a compressed string such that decoding is efficiently possible – while finding this compressed string may be inefficient.

One central property of perfect compression – unifying the core properties of both randomness condensers and resource-bounded Kolmogorov complexity [GS85; TVZ05] – is the ability to efficiently map bitstrings from some distribution to uniformly random-looking bitstrings while at the same time allowing the original input to be efficiently recovered.

CRYPTOGRAPHIC APPLICATIONS OF COMPRESSION. To demonstrate the usefulness of this property, consider the following scenario. Suppose, Alice wants to confidentially communicate some (random) sequence of seven English words to Bob.¹⁴ However, Alice and Bob only share a common password with *low entropy*, i. e., a password which can be efficiently guessed. Using the password as a secret key to encrypt Alice’s message allows an eavesdropper to try all (most likely) passwords until decryption with a candidate password reveals a sequence of seven English words.

This problem appears in the context of, e. g., password-authenticated key exchange [BM92], honey encryption [JR14] or subliminal communication and steganography [vHL05; CGOS07; HPRV19]. A natural solution is to use perfect compression as was done in [BMN01] to extend [BM92] for elliptic curve groups. Perfect compression maps samples from a distribution to random-looking strings. Compressing the message to a random-looking string before encrypting renders the above brute-force attack useless: decryption with a wrong password reveals a random-looking plaintext (by a standard assumption on the used secret-key encryption scheme) which is exactly the same behavior as decryption with the correct password.

However, perfect compression exists only for a very small number of message distributions. On the other hand, the core property of compression – that outputs are smaller than inputs – is not necessary for these applications. Indeed, efficiently encoding x into a (possibly longer) random-looking string suffices, provided that x can be efficiently recovered.

¹³ Entropy measures the uncertainty or amount of randomness of a distribution.

¹⁴ This message distribution is just a toy example. A more useful but more complex message distribution is the distribution modeling natural English language.

RELAXING COMPRESSION. Since compression (particularly perfect compression) only exists for a very small number of distributions, we vary the notion of compression as follows. We relax perfect compression by eliminating the requirement for shorter outputs and analyze several variants of the resulting notion, see also Figure 1.4 for an overview. Moreover, we consider variants assuming the presence of a trusted setup in the form of a common reference string.

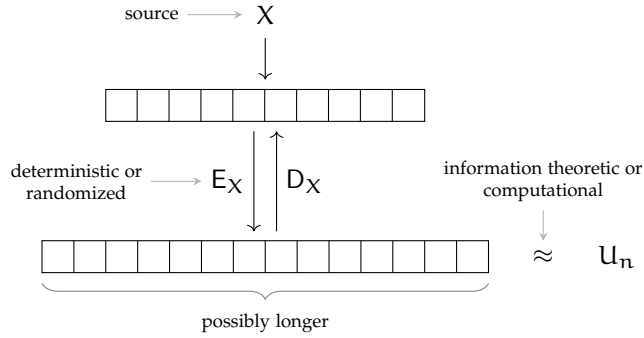


Figure 1.4: Illustration of our notion of pseudorandom encodings. Compared to perfect compression, we drop the requirement for shorter outputs, we consider deterministic and randomized encoding algorithms and we consider computational and information-theoretic indistinguishability from randomness. Note that we also study variants assuming a trusted setup which are not captured by this visual.

The resulting notion only preserves the properties that encoded samples can be efficiently recovered and that encoded samples appear to be random. For this reason, we refer to this notion as *pseudorandom encodings*. A pseudorandom encoding scheme for a distribution X is a tuple (E_X, D_X) , where E_X is an efficient (possibly randomized) encoding algorithm and D_X is an efficient deterministic decoding algorithm satisfying the following two properties.

CORRECTNESS. Decoding works with overwhelming probability. More formally,

$$\Pr [y \leftarrow X_\lambda : D_X(E_X(y)) = y]$$

is overwhelming.

PSEUDORANDOMNESS. An encoded sample is indistinguishable from true randomness. More formally,

$$E_X(X_\lambda) \approx U_{n(\lambda)},$$

where $n(\lambda)$ is the output length of E_X and “ \approx ” denotes some notion of indistinguishability.

Additionally, we consider distributions X which can accept an input $m \in L \subseteq \{0, 1\}^*$, i. e., distribution families $(X_m)_{m \in L}$. In this case, the above properties need to hold for all such inputs.

Our main focus is on the hypothesis that *all* efficiently samplable distributions can be pseudorandomly encoded. We refer to this hypothesis as the *pseudorandom encoding hypothesis* (PREH).

The above notion offers several degrees of freedom: the encoding algorithm can be deterministic or randomized and the indistinguishability requirement can be information-theoretical or computational. Additionally, we may assume a trusted setup, i. e., a globally known common reference string (**CRS**) which is set up honestly. If a distribution X admits an input $m \in L$, this further entails two different notions: a non-adaptive version, where correctness and pseudorandomness hold only for inputs m which are independent of the **CRS**, and an adaptive version, where correctness and pseudorandomness hold also for inputs m which are chosen depending on the **CRS**. See Figure 1.5 for an overview of these notions.

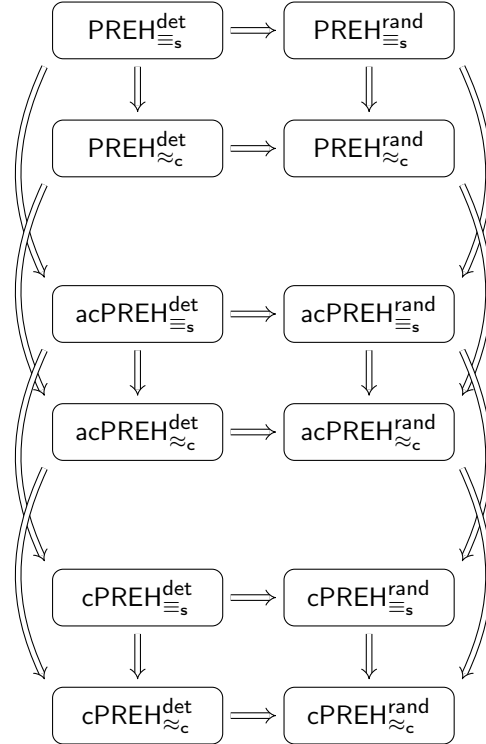


Figure 1.5: An overview of the different notions of pseudorandom encodings. The superscripts “det” and “rand” indicate whether the encoding algorithm is deterministic or allowed to be randomized. The subscripts “ \equiv_s ” and “ \approx_c ” indicate whether the guarantees are information-theoretical or computational. The prefixes “c” and “ac” indicate that a trusted setup is assumed where inputs are required to be chosen statically or adaptively, respectively.

In Part III of this thesis which is based on the publication [ACI+20], we provide a classification of pseudorandom encodings and identify computational, randomized pseudorandom encodings with common setup as the best possible and useful notion, and instantiate this building block from IO.

In other words, assuming that we can protect the computations inside some program from the user of the program (i. e., if IO exists), then pseudorandom encoding schemes exist for all efficiently samplable distributions in their weakest variant. However, this is not the only common ground between pseudorandom encodings and the protection of computations as we will elaborate in the following.

¹⁵ We consider the feasibility of MPC for all (possibly randomized) programs.

PROTECTING COMPUTATIONS. Surprisingly, the resulting novel notion of pseudorandom encodings turns out to be strongly connected to *fully adaptively* secure multi-party computation.¹⁵ Multi-party computation (MPC) allows parties to jointly compute any possibly randomized program on their inputs without revealing them. More precisely, while every party eventually learns the output, nothing beyond what follows from that party’s input and output can be inferred. In particular, no inputs from other parties can be learned.

Security in such a setting is defined by comparing the actual protocol execution (the real world) with an ideally secure protocol in the presence of an adversary who may take control of parties. If every adversarial behavior in the real world can be emulated in the ideally secure world, then the real world protocol is “at least as secure as” the ideally secure protocol [GMW87]. The greatest challenge pose *adaptive* adversaries who may adaptively decide to corrupt any number of participating parties at any time during the protocol execution. This behavior is particularly challenging to emulate.

In other words, multi-party computation allows for executing *protected computations* in settings, where many mutually distrustful parties are involved. This manner of protecting computations should be contrasted with the way code obfuscation protects computations. Code obfuscation conceals the way *how* a program derives outputs from entities who evaluate the program. In the setting of multi-party computation, on the other hand, every party may very well know the program which is to be jointly evaluated. However, multi-party computation guarantees that no party is able to infer a single intermediate value which does not already follow from that party’s input and output.

The notion of adaptively secure MPC was first considered in [BH93; CFGN96]. [BH93] relies on secure and trusted erasures. However, since erasing data is not verifiable and may be subject to hardware limitations, it is (arguably) not satisfying to rely on other parties to erase their internal data. Hence, we only consider a setting, where secure erasures are not assumed. There are several positive results on the feasibility of adaptive MPC in restricted settings such as the honest majority setting (where a majority of the parties remains uncorrupted) [CFGN96; Bea97], the all-but-one corruptions setting (where at least one party remains uncorrupted) [KO04; IPS08; GS12; HP14; DPR16], or restricting the class of supported functions to so-called “well-formed” ones (which do not hide their internal randomness) [CLOS02]. We consider the strongest notion of adaptive security for *all* randomized functionalities allowing adaptive corruptions of *all* parties [IKOS10; DKR15; CPV17]. A setting, where all parties are corrupted may seem unnatural. However, as outlined in [DKR15], a protocol running a sub-protocol among a strict subset of the parties should remain secure even if all parties running the sub-protocol are corrupted.

HOW TO EMULATE ADAPTIVE CORRUPTION OF ALL PARTIES. Consider the adversary who merely observes the entire protocol execution and decides to corrupt all parties after the protocol has terminated. Since secure erasures are not assumed, this adversary learns well-distributed random tapes which explain all messages sent by the individual parties. Recovering such valid-looking random tapes *after the output is fixed* is a very challenging task.

Naively speaking, invertible sampling constitutes the problem to invert randomized algorithms, i. e., to produce a valid-looking random tape that leads to a given output. However, there exist several randomized algorithms which cannot be inverse sampled to this extent. Consider, for instance, the algorithm which randomly samples two large prime numbers and outputs their product. Given our current understanding of computational number theory, this algorithm cannot be inverse sampled in the above sense. In fact, invertible sampling due to [CFGN96; DN00; GKM+00; IKOS10] only requires the existence of an inverse samplable algorithm which induces an output distribution similar to the output distribution of the actual sampler. This notion of invertible sampling is exactly what is both necessary and sufficient to emulate fully adaptive adversarial behavior. Ishai, Kumarasubramanian, Orlandi, and Sahai [IKOS10] establish an equivalence between invertible sampling and *fully adaptively secure* multi-party computation (MPC). More precisely, fully adaptively secure multi-party computation is possible if and only if invertible sampling is possible for all efficiently samplable distributions with input.

Our study of pseudorandom encodings reveals a surprising connection between invertible sampling and pseudorandom encodings: our novel notion of pseudorandom encodings emerges to be equivalent to invertible sampling. Consequently, pseudorandom encodings are equivalent to fully adaptively secure MPC – the strongest form of protecting computations in a distributed setting.

Our study of pseudorandom encodings unveils further unexpected connections between several fields within cryptography, like fully adaptive MPC, honey encryption and steganography. We refer the reader to Part III of this thesis for more details.

SUBSEQUENT WORK. In [ACI+20], we obtain partial evidence that indistinguishability obfuscation is not only sufficient for pseudorandom encodings but also *necessary*. Hence, we raised the question whether IO is indeed necessary for pseudorandom encodings. This open question transpired quite helpful for the development of plausible IO candidates. In fact, shortly after uploading a preliminary version of [ACI+20] to the IACR cryptology ePrint archive, Wee and Wichs [WW20] answer our question positively giving rise to a novel construction of IO from the assumption that pseudorandom encoding schemes exist for a very small class of distributions (and the learning with errors assumption).

1.4 STRUCTURE OF THIS THESIS

This thesis is divided into 3 parts. Each part contains an introduction including a short summary of the contribution presented in that part and a discussion setting the results in context to the state-of-the-art. These introductions are taken verbatim from the works [ACH20; AHK20; ACI+20] with minor changes.

Before we present our contributions in Parts I, II and III in detail, Chapter 2 introduces basic notation and definitions of cryptographic primitives which we will use in this thesis. Chapter 2 is in part taken verbatim from the works [ACH20; AHK20; ACI+20].

- In Part **I** of this thesis, we present our results regarding probabilistic indistinguishability obfuscation. We introduce a novel framework called “doubly-probabilistic indistinguishability obfuscation” which allows to avoid subexponentially secure **IO** in several contexts. Part **I** is taken verbatim from the publication [ACH20] with minor changes.
- In Part **II** of this thesis, we present our results on instantiating the algebraic group model from falsifiable assumptions. We introduce a novel framework called “algebraic wrapper” which can be instantiated in the standard model and allows to transport several proofs from the algebraic group model into the standard model. Part **II** is taken verbatim from the publication [AHK20] with minor changes.
- In Part **III** of this thesis, we introduce our results on pseudorandom encodings. We introduce the novel notion of pseudorandom encodings and provide a unified study including positive and negative results on the feasibility of several flavors of pseudorandom encodings. Furthermore, we demonstrate how pseudorandom encodings relate to various areas in cryptography and derive new insights for those fields. Part **III** is taken verbatim from the publication [ACI+20] with minor changes.

Finally, in Chapter 19, we provide a brief outlook on possible further research directions.

2 | PRELIMINARIES

In this chapter, we introduce notations and technical preliminaries which are common to the subsequent Parts I, II and III. The preliminaries are in part taken verbatim from the works [ACH20; AHK20; ACI+20].

2.1 NOTATIONS

For $n \in \mathbb{N}$, we denote by $[n]$ the set $\{1, \dots, n\}$. We denote the set of all functions mapping from set \mathcal{X} to \mathcal{Y} by $\text{maps}(\mathcal{X}, \mathcal{Y})$. The security parameter is denoted by λ and is implicitly given to all algorithms in unary. A probabilistic polynomial time (PPT) algorithm runs in time polynomial in the (implicit) security parameter λ . We consider non-uniform polynomial time adversaries, i. e., polynomial time adversaries receiving a polynomially bounded auxiliary input (also called “advice string”) depending only on the security parameter.¹⁶ A function $\text{negl}: \mathbb{N} \rightarrow \mathbb{R}$ is negligible in λ if for every constant $c \in \mathbb{N}$, there exists a bound $n_c \in \mathbb{N}$, such that for all $n \geq n_c$, $|\text{negl}(n)| \leq n^{-c}$. In asymptotic notation, the set of negligible functions is $\lambda^{-\omega(1)}$. A function g is overwhelming if $1 - g(\lambda)$ is a negligible function.

Given a finite set \mathcal{X} , the notation $x \leftarrow \mathcal{X}$ means a uniformly random assignment of an element of \mathcal{X} to the variable x . Given a probability distribution D , the notation $x \leftarrow D$ means sampling an element according to the distribution D and assigning that element to x . The support of a probability distribution D is the set of all elements with a non-zero probability to occur and is denoted by $\text{supp}(D)$. We denote the uniform distribution over bitstrings of length n by U_n . We use the notation $\{C_1, C_2, \dots, C_m : D\}$ to denote the distribution of D which is obtained as the result of the process defined by the sequence of instructions C_1, \dots, C_m .

Let E_1 and E_2 be events of some probabilistic process. We denote the probability that E_1 occurs by $\Pr[E_1]$, and the probability that E_1 occurs conditioned on the fact that the event E_2 occurred by $\Pr[E_1 \mid E_2]$.

Given a possibly randomized algorithm A , the notation $y \leftarrow A(x)$ means evaluation of A on input of x with fresh random coins r and assignment to the variable y . Sometimes we make the random coins r used by A explicit and write $A(x; r)$. The notation $\mathcal{A}^\mathcal{O}$ indicates that the algorithm \mathcal{A} is given oracle access to \mathcal{O} .

When describing security games, we always consider *stateful* adversaries unless explicitly stated otherwise. In some cases, we emphasize this by handling the adversarial state (usually denoted by st) explicitly.

In game based proofs, out_i denotes the output of game \mathbf{G}_i . Further, we will use **this notation** to highlight differences to previous hybrids.

¹⁶ Note that by a coin-fixing argument, it is sufficient to consider non-uniform deterministic adversaries. Most of the results in this thesis apply for uniform PPT adversaries as well. In case we make explicit use of the non-uniformity of the adversary, we remark this explicitly.

DEFINITION 2.1 (Statistical distance). The *statistical distance* Δ between two probability distributions X and Y (over some set \mathcal{X}) is

$$\Delta(X, Y) := \sum_{x \in \mathcal{X}} |\Pr[X = x] - \Pr[Y = x]|.$$

In other words, the statistical distance is the metric induced by the Manhattan norm (or ℓ_1 norm) on the probability space. We say that two distributions are statistically close if their statistical distance is negligible.

DEFINITION 2.2 (Min-entropy). The *min-entropy* of a distribution X is

$$H_\infty(X) = -\log \left(\max_{x \in \text{supp}(X)} \Pr[X = x] \right).$$

In other words, the probability of correctly guessing the outcome of a probability experiment with min-entropy μ is at most $2^{-\mu}$.

In many cases, some information Z that is correlated to the actual source X is known. Since for our purposes, the conditional part Z is not under adversarial control, we use the notion of *average* conditional min-entropy as used in [HLR07; DORS08].

DEFINITION 2.3 (Average min-entropy, [HLR07; DORS08]). Let (X, Z) be a joint distribution. The *average min-entropy* of X conditioned on Z is

$$\tilde{H}_\infty(X | Z) := -\log \left(\mathbb{E}_{z \leftarrow Z} \left[\max_x \Pr[X_z = x] \right] \right).$$

Extractors allow to obtain (almost) uniform randomness from samples from any weakly random source of entropy, as long as a short uniform random seed is available. Hence, even weakly random distributions can be useful for cryptography.

DEFINITION 2.4 (Extractor, [ILL89]). A function $\text{Ext}: \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is a (k, ϵ) -*extractor* if for all distributions X on $\{0, 1\}^n$ with min-entropy $H_\infty(X) \geq k$ (which are independent of the key K), the statistical distance

$$\Delta \left(\{K \leftarrow \mathcal{U}_d: (K, \text{Ext}(K, X))\}, \{K \leftarrow \mathcal{U}_d: (K, \mathcal{U}_m)\} \right)$$

is at most ϵ .

Initially, such extractors were used to reduce the randomness complexity of algorithms, [Sip88; San87; NZ96].

2.2 HARDNESS ASSUMPTIONS

Cryptography aims to provide an asymmetry between the computational cost for an honest user and the computational cost for adversaries. The honest user is provided with computationally affordable mechanisms to, e.g., protect information, whereas it should be computationally costly for an adversary to obtain this protected information. Asymptotic security is the complexity theoretic approach for achieving this. The computational costs for user and adversary are parametrized by a *security parameter* $\lambda \in \mathbb{N}$. Tasks that are computable in probabilistic polynomial-time in the security parameter are considered efficient, all other tasks are considered infeasible

(or intractable). The philosophy behind asymptotic security is that one is able to scale the security guarantee (by choosing a suitable value for λ) thus covering arbitrarily powerful adversaries. A security proof consists of a reduction¹⁷ of a computational problem to the task of breaking the security guarantee in question.

The validity of solutions for adversarial tasks in cryptography can be efficiently tested. In other words, the adversarial tasks of interest to cryptography are in \mathcal{NP} .¹⁸ Thus, a necessary condition for the existence of (most) cryptographic schemes is $\mathcal{BPP} \not\subseteq \mathcal{NP}$ (implying $\mathcal{P} \neq \mathcal{NP}$).

The complexity classes \mathcal{P} and \mathcal{NP} are only concerned with the worst-case complexity. However, worst-case hardness does not suffice for cryptography since finding hard instances must be tractable for the honest users (i. e., must be computable in probabilistic polynomial-time). Therefore, for cryptography, we assume the *average-case* hardness of problems in \mathcal{NP} .

The fundamental notion of average-case hardness was conceived by Levin [Lev86]. Average-case hardness against polynomial-time adversaries for computational and decisional problems is roughly defined as follows.

DEFINITION 2.5 (Distributional search problem (informal), [Lev86; IL90]). A *distributional search problem* is a tuple (S, μ) , where $S \subseteq \{0, 1\}^* \times \{0, 1\}^*$ and $\mu: \{0, 1\}^* \rightarrow [0, 1]$ is a distribution over inputs.

DEFINITION 2.6 (Distributional decision problem (informal), [Lev86; IL90]). A *distributional decision problem* is a tuple (D, μ) , where $D: \{0, 1\}^* \rightarrow \{0, 1\}$ and $\mu: \{0, 1\}^* \rightarrow [0, 1]$ is a distribution over inputs.

A decision problem (D, μ) is in distributional \mathcal{NP} if μ is computable in polynomial time (that is efficiently samplable) and D is an \mathcal{NP} -predicate. Furthermore, the class of distributional \mathcal{NP} contains complete problems.

For cryptography, a more direct way of formalizing average-case hardness emerged. Instead of considering one fixed input distribution, in cryptography, we consider a sequence of input distributions $(\mu_\lambda)_{\lambda \in \mathbb{N}}$ such that μ_λ is a distribution over $\{0, 1\}^\lambda$.

DEFINITION 2.7 (Average-case hardness (informal), [Lev86; GM82; BM82; Yao82]). A distributional decision problem $(D, (\mu_\lambda)_{\lambda \in \mathbb{N}})$ (resp., search problem $(S, (\mu_\lambda)_{\lambda \in \mathbb{N}})$) is *hard on the average* if $(\mu_\lambda)_{\lambda \in \mathbb{N}}$ is efficiently samplable (in λ) and for all \mathcal{PPT} adversaries, the probability to correctly solve D is at most negligibly better than guessing (resp., at most negligible), where the probability is also over the input choice from $(\mu_\lambda)_{\lambda \in \mathbb{N}}$.

SUBEXPONENTIAL HARDNESS. There are two common ways to define hardness against subexponential adversaries. The weaker variant requires the advantage of polynomial-time adversaries to decrease subexponentially in the security parameter, i. e., requires the advantage to be in $2^{-o(\lambda)}$.¹⁹ This definition is commonly used in the context of subexponentially secure \mathcal{IO} . The stronger variant requires the advantage of subexponential-time adversaries to decrease subexponentially in the security parameter.

NON-UNIFORM ADVERSARIES. A stronger model of computation considers so-called non-uniform adversaries. This model of computation originates from the complexity class \mathcal{P}/poly containing all languages which are decidable by polynomial-sized circuits. A non-uniform (\mathcal{PPT}) adversary is a tuple

¹⁷ A reduction from a problem A to a problem B is a probabilistic polynomial-time algorithm which solves problem A (with non-negligible probability) given oracle access to a problem solver for B .

¹⁸ For simplicity, we do not make a distinction between decision and search problems at this point.

¹⁹ That is, the advantage is at most $2^{-\lambda^\epsilon}$ for some $0 \leq \epsilon < 1$.

$(\mathcal{A}, (a_\lambda)_{\lambda \in \mathbb{N}})$, where \mathcal{A} is a (probabilistic) polynomial-time machine and $(a_\lambda)_{\lambda \in \mathbb{N}}$ is a sequence of polynomial-sized bitstrings. That is, a_λ provides \mathcal{A} with one additional advice string per input length. Note that by a coin-fixing argument, non-uniform PPT adversaries and non-uniform deterministic polynomial-time adversaries are equivalent.

2.2.1 Efficient Falsifiability

An assumption is falsifiable if an observation can show it to be false. Being able to assess the plausibility of hardness assumptions is an important task. For this purpose, Naor [Naor03] introduces a classification of hardness assumptions based on the *complexity of falsifying* them. Informally, a cryptographic assumption is falsifiable if the assumption can be formalized as an interaction between an *efficient* challenger and an adversary, such that the challenger can efficiently provide the adversary with a challenge and is able to efficiently verify the adversary’s solution attempt. Naor classifies assumptions into four classes. These are “efficiently falsifiable” assumptions, “falsifiable” assumptions, “somewhat falsifiable” assumptions²⁰ and “non-falsifiable assumptions” (ordered decreasing in plausibility). Assumptions which are at least somewhat falsifiable are acceptable for use in cryptography. This, however, is (arguably) not true for non-falsifiable assumptions.²¹

²⁰ The difference between efficient falsifiability, falsifiability and somewhat falsifiability is mostly the runtime to verify the solution.

²¹ Using non-falsifiable assumptions might be better than simply assuming “the scheme is secure”, but a security reduction based on such assumptions is (arguably) not satisfying.

Note that falsifiability is not closed under security reductions. There are security properties that are non-falsifiable (when formulated as an assumption) but can be instantiated from falsifiable assumptions. For instance, the soundness property of non-interactive zero-knowledge (NIZK) proof systems (see Section 2.8) is one such property. In this thesis, we only use falsifiable assumptions and security properties which can be based on falsifiable assumptions.

2.3 ONE-WAY FUNCTIONS

Before we define one-way functions, we define the notion of efficiently computable function family ensembles.

DEFINITION 2.8 (Function family ensemble). Let $n(\lambda), m(\lambda)$ be polynomially bounded length functions. A *function family ensemble* is a tuple of PPT algorithms $\mathcal{F} = (\text{Gen}, \text{Eval})$ such that

$$\left\{ \text{Eval}(f, \cdot): \{0, 1\}^{n(\lambda)} \rightarrow \{0, 1\}^{m(\lambda)} \mid f \in \text{supp}(\text{Gen}(1^\lambda)) \right\}_{\lambda \in \mathbb{N}}$$

is an efficiently computable family of functions.

For ease of notation, we often write $f(x)$ instead of $\mathcal{F}.\text{Eval}(f, x)$. A one-way function is a function on bitstrings which can be efficiently evaluated but which is hard to invert given the image of a random input.

DEFINITION 2.9 (One-way function, [DH76]). A *one-way function (OWF)* family ensemble is a function family ensemble $\mathcal{F} = (\text{Gen}, \text{Eval})$, such that for all PPT adversaries \mathcal{A} ,

$$\text{Adv}_{\mathcal{F}, \mathcal{A}}^{\text{ow}}(\lambda) := \Pr \left[\begin{array}{l} f \leftarrow \text{Gen}(1^\lambda) \\ x \leftarrow \mathcal{U}_{n(\lambda)} \quad : f(x') = f(x) \\ x' \leftarrow \mathcal{A}(f, f(x)) \end{array} \right]$$

is negligible in λ .

The existence of one-way functions is the most fundamental hardness assumption in cryptography.

2.4 PUNCTURABLE PSEUDORANDOM FUNCTIONS

Pseudorandom functions are functions on bitstrings which are computationally indistinguishable from truly random functions as long as the function key remains secret.

DEFINITION 2.10 (Pseudorandom function, [GGM86]). A family of *pseudorandom functions (PRFs)* is a tuple of PPT algorithms $F = (\text{KGen}, \text{Eval})$ and two computable functions $n(\lambda), m(\lambda)$ such that

- $\text{KGen}(1^\lambda)$ outputs a key K , and
- $\text{Eval}(K, \cdot)$, given a key $K \in \text{supp}(\text{KGen}(1^\lambda))$, defines a function mapping bitstrings of length $n(\lambda)$ to bitstrings of length $m(\lambda)$.

A PRF F is required to satisfy the following property.

$$\text{Adv}_{F, \mathcal{A}}^{\text{prf}}(\lambda) := \left| \Pr \left[K \leftarrow \text{KGen}(1^\lambda) : \mathcal{A}^{F, \text{Eval}(K, \cdot)}(1^\lambda) = 1 \right] - \Pr \left[R \leftarrow \text{maps}(\{0, 1\}^{n(\lambda)}, \{0, 1\}^{m(\lambda)}) : \mathcal{A}^{R(\cdot)}(1^\lambda) = 1 \right] \right|$$

is a negligible function in λ .

For ease of notation we often write $F(K, x)$ to denote $F.\text{Eval}(K, x)$. Pseudorandom functions can be constructed from one-way functions [GGM86].

For some applications, it will be useful to generate PRF keys which can be used to evaluate the PRF everywhere except for a pre-defined subset of pre-images. This procedure is called “puncturing”. A punctured PRF key is not only insufficient to evaluate the PRF at a punctured point, but it also does not leak any information about the actual PRF images at punctured points.

DEFINITION 2.11 (Puncturable PRF, [GGM86; BW13; BGI14; KPTZ13]). A family of *puncturable PRFs (pPRFs)* is a tuple of PPT algorithms $F = (\text{KGen}, \text{Punct}, \text{Eval})$ and two computable functions $n(\lambda), m(\lambda)$ such that

- $\text{KGen}(1^\lambda)$ outputs a key K ,
- $\text{Punct}(K, T)$, on input of a key K , and a subset of the input space $T \subset \{0, 1\}^{n(\lambda)}$, outputs a punctured key $K\{T\}$,
- $\text{Eval}(K', \cdot)$, on input of a key K' defines a function mapping bitstrings of length $n(\lambda)$ to bitstrings of length $m(\lambda)$. $\text{Eval}(K\{T\}, x)$ on input of a punctured key $K\{T\}$ and an $x \in T$, outputs a special symbol \perp .

A **pPRF** is required to satisfy the following properties.

FUNCTIONALITY PRESERVED UNDER PUNCTURING. For all **PPT** adversaries \mathcal{A} which output a set $T \subseteq \{0, 1\}^{n(\lambda)}$, we have that

$$\Pr \left[\begin{array}{l} T \leftarrow \mathcal{A}(1^\lambda) \\ K \leftarrow \text{KGen}(1^\lambda) \\ K\{T\} \leftarrow \text{Punct}(K, T) \end{array} : \forall x \in \{0, 1\}^{n(\lambda)} \setminus T: \text{Eval}(K, x) = \text{Eval}(K\{T\}, x) \right] = 1.$$

PSEUDORANDOM AT PUNCTURED POINTS. For all **PPT** adversaries \mathcal{A} ,

$$\text{Adv}_{F, \mathcal{A}}^{\text{s-prpf}}(\lambda) := \left| \Pr \left[\text{Exp}_{F, \mathcal{A}, 0}^{\text{s-prpf}}(\lambda) = 1 \right] - \Pr \left[\text{Exp}_{F, \mathcal{A}, 1}^{\text{s-prpf}}(\lambda) = 1 \right] \right|$$

is negligible, where $\text{Exp}_{F, \mathcal{A}, b}^{\text{s-prpf}}(\lambda)$ is defined in Figure 2.1.

```

Exp_{F, \mathcal{A}, b}^{\text{s-prpf}}(\lambda)
(T, st) \leftarrow \mathcal{A}(1^\lambda)
K \leftarrow F.\text{KGen}(1^\lambda)
K\{T\} \leftarrow F.\text{Punct}(K, T)
y_0 \leftarrow F.\text{Eval}(K, T)
y_1 \leftarrow \{0, 1\}^{m(\lambda)}
return \mathcal{A}(st, K\{T\}, y_b)

```

Figure 2.1: Selective security game $\text{Exp}_{F, \mathcal{A}, b}^{\text{s-prpf}}(\lambda)$ for puncturable pseudorandom functions.

[[BW13](#); [BGI14](#); [KPTZ13](#)] observed that the original tree-based PRF construction from [[GGM86](#)] already satisfies this definition. In this thesis, we will additionally be concerned with a special type of **pPRFs** which allow to extract (almost) uniformly random strings from high-entropy inputs, even when the function key is public (or not guaranteed to be hidden, as is the case in many applications involving indistinguishability obfuscation). That is, **pPRFs** which also act as randomness extractors (cf. Definition 2.4).

DEFINITION 2.12 (Extracting **pPRF**, [[SW14](#)]). An *extracting* puncturable **PRF** family with error $\epsilon(\cdot)$ for min-entropy $k(\cdot)$ is a family of puncturable **PRFs** F mapping $n(\lambda)$ bits to $m(\lambda)$ bits such that for all $\lambda \in \mathbb{N}$, if X is a distribution over $\{0, 1\}^{n(\lambda)}$ with min-entropy at least $k(\lambda)$, then the statistical distance

$$\Delta \left(\{K \leftarrow \text{KGen}(1^\lambda): (K, F(K, X))\}, \{K \leftarrow \text{KGen}(1^\lambda): (K, U_{m(\lambda)})\} \right)$$

is at most $\epsilon(\lambda)$.

Due to [[SW14](#)] (full version, Theorem 3), extracting **pPRFs** exist assuming one-way functions. On a high level, Sahai and Waters [[SW14](#)] show that composing a puncturable **PRF** (which is almost always injective) with a strong extractor yields an extracting **pPRF**.

2.5 OBFUSCATION

Code obfuscation has been formally introduced in [Hado0; BGI+01; BGI+12]. An obfuscator is an efficient compiler which transforms a given program into an unintelligible one while preserving its original functionality. Several notions of “unintelligibility” for program obfuscation have been studied. The strongest conceivable notion is virtual black-box (VBB) obfuscation. This notion requires that the obfuscated program is a “virtual black-box”, i. e., does not leak anything about the original program except for its input-output behavior. More formally, everything being efficiently computable given the obfuscated program, can also be efficiently computed given mere oracle access to the original program. However, Barak et al. [BGI+01; BGI+12] show that if one-way functions exist, there are inherently unobfuscatable programs. On the other hand, the existence of a VBB obfuscator (for a relatively small class of programs) implies the existence of one-way functions. Hence, VBB obfuscation is unconditionally impossible.

Barak et al. [BGI+01; BGI+12] also introduce a weaker notion of obfuscation called “indistinguishability obfuscation” (IO). An indistinguishability obfuscator only guarantees that if two programs are functionally equivalent, then their obfuscations are computationally indistinguishable. This definition avoids the black-box paradigm along with the impossibility results for VBB obfuscation. However, intuitively, this notion is rather limited compared to VBB since there are no guarantees for circuits which compute different functions. Barak et al. [BGI+01; BGI+12] show that an inefficient indistinguishability obfuscator exists unconditionally. Namely, the obfuscator which on input of a Boolean circuit C outputs the lexicographically first circuit which computes the same function as C and has size $|C|$ is an inefficient indistinguishability obfuscator. If $\mathcal{P} = \mathcal{NP}$, the polynomial hierarchy collapses and finding such an obfuscation is efficiently possible. Hence, the existence of an indistinguishability obfuscator cannot imply one-way functions.

Later, Goldwasser and Rothblum [GR07] show that indistinguishability obfuscation is equivalent to “best-possible” obfuscation. An obfuscator is a best-possible obfuscator if any information that is not hidden by the obfuscated circuit, is also not hidden by any other circuit of similar size which computes the same function. This is formalized via a simulator. Namely, a best-possible obfuscator \mathcal{O} guarantees that for every adversary (or “learner”) \mathcal{L} , there exists a simulator Sim such that for any pair of functionally equivalent circuits C_1, C_2 (of identical length), the distributions $\mathcal{L}(\mathcal{O}(C_1))$ and $\text{Sim}(C_2)$ are computationally indistinguishable. That is, no obfuscator can hide more information than a best-possible obfuscator. Further, Goldwasser and Rothblum [GR07] show that indistinguishability obfuscation is actually equivalent to best-possible obfuscation.

INDISTINGUISHABILITY OBFUSCATION. Let $\mathcal{C} = (\mathcal{C}_\lambda)_{\lambda \in \mathbb{N}}$ be an ensemble of sets \mathcal{C}_λ of Boolean circuits. The set \mathcal{C}_λ contains circuits of size λ with input length $n(\lambda)$.

DEFINITION 2.13 (Indistinguishability obfuscation, [GGH+13]). A uniform PPT algorithm $i\mathcal{O}$ is an *indistinguishability obfuscator* (IO) for a circuit class $\mathcal{C} = (\mathcal{C}_\lambda)_{\lambda \in \mathbb{N}}$ if the following conditions are satisfied.

CORRECTNESS. For all security parameters $\lambda \in \mathbb{N}$, all circuits $C \in \mathcal{C}_\lambda$, all inputs x , we have that

$$\Pr[\Lambda \leftarrow \text{iO}(C): \Lambda(x) = C(x)] = 1.$$

SECURITY. For every (possibly non-uniform) **PPT** distinguisher \mathcal{A} , for all pairs of circuits $(C_0, C_1) \in \mathcal{C}_\lambda \times \mathcal{C}_\lambda$ such that for all inputs x , $C_0(x) = C_1(x)$, we have that

$$\text{Adv}_{\text{iO}, \mathcal{A}}^{\text{io-ind}}(\lambda) := \left| \Pr[\mathcal{A}(1^\lambda, \text{iO}(C_0)) = 1] - \Pr[\mathcal{A}(1^\lambda, \text{iO}(C_1)) = 1] \right|$$

is negligible in λ .

Definition 2.13 can also be formulated with respect to an efficient probabilistic circuit sampler D for \mathcal{C} . Intuitively, this corresponds to a “legitimate” adversary choosing the circuit pair (C_0, C_1) in a “first phase”.

A circuit sampler for \mathcal{C} is defined as a family of efficiently samplable distributions $D = (D_\lambda)_{\lambda \in \mathbb{N}}$, where D_λ is a distribution over triplets (C_0, C_1, z) with $C_0, C_1 \in \mathcal{C}_\lambda$ such that C_0 and C_1 take inputs of the same length and $z \in \{0, 1\}^{\text{poly}(\lambda)}$ is some state information. This state information z allows to pass information which the circuit sampler D generates alongside (C_0, C_1) to the adversary \mathcal{A} . We write $D(1^\lambda)$ to denote the efficient sampling process according to the distribution D_λ .

DEFINITION 2.14 (Functionally equivalent circuit sampler). We call a circuit sampler $D = (D_\lambda)_{\lambda \in \mathbb{N}}$ a *functionally equivalent circuit sampler* if D guarantees that the circuits C_0 and C_1 are functionally equivalent in the sense that for all inputs x , we have $C_0(x) = C_1(x)$.

The alternative definition of indistinguishability obfuscation relative to functionally equivalent circuit samplers is as follows.

DEFINITION 2.15 (Indistinguishability obfuscation (alternative), [GGH+13]). A uniform **PPT** algorithm iO is an *indistinguishability obfuscator (IO)* for a circuit class $\mathcal{C} = (\mathcal{C}_\lambda)_{\lambda \in \mathbb{N}}$ if the following conditions are satisfied.

CORRECTNESS. For all security parameters $\lambda \in \mathbb{N}$, all circuits $C \in \mathcal{C}_\lambda$, all inputs x , we have that

$$\Pr[\Lambda \leftarrow \text{iO}(C): \Lambda(x) = C(x)] = 1.$$

SECURITY. For every (possibly non-uniform) **PPT** distinguisher (D, \mathcal{A}) , where D is a functionally equivalent circuit sampler for \mathcal{C} , we have that

$$\begin{aligned} \text{Adv}_{\text{iO}, D, \mathcal{A}}^{\text{io-ind}}(\lambda) := & \\ & \left| \Pr[(C_0, C_1, z) \leftarrow D_\lambda: \mathcal{A}(1^\lambda, C_0, C_1, z, \text{iO}(C_0)) = 1] \right. \\ & \left. - \Pr[(C_0, C_1, z) \leftarrow D_\lambda: \mathcal{A}(1^\lambda, C_0, C_1, z, \text{iO}(C_1)) = 1] \right| \end{aligned}$$

is negligible in λ .

For non-uniform adversaries, an indistinguishability obfuscator according to Definition 2.13 for a circuit class \mathcal{C} is also an indistinguishability obfuscator according to Definition 2.15 for the circuit class \mathcal{C} . This can be realized by a coin-fixing argument. We can fix a worst-case choice of random

coins r for D_λ . Given random coins r , D_λ produces concrete (C_0, C_1, z) for which \mathcal{A} 's advantage is still negligible due to Definition 2.13. For this thesis, Definition 2.15 will be sufficient.

[JLS20] provides the first instantiation of **IO** based on well-founded falsifiable assumptions.

IO IS EFFICIENTLY FALSIFIABLE. Definition 2.15 enjoys efficient falsifiability due to Naor [Na003] as opposed to Definition 2.13. The existence of an indistinguishability obfuscator according to the original definition from [BGI+12] (see Definition 2.13) falls within the category of non-falsifiable assumptions. Definition 2.13 can be modeled as a game between an adversary and a challenger. The adversary sends a pair of circuits (C_0, C_1) to the challenger who obfuscates a random one of them and returns this obfuscation to the adversary who then has to guess which one was obfuscated. However, the challenger is unable to efficiently test whether the adversarially chosen circuits C_0 and C_1 are functionally equivalent, i. e., if for all inputs x , $C_0(x) = C_1(x)$. This would require the challenger to solve the SAT instance $C_0(x) \oplus C_1(x)$.

The alternative formulation from Definition 2.15 requires that no adversary can distinguish obfuscations of a circuit pair (C_0, C_1) which is (honestly) sampled from a functionally equivalent circuit sampler (cf. Definition 2.14). Hence, Definition 2.15 does satisfy efficient falsifiability.

DIFFERING-INPUTS OBFUSCATION. The notion of differing-inputs obfuscation (**dIO**) was defined in [BGI+12; BCP14; ABG+13] as a generalization of **IO**. The notion of **dIO** is defined very similarly to Definition 2.15 except that security is required to hold not only for functionally equivalent circuit samplers (see Definition 2.14), but for arbitrary circuit samplers. This definition clearly leads to problems if the input-output behavior of the sampled circuits C_0 and C_1 is easily distinguishable. Therefore, [ABG+13] restrict the definition of **dIO** to classes of circuits, where given (C_0, C_1, z) sampled from D , it is infeasible to find an input x , where C_0 and C_1 differ. Another possibility is to require that for every adversary, there exists a suitable extractor recovering an input x , where the sampled circuits differ [BCP14]. The authors of [BCP14; ABG+13; BP15b] provide several applications of **dIO**, e. g., functional witness encryption, adaptively secure functional encryption, multiparty non-interactive key-exchange (without setup and only logarithmic-size messages), perfect zero knowledge succinct non-interactive arguments (**SNARGs**), and many more. However, there is evidence that **dIO** for all circuits is implausible [GGHW17; BP15b].

2.5.1 Probabilistic Indistinguishability Obfuscation

Indistinguishability obfuscation for probabilistic circuits, or probabilistic indistinguishability obfuscation (**pIO**), was introduced in [CLTV15]. This notion generalizes the classical notion of **IO** from [Hadoo; BGI+12]. Similar to Definition 2.15, **pIO** is defined relative to a class of circuit samplers \mathcal{C} . A circuit sampler is defined as before, see Section 2.5. Informally, a probabilistic indistinguishability obfuscator allows to obfuscate probabilistic circuits and is secure relative to a circuit sampler D if obfuscations $\text{piO}(C_0)$ and $\text{piO}(C_1)$ of circuits $(C_0, C_1) \leftarrow D$ are indistinguishable. Let $\mathcal{C} = (\mathcal{C}_\lambda)_{\lambda \in \mathbb{N}}$ be an ensemble

of sets \mathcal{C}_λ of probabilistic circuits. The set \mathcal{C}_λ contains circuits of size λ with input length $n(\lambda)$ expecting (at most) $m(\lambda)$ random bits.

DEFINITION 2.16 (Probabilistic indistinguishability obfuscation for a class of samplers \mathcal{C} , [CLTV15]). A *probabilistic indistinguishability obfuscator* (pIO) for a class of samplers \mathcal{C} over the probabilistic circuit family $\mathcal{C} = (\mathcal{C}_\lambda)_{\lambda \in \mathbb{N}}$ is a uniform PPT algorithm piO , such that the following properties hold.

CORRECTNESS. For every PPT adversary \mathcal{A} , every $C \in \mathcal{C}_\lambda$, the advantage of \mathcal{A} , given the description of the circuit C , to distinguish oracle access to the original randomized circuit C from oracle access to $\text{piO}(C)$ without querying the same input more than once, is negligible.

SECURITY WITH RESPECT TO \mathcal{C} . For all circuit samplers $D \in \mathcal{C}$, for all PPT adversaries \mathcal{A} , the advantage

$$\begin{aligned} \text{Adv}_{\text{piO}, D, \mathcal{A}}^{\text{pio-ind}}(\lambda) := & \\ & \left| \Pr \left[(C_0, C_1, z) \leftarrow D_\lambda : \mathcal{A}(1^\lambda, C_0, C_1, z, \text{piO}(1^\lambda, C_0)) = 1 \right] \right. \\ & \left. - \Pr \left[(C_0, C_1, z) \leftarrow D_\lambda : \mathcal{A}(1^\lambda, C_0, C_1, z, \text{piO}(1^\lambda, C_1)) = 1 \right] \right| \end{aligned}$$

is negligible.

For the purpose of a versatile notion of probabilistic indistinguishability obfuscation, we require security to hold for sampler classes which are broader than the restrictive class of functionally equivalent samplers from Definition 2.14. In fact, the power of pIO lies in the generality of feasible circuit sampler classes.

The broadest class of circuit samplers we consider is the class of dynamic-input indistinguishable circuit samplers, introduced in [CLTV15]. Informally, a circuit sampler D is a dynamic-input indistinguishable circuit sampler if D produces circuit pairs (C_0, C_1) such that even for adversarially chosen inputs x , the distributions induced by $C_0(x)$ and $C_1(x)$ are computationally indistinguishable.

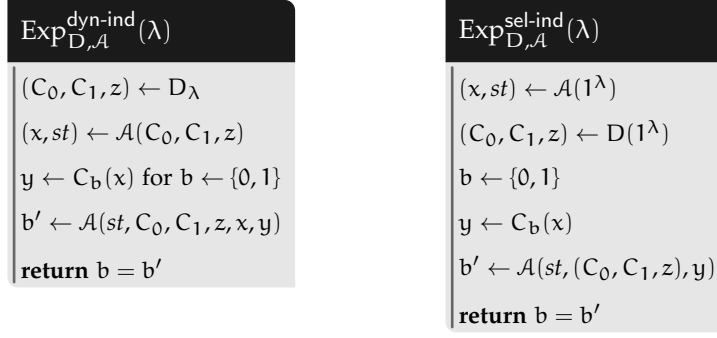
This class of circuit samplers is the broadest conceivable class which still results in a feasible notion of pIO. This follows from the observation that if the output distributions $C_0(x)$ and $C_1(x)$ are efficiently distinguishable for some input x , then so are their obfuscations.

DEFINITION 2.17 (Dynamic-input indistinguishable samplers, [CLTV15]). The class $\mathcal{C}^{\text{dyn-ind}}$ of *dynamic-input indistinguishable samplers* for a circuit family \mathcal{C} contains all circuit samplers $D = \{D_\lambda\}_{\lambda \in \mathbb{N}}$ for \mathcal{C} such that for every (non-uniform) PPT adversary \mathcal{A} , the advantage

$$\text{Adv}_{D, \mathcal{A}}^{\text{dyn-ind}}(\lambda) := \left| \Pr \left[\text{Exp}_{D, \mathcal{A}}^{\text{dyn-ind}}(\lambda) = 1 \right] - \frac{1}{2} \right|$$

is negligible, where the game $\text{Exp}_{D, \mathcal{A}}^{\text{dyn-ind}}(\lambda)$ is defined in Figure 2.2a.

As observed in [CLTV15], restricting pIO for dynamic-input indistinguishable samplers to the circuit class of deterministic circuits results in differing-inputs obfuscation. More precisely, for every $D \in \mathcal{C}$, given $(C_0, C_1, z) \leftarrow D$, it must be computationally hard to find an input x such that $C_0(x) \neq C_1(x)$. Hence, pIO for dynamic-input indistinguishable samplers inherits the implausibility result from [GGHW17].



(a) Game $\text{Exp}_{D,\mathcal{A}}^{\text{dyn-ind}}(\lambda)$ for the indistinguishability property of dynamic-input samplers. (b) Game $\text{Exp}_{D,\mathcal{A}}^{\text{sel-ind}}(\lambda)$ for the X-indistinguishability property of static-input samplers.

Figure 2.2: Definitions of the security games for dynamic-input indistinguishable samplers (2.2a) and (static-input) X-ind samplers (2.2b).

STATIC-INPUT PIO. The above implausibility suggests the study of a weaker notion. Canetti et al. [CLTV15] introduce a corresponding “static-inputs” notion, see Figure 2.2b for the corresponding game. At first, this definition renders the resulting class of circuit samplers even broader (resulting in an impossible notion). However, Canetti et al. restrict this class of circuit samplers accounting for the number of inputs, where the sampled circuits actually behave differently as follows. Informally, let X be (a lower bound of) the number of inputs, where circuits sampled from D behave differently. Then, a circuit sampler D is in the class of X-ind samplers if for every adversary, the distinguishing advantage in the game defined in Figure 2.2b is upper bounded by $\text{negl}(\lambda) \cdot X^{-1}(\lambda)$. Due to [CLTV15], **pio** for this class of samplers is achievable from **io** and one-way functions.

DEFINITION 2.18 (X-ind sampler, [CLTV15]). Let $X(\lambda)$ be a function upper bounded by 2^λ . The class $\mathcal{C}^{\text{X-ind}}$ of X-ind samplers for a circuit family \mathcal{C} contains all circuit samplers $D = (D_\lambda)_{\lambda \in \mathbb{N}}$ for \mathcal{C} such that for all $\lambda \in \mathbb{N}$, there exists a set $\mathcal{X}_\lambda \subseteq \{0, 1\}^*$ with $|\mathcal{X}_\lambda| \leq X(\lambda)$, such that

X-DIFFERING INPUTS. With overwhelming probability over the choice of $(C_0, C_1, z) \leftarrow D_\lambda$, for every $x \notin \mathcal{X}_\lambda$, for all $r \in \{0, 1\}^{m(\lambda)}$, we have that $C_0(x; r) = C_1(x; r)$.

X-INDISTINGUISHABILITY. For all (non-uniform) adversaries \mathcal{A} , the advantage

$$X(\lambda) \cdot \left(\Pr \left[\text{Exp}_{D,\mathcal{A}}^{\text{sel-ind}}(\lambda) = 1 \right] - \frac{1}{2} \right)$$

is negligible, where $\text{Exp}_{D,\mathcal{A}}^{\text{sel-ind}}(\lambda)$ is defined in Figure 2.2b.

On a high level, the construction of **pio** for the sampler class $\mathcal{C}^{\text{X-ind}}$ from [CLTV15] works as follows. Given a probabilistic circuit C with size at most λ , $\text{pio}(C)$ samples a **PRF** key K and obfuscates the deterministic circuit \bar{C} which on input of x evaluates the deterministic circuit $C(x; F(K, x))$.

This construction respects the support of the original randomized circuit, i. e., the output of the obfuscated circuit is always in the support of the original circuit (on the same inputs). More formally, for all circuits $C \in \mathcal{C}_\lambda$,

all inputs $x \in \{0,1\}^*$ (of matching length), all $\Lambda \in \text{supp}(\text{piO}(C))$, we have $\Lambda(x) \in \text{supp}(C(x))$.

NOTIONS OF CORRECTNESS. Defining correctness for probabilistic obfuscation is slightly more tricky than for deterministic obfuscation. This is because an obfuscated program is defined to be *deterministic* whereas the original program can use randomness. In [CLTV15], correctness is defined by requiring that evaluations of the obfuscated program must be indistinguishable from evaluations of the original randomized program, where the randomized program must not be called twice on the same input (to exclude trivial attacks). This definition of correctness is inspired by the security definition a pseudorandom function, see Definition 2.10, which is used by [CLTV15] to derive the randomness for the obfuscated circuit.

Another notion of correctness is to require that the obfuscated program *respects the support* of the original randomized program. This notion is useful because, in many cases, it allows to avoid a reduction to the correctness property above which can simplify proofs.

Dodis et al. [DHRW16] introduce another notion of correctness which guarantees that a one-time evaluation of the obfuscated program is distributed *identically* to the distribution induced by the original randomized circuit (on the same input), where the distribution is also over the randomness of the obfuscator. Clearly, this notion of correctness implies that the obfuscated circuit respects the support of the original program.

In Part I, we elaborate on these notions in more detail and, in Part II, we introduce a novel notion of *statistical* correctness which proves to be quite useful.

2.6 ASSUMING ADVERSARIAL KNOWLEDGE

To make assumptions on adversarial knowledge is a common paradigm used in cryptography. Such assumptions state that if an adversary is able to derive some output, then he *must know* how he computed that output. This kind of “adversarial knowledge” is modeled by the ability to algorithmically extract certain information from the adversary’s code (and random tape).

With the knowledge-of-exponent assumption (KEA), Damgård [Dam92] introduced the first assumption about adversarial knowledge. The KEA is an assumption on cryptographic groups and informally states that given two random group generators g and h , then the only way to produce (A, B) such that $\text{dlog}_A(B) = \text{dlog}_g(h)$ is to compute (g^c, h^c) for some $c \in \mathbb{Z}_p$. That is, every adversary producing such (A, B) must know c such that $(A, B) = (g^c, h^c)$.

DEFINITION 2.1g (Knowledge-of-exponent assumption, [Dam92]). Let G be a cyclic group with group parameter generation algorithm $G\text{Gen}$ producing triplets of the form $pp_G := (G, p, g)$. The knowledge-of-exponent assumption

(KEA) holds relative to GGen if for every PPT adversary \mathcal{A} , there exists a PPT extractor \mathcal{E} such that

$$\Pr \left[\begin{array}{l} pp_G \leftarrow \text{GGen}(1^\lambda) \\ x \leftarrow \mathbb{Z}_p \\ (A, B) \leftarrow \mathcal{A}(pp_G, g^x; r) \\ c \leftarrow \mathcal{E}(pp_G, g^x, r) \end{array} : A^x = B \wedge g^c \neq A \right]$$

is a negligible function in λ .

Note that KEA is only nontrivial if computing discrete logarithms is hard relative to GGen.

This paradigm of knowledge-type assumptions was later generalized as extractable functions by Canetti and Dakdouk [CD08; CD09]. These are (one-way) function families such that every adversary who (given a function description f) outputs an image $y \in \text{image}(f)$ must know a pre-image of y under f .

DEFINITION 2.20 (Extractable one-way functions without auxiliary information, [BCPR16]). An *extractable one-way function (EOWF) family ensemble without auxiliary information* is a function family ensemble $\mathcal{F} = (\text{Gen}, \text{Eval})$ (cf. Definition 2.8) if the following two properties are satisfied.

ONE-WAYNESS. For every PPT adversary \mathcal{A} ,

$$\Pr \left[\text{Exp}_{\mathcal{F}, \mathcal{A}}^{\text{ow}}(\lambda) = 1 \right]$$

is negligible, where $\text{Exp}_{\mathcal{F}, \mathcal{A}}^{\text{ow}}(\lambda)$ is defined in Figure 2.3a.

EXTRACTABILITY. For every PPT adversary \mathcal{X} (using a random tape of length $p(\lambda)$), there exists a PPT algorithm \mathcal{E}_X such that

$$\Pr \left[\text{Exp}_{\mathcal{F}, \mathcal{X}, \mathcal{E}_X}^{\text{ext}}(\lambda) = 1 \right]$$

is overwhelming, where $\text{Exp}_{\mathcal{F}, \mathcal{X}, \mathcal{E}_X}^{\text{ext}}(\lambda)$ is defined in Figure 2.3b.

$\text{Exp}_{\mathcal{F}, \mathcal{A}}^{\text{ow}}(\lambda)$	$\text{Exp}_{\mathcal{F}, \mathcal{X}, \mathcal{E}_X}^{\text{ext}}(\lambda)$
$f \leftarrow \text{Gen}(1^\lambda)$	$f \leftarrow \text{Gen}(1^\lambda)$
$x \leftarrow \{0, 1\}^{n(\lambda)}$	$r_X \leftarrow \{0, 1\}^{p(\lambda)}$
$y := f(x)$	$y \leftarrow \mathcal{X}(f; r_X)$
$x' \leftarrow \mathcal{A}(f, y)$	$x \leftarrow \mathcal{E}_X(f, r_X)$
return $f(x') = y$	return $(f(x) = y) \vee (\forall x': f(x') \neq y)$

(a) The one-way game.

(b) The extraction game.

Figure 2.3: One-way and extraction games for EOWFs.

DEFINITION 2.21 (Extractable one-way function family ensembles with common auxiliary information, [BCPR16]). An *extractable one-way function (EOWF) family ensemble with common auxiliary information* is a function family ensemble $\mathcal{F} = (\text{Gen}, \text{Eval})$ (cf. Definition 2.8) if the following properties are satisfied.

ONE-WAYNESS. As in Definition 2.20.

ONE-WAYNESS (STRONGER). For every PPT adversary \mathcal{A} , for every polynomial b and for every $z \in \{0, 1\}^{b(\lambda)}$,

$$\Pr \left[\text{Exp}_{\mathcal{F}, \mathcal{A}, z}^{\text{ow-aux}}(\lambda) = 1 \right]$$

is negligible, where $\text{Exp}_{\mathcal{F}, \mathcal{A}, z}^{\text{ow-aux}}(\lambda)$ is defined in Figure 2.4a.

EXTRACTABILITY. For every PPT adversary \mathcal{X} , there exists a PPT algorithm $\mathcal{E}_{\mathcal{X}}$ such that for every polynomial b and every $z \in \{0, 1\}^{b(\lambda)}$,

$$\Pr \left[\text{Exp}_{\mathcal{F}, \mathcal{X}, \mathcal{E}_{\mathcal{X}}, z}^{\text{ext-aux}}(\lambda) = 1 \right]$$

is overwhelming, where $\text{Exp}_{\mathcal{F}, \mathcal{X}, \mathcal{E}_{\mathcal{X}}, z}^{\text{ext-aux}}(\lambda)$ is defined in Figure 2.4b.

$\text{Exp}_{\mathcal{F}, \mathcal{A}, z}^{\text{ow-aux}}(\lambda)$	$\text{Exp}_{\mathcal{F}, \mathcal{X}, \mathcal{E}_{\mathcal{X}}, z}^{\text{ext-aux}}(\lambda)$
$f \leftarrow \text{Gen}(1^\lambda)$	$f \leftarrow \text{Gen}(1^\lambda)$
$x \leftarrow \{0, 1\}^{n(\lambda)}$	$r_{\mathcal{X}} \leftarrow \{0, 1\}^{p(\lambda)}$
$y := f(x)$	$y \leftarrow \mathcal{X}(f, z; r_{\mathcal{X}})$
$x' \leftarrow \mathcal{A}(f, y, z)$	$x \leftarrow \mathcal{E}_{\mathcal{X}}(f, z, r_{\mathcal{X}})$
return $f(x') = y$	return $(f(x) = y) \vee (\forall x': f(x') \neq y)$

(a) The one-way game.

(b) The extraction game.

Figure 2.4: One-way and extraction game for EOWFs with common auxiliary input.

We note that the weaker notion of one-wayness without auxiliary input is sufficient for us. Assuming non-uniform adversaries, one-wayness without auxiliary input and one-wayness with auxiliary input are equivalent. We emphasize that in contrast to one-wayness, extractability with common auxiliary input is strictly stronger than extractability without common auxiliary input (even when considering non-uniform adversaries).

DEFINITION 2.22 (Extractable one-way function family ensembles with b -bounded common auxiliary information, [BCPR16]). Like Definition 2.21 but with a fixed polynomial b determining the length of the common auxiliary information.

In Section 17.2.1 we define further variants of EOWFs.

²² Note that Barak [Bar01] introduces a non-black-box technique that does not rely on knowledge-type assumptions. However, it is not known whether this technique extends to settings other than zero-knowledge proofs.

ON THE NEED FOR NON-BLACK-BOX TECHNIQUES. The knowledge-of-exponent assumption [HT98; BP04] as well as extractable one-way functions [CD09; BCC+17] imply 3-message zero-knowledge protocols. However, this notion of zero-knowledge protocols is known to be impossible with respect to black-box simulation [GK96; Kato08; FGJ18]. Hence, constructing EOWFs (or groups where KEA holds) requires non-black-box techniques. So far, no such constructions are known.²²

FALSIFIABILITY AND PLAUSIBILITY. Furthermore, knowledge-type assumptions such as the knowledge-of-exponent assumption or the existence of extractable one-way functions are not falsifiable in the sense of Naor [Nao03],

see Section 2.2.1. Hence, it is difficult to reason about their plausibility. While assuming KEA or the existence of EOWFs might be better than simply assuming “the scheme is secure”, a proof based on such assumptions is not satisfying.

In Part III of this thesis, we use knowledge-type assumptions “negatively” in the sense that we prove that the existence of certain knowledge-type assumptions conflicts with other notions. Furthermore, in Part II of this thesis, we demonstrate how certain knowledge-type assumptions can be avoided.

We note that the notion of extractable one-way functions with unbounded common auxiliary information as per Definition 2.21 conflicts with indistinguishability obfuscation [BCPR16] since the common auxiliary information can contain a circuit which, given a function description f , outputs $f(F(K, f))$ thus forcing any extractor to break the pseudorandom function F . More precisely, consider the universal adversary which simply executes all commands in its auxiliary input. Then, there needs to be one “universal” extractor which works for *all* adversarial strategies. Suppose the auxiliary input contains an *obfuscated* program which computes $f(F(K, f))$ as above, then the universal extractor needs to either break the pseudorandom function F or the security of the used obfuscator. Hence, it is unlikely that such EOWFs exist (even heuristically).

2.7 PUBLIC-KEY ENCRYPTION AND ITS VARIANTS

Encryption enables users to confidentially communicate messages over insecure channels. Secret-key encryption (SKE) assumes both users know a common secret key which is used for both encryption and decryption. SKE is known to exist from OWFs [GM84]. In contrast to classical secret-key encryption, public-key encryption (PKE) schemes use two keys; a public key for encryption and a secret key for decryption. Currently, it is not known whether PKE schemes can be constructed from OWFs but there is evidence, that this might not be possible [IR89; Dac16]. However, if IO exists, then PKE schemes can be constructed from OWFs.

DEFINITION 2.23 (Public-key encryption, [ElG85; Yao82]). A *public-key encryption (PKE) scheme* for message space $\mathcal{M} = \{\mathcal{M}_\lambda\}_{\lambda \in \mathbb{N}}$ is a tuple of algorithms $E = (\text{KGen}, \text{Enc}, \text{Dec})$ such that the following properties are satisfied.

PERFECT CORRECTNESS. For all $\lambda \in \mathbb{N}$, all $(pk, sk) \in \text{supp}(\text{KGen}(1^\lambda))$, all $m \in \mathcal{M}_\lambda$,

$$\Pr [\text{Dec}(sk, \text{Enc}(pk, m)) = m] = 1.$$

IND-CPA SECURITY. For all legitimate PPT adversaries \mathcal{A} ,

$$\text{Adv}_{E, \mathcal{A}}^{\text{ind-cpa}}(\lambda) := \Pr \left[\begin{array}{l} (pk, sk) \leftarrow \text{KGen}(1^\lambda) \\ (m_0, m_1, st) \leftarrow \mathcal{A}(pk) \\ c^* \leftarrow \text{Enc}(pk, m_0) \end{array} : \mathcal{A}(pk, c^*, st) = 1 \right]$$

$$- \Pr \left[\begin{array}{l} (pk, sk) \leftarrow \text{KGen}(1^\lambda) \\ (m_0, m_1, st) \leftarrow \mathcal{A}(pk) \\ c^* \leftarrow \text{Enc}(pk, m_1) \end{array} : \mathcal{A}(pk, c^*, st) = 1 \right]$$

is negligible, where legitimate means that \mathcal{A} always outputs two messages $m_0, m_1 \in \mathcal{M}$ of identical length.

Without loss of generality, we assume that sk is the random tape used for key generation. Therefore, making the random tape of KGen explicit, we write $(pk, sk) = \text{KGen}(1^\lambda; sk)$.

Note that **IND-CPA** security as defined in Definition 2.23 is not the de facto security for use in practice. However, it suffices as building block for our applications.

2.7.1 Fully Homomorphic Encryption

Equipping public-key encryption schemes with the additional property that one can publicly evaluate programs on the contents of ciphertexts was a long standing open problem since [RAD78] and was resolved by Gentry [Gen09]. This variant of public-key encryption has many applications in theory and practice. For instance, fully homomorphic encryption allows to outsource computations on private data without compromising its confidentiality.

Let $\mathcal{P} = (\mathcal{P}_\lambda)_{\lambda \in \mathbb{N}}$ be a family of sets of polynomial sized circuits of arity $\alpha(\lambda)$, i. e., the set \mathcal{P}_λ contains circuits of polynomial size in λ . We assume that for any $\lambda \in \mathbb{N}$ the circuits in \mathcal{P}_λ share the common input domain $(\{0, 1\}^{\text{poly}(\lambda)})^{\alpha(\lambda)}$ for a fixed polynomial $\text{poly}(\lambda)$. A homomorphic encryption scheme enables evaluation of circuits on encrypted data.

DEFINITION 2.24 (Fully homomorphic encryption, [Gen09]). A *fully homomorphic (public-key) encryption (FHE) scheme* with message space $\mathcal{M} = \{\mathcal{M}_\lambda\}_{\lambda \in \mathbb{N}}$ for a deterministic circuit family $\mathcal{P} = (\mathcal{P}_\lambda)_{\lambda \in \mathbb{N}}$ of arity $\alpha(\lambda)$ and input domain $(\{0, 1\}^{\text{poly}(\lambda)})^{\alpha(\lambda)}$ is a tuple of **PPT** algorithms $\text{FHE} = (\text{KGen}, \text{Enc}, \text{Dec}, \text{Eval})$ such that $(\text{KGen}, \text{Enc}, \text{Dec})$ is a perfectly correct **IND-CPA** secure public-key encryption scheme and the following properties are met.

PERFECT CORRECTNESS. For all $\lambda \in \mathbb{N}$, all $m_1, \dots, m_{\alpha(\lambda)} \in \mathcal{M}_\lambda$, $P \in \mathcal{P}_\lambda$,

$$\Pr \left[\begin{array}{l} (pk, sk) \leftarrow \text{KGen}(1^\lambda) \\ c_i \leftarrow \text{Enc}(pk, m_i) \\ c \leftarrow \text{Eval}(pk, P, c_1, \dots, c_{\alpha(\lambda)}) \end{array} : \begin{array}{l} \text{Dec}(sk, c) \\ = \\ P(m_1, \dots, m_{\alpha(\lambda)}) \end{array} \right] = 1.$$

COMPACTNESS. The size of the output of Eval is polynomial in λ and independent of the size of the circuit P .

The first **FHE** scheme based on standard lattice assumptions (additionally assuming circular security) is due to [BV14a; BV11; BV14b].²³

Further, due to [CLTV15], probabilistic indistinguishability obfuscation in conjunction with (slightly super-polynomially secure) perfectly correct and perfectly re-randomizable public-key encryption yields a perfectly correct and perfectly re-randomizable fully homomorphic encryption scheme (without using circular security assumptions).

²³ Circular security means that it is assumed that the scheme remains secure even if an encryption of the secret key is known.

2.8 NON-INTERACTIVE ZERO-KNOWLEDGE PROOF SYSTEMS

Zero-knowledge protocols [GMR85] are a fundamental notion in cryptography. They allow to prove the validity of statements without revealing any information beyond that fact. One-way functions are both sufficient [GMR85] and necessary [OW93] for the existence of interactive zero-knowledge protocols for languages in \mathcal{NP} .

Non-interactive zero-knowledge (NIZK) proof systems [BFM88] are the non-interactive analogue of zero-knowledge protocols, where only one message between the prover and the verifier is exchanged. In contrast to interactive zero-knowledge protocols, NIZK proof systems in the plain model are only possible for languages in \mathcal{BPP} [GO94]. However, the notion of zero-knowledge is not meaningful if restricted to \mathcal{BPP} since a verifier does not need the prover to verify the validity of a statement. Hence, NIZK proof systems for \mathcal{NP} only exist in stronger models such as the common reference string model, where a reference string is generated by a trusted party and made available to all parties.

NIZK proof systems find use in various areas of cryptography. A NIZK proof system for a language L with witness relation \mathcal{R}_L enables to prove in a non-interactive manner that some statements are in L without leaking information about corresponding witnesses.

DEFINITION 2.25 (Non-interactive zero-knowledge proof system [GOS06]). A non-interactive zero-knowledge (NIZK) proof system for a language $L \in \mathcal{NP}$ (with witness relation \mathcal{R}_L) is a tuple of PPT algorithms $\text{NIZK} = (\text{Setup}, \text{Prove}, \text{Verify})$ such that Setup is a common reference string generation algorithm, Prove is a proving algorithm and Verify is a (deterministic) verification algorithm.

- $\text{Setup}(1^\lambda)$ outputs a common reference string σ .
- $\text{Prove}(\sigma, x, w)$, on input of σ , a statement x and a witness w , outputs a proof π .
- $\text{Verify}(\sigma, x, \pi)$, on input of σ , a statement x and a proof π , deterministically outputs either 1 or 0.

A NIZK proof system is required to meet the following properties.

PERFECT COMPLETENESS. For every $\lambda \in \mathbb{N}$, every $(x, w) \in \mathcal{R}_L$ with $|x| = \lambda$, we have that

$$\Pr \left[\begin{array}{l} \sigma \leftarrow \text{Setup}(1^\lambda) \\ \pi \leftarrow \text{Prove}(\sigma, x, w) \end{array} : \text{Verify}(\sigma, x, \pi) = 1 \right] = 1,$$

where the probability is over the random coins of Setup and Prove.

STATISTICAL SOUNDNESS. For every $\lambda \in \mathbb{N}$, every $x \notin L$ (of polynomial length) and every (possibly unbounded) adversary \mathcal{A} ,

$$\text{Adv}_{\text{NIZK}, \mathcal{A}}^{\text{sound}}(\lambda) := \Pr \left[\begin{array}{l} \sigma \leftarrow \text{Setup}(1^\lambda) \\ \pi \leftarrow \mathcal{A}(\sigma, x) \end{array} : \text{Verify}(\sigma, x, \pi) = 1 \right] < 2^{-\lambda},$$

where the probability is over the random coins of Setup and \mathcal{A} .

COMPUTATIONAL ZERO-KNOWLEDGE. There exists a **PPT** algorithm $\text{Sim} = (\text{Sim}_0, \text{Sim}_1)$ such that for every **PPT** adversary \mathcal{A} ,

$$\text{Adv}_{\text{NIZK}, \text{Sim}, \mathcal{A}}^{\text{zk}}(\lambda) := \left| \Pr \left[\sigma \leftarrow \text{Setup}(1^\lambda): \mathcal{A}^{\text{Prove}(\sigma, \cdot)}(\sigma) = 1 \right] - \Pr \left[(\sigma, \tau_\Pi) \leftarrow \text{Sim}_0(1^\lambda): \mathcal{A}^{\text{Sim}'_1(\sigma, \tau_\Pi, \cdot)}(\sigma) = 1 \right] \right|$$

is negligible in λ , where $\text{Sim}'_1(\sigma, \tau_\Pi, x, w)$ operates like $\text{Sim}_1(\sigma, \tau_\Pi, x)$ only if $(x, w) \in \mathcal{R}_\perp$ and returns \perp otherwise.

REMARK 2.1. The above definition of **NIZK** proof systems can be strengthened such that NIZK.Setup is required to produce uniformly random common reference strings from $\{0, 1\}^{n_\Pi(\lambda)}$, so-called common *random* strings.

For interactive zero-knowledge proof systems, one-way functions are sufficient [**GMW86**] (even in the plain model, i. e., without a common reference string). This is not the case for non-interactive zero-knowledge proof systems. To date, **NIZK** proof systems can be built from a variety of assumptions including assumptions on pairing-friendly groups [**GOS06**; **CH19**], certain trapdoor permutations (and, hence, RSA) [**FLS90**], the learning-with-errors (**LWE**) assumption [**PS19**] and indistinguishability obfuscation (in conjunction with one-way functions) [**BP15a**].

Part I

DOUBLY-PROBABILISTIC
INDISTINGUISHABILITY OBFUSCATION

I	DOUBLY-PROBABILISTIC INDISTINGUISHABILITY OBFUSCATION	
3	INTRODUCTION	41
3.1	Technical Overview	43
4	PRELIMINARIES	47
4.1	Perfect Puncturable PRFs	47
4.2	Extremely Lossy Functions	47
4.2.1	Instantiating Extremely Lossy Functions	48
5	DOUBLY-PROBABILISTIC IO	51
6	CONSTRUCTION	59
6.1	Overview	59
6.2	Constructing Doubly-Probabilistic IO	60
7	LEVELED HOMOMORPHIC ENCRYPTION	73

3

INTRODUCTION

In this part, we present the results of [ACH20]. This part is taken verbatim from [ACH20] with minor changes.

OBFUSCATION. Code obfuscation has been formalized already in the early 2000s as a cryptographic building block, by Hada [Hadoo] and Barak et al. [BGI+01], along with a number of early positive [Can97; LPS04; Wee05; HRsV07; HMS07] and negative [BGI+01; GK05; Wee05] results. However, prior to the candidate obfuscation scheme of Garg et al. [GGH+13], only relatively few positive results on obfuscation were known.

The first candidate obfuscator from [GGH+13] changed things. Their work identified indistinguishability obfuscation (IO, cf. [BGI+01; GR07]) as an achievable *and* useful general notion of obfuscation: it presented a candidate indistinguishability obfuscator, along with a first highly nontrivial application (functional encryption). Since then, a vast number of applications have been proposed, ranging from functional [GGH+13], deniable [SW14], and fully homomorphic [CLTV15] encryption, over multi-party computation (e. g., [GGHR14]), to separation results (e. g., [HRW16]). In the process, powerful techniques like “puncturing” [SW14] have been discovered, which have found applications even beyond obfuscation (e. g., in multi-party computation [BL18a; GS18], instantiating the Fiat-Shamir paradigm [CCRR18], and verifiable random functions [Bit17; GHKW17]). Besides, the notion of IO itself has been refined, and related to other notions of obfuscation [ABG+13; BP13; BCP14; BCKP14; CLTV15; IPS15], and various different constructions of obfuscators have been presented [PST14; Zim15; AJ15; BV15; AS17; Lin17; LT17; CVW18; GJK18; CHVW19; AJL+19; Agr19; JLMS19; BDGM20a; GJLS20; JLS20; GP20; AP20; BIJ+20; BDGM20b; WW20].

SUBEXPONENTIAL ASSUMPTIONS. It is currently hard to find a cryptographic primitive that can *not* be constructed from IO (in combination with another mild assumption such as the existence of one-way functions). However, some of the known IO-based constructions come only with *subexponential* reductions to IO. For instance, the only known IO-based constructions of fully homomorphic encryption [CLTV15], spooky encryption [DHRW16], and graded encoding schemes [FHHL18] suffer from reductions with a subexponential loss.

Hence, while IO has generally been recognized as an extremely powerful primitive (even to the extent being called a “central hub” for cryptography [SW14]), it is not at all clear if this also holds for *polynomially* secure IO. Indeed, it is conceivable that only polynomially secure IO exists, in which case much of IO’s power stands in question.

More generally, subexponential reductions (in particular to IO) are undesirable. Namely, the security of existing IO constructions is still not well-understood, and in particular current state-of-the-art constructions of IO schemes (such as [AS17; Lin17; LT17]) already require subexponential com-

putational assumptions themselves. Hence, assuming subexponential IO is a particularly risky bet. This suspicion is confirmed in part by Pass and Shelat [Ps16] who separate polynomial and subexponential security for virtual black-box obfuscation.

Removing subexponential assumptions in general and from IO -based constructions in particular has already explicitly been considered in [LM16; GS16] and [GPS16; GPSZ17; LZ17] respectively. These works offer general techniques and ideas to turn subexponential reductions into polynomial ones. For instance, [GPSZ17; LZ17] offer ways to replace (subexponential) IO -based constructions with (polynomial) constructions based on functional encryption. Of course, this requires a special structure of the primitive to be implemented, and is demonstrated for several primitives, including non-interactive key exchange and short signature schemes.

OUR CONTRIBUTION. In this part, we are also concerned with substituting subexponential with polynomial reductions in IO -based constructions. Unlike [GPSZ17; LZ17], however, we do not follow the approach of using functional encryption directly in place of IO , but instead will employ extremely lossy functions (ELFs) [Zha16] to “absorb” subexponential complexity.²⁴

²⁴ That means that our final schemes depend on ELFs, which are currently only known to be instantiable from exponential assumptions. However, we stress that ELFs can be built from exponential variants of very standard assumptions, such as the decisional Diffie-Hellman (DDH) assumption.

We will implement a variant of probabilistic indistinguishability obfuscation (pIO , introduced in [CLTV15]) using polynomially secure IO (and ELFs). A pIO scheme can be used to obfuscate *probabilistic* (i. e., randomized) programs, and are currently the only way to obtain, e. g., fully homomorphic encryption (FHE) schemes without circular security assumptions [CLTV15]. However, the only previous construction of pIO schemes required subexponentially secure IO [CLTV15]. Hence, our construction yields the first FHE scheme from polynomially secure IO (and ELFs). Similarly, we can turn the assumption of subexponentially secure IO into polynomially secure IO (plus ELFs) in the construction of spooky encryption from [DHRW16].

Both FHE and spooky encryption are quite powerful primitives, and we obtain several “spin-off results” by revisiting their implications. For instance, when instantiating the pIO -based FHE construction of [CLTV15] with our pIO scheme and a suitable public-key encryption scheme, we obtain a fully key-dependent message (KDM) secure public-key encryption scheme from (polynomially secure) IO and the exponentially secure DDH assumption (and no further assumptions). Under the same assumptions, we obtain multi-key FHE with threshold decryption and function secret sharing schemes from the spooky encryption construction from [DHRW16].

ON THE PLAUSIBILITY OF ELFS. One could argue that we trade one exponential assumption for another, and it is not clear that assuming polynomial IO and exponential DDH is any better than assuming only subexponential IO in the first place. Seconding Zhandry [Zha16] here, we think that exponential DDH is a realistic assumption that is far more popular, better-investigated, and arguably more plausible than subexponential IO . Much of the currently deployed cryptography relies on (in fact a strong variant of) exponential DDH, because parameters are almost always chosen according to the best known attacks.

ON THE NUMBER OF ASSUMPTIONS. Another natural observation is that IO for general circuits is already an exponential family of assumptions in

itself (one for each obfuscated circuit) [PST14]. It might seem that this lets the challenge of relying on polynomially secure IO instead of subexponentially secure IO appear less appealing. We make two comments on that.

- First, being an exponential family of assumptions and assuming resistance against subexponential adversaries are orthogonal issues. Many cryptographic assumptions have several dimensions of strengths, and relaxing the assumption in any of these dimensions is desirable.²⁵ In this part, we make progress in one important dimension. By replacing subexponential IO by polynomial IO plus exponential DDH, we effectively trade an *exponential* number of subexponential hardness assumptions in exchange for a *single* (plausible, well-studied) exponential hardness assumption (plus an exponential family of polynomial hardness assumptions).
- Second, IO being an exponential family of assumptions can be considered an artificial consequence of working on the general notion of IO for *arbitrary circuits*. When using IO in concrete constructions (e. g., in all the constructions described in this part), one almost never needs to assume IO for all circuits. It usually suffices to assume IO for a constant number of specific circuits (namely those being obfuscated in the construction and the analysis). Hence, IO is a small number of assumptions when used for building a cryptographic primitive.

²⁵ For example, if a protocol relies on the subexponential hardness of LWE with exponential modulus-to-noise ratio, it would be desirable to achieve the same while relying only on polynomially secure LWE, even if the modulus-to-noise ratio remains exponential.

3.1 TECHNICAL OVERVIEW

THE PIO CONSTRUCTION OF CANETTI ET AL. To describe our ideas, it will be helpful to briefly review the work of Canetti et al. [CLTV15]. In a nutshell, they define the notion of pIO as a way to obfuscate probabilistic programs, and show how to use pIO to implement the first FHE scheme without any circular security assumption. Intuitively, where the notion of IO captures that the obfuscation $iO(P)$ of a *deterministic* program P does not leak anything beyond the functionality of P , pIO captures the same for *probabilistic* programs P .²⁶

They also show how to implement pIO with an indistinguishability obfuscator iO and a pseudorandom function (PRF) F . Namely, in order to obfuscate a probabilistic program P , Canetti et al. obfuscate the *deterministic* program P' that, on input x , runs $P(x)$ with random coins $r = F(K, x)$. Here, K is a PRF key hard-coded into P' . The security proof uses “puncturing” techniques [SW14] and a hybrid argument over all possible P -inputs x . More specifically, for each P -input x , separate reductions to the security of iO and F show that the execution of $P'(x)$ is secure.²⁷

This proof strategy is very general and does not need to make any specific assumptions about the structure of P . (In fact, this strategy can be viewed as a specific form of “complexity leveraging”, technically similar to the conversion of selective security into adaptive security, e. g., [BB04a].) However, the price to pay is a reduction loss which is linear in the size of the input domain (which usually is exponentially large). In particular, even after scaling security parameters suitably, Canetti et al. still require subexponentially secure IO and PRFs.

²⁶ This is of course an oversimplification. Also, [CLTV15] define several types of pIO security that provide a tradeoff between security and achievability.

²⁷ Again, we are not very specific about the form of desired or assumed security. However, we believe that for this exposition, these specifics do not matter.

MORE ON PREVIOUS WORKS TO REMOVE SUBEXPONENTIALITY. There are a number of known ways to deal with subexponential reduction losses due to complexity leveraging (or related techniques). For instance, various semi-generic (pre-IO) techniques seek to achieve adaptive security (for different primitives) by establishing an algebraic or combinatorial structure on the used inputs [BB04b; Wat05; HK08; HW09], and can sometimes be adapted to the IO setting [HSW14]. But like the already-mentioned, somewhat more general approaches [GPSZ17; LZ17], these works make specific assumptions about the structure of the involved computations.

A somewhat more general approach (that works for more general classes of programs) was outlined by Zhandry [Zha16], who introduces the notion of “extremely lossy functions” (ELFs). Intuitively, an ELF is an injective function G that can be switched into an “extremely lossy mode”, in which its range is polynomially small. Such an ELF can sometimes be used to “preprocess” inputs in a cryptographic scheme, with the following benefit: a security reduction can switch the ELF to extremely lossy mode, so that only a polynomial number of (preprocessed) inputs $G(x)$ need to be considered. This simplifies a potential hybrid argument over all (preprocessed) inputs $G(x)$, and can lead to a polynomial (instead of a subexponential) reduction.

However, trying to apply this strategy to the construction and reduction of Canetti et al. (as sketched above) directly fails. Namely, in their application, inputs to the cryptographic scheme will be inputs x to an arbitrary (probabilistic) program P ; preprocessing them with an ELF will destroy their structure, and it is not clear how to run P on ELF-preprocessed inputs $G(x)$. Indeed, applying ELFs to realize pIO requires fundamentally different techniques.

MAIN IDEA: PIO WITH SPARSIFIABLE INPUTS. Instead, we will restrict ourselves to programs P that take as input an element x from a small number of (arbitrary but efficiently samplable) distributions. In other words, all possible inputs x need to be in the range of one of a small number of efficient samplers S_i . As an example, for $i \in \{0, 1\}$, sampler S_i could sample ciphertexts C that encrypt plaintext i . Moreover, we require that all inputs to a program P to be obfuscated are at some point actually sampled from some S_i according to a certain process.

Obfuscating a given probabilistic program P (that takes as inputs one or more x as above) now consists of two steps:

1. First, we *encode* all inputs x , in the sense that we compile S_i to attach a “certificate” aux to x . This certificate aux guarantees that x has really been sampled using S_i . Furthermore, the compiled sampler S_i uses preprocessed random coins of the form $G(r)$ (instead of r) for an ELF G . (When G is in injective mode, this does not affect the distribution of sampled x .) The certificate aux additionally guarantees this choice of random coins.²⁸
2. Second, we produce the actual obfuscation of the probabilistic program P as follows. We use an indistinguishability obfuscator iO to obfuscate the following (deterministic) variant P' of P : on inputs x_1, \dots, x_ℓ with certificates aux_1, \dots, aux_ℓ , P' first checks the certificates aux_i and aborts if one of them is invalid. Next, P' runs $P(x_1, \dots, x_\ell)$, with random

²⁸ Looking ahead, this “certificate” will be implemented using a NIZK in our construction.

coins $F(K, (x_i)_{i=1}^{\ell})$ for a PRF F and a hard-coded PRF key K . Finally, P' outputs P 's output.

Maybe the most important property of this setup is that now the sets of inputs x_i are “sparsifiable” in the following sense. If we set G to extremely lossy mode, then only a polynomial number of different random coins r can occur. Hence, each S_i will output one of only a small number of possible samples (e. g., encryptions C generated with random coins from a small set). In that sense, the set of possible inputs x_i to P has been “sparsified”, and a hybrid argument over all possible inputs as in [CLTV15] is possible with polynomial loss.

We stress that our technique of applying ELFs fundamentally differs from [Zha16]. In [Zha16], the constructed primitive itself ensures that G is applied on all inputs. When approaching the challenge of constructing pIO, however, the input to the primitive must externally be sampled using random coins that are preprocessed with G . This process is not under the control of the primitive and therefore requires a mechanism certifying that inputs are generated according to this specific process. We implement this mechanism using the combination of compiling the sampler for the input distribution into a “certifying sampler” (step 1) and restricting correctness of the obfuscated program (step 2).

Surprisingly, our pIO scheme achieves the notion of “dynamic-input pIO” [CLTV15], a very strong variant of pIO security. On a high level, dynamic-input pIO guarantees indistinguishability between obfuscations of probabilistic programs as long as their output distributions on adversarially chosen inputs are indistinguishable. This constitutes a very strong requirement and, in fact, implies differing-inputs obfuscation [BGI+01; ABG+13], a notion for which strong impossibility results exist [GGHW17; BSW16]. However, our obfuscator produces circuits which are only required to work on inputs certifiably generated according to a specific process. Hence, our pIO scheme enjoys a restricted form of correctness. This enables us to circumvent the impossibility results [GGHW17; BSW16].

APPLICATIONS. One obvious question is of course how restrictive our assumption on input domains really is. We show that our assumptions apply to two existing pIO-based constructions, with a number of interesting consequences.

First, we revisit the pIO-based construction of fully homomorphic encryption from [CLTV15]. Here, pIO is used to obfuscate the FHE evaluation algorithm that takes two ciphertexts (say, of two bit plaintexts b_0 and b_1) as input, and outputs a ciphertext of the NAND of the two plaintexts (i. e., $b_0 \bar{\wedge} b_1$). If we set S_b to be a sampler that samples an encryption of b , this setting perfectly fits our scheme. Hence, we obtain first a leveled homomorphic encryption (LHE) scheme, and from this an FHE scheme using the high-level strategy from [CLTV15]. Hence, putting this together with our pIO construction, we obtain an FHE scheme from polynomially secure IO and an ELF (and no further assumptions).

We note that the above FHE scheme is also fully key-dependent message (KDM, see [BRSo3]) secure when implemented with a suitable basic public-key encryption scheme (such as the DDH-based scheme of [BHHO08]). In that case, the FHE is secure even when an encryption of its own secret key $C_{sk} = \text{Enc}(pk, sk)$ is public. However, such an encryption C_{sk} can be

transformed into an encryption $\text{Enc}(pk, f(sk))$ of an arbitrary function of sk thanks to the fully homomorphic properties of the **FHE** scheme. This leads to a conceptually very simple fully **KDM**-secure encryption scheme from polynomial assumptions (and **ELFs**). (We stress that we do not claim novelty for this observation. The connection between **FHE** and **KDM** security has already been observed in [BHHI10], and [DHRW16] have observed that the **FHE** construction of Canetti et al. preserves interesting properties of the underlying encryption scheme. However, [DHRW16] do not explicitly mention **KDM** security, and we find these consequences interesting enough to point out.)

As a second application, we consider spooky encryption (with **CRS**) introduced by Dodis et al. [DHRW16]. Note that this is not the contribution of this author and is hence not included in this thesis. We refer the reader to the original publication [ACH20] for more details.

We believe that our new notion of obfuscation will prove useful in other applications; for example, it would likely allow to improve the recent result [CRRV17], which constructed **IND-CCA₁**-secure **FHE** from subexponentially secure **IO**.

FOLLOW-UP WORK. In the recent work [DN18], Döttling and Nishimaki define the notion universal proxy re-encryption (**UPRE**). **UPRE** schemes allow a proxy to convert any ciphertext under any public key of any existing **PKE** scheme into a ciphertext under any public key of any possibly different existing **PKE** scheme. [DN18] instantiate **UPRE** based on probabilistic **IO** due to [CLTV15]. **UPRE** for all **PKE** schemes (including non re-randomizable ones) requires dynamic-input **pIO**, which implies differing-inputs obfuscation. However, [DN18] observe that our notion of doubly-probabilistic **IO** suffices which yields an instantiation of **UPRE** for all **PKE** schemes based on polynomial **IO** and exponential **DDH**.

4

PRELIMINARIES

In this chapter, we introduce the necessary preliminaries for this part. In Section 4.1, we introduce a variant of pseudorandom functions and in Section 4.2 we introduce extremely lossy functions – the core tool for our construction in Chapter 6.

4.1 PERFECT PUNCTURABLE PRFS

In this part, we use puncturable PRFs with the additional property that a one-time evaluation is uniformly distributed over the output space.

DEFINITION 4.1 (Perfect pPRF, [DHRW16]). A *perfect puncturable PRF* is a family of puncturable PRFs F mapping $n(\lambda)$ bits to $m(\lambda)$ bits such that for all $\lambda \in \mathbb{N}$ and all inputs $x \in \{0, 1\}^{n(\lambda)}$, the distribution

$$\{K \leftarrow F.\text{KGen}(1^\lambda) : F.\text{Eval}(K, x)\}$$

is identical to the uniform distribution over $\{0, 1\}^{m(\lambda)}$.

Given any puncturable PRF F' , we can build a perfect puncturable PRF F by sampling two keys $K_1 \leftarrow F'.\text{KGen}(1^\lambda)$ and $K_2 \leftarrow \{0, 1\}^{m(\lambda)}$ in the key generation algorithm and defining the evaluation algorithm to output $F'.\text{Eval}(K_1, x) \oplus K_2$ on input of x , see [DHRW16].

4.2 EXTREMELY LOSSY FUNCTIONS

In this section, we present extremely lossy functions (ELFs) introduced in [Zha16]. Extremely lossy functions are a highly powerful primitive for complexity absorption allowing to replace subexponential or even exponential security assumptions with polynomial ones.

Informally, an ELF is a function that can be generated in two different modes, an injective mode and an extremely lossy mode. In injective mode, the range of the ELF has exponential size whereas the range comprises only polynomially many elements in extremely lossy mode. However, an adversary running in time $\mathcal{O}(\sqrt{r})$, where r denotes the cardinality of the range in extremely lossy mode, can distinguish the injective from the extremely lossy mode by evaluating the function on \sqrt{r} inputs until he finds a collision. Hence, classic security against arbitrary polynomial time distinguishers between injective and lossy mode is impossible. However, security against all adversaries running in time at most r^c for $c < 1/2$ is still achievable. Consequently, the parameter r is chosen *depending* on the adversary \mathcal{A} , such that for any polynomial time adversary, there exists a polynomial r , such that \mathcal{A} cannot distinguish between injective and lossy mode.

DEFINITION 4.2 (Extremely lossy function, [Zha16]). An *extremely lossy function (ELF)* is a tuple of PPT algorithms $\text{ELF} = (\text{Gen}, \text{Eval})$ such that

- Gen on input of (M, r) , where M is an integer and $r \in [M]$, outputs a function description $G: [M] \rightarrow [N]$.
- If $r = M$, G is injective with overwhelming probability (in $\log M$) over the randomness of $\text{Gen}(M, M)$.
- For every $r \in [M]$, $|G([M])| < r$ with overwhelming probability (in $\log M$) over the randomness of $\text{Gen}(M, r)$.
- $\text{Eval}(G, \cdot)$ evaluates the function $G: [M] \rightarrow [N]$ and runs in polynomial time in $\log M$.

An extremely lossy function ELF is required to satisfy the following property.

INDISTINGUISHABILITY. For every large enough M , every polynomial P , and every inverse polynomial function δ , there exists a polynomial Q such that for every adversary \mathcal{A} running in time at most $P(\log M)$ and every $r \in [Q(\log M), M]$, the advantage of \mathcal{A}

$$\begin{aligned} \text{Adv}_{\text{ELF}, (M, r), \mathcal{A}}^{\text{elf}}(\log M) := \\ \left| \Pr [G \leftarrow \text{ELF.Gen}(M, M): \mathcal{A}(\log M, G) = 1] \right. \\ \left. - \Pr [G \leftarrow \text{ELF.Gen}(M, r): \mathcal{A}(\log M, G) = 1] \right| \end{aligned}$$

is upper bounded by $\delta(\log M)$.

Note that the above definition does not explicitly use a security parameter λ . Instead, the bitlength of M (i. e., $\log M$) acts as security parameter. As a shorthand notation, we often write $G(x)$ to denote $\text{ELF.Eval}(G, x)$. For our purpose, we consider extremely lossy functions satisfying additional properties.

DEFINITION 4.3 (Strong regularity, [Zha16]). An extremely lossy function $\text{ELF} = (\text{Gen}, \text{Eval})$ is *strongly regular* if for every polynomial r , with overwhelming probability over the choice of $G \leftarrow \text{Gen}(M, r)$, the distribution $\{x \leftarrow [M]: G(x)\}$ is statistically close to the uniform distribution over $G([M])$.

Note that if an ELF is strongly regular, it is possible to efficiently enumerate its image. The set of values obtained by evaluating an ELF on $\log(M) \cdot r \cdot \log(r)$ random inputs, where r is an upper bound on the size of its image, contains the entire image of the ELF with overwhelming probability.

4.2.1 Instantiating Extremely Lossy Functions

In [Zha16], Zhandry provides a construction of strongly regular extremely lossy functions based on the exponential hardness of the decisional Diffie-Hellman assumption **eDDH** (or any of its variants, such as the decision linear assumption). Let $G\text{Gen}$ be a group parameter generation algorithm. That is, $G\text{Gen}$ samples tuples of the form (G, p, g) such that G is a description of a cyclic group, p is its order and g is a generator of G .

DEFINITION 4.4 (Exponential **DDH** assumption, [Zha16]). The *exponential decisional Diffie-Hellman (eDDH) assumption* is true relative to $G\text{Gen}$ if there exists a polynomial q such that the following is true. For every time bound

t and probability ϵ , let $\lambda := \log q(t, 1/\epsilon)$. For every adversary \mathcal{A} running in time at most t , the advantage

$$\left| \Pr \left[\begin{array}{l} (\mathbb{G}, p, g) \leftarrow \text{GGen}(1^\lambda) \\ (a, b, c) \leftarrow \mathbb{Z}_p^3 \end{array} : \mathcal{A}(\mathbb{G}, g, g^a, g^b, g^c) = 1 \right] - \Pr \left[\begin{array}{l} (\mathbb{G}, p, g) \leftarrow \text{GGen}(1^\lambda) \\ (a, b) \leftarrow \mathbb{Z}_p^2 \end{array} : \mathcal{A}(\mathbb{G}, g, g^a, g^b, g^{ab}) = 1 \right] \right| \leq \epsilon.$$

As noted in [Zha16], groups based on elliptic curves are plausible candidates where this assumption holds. In practical instantiations of DDH over elliptic curves, the size of the group is chosen assuming that the best attack requires time in $\mathcal{O}(\sqrt{p})$, hence disproving eDDH (which amounts to showing that there is an attack which takes time less than p^c for any constant c) would have considerable practical implications. On the other hand, relying on some form of exponential hardness assumption seems necessary, since an instantiation from polynomial hardness would have surprising complexity-theoretic implications. More precisely, given access to only some super-logarithmic amount of non-determinism (i. e., $\omega(\log \log M)$ bits, where $[M]$ is the domain of the ELF), it is easy to distinguish between injective and lossy mode of the ELF. This is due to the fact that in lossy mode, the image of G has only polynomial size which means that the restriction of G to the set $D = [2^{\omega(\log \log M)}]$ (having super-polynomial cardinality) is guaranteed to have a collision (which is not the case in injective mode), and using only $\omega(\log \log M)$ bits of non-determinism this collision can be guessed.

NON-BLACK-BOX TECHNIQUES. Using extremely lossy function in a proof of security requires to explicitly use the size of the adversary as a circuit. Hence, the resulting proof does not treat the adversary as a black-box but is “slightly non-black-box”. Consequently, in some cases, extremely lossy functions can be used to avoid black-box impossibility results.

5

DOUBLY- PROBABILISTIC IO

In this chapter, we define our novel variant of indistinguishability obfuscation for probabilistic circuits, which takes into account the fact that in many applications, obfuscated (probabilistic) circuits might only have to be evaluated on inputs coming from specific distributions.

We leverage the fact that obfuscated circuits only have to work when fed with certain input distributions as follows. By “compiling” the input sampler S^{in} into a program which additionally produces auxiliary material *aux*, we introduce the possibility to verify that a given input was indeed produced according to S^{in} (on some input $m \in L$). We also refer to this auxiliary material as “certificate”. As a result, an obfuscated circuit can always verify whether its input was sampled “honestly”. If some input lacks a valid certificate, the obfuscated program simply outputs some error symbol \perp . In other words, we *restrict the correctness* of the obfuscated circuit to only hold for such well-formed inputs.

However, this approach faces two issues. First, the inputs to an obfuscated circuit might not be sampled “all at once” from a single distribution; rather, they can come from different and independent sources. We capture this behavior by defining *ℓ-source obfuscation*, to account for the fact that different inputs might have been sampled independently. Second, when inputs are sampled by different parties, there might still be interdependencies which must be accounted for. For example, a party might sample an input (e. g., a public key of an encryption scheme), pass it to a second party, who then samples a second input from a distribution that is parametrized by the first input (e. g., a ciphertext under that public key). We handle this possibility by ordering the ℓ inputs to the obfuscated circuit, and by considering a *stateful* input sampler S^{in} : when S^{in} is used to generate the i 'th sample y_i , it receives in addition to its input a state $\text{stf}(y_1, \dots, y_{i-1})$, where stf is some fixed efficiently computable *state function* (which depends on the particular application), and the y_j (for $j \in [i-1]$) are outputs sampled by the first $i-1$ sources. The state function captures the fact that a particular application might define an arbitrary communication pattern, and specifies which samples a party should have access to when generating its sample.

In the following, we consider a circuit class $\mathcal{C} = \{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$ of (probabilistic) circuits expecting ℓ inputs of length $n(\lambda)$, such that \mathcal{C}_λ only contains circuits of size λ .

Additionally, we admit the possibility that a sampler produces some additional correlated output, that will not serve as input to an obfuscated circuit. Hence, there is no need to “certify” this input using the auxiliary information, and we call this output unauthenticated output. Continuing the use case from above, given a sampler producing some public key, the unauthenticated part of that sampler’s output could be a corresponding secret key. Figure 5.1 visualizes the typical use case of a 1-source [dpIO](#) scheme (Figure 5.1a) and of an ℓ -source [dpIO](#) scheme (Figure 5.1b).

We now define input samplers more formally.

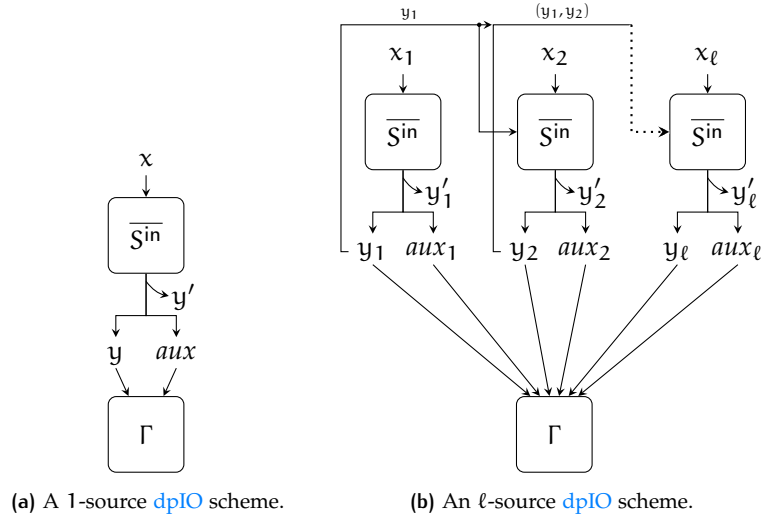


Figure 5.1: Overview of our formalization for 1-source **dpIO** and ℓ -source **dpIO**. S^{in} denotes an input sampler, $\overline{S^{\text{in}}}$ denotes an encoded input sampler and Γ denotes the obfuscation of a circuit C with respect to the input sampler S^{in} . Figure 5.1a depicts a 1-source **dpIO** scheme. Hence, the input sampler is stateless. Figure 5.1b depicts an ℓ -source **dpIO** scheme. The state for the i -th sample of S^{in} may depend on all previous (authenticated) outputs y_1, \dots, y_{i-1} . Formally, the state function derives the state from y_1, \dots, y_{i-1} , but is omitted for simplicity.

DEFINITION 5.1 (Input sampler). An *input sampler* with input space L and state space \mathcal{T} is a probabilistic algorithm S^{in} such that $S^{\text{in}}(\text{state}, x)$ produces outputs of the form $(y; y')$ with $y \in \{0, 1\}^{n(|x|)}$ (for $\text{state} \in \mathcal{T}$ and $x \in L$). A class of input samplers \mathcal{S}^{in} is an ensemble of input samplers sharing the same input space and state space.

We will treat $S^{\text{in}}(\cdot, \cdot)$ as a Boolean circuit. Looking ahead, the first part y will be used as input to the probabilistic circuit to be obfuscated, and hence needs to be authenticated. Since the second part y' is not meant as input to the probabilistic circuit, y' will not be authenticated. A state function stf efficiently maps $(\ell - 1)$ -tuples from $(\{0, 1\}^{n(\lambda)} \cup \{\perp\})^{\ell-1}$ to states from \mathcal{T} .

We are now prepared to define our novel notion of probabilistic indistinguishability obfuscation. Since there are two probabilistic processes involved (the input sampling and the evaluation of the (obfuscated) probabilistic circuit), we call this notion doubly-probabilistic indistinguishability obfuscation.

DEFINITION 5.2 (Doubly-probabilistic indistinguishability obfuscation). Let ℓ be an integer. Let $\text{stf}: (\{0, 1\}^{n(\lambda)} \cup \{\perp\})^{\ell-1} \rightarrow \mathcal{T}$ be an efficiently computable state function. Let \mathcal{S}^{in} be a class of input samplers with input space L and state space \mathcal{T} . Let $\mathcal{C} = \{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$ be a family of (probabilistic) circuits expecting ℓ inputs of length $n(\lambda)$, and let \mathcal{C} be a class of circuit samplers over \mathcal{C} . An ℓ -source doubly-probabilistic indistinguishability obfuscation (**dpIO**) scheme for $(\text{stf}, \mathcal{S}^{\text{in}}, \mathcal{C}, \mathcal{C})$ is a triple of **PPT** algorithms $\text{dpiO} = (\text{Setup}, \text{Encode}, \text{Obfuscate})$ such that

- $\text{Setup}(1^\lambda)$ outputs public parameters pp ,

- $\text{Encode}(pp, S^{\text{in}})$, on input of pp and an input sampler $S^{\text{in}} \in \mathcal{S}^{\text{in}}$, the deterministic algorithm Encode outputs an encoded input sampler \overline{S}^{in} ,
- $\text{Obfuscate}(pp, S^{\text{in}}, C)$, on input of pp , an input sampler $S^{\text{in}} \in \mathcal{S}^{\text{in}}$ and a circuit $C \in \mathcal{C}_{\ell, \lambda}$, Obfuscate outputs a circuit Γ of size $\text{poly}(\lambda, |C|)$. We call Γ an *obfuscation* of C with respect to S^{in} .

A **dpIO** scheme is required to satisfy *simulatability of encodings* (Definition 5.3), *restricted correctness* (Definition 5.4) and *security with respect to the class of circuit samplers* \mathcal{C} (Definition 5.5).

Given the definition above, some remarks are in order. A **dpIO** scheme is parametrized not only by a circuit class \mathcal{C} and a circuit sampler class \mathcal{C} , but also by an input sampler class \mathcal{S}^{in} and a state function stf . Circuits are always obfuscated with respect to some input sampler S^{in} . As a consequence, the obfuscated circuit expects inputs from the corresponding encoded input sampler \overline{S}^{in} . Furthermore, the obfuscation procedure takes circuits $C \in \mathcal{C}_{\lambda}$ as input. That is, C expects ℓ inputs of length $n(\lambda)$ and $|C| = \lambda$.

Informally, simulatability of encodings ensures that, on any (adversarially chosen) input x , state $state$, and input sampler S^{in} , the encoded sampler \overline{S}^{in} , on input of $(state, x)$, outputs samples of the form $(y, aux; y')$, where $(y; y')$ is distributed as outputs from $S^{\text{in}}(state, x)$, and aux does not leak any nontrivial information about the inputs. This is formalized by requiring the existence of a simulator that can simulate aux given only y .

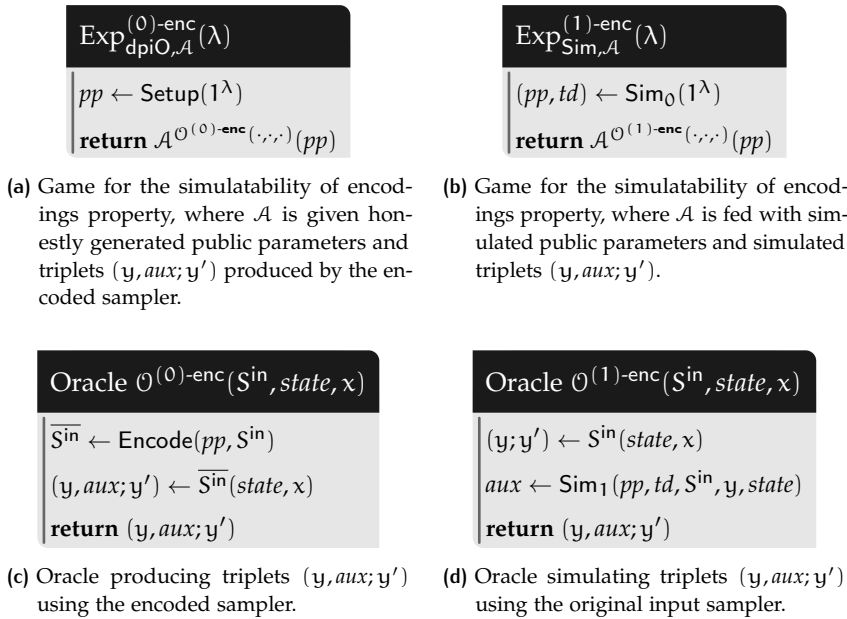


Figure 5.2: Definitions of the security games for the simulatability of encodings property of **dpIO** schemes. The **PPT** algorithm \mathcal{A} can interact polynomially many times with either $\mathcal{O}^{(0)\text{-enc}}$ (Figure 5.2c) or $\mathcal{O}^{(1)\text{-enc}}$ (Figure 5.2d). Note that \mathcal{A} provides the input sampler S^{in} as a Boolean circuit.

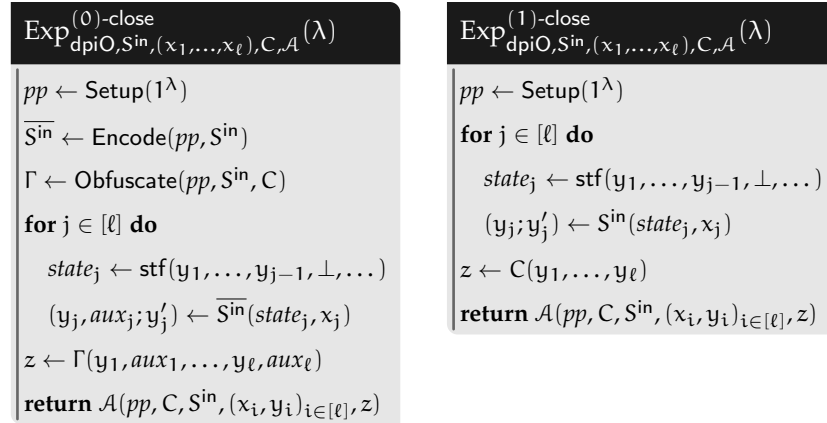
DEFINITION 5.3 (Simulatability of encodings). An ℓ -source **dpIO** scheme dpiO for $(\text{stf}, \mathcal{S}^{\text{in}}, \mathcal{C}, \mathcal{C})$ satisfies *simulatability of encodings* if for every **PPT** adversary \mathcal{A} , there exists a **PPT** simulator $\text{Sim} = (\text{Sim}_0, \text{Sim}_1)$ such that

$$\text{Adv}_{\text{dpiO}, \text{Sim}, \mathcal{A}}^{\text{enc}}(\lambda) := \left| \Pr \left[\text{Exp}_{\text{dpiO}, \mathcal{A}}^{(0)\text{-enc}}(\lambda) = 1 \right] - \Pr \left[\text{Exp}_{\text{Sim}, \mathcal{A}}^{(1)\text{-enc}}(\lambda) = 1 \right] \right|$$

is negligible, where the games $\text{Exp}_{\text{dpiO}, \mathcal{A}}^{(0)\text{-enc}}(\lambda)$ and $\text{Exp}_{\text{Sim}, \mathcal{A}}^{(1)\text{-enc}}(\lambda)$ are defined in Figure 5.2.

Intuitively, restricted correctness ensures that in an “honest scenario”, i. e., when the inputs (y_1, \dots, y_ℓ) are sampled using the input sampler S^{in} , then the obfuscated circuit is guaranteed to behave “correctly”. More precisely, restricted correctness guarantees that the support of the original circuit is respected, that the outputs of the obfuscated (deterministic) circuit are computationally indistinguishable from once-per-input evaluations of the original (probabilistic) circuit, and that the output distribution of the obfuscated circuit (taken over the random coins of the obfuscator) is one-time statistically close to the output distribution of the original circuit (taken over its internal random coins). Note that the statistical indistinguishability does not extend to multiple evaluations.

To make this definition meaningful, we need a way to let the obfuscated circuit verify that the inputs are well-formed. We do not require that the inputs were generated through S^{in} with uniformly random coins, but only that they were generated through S^{in} with *some* random coins (and some input). To enable verification, we let the parties generate their input using the encoded sampler $\overline{S^{\text{in}}}$ instead. An encoded sampler produces inputs as S^{in} , and in addition produces auxiliary information which can be used by the obfuscated program to verify that the inputs were honestly constructed. More formally, for a given input y , the obfuscated circuit verifies that there exists an input x , random coins r , and an unauthenticated part y' such that $(y; y') = S^{\text{in}}(x; r)$.



(a) Closeness game, where \mathcal{A} receives a single output from the obfuscated circuit.

(b) Closeness game, where \mathcal{A} receives a single output from the original randomized circuit.

Figure 5.3: Games $\text{Exp}_{\text{dpiO}, S^{\text{in}}, (x_1, \dots, x_\ell), C, \mathcal{A}}^{(b)\text{-close}}(\lambda)$ (for $b \in \{0, 1\}$) modeling the statistical closeness property of an ℓ -source dpiO scheme.

A small technicality is that we must allow the input sampler to depend on state information, to capture interdependencies between the inputs. This means that the auxiliary information will have to certify that an input was generated correctly, with respect to some state that the obfuscated circuit might not have access to (which would prevent it from verifying the certificate). However, this issue disappears by restricting the interdependencies to only involve a state computed from the *previous samples* (as opposed to more

complex interdependencies which would involve, for example, the coins used to produce these samples). See also Figure 5.1b for an overview. In this case, the obfuscated circuit can test the certificates in an incremental way. The first test verifies that y_1 was correctly constructed with respect to the state $\text{stf}(\perp, \dots, \perp)$, the second test verifies that y_2 was correctly constructed with respect to the state $\text{stf}(y_1, \perp, \dots, \perp)$, and the i -th test verifies that y_i was correctly constructed with respect to the state $\text{stf}(y_1, \dots, y_{i-1}, \perp, \dots, \perp)$.

DEFINITION 5.4 (Restricted correctness). An ℓ -source **dpIO** scheme dpiO for $(\text{stf}, \mathcal{S}^{\text{in}}, \mathcal{C}, \mathcal{C})$ satisfies *restricted correctness* if the following properties are satisfied.

STATISTICAL CLOSENESS. One-time evaluation of the encoded input sampler and the obfuscated circuit (on these inputs) is statistically close to one-time evaluation of the original input sampler and the original circuit (on these inputs). That is, for every $S^{\text{in}} \in \mathcal{S}^{\text{in}}$, $(x_1, \dots, x_\ell) \in \mathbb{L}^\ell$, and $C \in \mathcal{C}_\lambda$, for every (possibly unbounded) adversary \mathcal{A} , the advantage

$$\begin{aligned} \text{Adv}_{\text{dpiO}, S^{\text{in}}, (x_1, \dots, x_\ell), C, \mathcal{A}}^{\text{close}}(\lambda) := & \\ & \left| \Pr \left[\text{Exp}_{\text{dpiO}, S^{\text{in}}, (x_1, \dots, x_\ell), C, \mathcal{A}}^{(0)\text{-close}}(\lambda) = 1 \right] \right. \\ & \left. - \Pr \left[\text{Exp}_{\text{dpiO}, S^{\text{in}}, (x_1, \dots, x_\ell), C, \mathcal{A}}^{(1)\text{-close}}(\lambda) = 1 \right] \right| \end{aligned}$$

is negligible, where $\text{Exp}_{\text{dpiO}, S^{\text{in}}, (x_1, \dots, x_\ell), C, \mathcal{A}}^{(b)\text{-close}}(\lambda)$ (for $b \in \{0, 1\}$) is defined in Figure 5.3.

COMPUTATIONAL CORRECTNESS. Evaluations of an obfuscated program are computationally indistinguishable from evaluations of the original randomized program on well-formed inputs. That is, for every $S^{\text{in}} \in \mathcal{S}^{\text{in}}$ and $C \in \mathcal{C}_\lambda$, for every **PPT** adversary \mathcal{A} , the advantage

$$\begin{aligned} \text{Adv}_{\text{dpiO}, C, \mathcal{A}}^{\text{rcorr}}(\lambda) := & \left| \Pr \left[\text{Exp}_{\text{dpiO}, C, \mathcal{A}}^{(0)\text{-rcorr}}(\lambda) = 1 \right] \right. \\ & \left. - \Pr \left[\text{Exp}_{\text{dpiO}, C, \mathcal{A}}^{(1)\text{-rcorr}}(\lambda) = 1 \right] \right| \end{aligned}$$

is negligible, where $\text{Exp}_{\text{dpiO}, C, \mathcal{A}}^{(b)\text{-rcorr}}(\lambda)$ (for $b \in \{0, 1\}$) is defined in Figure 5.4 and the oracles $\mathcal{O}^{(0)\text{-rcorr}}$ and $\mathcal{O}^{(1)\text{-rcorr}}$ must not be called twice on the same input.

SUPPORT RESPECTING. The encoded sampler $\overline{S^{\text{in}}}$ and the obfuscated circuit Γ respect the support of the original sampler S^{in} and the original circuit C , respectively. That is, for all $pp \in \text{supp}(\text{Setup}(1^\lambda))$, $\overline{S^{\text{in}}} := \text{Encode}(pp, S^{\text{in}})$ and all $\Gamma \in \text{supp}(\text{Obfuscate}(pp, S^{\text{in}}, C))$, we have that

- for all $x \in \mathbb{L}$, $\text{state} \in \mathcal{T}$, $(y, \text{aux}; \cdot) \in \text{supp}(\overline{S^{\text{in}}}(\text{state}, x))$, we have that $(y; \cdot) \in \text{supp}(S^{\text{in}}(\text{state}, x))$, and
- for all inputs $x_1, \dots, x_\ell \in \mathbb{L}$, all $j \in [\ell]$, all inputs $(y_j, \text{aux}_j; \cdot) \in \text{supp}(\overline{S^{\text{in}}}(\text{state}_j, x_j))$, where $\text{state}_j := \text{stf}(y_1, \dots, y_{j-1}, \perp, \dots, \perp)$, we have that

$$\Gamma(y_1, \text{aux}_1, \dots, y_\ell, \text{aux}_\ell) \in \text{supp}(C(y_1, \dots, y_\ell)).$$

Note that closeness implies that samples produced with the original input sampler S^{in} and samples produced with the encoded input sampler $\overline{S^{\text{in}}}$ are statistically close.

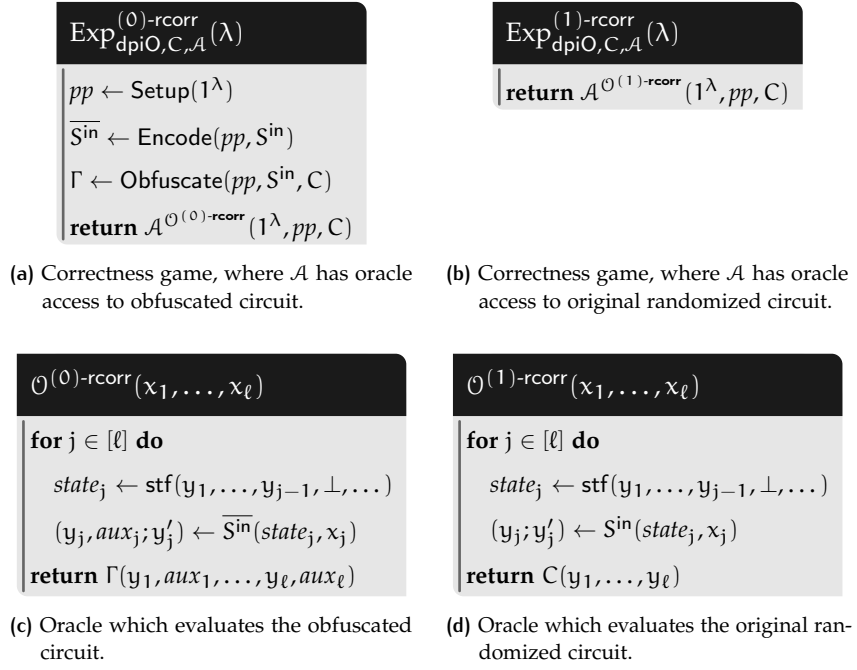


Figure 5.4: Experiments $\text{Exp}_{\text{dpiO}, \mathcal{C}, \mathcal{A}}^{(0)\text{-rcorr}}(\lambda)$ and $\text{Exp}_{\text{dpiO}, \mathcal{C}, \mathcal{A}}^{(1)\text{-rcorr}}(\lambda)$ modeling the computational correctness property an ℓ -source **dpiO** scheme.

Our notion of security is close in spirit to the standard indistinguishability notion for obfuscation of probabilistic circuits of [CLTV15] (see Definition 2.16). However, in our scenario, the security notion must account for the fact that public parameters pp are generated in a setup phase; the indistinguishability property of obfuscated circuits must therefore hold when many circuits are obfuscated with respect to the *same* public parameters. This suggests an oracle-based security notion.

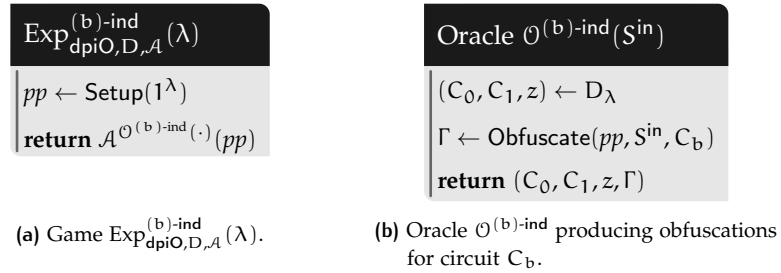


Figure 5.5: Definition of the games $\text{Exp}_{\text{dpiO}, \mathcal{D}, \mathcal{A}}^{(b)\text{-ind}}(\lambda)$ (for $b \in \{0, 1\}$) modeling the indistinguishability property of a **dpiO** scheme. The adversary \mathcal{A} can interact polynomially many times with the oracle $\mathcal{O}^{(b)\text{-ind}}$. Given a description of an input sampler S^{in} , the oracle $\mathcal{O}^{(b)\text{-ind}}$ samples (C_0, C_1, z) from the circuit sampler and obfuscates C_b with respect to S^{in} .

DEFINITION 5.5 (Security with respect to \mathcal{C}). An ℓ -source **dpiO** scheme dpiO for $(\text{stf}, S^{\text{in}}, \mathcal{C}, \mathcal{C})$ satisfies *security with respect to* \mathcal{C} if for every circuit sampler $\mathcal{D} = \{\mathcal{D}_\lambda\}_{\lambda \in \mathbb{N}} \in \mathcal{C}$, for every **PPT** adversary \mathcal{A} , the advantage

$$\text{Adv}_{\text{dpiO}, \mathcal{D}, \mathcal{A}}^{\text{ind}}(\lambda) := \left| \Pr \left[\text{Exp}_{\text{dpiO}, \mathcal{D}, \mathcal{A}}^{(0)\text{-ind}}(\lambda) = 1 \right] - \Pr \left[\text{Exp}_{\text{dpiO}, \mathcal{D}, \mathcal{A}}^{(1)\text{-ind}}(\lambda) = 1 \right] \right|$$

is negligible, where the games $\text{Exp}_{\text{dpiO},D,\mathcal{A}}^{(0)\text{-ind}}(\lambda)$ and $\text{Exp}_{\text{dpiO},D,\mathcal{A}}^{(1)\text{-ind}}(\lambda)$ are defined in Figure 5.5.

6

CONSTRUCTION

Before describing our construction, we recall some notation for this chapter. Let $\ell \in \mathbb{N}$ be some constant. Let $\mathcal{C} = (\mathcal{C}_\lambda)_{\lambda \in \mathbb{N}}$ be a circuit class such that \mathcal{C}_λ contains (arbitrary) circuits of size λ . Recall that $\mathcal{C} = \{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$ is a class of (probabilistic) circuits expecting ℓ inputs of length $n(\lambda)$, such that \mathcal{C}_λ only contains circuits of size λ .

Let \mathcal{S}^{in} be a class of (arbitrary) input samplers with *polynomially sized input space* L , that is $|L \cap \{0, 1\}^\lambda| = \text{poly}(\lambda)$ for some polynomial $\text{poly}(\lambda)$. Note that such input samplers can produce exponentially many y and, in fact, these are the cases of interest for our applications.

In this chapter, we construct an ℓ -source **dpIO** scheme, for the input sampler class \mathcal{S}^{in} , and dynamic-input indistinguishable circuit samplers for \mathcal{C} . Our construction relies on polynomially secure indistinguishability obfuscation (Definition 2.15), polynomially secure (perfect) puncturable pseudo-random functions (Definition 4.1), statistically sound non-interactive zero-knowledge proof systems (Definition 2.25), and an extremely lossy function (Definition 4.2).

6.1 OVERVIEW

We start by providing a high-level overview of our construction. The setup algorithm **Setup** generates parameters G for the extremely lossy function (in injective mode) and a common reference string σ for the non-interactive zero-knowledge proof system.

Given an input sampler $S^{\text{in}} \in \mathcal{S}^{\text{in}}$, the encode algorithm **Encode** compiles S^{in} to an encoded input sampler \overline{S}^{in} as follows. On input of $(state, x; r)$, \overline{S}^{in} samples $(y; y')$ using the input sampler $S^{\text{in}}(state, x; G(r))$, certifies y with a **NIZK** proof $\pi = aux$ produced via $\text{NIZK.Prove}(\sigma, y, L_{state}^{G, S^{\text{in}}}, (y', x, r))$, and outputs $(y, aux; y')$. The language $L_{state}^{G, S^{\text{in}}}$ contains all values y for which there exists (y', x, r) such that $(y; y') = S^{\text{in}}(state, x, G(r))$. We refer to values $y \in L_{state}^{G, S^{\text{in}}}$ as well-formed inputs. Note that if G is in injective mode, $L_{state}^{G, S^{\text{in}}}$ is, in general, a trivial language, i. e., the union over all $x \in L$ of the first component of $\text{supp}(S^{\text{in}}(state, x))$. Informally speaking, simulatability of the encodings directly follows from the injectivity of G , and the zero-knowledge property of the proof system.

For the sake of simplicity, we describe the intuition behind our construction for a 1-source **dpIO** scheme. The obfuscator **Obfuscate**, on input of an input sampler S^{in} and a circuit C , obfuscates C (with respect to S^{in}) as follows. First, **Obfuscate** samples a **PRF** key K for F . Then, it constructs a Boolean circuit \overline{C} which, on input of (y, aux) , verifies the certificate aux with NIZK.Verify . If the test passes, \overline{C} evaluates the circuit C on input of y and random coins $F(K, y)$. If the test fails, \overline{C} outputs an error symbol \perp . Informally, restricted correctness follows from the correctness properties of all used components.

Recall that our main concern in this part is to avoid a subexponential reduction to indistinguishability obfuscation. Hence, the tricky part is to achieve security with respect to a preferably large class \mathcal{C} of circuit samplers. Interestingly, we are able to prove our **dpIO** scheme secure with respect to the class of dynamic-input indistinguishable samplers $\mathcal{C}^{\text{dyn-ind}}$ over the class of all polynomial sized circuits. This even exceeds the capabilities of the **pIO** construction due to [CLTV15], since general **pIO** for the class of dynamic-input indistinguishable samplers falls victim to the implausibility result due to [GGHW17]. We avoid this implausibility by restricting the correctness of obfuscated programs to only hold for “certified inputs”.

Therefore, to prove indistinguishability between the obfuscations of two dynamic-input indistinguishable circuits (C_0, C_1) , we resort to the standard puncturing strategy of [CLTV15]. That is, we proceed through a sequence of hybrids, successively modifying of the obfuscated circuit, an input at a time. More precisely, starting with an obfuscation of C_0 , we program the obfuscated circuit to behave like $C_1(y; r)$ instead of $C_0(y; r)$, for each input y . Using this naive approach directly, each such replacement relies on the security of the underlying **IO** scheme (and the security of the **PRF**, and the dynamic-input indistinguishability of C_0 and C_1).

The main issue of this approach is that the number of possible inputs y (hence the number of hybrids) is exponential – indeed, this is the reason why the **pIO** scheme of [CLTV15] requires subexponentially secure primitives (**iO** and **F**). To get around this issue, we leverage the origin of inputs to circuits in conjunction with the power of the extremely lossy function G . We first switch G to an appropriate extremely lossy mode, which by definition cannot be distinguished from the injective mode. Now, the soundness of the **NIZK** proof system ensures that all valid inputs y originate from $S^{\text{in}}(\text{state}, x; G(r))$ for some (x, r) (omitting y' for simplicity). For a given *state*, the quantity of such values is bounded by the size of the range of G – which is merely polynomial – times the size of the input domain (of the input sampler) $L \cap \{0, 1\}^\lambda$. Therefore, in all applications where the inputs to the obfuscated circuit are sampled using sampler inputs from a small domain, we can base security on *polynomially secure* **iO**. Note that for $\ell > 1$, the state input *state* for the input sampler S^{in} depends in a deterministic manner (the state function) from previous samples, which exist only in polynomial quantity. Therefore, the above approach generalizes to ℓ -source **dpIO**.

6.2 CONSTRUCTING DOUBLY-PROBABILISTIC IO

For our construction, we employ a perfectly sound **NIZK** proof system for the following (parametrized) language

$$L_{\text{state}}^{G, S^{\text{in}}} := \left\{ y \mid \begin{array}{l} \exists y' \in \{0, 1\}^* \\ \exists x \in L \cap \{0, 1\}^\lambda : (y; y') = S^{\text{in}}(\text{state}, x; G(r)) \\ \exists r \in \{0, 1\}^{k(\lambda)} \end{array} \right\}.$$

Note that $L_{\text{state}}^{G, S^{\text{in}}}$ is well-defined for arbitrary values of *state*, i. e., even for state information which cannot occur during honest use. Looking ahead,

validity of the state information is ensured independently by the obfuscated circuit.

For simplicity in the analysis, we use a **NIZK** proof system that satisfies the following property. With overwhelming probability over the coins of $\text{Setup}(1^\lambda)$, there does not exist any pair (x, π) such that $x \notin L$ and $\text{NIZK.Verify}(\sigma, x, \pi) = 1$. We call a **NIZK** that satisfies this property *almost perfectly sound*. We note that there is a simple folklore method which allows to construct an almost perfectly sound **NIZK** proof system starting from any statistically sound **NIZK** proof system. Consider a $2^{-\lambda}$ -statistically sound **NIZK** proof system, for statements $x \in \{0, 1\}^n$, for some polynomial $n = n(\lambda)$. Using parallel repetitions, the soundness of the proof system can be amplified to $2^{-\lambda-n}$. That is, for any statement $x \notin L$, the probability

$$\Pr_{\sigma} [\exists \pi: \text{Verify}(\sigma, x, \pi) = 1] \leq 2^{-\lambda-n}.$$

Then, it necessarily holds that for all possible σ except a $2^{-\lambda}$ fraction of them, there does not exist any pair (x, π) such that $x \notin L$ and π is an accepting proof. To realize this, let E_x^{σ} denote the event that there exists a proof π such that $\text{Verify}(\sigma, x, \pi) = 1$. Then, by a union bound argument,

$$\Pr_{\sigma} [\exists x \in \{0, 1\}^n \setminus L: E_x^{\sigma}] \leq \sum_{x \in \{0, 1\}^n \setminus L} \Pr_{\sigma} [E_x^{\sigma}] \leq 2^n \cdot 2^{-\lambda-n}.$$

Hence, the **NIZK** proof system obtained via parallel repetitions is almost perfectly sound.

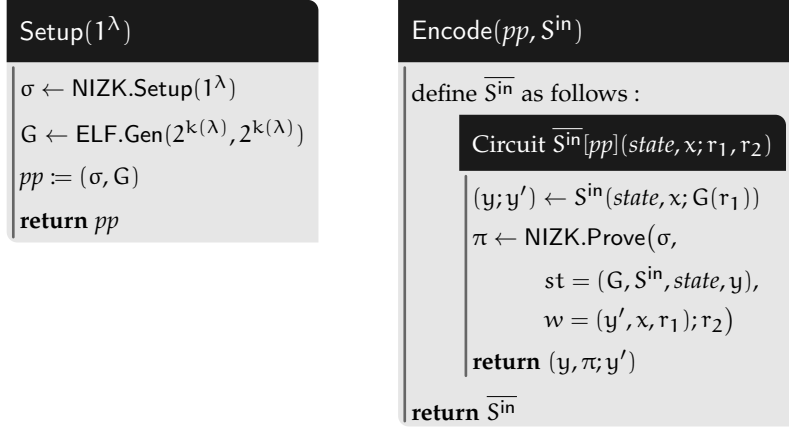
Let $\ell \in \mathbb{N}$ be a constant, let $\text{stf}: (\{0, 1\}^{n(\lambda)} \cup \{\perp\})^{\ell-1} \rightarrow \mathcal{T}$ be an efficiently computable state function, and let $\mathcal{C} = \{\mathcal{C}_{\lambda}\}_{\lambda}$ be a family of (randomized) circuits taking ℓ inputs of size $n(\lambda)$ such that \mathcal{C}_{λ} only contains circuits of size λ . Let \mathcal{S}^{in} be a class of (arbitrary) input samplers with randomness space $\{0, 1\}^{k(\lambda)}$ polynomially sized input space L , that is $|L \cap \{0, 1\}^{\lambda}| = \text{poly}(\lambda)$ for some polynomial $\text{poly}(\lambda)$. Further, let $\mathcal{C}^{\text{dyn-ind}}$ be the class of dynamic-input indistinguishable samplers (over \mathcal{C}).

THEOREM 6.1. *Let ELF be a strongly regular extremely lossy function, let iO be a perfectly correct polynomially secure indistinguishability obfuscator, let F be a polynomially secure (perfect) puncturable **PRF**, and let **NIZK** be a statistically sound and computationally zero-knowledge **NIZK** proof system for the family of languages $\{\Gamma_{\text{state}, G, \mathcal{S}^{\text{in}}}\}_{\text{state}, G, \mathcal{S}^{\text{in}}}$. Then, the **dpIO** scheme $\text{dpiO} = (\text{Setup}, \text{Encode}, \text{Obfuscate})$ defined in Figure 6.1 is an ℓ -source **dpIO** scheme for $(\text{stf}, \mathcal{S}^{\text{in}}, \mathcal{C}, \mathcal{C}^{\text{dyn-ind}})$.*

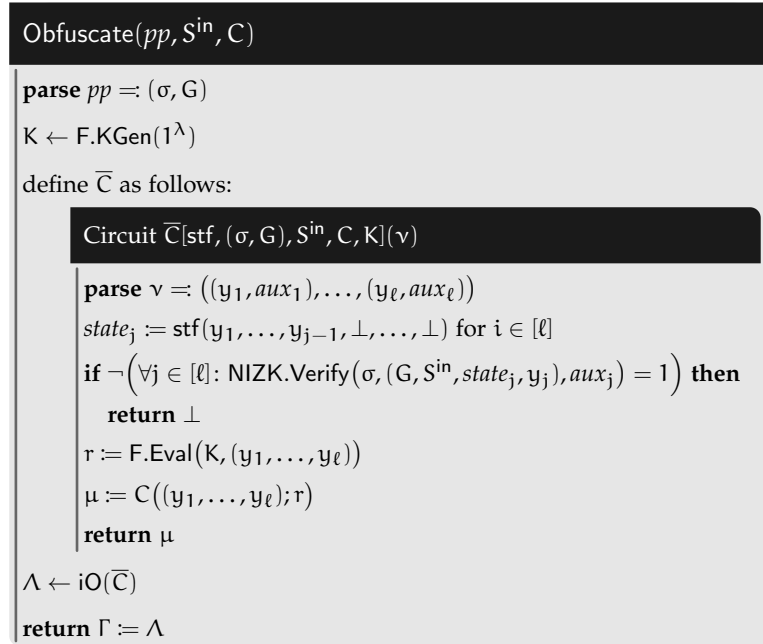
Statistically sound perfectly correct **NIZK** proof systems for \mathcal{NP} can be constructed from polynomially secure indistinguishability obfuscation and extremely lossy functions [BP15a]. Further, **ELFs** imply the existence of one-way functions, hence of (perfect) puncturable **PRFs** [GGM84a; HILL99]. Together with Theorem 6.1, we obtain the following corollary.

COROLLARY 6.1. *Assuming polynomially secure indistinguishability obfuscation and extremely lossy functions, for every constant ℓ , for every class of input samplers \mathcal{S}^{in} with a polynomial sized input domain, for every state function stf and for every probabilistic circuit class \mathcal{C} , there exists an ℓ -source doubly-probabilistic indistinguishability obfuscation scheme for $(\text{stf}, \mathcal{S}^{\text{in}}, \mathcal{C}, \mathcal{C}^{\text{dyn-ind}})$, where $\mathcal{C}^{\text{dyn-ind}}$ is the class of dynamic-input circuit samplers for \mathcal{C} .*

Proof of Theorem 6.1. We prove that dpiO as defined in Figure 6.1 satisfies simulatability of encodings (cf. Definition 5.3), statistical restricted correctness (cf. Definition 5.4), and indistinguishability (cf. Definition 5.5).



(a) Setup algorithm producing public parameters.

(b) Encoding algorithm which compiles input samplers S^{in} to so-called encoded input samplers $\overline{S^{\text{in}}}$.(c) Obfuscate algorithm producing obfuscations Γ of circuits C with respect to an input sampler S^{in} .Figure 6.1: Our construction of an ℓ -source **dpIO** scheme $\text{dpiO} = (\text{Setup}, \text{Encode}, \text{Obfuscate})$.

SIMULATABILITY OF ENCODINGS. We prove that there exists a **PPT** simulator $\text{Sim} = (\text{Sim}_0, \text{Sim}_1)$ such that for every **PPT** adversary \mathcal{A} , the advantage $\text{Adv}_{\text{dpiO}, \text{Sim}, \mathcal{A}}^{\text{enc}}(\lambda)$ is negligible. By the zero-knowledge property of **NIZK**, there exists a simulator $\text{NIZK.Sim} = (\text{NIZK.Sim}_0, \text{NIZK.Sim}_1)$ such that the advantage $\text{Adv}_{\text{NIZK}, \text{NIZK.Sim}, \mathcal{B}}^{\text{zk}}(\lambda)$ is negligible for all **PPT** adversaries \mathcal{B} . In Figure 6.2, we construct a simulator $\text{Sim} = (\text{Sim}_0, \text{Sim}_1)$.

Let \mathcal{A} be a **PPT** distinguisher between $\text{Exp}_{\text{dpiO}, \mathcal{A}}^{(0)\text{-enc}}(\lambda)$ and $\text{Exp}_{\text{Sim}, \mathcal{A}}^{(1)\text{-enc}}(\lambda)$ (making at most Q oracle queries for some polynomial Q). We prove indistinguishability between the real and the ideal distribution via a series of hybrids starting from the ideal distribution $\text{Exp}_{\text{Sim}, \mathcal{A}}^{(1)\text{-enc}}(\lambda)$.

GAME G_0 . This game is identical to $\text{Exp}_{\text{Sim}, \mathcal{A}}^{(1)\text{-enc}}(\lambda)$. We remark that in this game, the tuple $(y; y')$ is produced using the adversarially chosen sampler S^{in} on input of the adversarially chosen state $state$ and input x , where S^{in} is supplied with uniform randomness from $\{0, 1\}^{k(\lambda)}$.

GAME G_1 . This game is identical to G_0 except for the fact that for each query $(S^{\text{in}}, state, x)$, the sampler S^{in} is supplied with randomness $G(r)$ for uniform r (instead of true randomness). Due to the strong regularity of G and by a standard hybrid argument over all queries, the statistical distance between G_0 and G_1 is negligible. More precisely, $|\Pr[out_1 = 1] - \Pr[out_0 = 1]| \leq Q \cdot \epsilon_{\text{reg}}(\lambda)$, where Q is a polynomial upper bound for the number of oracle queries of \mathcal{A} .

$\text{Sim}_0(1^\lambda)$	$\text{Sim}_1(pp, td, S^{\text{in}}, y, state)$
$(\sigma, \tau_{\Pi}) \leftarrow \text{NIZK.Sim}_0(1^\lambda)$ $G \leftarrow \text{ELF.Gen}(2^{k(\lambda)}, 2^{k(\lambda)})$ $pp := (\sigma, G)$ $td := \tau_{\Pi}$ return (pp, td)	parse $pp = (\sigma, G)$ parse $td = \tau_{\Pi}$ $\pi \leftarrow \text{NIZK.Sim}_1(\sigma, \tau_{\Pi}, (G, S^{\text{in}}, state, y))$ $aux := \pi$ return aux

(a) Definition of the first phase of the simulator Sim_0 . Sim_0 generates a simulated CRS together with a corresponding trapdoor and sets up the parameters for the ELF in injective mode.

(b) Definition of the second phase of the simulator Sim_1 . Sim_1 uses the NIZK trapdoor τ_{Π} to simulate a proof π for y .

Figure 6.2: Definition of the simulator $\text{Sim} = (\text{Sim}_0, \text{Sim}_1)$ in Figures 6.2a and 6.2b, respectively, for the proof of the simulatability of encodings property of dpiO .

GAME G_2 . This game is the same as G_1 with the difference that σ is produced honestly using $\text{NIZK.Setup}(1^\lambda)$. Additionally, for each adversarial query $(S^{\text{in}}, state, x)$, the proof π is produced honestly by $\text{NIZK.Prove}(\sigma, (G, S^{\text{in}}, state, y), (y', x, r))$, where $G(r)$ are the random coins supplied to the sampler S^{in} . The view of \mathcal{A} in game G_2 is distributed exactly as in the real game $\text{Exp}_{\text{dpiO}, \mathcal{A}}^{(0)\text{-enc}}(\lambda)$.

We construct a PPT adversary \mathcal{B}_2 for the zero-knowledge property of NIZK. Given a common reference string σ , \mathcal{B}_2 samples an ELF G in injective mode via $\text{ELF.Gen}(2^{k(\lambda)}, 2^{k(\lambda)})$ and invokes \mathcal{A} on input of $pp := (\sigma, G)$. Each time \mathcal{A} queries his oracle on $(S^{\text{in}}, state, x)$, \mathcal{B}_2 draws random coins $r \in \{0, 1\}^{k(\lambda)}$ and invokes the sampler S^{in} on input of $(state, x)$ with random coins $G(r)$ to obtain $(y; y')$. In order to produce π , \mathcal{B}_2 calls his prove oracle on input of $(G, S^{\text{in}}, state, y)$ with witness (y', x, r) . Therefore, if \mathcal{B}_2 is supplied with an honest CRS and honestly generated proofs, \mathcal{B}_2 perfectly simulates G_2 for \mathcal{A} . Otherwise, if \mathcal{B}_2 is supplied with a simulated CRS and simulated proofs, \mathcal{B}_2 perfectly simulates G_1 . Hence, $|\Pr[out_2 = 1] - \Pr[out_1 = 1]| \leq \text{Adv}_{\text{NIZK}, \text{NIZK.Sim}, \mathcal{B}_2}^{\text{zk}}(\lambda)$.

Thus, we have

$$\begin{aligned} \text{Adv}_{\text{dpiO}, \text{Sim}, \mathcal{A}}^{(0)\text{-enc}}(\lambda) &= |\Pr[out_2 = 1] - \Pr[out_0 = 1]| \\ &\leq Q \cdot \epsilon_{\text{reg}}(\lambda) + \text{Adv}_{\text{NIZK}, \text{NIZK.Sim}, \mathcal{B}}^{\text{zk}}(\lambda) \end{aligned}$$

concluding the proof since both terms are negligible in λ .

RESTRICTED CORRECTNESS. Let $S^{\text{in}} \in \mathcal{S}^{\text{in}}$ be an input sampler, let x_1, \dots, x_ℓ be arbitrary inputs from $L \cap \{0, 1\}^\lambda$, and let C be a circuit from the family \mathcal{C}_λ . To prove that dpiO satisfies Definition 5.4, we prove that dpiO satisfies statistical closeness, computational correctness and is support respecting.

Statistical closeness. We begin with statistical closeness and proceed over a series of hybrids for that purpose.

GAME G_0 . This game is the ideal game $\text{Exp}_{\text{dpiO}, S^{\text{in}}, (x_1, \dots, x_\ell), C, \mathcal{A}}^{(1)\text{-close}}(\lambda)$. Note that the sampler S^{in} is called using true randomness whereas the experiment $\text{Exp}_{\text{dpiO}, S^{\text{in}}, (x_1, \dots, x_\ell), C, \mathcal{A}}^{(0)\text{-close}}(\lambda)$ generates samples using $G(r)$, where r is truly random.

GAME G_1 . This game is identical to G_0 with the difference that each call of the sampler S^{in} is supplied with $G(r)$ as randomness (where r is sampled uniformly for each call). Due to the strong regularity of G , and by a hybrid argument over all ℓ calls of S^{in} , the statistical distance between G_1 and G_0 is negligible.

GAME G_2 . This game is the real game $\text{Exp}_{\text{dpiO}, S^{\text{in}}, (x_1, \dots, x_\ell), C, \mathcal{A}}^{(0)\text{-close}}(\lambda)$. We now argue that the view of \mathcal{A} in G_1 is distributed identically to his view in G_2 . G_2 samples public parameters pp via $\text{Setup}(1^\lambda)$ and $\overline{S^{\text{in}}}$ via $\overline{S^{\text{in}}} \leftarrow \text{Encode}(pp, S^{\text{in}})$. Further, (y_j, aux_j) are sampled as $state_j \leftarrow \text{stf}(y_1, \dots, y_{j-1}, \perp, \dots, \perp)$ and $(y_j, aux_j, y'_j) \leftarrow \overline{S^{\text{in}}}(state_j, x_j)$, for $j \in [\ell]$. Let Γ be the obfuscation $\Gamma \leftarrow \text{Obfuscate}(pp, S^{\text{in}}, C)$ of the circuit C with respect to sampler S^{in} . Due to the perfect correctness of iO , Γ computes the same functionality as $\overline{C}[\text{stf}(\sigma, G), S^{\text{in}}, C, K]$, where K is a freshly generated key for the PRF F . Hence, by the perfect completeness of NIZK, on input of $((y_1, aux_1), \dots, (y_\ell, aux_\ell))$, Γ evaluates the circuit C on input of (y_1, \dots, y_ℓ) with random coins $F(K, (y_1, \dots, y_\ell))$. Therefore, the view of \mathcal{A} in the games G_1 and G_2 only differs in the fact that G_1 supplies C with true random coins whereas G_2 supplies C with $F(K, (y_1, \dots, y_\ell))$ as randomness. As F is a perfect PRF (see Definition 4.1), the distribution $\{K \leftarrow F.\text{KGen}(1^\lambda) : F(K, (y_1, \dots, y_\ell))\}$ is identical to the uniform distribution over the image of F . Therefore, the view of \mathcal{A} in G_1 and G_2 is distributed identically.

Computational correctness. Computational correctness follows from the security of the PRF. Let C be a probabilistic circuit from \mathcal{C}_λ . Consider the following series of hybrids.

GAME G_0 . This game is identical to $\text{Exp}_{\text{dpiO}, C, \mathcal{A}}^{(1)\text{-rcorr}}(\lambda)$.

GAME G_1 . Game G_1 is identical to G_0 except that initially, a PRF key K is sampled and the random coins for circuit C are computed as $F(K, (y_1, \dots, y_\ell))$. Due to the security of F , we have

$$|\Pr[out_1 = 1] - \Pr[out_0 = 1]| \leq \text{Adv}_{F, \mathcal{A}}^{\text{prf}}(\lambda).$$

GAME G_2 . G_2 is identical to $\text{Exp}_{\text{dpiO}, C, \mathcal{A}}^{(0)\text{-rcorr}}(\lambda)$. Due to perfect correctness of iO and perfect completeness of NIZK, we have that $\Pr[out_1 = 1] = \Pr[out_2 = 1]$.

Support respecting. Since an encoded input sampler \overline{S}^{in} produces y by evaluating S^{in} with random coins $G(r) \subseteq \{0, 1\}^{k(\lambda)}$, we have that for all $pp \in \text{supp}(\text{Setup}(1^\lambda))$, all $\overline{S}^{\text{in}} \leftarrow \text{Encode}(pp, S^{\text{in}})$, all $x \in L$, all $state \in \mathcal{T}$, and all $(y, aux; \cdot) \in \text{supp}(\overline{S}^{\text{in}}(state, x))$, $(y; \cdot) \in \text{supp}(S^{\text{in}}(state, x))$.

Furthermore, by construction, perfect completeness of NIZK and perfect correctness of iO , for all $\Gamma \leftarrow \text{Obfuscate}(pp, S^{\text{in}}, C)$, all $j \in [\ell]$, all inputs $(y_j, aux_j; \cdot) \in \text{supp}(\overline{S}^{\text{in}}(state_j, x_j))$, where $state_j := \text{stf}(y_1, \dots, y_{j-1}, \perp, \dots, \perp)$, we have that $\Gamma(y_1, aux_1, \dots, y_\ell, aux_\ell) \in \text{supp}(C(y_1, \dots, y_\ell))$.

HYBRID ARGUMENTS. In the following, we heavily use a standard technique called hybrid arguments. Computational indistinguishability of random variables can easily be seen to satisfy transitivity. That is, if the random variables X_1, X_2 and X_2, X_3 are pairwise computationally indistinguishable, then X_1 and X_3 are computationally indistinguishable. Basically, this follows from the triangle inequality. More precisely, for all PPT adversaries \mathcal{A} , $\Pr[\mathcal{A}(X_1) = 1] - \Pr[\mathcal{A}(X_2) = 1] \leq \epsilon_1(\lambda)$ and $\Pr[\mathcal{A}(X_2) = 1] - \Pr[\mathcal{A}(X_3) = 1] \leq \epsilon_2(\lambda)$ for some negligible functions $\epsilon_1(\lambda), \epsilon_2(\lambda)$. Then,

$$\begin{aligned} & \left| \Pr[\mathcal{A}(X_1) = 1] - \Pr[\mathcal{A}(X_3) = 1] \right| \\ &= \left| \Pr[\mathcal{A}(X_1) = 1] - \Pr[\mathcal{A}(X_2) = 1] + \Pr[\mathcal{A}(X_2) = 1] - \Pr[\mathcal{A}(X_3) = 1] \right| \\ &\leq \underbrace{\left| \Pr[\mathcal{A}(X_1) = 1] - \Pr[\mathcal{A}(X_2) = 1] \right|}_{\leq \epsilon_1(\lambda)} + \underbrace{\left| \Pr[\mathcal{A}(X_2) = 1] - \Pr[\mathcal{A}(X_3) = 1] \right|}_{\leq \epsilon_2(\lambda)}. \end{aligned}$$

However, this transitivity does not immediately generalize to polynomially many “steps”. This is because negligible functions can be somewhat contrived such that the sum of polynomially many negligible functions does not turn out to be negligible. For instance, the functions

$$\epsilon_i(\lambda) := \begin{cases} 1 & \text{if } \lambda \leq i \\ 0 & \text{otherwise} \end{cases}$$

are negligible (for every $i \in \mathbb{N}$), whereas the function $f(\lambda) := \sum_{i=1}^{\lambda} \epsilon_i(\lambda)$ is not negligible. In order to prove indistinguishability of polynomially many random variables (or security games), we use a standard proof technique called hybrid arguments. Basically, the above intuition to generalize the transitivity of indistinguishability to polynomially many steps works if each pairwise indistinguishability can be upper bounded by the *same* negligible function $\epsilon(\lambda)$. Then, intuitively, the triangle inequality yields

$$\begin{aligned} & \left| \Pr[\mathcal{A}(X_1) = 1] - \Pr[\mathcal{A}(X_n) = 1] \right| \\ &\leq \sum_{i=1}^{n-1} \underbrace{\left| \Pr[\mathcal{A}(X_i) = 1] - \Pr[\mathcal{A}(X_{i+1}) = 1] \right|}_{\leq \epsilon(\lambda)} \leq (n-1)\epsilon(\lambda). \end{aligned}$$

However, ensuring that all $\left| \Pr[\mathcal{A}(X_i) = 1] - \Pr[\mathcal{A}(X_{i+1}) = 1] \right|$ can be upper bounded by the same negligible function actually requires some extra work.

Suppose the indistinguishability between X_i and X_{i+1} can be based on the same decisional assumption such that a successful distinguisher between X_1 and X_n can be used to construct a single adversary \mathcal{B} on that assumption. Suppose this decisional assumption states that the distributions Y_0 and Y_1 are indistinguishable. The distinguisher \mathcal{B} , on input of a challenge y from

Y_0 or Y_1 , guesses an index $i \leftarrow [n-1]$, uses his y to sample x from the distribution X_i (if y is distributed according to Y_0) or from X_{i+1} (if y is distributed according to Y_1), and calls \mathcal{A} on x . The advantage of \mathcal{B} is

$$\begin{aligned}
& \Pr [\mathcal{B}(Y_0) = 1] - \Pr [\mathcal{B}(Y_1) = 1] \\
&= \sum_{j=1}^{n-1} \left(\Pr [\mathcal{B}(Y_0) = 1 \wedge i = j] - \Pr [\mathcal{B}(Y_1) = 1 \wedge i = j] \right) \\
&= \frac{1}{n-1} \cdot \sum_{j=1}^{n-1} \left(\Pr [\mathcal{B}(Y_0) = 1 \mid i = j] - \Pr [\mathcal{B}(Y_1) = 1 \mid i = j] \right) \\
&= \frac{1}{n-1} \cdot \sum_{j=1}^{n-1} \left(\Pr [\mathcal{A}(X_j) = 1] - \Pr [\mathcal{A}(X_{j+1}) = 1] \right) \\
&= \frac{1}{n-1} \cdot \left(\Pr [\mathcal{A}(X_1) = 1] - \Pr [\mathcal{A}(X_n) = 1] \right).
\end{aligned}$$

This formally proves that X_1 and X_n are indistinguishable.

In order to avoid notational overhead, we will not make each reduction explicit.

SECURITY. Let $D \in \mathcal{C}^{\text{dyn-ind}}$ be an arbitrary dynamic-input indistinguishable circuit sampler over \mathcal{C} . Let $p_{\mathcal{C}}$ be an upper bound on the size of the circuits \bar{C} , \bar{C}' and \bar{C}'' (including all hard-coded variables). We assume that all circuits to be obfuscated using iO are padded to size $p_{\mathcal{C}}(\lambda)$.

To prove that dpiO satisfies indistinguishability (Definition 5.5), we proceed over a series of hybrids. Toward contradiction, we assume that there exists an adversary \mathcal{A} of polynomial size s who distinguishes $\text{Exp}_{\text{dpiO}, D, \mathcal{A}}^{(0)\text{-ind}}(\lambda)$ from $\text{Exp}_{\text{dpiO}, D, \mathcal{A}}^{(1)\text{-ind}}(\lambda)$ with non-negligible advantage after making a polynomial number Q of queries to the oracle.

GAME G_0 . This game samples a random bit $b \leftarrow \{0, 1\}$, and simulates the experiment $\text{Exp}_{\text{dpiO}, D, \mathcal{A}}^{(b)\text{-ind}}(\lambda)$. More precisely, \mathcal{A} has access to the public parameters pp and an oracle $\mathcal{O}^{(b)\text{-ind}}(D)$, which on input of an input sampler S^{in} , draws a sample (C_0, C_1, z) from D and outputs (C_0, C_1, z) together with an obfuscation $\text{Obfuscate}(pp, S^{\text{in}}, C_b)$. \mathcal{A} outputs a guess b' and the game returns 1 if and only if $b' = b$. By assumption, the advantage of \mathcal{A} in distinguishing $\text{Exp}_{\text{dpiO}, D, \mathcal{A}}^{(0)\text{-ind}}(\lambda)$ from $\text{Exp}_{\text{dpiO}, D, \mathcal{A}}^{(1)\text{-ind}}(\lambda)$ is non-negligible. Hence,

$$\begin{aligned}
\epsilon(\lambda) &:= \left| \Pr [\text{out}_0 = 1] - \frac{1}{2} \right| \\
&= \frac{1}{2} \cdot \left| \Pr \left[\text{Exp}_{\text{dpiO}, D, \mathcal{A}}^{(0)\text{-ind}}(\lambda) = 1 \right] - \Pr \left[\text{Exp}_{\text{dpiO}, D, \mathcal{A}}^{(1)\text{-ind}}(\lambda) = 1 \right] \right|
\end{aligned}$$

is a non-negligible function in λ .

GAME G_1 . This game is identical to G_0 except for the fact that the **ELF** G is sampled in *extremely lossy mode*. More precisely, G is sampled via $\text{ELF.Gen}(2^{k(\lambda)}, t)$, where t is a polynomial such that every **PPT** algorithm of (circuit) size s has advantage at most $\epsilon/2$ in distinguishing $\text{ELF.Gen}(2^{k(\lambda)}, 2^{k(\lambda)})$ from $\text{ELF.Gen}(2^{k(\lambda)}, t)$. The advantage of \mathcal{A} in G_1 is therefore lower bounded by $\epsilon/2$, that is

$$\left| \Pr [\text{out}_1 = 1] - \frac{1}{2} \right| \geq \frac{\epsilon}{2}. \tag{6.1}$$

Otherwise, \mathcal{A} has advantage greater than $\epsilon/2$ in distinguishing $\text{ELF.Gen}(2^{k(\lambda)}, 2^{k(\lambda)})$ from $\text{ELF.Gen}(2^{k(\lambda)}, t)$.

GAME \mathbf{G}_2 . This game proceeds exactly as \mathbf{G}_1 , except that after sampling $b \leftarrow \{0, 1\}$, the game always sets up the experiment $\text{Exp}_{\text{dpiO}, D, \mathcal{A}}^{(1)\text{-ind}}(\lambda)$, but still returns 1 if and only if \mathcal{A} 's guess b' equals b . Clearly,

$$\Pr[out_2 = 1] = \frac{1}{2} \quad (6.2)$$

and the advantage of \mathcal{A} in \mathbf{G}_2 equals 0.

In the following, we prove that

$$|\Pr[out_2 = 1] - \Pr[out_1 = 1]| \leq Q \cdot \text{negl}(\lambda), \quad (6.3)$$

where negl is a negligible function and Q is a polynomial in λ . Equation (6.3) together with Equations (6.1) and (6.2) yields a contradiction to our assumption that ϵ is non-negligible as a function in λ .

In the following, S_q^{in} denotes the input sampler queried in the q -th oracle query. The circuits $C_{q,0}$ and $C_{q,1}$ denote the circuits sampled in the q -th oracle query using the circuit sampler D . Towards proving Equation (6.3), we define a sequence of hybrids as follows.

GAME $\mathbf{G}_{1,q}$. This game is identical to \mathbf{G}_1 except for the fact that the first $q-1$ oracle queries are answered using an obfuscation Γ_k of $C_{k,1}$ instead of $C_{k,b}$. Clearly, we have that

$$\Pr[out_1 = 1] = \Pr[out_{1,1} = 1] \quad (6.4)$$

$$\Pr[out_2 = 1] = \Pr[out_{1,Q+1} = 1] \quad (6.5)$$

Furthermore, we have that

$$\left| \Pr[out_{1,Q+1} = 1] - \Pr[out_{1,1} = 1] \right| \quad (6.6)$$

$$\leq \sum_{q=1}^Q \left| \Pr[out_{1,q} = 1] - \Pr[out_{1,q+1} = 1] \right|. \quad (6.7)$$

Due to Equations (6.4), (6.5) and (6.7), it suffices to upper bound the distinguishing gap between $\mathbf{G}_{1,q}$ and $\mathbf{G}_{1,q+1}$ in order to prove Equation (6.3).

Due to the statistical soundness of NIZK, with overwhelming probability over the random coins of Setup, for all $q \in [Q]$, there is no word $w \in \{0, 1\}^* \setminus L_{\text{state}}^{\mathbf{G}, S_q^{\text{in}}}$ such that a valid proof π for the statement $w \in L_{\text{state}}^{\mathbf{G}, S_q^{\text{in}}}$ exists. Hence, every obfuscated circuit Γ_q , produced by the oracle on input of an input sampler S_q^{in} , simulates the randomized computation of the circuit $C_{q,b}$ only on well-formed inputs, i. e., on outputs of S_q^{in} using random coins from the range of G . Consequently, we only need to perform the hybrid argument over all circuit inputs from [CLTV15] restricted to well-formed inputs. Since ELF is in extremely lossy mode, this set of well-formed inputs is *extremely sparse*. Let

$$B_{q,j} := \left\{ S_q^{\text{in}}(\text{stf}(y_1, \dots, y_{j-1}), x; G(r)) \mid \begin{array}{l} x \in L \quad , \quad r \in \{0, 1\}^{k(\lambda)} \\ k \in [j-1], y_k \in B_{q,k} \end{array} \right\}$$

be this set of all well-formed inputs for input position $j \in [\ell]$. The set $B_{q,j}$ contains at most $|L| \cdot t^{j-1}$ elements. Further, let

$$\gamma_{q,1} < \dots < \gamma_{q,\bar{t}}$$

be the ordered enumeration of all ℓ -tuples in $B_q := \prod_{j=1}^{\ell} B_{q,j}$, where \bar{t} denotes the size of B_q .²⁹ Hence, the total number of well-formed inputs

$$\bar{t} := \prod_{j=1}^{\ell} |B_{q,j}| \leq \left(|L| \cdot t^{\ell-1}\right)^{\ell} \leq |L|^{\ell} \cdot t^{(\ell^2)}$$

is polynomial in λ (given that ℓ is a constant, and $|L|$ and t are polynomial).

Towards proving indistinguishability between $G_{1,q}$ and $G_{1,q+1}$, we conduct a hybrid argument over all well-formed inputs for the obfuscation Γ_q and gradually replace the evaluation of circuit $C_{q,b}$ with $C_{q,1}$. From here on, our proof strategy is similar to the one employed in [CLTV15]. However, we only need to consider polynomially many hybrids, hence, this strategy only incurs a polynomial reduction loss relative to the underlying assumptions.

GAME $G_{1,q,i}$. The game $G_{1,q,i}$ is identical to game $G_{1,q}$ except that the oracle answers the q -th query differently using an obfuscation of the circuit

$$\overline{C}'[\text{stf}, (\sigma, G), S_q^{\text{in}}, C_{q,b}, C_{q,1}, K_q, \gamma_{q,i}],$$

defined in Figure 6.3, using iO .

```

Circuit  $\overline{C}'[\text{stf}, (\sigma, G), S_q^{\text{in}}, C_0, C_1, K, \gamma_i](x)$ 
parse  $x = ((y_1, aux_1), \dots, (y_{\ell}, aux_{\ell}))$ 
 $state_j := \text{stf}(y_1, \dots, y_{j-1}, \perp, \dots, \perp)$ 
if  $\neg (\forall j \in [\ell]: \text{NIZK.Verify}(\sigma, (G, S_q^{\text{in}}, state_j, y_j), aux_j) = 1)$  then
  return  $\perp$ 
 $\gamma := (y_1, \dots, y_{\ell})$ 
if  $\gamma < \gamma_i$  then  $r := F(K, \gamma)$ ; return  $C_1(\gamma; r)$ 
if  $\gamma = \gamma_i$  then  $r := F(K, \gamma)$ ; return  $C_b(\gamma; r)$ 
if  $\gamma > \gamma_i$  then  $r := F(K, \gamma)$ ; return  $C_b(\gamma; r)$ 

```

Figure 6.3: Definition of the circuit \overline{C}' .

The circuits

$$\overline{C}[\text{stf}, (\sigma, G), S_q^{\text{in}}, C_{q,b}, K_q] \text{ and } \overline{C}'[\text{stf}, (\sigma, G), S_q^{\text{in}}, C_{q,b}, C_{q,1}, K_q, \gamma_{q,1}]$$

are functionally equivalent. On input of $x = ((y_1, aux_1), \dots, (y_{\ell}, aux_{\ell}))$, both return $C_{q,b}(y_1, \dots, y_{\ell})$ with randomness $r = F(K_q, (y_1, \dots, y_{\ell}))$. Hence, this game hop is justified by the security of iO . More formally, for every PPT adversary \mathcal{A} , there exists a functionally equivalent circuit sampler $D_{1,q,1}$ producing circuits as above and a PPT adversary $\mathcal{B}_{1,q,1}$ such that

$$\left| \Pr[out_{1,q} = 1] - \Pr[out_{1,q,1} = 1] \right| \leq \text{Adv}_{iO, D_{1,q,1}, \mathcal{B}_{1,q,1}}^{\text{io-ind}}(\text{pC}(\lambda)). \quad (6.8)$$

²⁹ We remark that the values of each set $B_{q,j}$ can be computed efficiently by evaluating S^{in} on all possible inputs and states from $L \times (\prod_{k=1}^{j-1} B_{q,k})$ and all possible random coins in the range of G . Note that since ELF is strongly regular, it is possible to enumerate the image of G in polynomial time.

Furthermore, let $\gamma_{q,\bar{t}+1} > \gamma_{q,\bar{t}}$, then the circuits

$$\begin{aligned} &\bar{C}' [\text{stf}, (\sigma, G), S_q^{\text{in}}, C_{q,b}, C_{q,1}, K_q, \gamma_{q,\bar{t}+1}] \text{ and} \\ &\bar{C} [\text{stf}, (\sigma, G), S_q^{\text{in}}, C_{q,1}, K_q] \end{aligned}$$

are functionally equivalent. Hence, for every PPT adversary \mathcal{A} , there exists a functionally equivalent circuit sampler $D_{1,q+1}$ producing circuits as above and a PPT adversary $\mathcal{B}_{1,q+1}$, such that

$$\begin{aligned} &\left| \Pr \left[\text{out}_{1,q,\bar{t}+1} = 1 \right] - \Pr \left[\text{out}_{1,q+1} = 1 \right] \right| \\ &\leq \text{Adv}_{\text{IO}, D_{1,q+1}, \mathcal{B}_{1,q+1}}^{\text{io-ind}}(\text{pC}(\lambda)). \end{aligned} \quad (6.9)$$

Furthermore, we have that

$$\begin{aligned} &\left| \Pr \left[\text{out}_{1,q,\bar{t}+1} = 1 \right] - \Pr \left[\text{out}_{1,q,1} = 1 \right] \right| \\ &\leq \sum_{i=1}^{\bar{t}} \left| \Pr \left[\text{out}_{1,q,i+1} = 1 \right] - \Pr \left[\text{out}_{1,q,i} = 1 \right] \right| \end{aligned} \quad (6.10)$$

Due to Equations (6.8), (6.9) and (6.10), it suffices to upper bound the distinguishing gap between $\mathbf{G}_{1,q,i}$ and $\mathbf{G}_{1,q,i+1}$ in order to prove indistinguishability between $\mathbf{G}_{1,q}$ and $\mathbf{G}_{1,q+1}$.

For this purpose, we aim to exploit the dynamic-input indistinguishability of the circuit sampler D . Therefore, we first need to supply the evaluation of $C_{q,b}$ (and $C_{q,1}$) at input of γ_i with true randomness. Hence, we define another series of hybrids between $\mathbf{G}_{1,q,i}$ and $\mathbf{G}_{1,q,i+1}$.

GAME $\mathbf{G}_{1,q,i,1}$. This game is identical to $\mathbf{G}_{1,q,i}$ except for the fact that we use a punctured PRF key $K_q\{\gamma_{q,i}\}$ produced via $F.\text{Punct}(K_q, \gamma_{q,i})$ and obfuscate the circuit

$$\bar{C}'' [\text{stf}, (\sigma, G), C_{q,b}, C_{q,1}, K_q\{\gamma_{q,i}\}, Y := C_{q,b}(\gamma_{q,i}; F(K_q, \gamma_{q,i})), \gamma_{q,i}]$$

as defined in Figure 6.4 using iO .

As F preserves the functionality under punctured keys, the circuits

$$\begin{aligned} &\bar{C}' [\text{stf}, (\sigma, G), S_q^{\text{in}}, C_{q,b}, C_{q,1}, K_q, \gamma_{q,i}] \text{ and} \\ &\bar{C}'' [\text{stf}, (\sigma, G), S_q^{\text{in}}, C_{q,b}, C_{q,1}, K_q\{\gamma_{q,i}\}, \underbrace{Y}_{:= C_{q,b}(\gamma_{q,i}; F(K_q, \gamma_{q,i}))}, \gamma_{q,i}] \end{aligned}$$

are functionally equivalent. On input of $\gamma_{q,i}$, the circuits used in $\mathbf{G}_{1,q,i}$ outputs $C_{q,b}(\gamma_{q,i})$ using randomness $r := F(K_q, \gamma_{q,i})$, and the circuit used in $\mathbf{G}_{1,q,i,1}$ outputs Y which is computed as $Y := C_{q,b}(\gamma_{q,i}; F(K_q, \gamma_{q,i}))$. On all remaining inputs, these circuits also behave identically as F preserves functionality under punctured keys. Hence, for all PPT adversaries \mathcal{A} , there exists a functionally equivalent circuit sampler $D_{1,q,i,1}$ producing circuits as above distributed as in $\mathbf{G}_{1,q,i}$ and $\mathbf{G}_{1,q,i,1}$, respectively, and a PPT adversary $\mathcal{B}_{1,q,i,1}$ such that

$$\left| \Pr \left[\text{out}_{1,q,i} = 1 \right] - \Pr \left[\text{out}_{1,q,i,1} = 1 \right] \right| \leq \text{Adv}_{\text{IO}, D_{1,q,i,1}, \mathcal{B}_{1,q,i,1}}^{\text{io-ind}}(\text{pC}(\lambda)). \quad (6.11)$$

```

Circuit  $\bar{C}''$  [stf, ( $\sigma, G$ ),  $S^{\text{in}}$ ,  $C_0, C_1, K\{\gamma_i\}, Y, \gamma_i](x)$ 
parse  $x := ((y_1, aux_1), \dots, (y_\ell, aux_\ell))$ 
 $state_j := \text{stf}(y_1, \dots, y_{j-1}, \perp, \dots, \perp)$ 
if  $\neg(\forall j \in [\ell]: \text{NIZK.Verify}(\sigma, (G, S^{\text{in}}, state_j, y_j), aux_j) = 1)$  then
  return  $\perp$ 
 $\gamma := (y_1, \dots, y_\ell)$ 
if  $\gamma < \gamma_i$  then  $r := F(K\{\gamma_i\}, \gamma)$ ; return  $C_1(\gamma; r)$ 
if  $\gamma = \gamma_i$  then return  $Y$ 
if  $\gamma > \gamma_i$  then  $r := F(K\{\gamma_i\}, \gamma)$ ; return  $C_b(\gamma; r)$ 

```

Figure 6.4: Definition of the circuit \bar{C}'' .

We note that the view of \mathcal{A} in game $\mathbf{G}_{1,q,i,1}$ does not depend on the PRF key K_q (only on the punctured PRF key $K_q\{\gamma_{q,i}\}$). This enables a reduction to the selective security of F .

GAME $\mathbf{G}_{1,q,i,2}$. In this game we replace the randomness $F(K_q, (\gamma_{q,i}))$ by true randomness, i. e., we produce Y as $Y := C_{q,b}(\gamma_{q,i}; R)$ for uniformly random $R \leftarrow \{0, 1\}^{k(\lambda)}$. This game hop is justified by the selective PRF property. More formally, for every PPT adversary \mathcal{A} , there exists a PPT adversary $\mathcal{B}_{1,q,i,2}$ such that

$$\left| \Pr \left[\text{out}_{1,q,i,2} = 1 \right] - \Pr \left[\text{out}_{1,q,i,1} = 1 \right] \right| \leq \text{Adv}_{F, \mathcal{B}_{1,q,i,2}}^{\text{s-prf}}(\lambda). \quad (6.12)$$

GAME $\mathbf{G}_{1,q,i,3}$. This game is the same as $\mathbf{G}_{1,q,i,2}$ except for the fact that Y is produced using the circuit $C_{q,1}$, i. e., $Y := C_{q,1}(\gamma_{q,i}; R)$. This game hop is justified by the fact that the circuit sampler D is a dynamic-input indistinguishable sampler. More precisely, for every PPT adversary \mathcal{A} there exists a PPT adversary $\mathcal{B}_{1,q,i,3}$ such that

$$\left| \Pr \left[\text{out}_{1,q,i,3} = 1 \right] - \Pr \left[\text{out}_{1,q,i,2} = 1 \right] \right| \leq \text{Adv}_{D, \mathcal{B}_{1,q,i,3}}^{\text{dyn-ind}}(\lambda). \quad (6.13)$$

GAME $\mathbf{G}_{1,q,i,4}$. This game is the same as $\mathbf{G}_{1,q,i,3}$ with the difference that we again use pseudorandom coins to compute Y , i. e., $Y := C_{q,1}(\gamma_{q,i}; F(K_q, \gamma_{q,i}))$. For every PPT adversary \mathcal{A} , there exists a PPT adversary $\mathcal{B}_{1,q,i,4}$ such that

$$\left| \Pr \left[\text{out}_{1,q,i,4} = 1 \right] - \Pr \left[\text{out}_{1,q,i,3} = 1 \right] \right| \leq \text{Adv}_{F, \mathcal{B}_{1,q,i,4}}^{\text{s-prf}}(\lambda). \quad (6.14)$$

Since F preserves functionality under punctured keys, the two circuits

$$\begin{aligned} & \bar{C}'' \left[\text{stf}, (\sigma, G), S_q^{\text{in}}, C_{q,b}, C_{q,1}, K_q\{\gamma_{q,i}\}, \underbrace{Y}_{:= C_{q,1}(\gamma_{q,i}; F(K_q, \gamma_{q,i}))}, \gamma_{q,i} \right] \\ & \bar{C}' \left[\text{stf}, (\sigma, G), S_q^{\text{in}}, C_{q,b}, C_{q,1}, K_q, \gamma_{q,i+1} \right] \end{aligned}$$

are functionally equivalent. Therefore, for all PPT adversaries \mathcal{A} , there exists a functionally equivalent circuit sampler $D_{1,q,i+1}$ producing circuits as above

distributed as in $\mathbf{G}_{1,q,i.4}$ and $\mathbf{G}_{1,q,i+1}$, respectively, and a PPT adversary $\mathcal{B}_{1,q,i+1}$ such that

$$\begin{aligned} & \left| \Pr \left[\text{out}_{1,q,i+1} = 1 \right] - \Pr \left[\text{out}_{1,q,i.4} = 1 \right] \right| \\ & \leq \text{Adv}_{\text{IO}, \mathcal{D}_{1,q,i+1}, \mathcal{B}_{1,q,i+1}}^{\text{io-ind}}(\rho_{\mathbf{C}}(\lambda)). \end{aligned} \quad (6.15)$$

Combining Equations (6.11), (6.12), (6.13), (6.14) and (6.15), we obtain

$$\begin{aligned} & \left| \Pr \left[\text{out}_{1,q,i+1} = 1 \right] - \Pr \left[\text{out}_{1,q,i} = 1 \right] \right| \\ & \leq 2 \cdot \text{Adv}_{\text{IO}, \mathcal{D}', \mathcal{B}'}^{\text{io-ind}}(\rho_{\mathbf{C}}(\lambda)) + 2 \cdot \text{Adv}_{\mathbf{F}, \mathcal{B}''}^{\text{s-prpf}}(\lambda) + \text{Adv}_{\mathcal{D}, \mathcal{B}'''}^{\text{dyn-ind}}(\lambda) \end{aligned} \quad (6.16)$$

for PPT adversaries \mathcal{B}' , \mathcal{B}'' , \mathcal{B}''' and a functionally equivalent circuit sampler \mathcal{D}' . Since IO is a secure indistinguishability obfuscator, \mathbf{F} is a puncturable PRF and \mathcal{D} is a dynamic-input indistinguishable sampler, the above quantities are negligible in λ .

Combining Equations (6.8), (6.9), (6.10) and (6.16), the advantage in distinguishing $\mathbf{G}_{1,q}$ and $\mathbf{G}_{1,q+1}$ can be upper bounded by

$$\begin{aligned} & \left| \Pr \left[\text{out}_{1,q+1} = 1 \right] - \Pr \left[\text{out}_{1,q} = 1 \right] \right| \\ & \leq |\mathbb{L}|^\ell \cdot t^{\ell^2} \cdot \text{negl}(\lambda) + 2 \cdot \text{Adv}_{\text{IO}, \mathcal{D}', \mathcal{B}'}^{\text{io-ind}}(\rho_{\mathbf{C}}(\lambda)) \end{aligned} \quad (6.17)$$

for some PPT adversary \mathcal{B}' and some functionally equivalent circuit sampler \mathcal{D}' , which is negligible.

Hence, combining Equations (6.4), (6.5), (6.7) and (6.17), the advantage to distinguish \mathbf{G}_1 and \mathbf{G}_2 is upper bounded by

$$\left| \Pr \left[\text{out}_2 = 1 \right] - \Pr \left[\text{out}_1 = 1 \right] \right| \leq Q \cdot \text{negl}'(\lambda) + 2 \cdot \text{Adv}_{\text{IO}, \mathcal{D}', \mathcal{B}'}^{\text{io-ind}}(\rho_{\mathbf{C}}(\lambda)),$$

for some negligible function negl' , for some PPT adversary \mathcal{B}' and some functionally equivalent circuit sampler \mathcal{D}' , which is negligible. This proves Equation (6.3).

Finally, together with Equations (6.1) and (6.2), we reach a contradiction concluding the proof. \square

EXTENSION. We note that our notion of doubly-probabilistic IO can be extended to support a set of independent public parameters such that encoded input samplers may depend on different public parameters. Suppose $(\text{Setup}, \text{Encode}, \text{Obfuscate})$ is a 5-source dpIO scheme. Then, the obfuscation $\text{Obfuscate}(pp_1[1-3], pp_2[4, 5], S^{\text{in}}, C)$ of the circuit C requires the first three inputs to be sampled with an encoded sampler with respect to pp_1 and the last two inputs to be sampled with an encoded sampler with respect to pp_2 .

The above prove can be directly extended to this setting by choosing each extremely-lossy mode such that \mathcal{A} has advantage at most $\epsilon/2\ell$ in distinguishing the lossy from the injective mode. For the applications covered in this thesis, we do not need this extension and refer the reader to the original publication [ACH20] for more details.

7

LEVELED HOMOMORPHIC ENCRYPTION

In this chapter, we demonstrate that our notion of **dpIO** from Chapter 5 can be applied to construct leveled homomorphic encryption in a similar way as in [CLTV15]. This construction leads to a transformation which operates on an encryption scheme E , satisfying **IND-CPA** security (and possibly other security properties, e.g., **KDM** security), and produces a leveled homomorphic encryption scheme that retains the security properties of E . We emphasize that since our notion of **dpIO** achieves security against dynamic-input indistinguishable samplers, we do not need to assume any additional properties for E . This is a further improvement compared to [CLTV15], where additional structural properties such as lossiness or re-randomizability are required.³⁰

Let $\mathcal{P} = (\mathcal{P}_{\lambda, \tilde{L}})_{\lambda, \tilde{L} \in \mathbb{N}}$ be a family of sets of polynomial sized depth- \tilde{L} circuits of arity k , i. e., the set $\mathcal{P}_{\lambda, \tilde{L}}$ contains circuits of polynomial size in λ of depth \tilde{L} . Further, for any $\lambda, \tilde{L} \in \mathbb{N}$ the circuits in $\mathcal{P}_{\lambda, \tilde{L}}$ share the common input domain $\{0, 1\}^k$.

DEFINITION 7.1 (Leveled homomorphic encryption [Gen09]). A leveled homomorphic encryption (**LHE**) scheme for message space $\{0, 1\}$ is a family of encryption schemes $\text{LHE} := \{E_{\tilde{L}} = (\text{KGen}_{\tilde{L}}, \text{Enc}_{\tilde{L}}, \text{Dec}_{\tilde{L}}, \text{Eval}_{\tilde{L}})\}_{\tilde{L} \in \mathbb{N}}$ such that each $E_{\tilde{L}}$ is an **IND-CPA** secure public-key encryption scheme for message space $\{0, 1\}$ which is homomorphic for all polynomial-size depth- \tilde{L} circuits. That is, each $E_{\tilde{L}}$ allows to homomorphically and compactly evaluate any polynomial-size depth- \tilde{L} circuit via an evaluation algorithm $\text{Eval}_{\tilde{L}}$ of size polynomial in (λ, \tilde{L}) , and all the $E_{\tilde{L}}$ use the same decryption circuit $\text{Dec}_{\tilde{L}} = \text{Dec}$. An **LHE** scheme is required to satisfy the following notion of correctness.

CORRECTNESS. For all $\lambda \in \mathbb{N}$, all $\tilde{L} \in \mathbb{N}$, all $m_1, \dots, m_k \in \mathcal{M}_\lambda$, $P \in \mathcal{P}_{\lambda, \tilde{L}}$,

$$\Pr \left[\begin{array}{l} (pk, ek, sk) \leftarrow \text{KGen}_{\tilde{L}}(1^\lambda) \\ c_i \leftarrow \text{Enc}_{\tilde{L}}(pk, m_i) \\ c \leftarrow \text{Eval}_{\tilde{L}}(ek, P, c_1, \dots, c_k) \end{array} : \begin{array}{l} \text{Dec}(sk, c) \\ = \\ P(m_1, \dots, m_k) \end{array} \right] = 1.$$

IND-CPA SECURITY. For all $\tilde{L} \in \mathbb{N}$, all **PPT** adversaries \mathcal{A} ,

$$\text{Adv}_{\text{LHE}, \mathcal{A}}^{\text{ind-cpa}}(\lambda) := \left| \Pr \left[\begin{array}{l} (pk, ek, sk) \leftarrow \text{KGen}_{\tilde{L}}(1^\lambda) \\ c^* \leftarrow \text{Enc}_{\tilde{L}}(pk, 0) \end{array} : \mathcal{A}(pk, ek, c^*) = 1 \right] - \Pr \left[\begin{array}{l} (pk, ek, sk) \leftarrow \text{KGen}_{\tilde{L}}(1^\lambda) \\ c^* \leftarrow \text{Enc}_{\tilde{L}}(pk, 1) \end{array} : \mathcal{A}(pk, ek, c^*) = 1 \right] \right|$$

is negligible in λ .

³⁰ It is not known whether such structural properties can be achieved from **IO** in conjunction with unstructured assumptions, such as one-way functions.

COMPACTNESS. The size of the output of $\text{Eval}_{\tilde{C}}$ is polynomial in (λ, \tilde{L}) and independent of the size of the circuit P .

Note that Definition 7.1 distinguishes between an evaluation key ek and a public (encryption) key pk . The evaluation key can equivalently be defined to be part of the public key pk as in Definition 2.24.

REMARK 7.1. Note that a leveled (or fully) homomorphic encryption scheme for message space $\{0, 1\}$ implies a leveled (or fully) homomorphic encryption scheme for message space $\{0, 1\}^\lambda$.

Let stf be the trivial state function, i. e., $\text{stf}: (y_1, y_2) \mapsto \perp$ for each tuple $(y_1, y_2) \in (\{0, 1\}^{n(\lambda)} \cup \{\perp\})^2$. Let $E = (\text{KGen}, \text{Enc}, \text{Dec})$ be an arbitrary IND-CPA secure public-key encryption scheme. Let the input sampler class \mathcal{S}^{in}

$$\mathcal{S}^{\text{in}} := \left\{ S_{pk}^{\text{in}} \mid (pk, \cdot) \in \text{supp}(E.\text{KGen}(1^\lambda)) \right\},$$

where S_{pk}^{in} is defined in Figure 7.1. For technical reasons, the input length of S_{pk}^{in} is required to match the security parameter λ (see Definition 5.1). Hence, the input domain L of the input samplers in \mathcal{S}^{in} equals $\{0^\lambda, 1^\lambda\}$. That is, \mathcal{S}^{in}

```

 $S_{pk}^{\text{in}}(\text{state}, x)$ 
y :=  $\perp$ , y' :=  $\perp$ 
y  $\leftarrow$   $\begin{cases} E.\text{Enc}(pk, 0) & , \text{if } x = 0^\lambda \\ E.\text{Enc}(pk, 1) & , \text{if } x = 1^\lambda \end{cases}$ 
return (y; y')
```

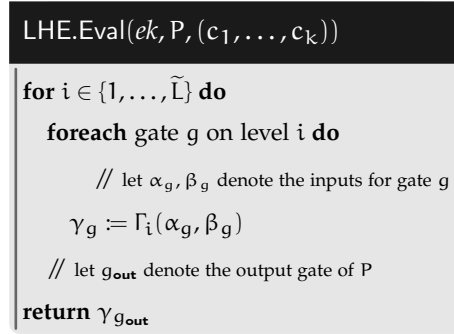
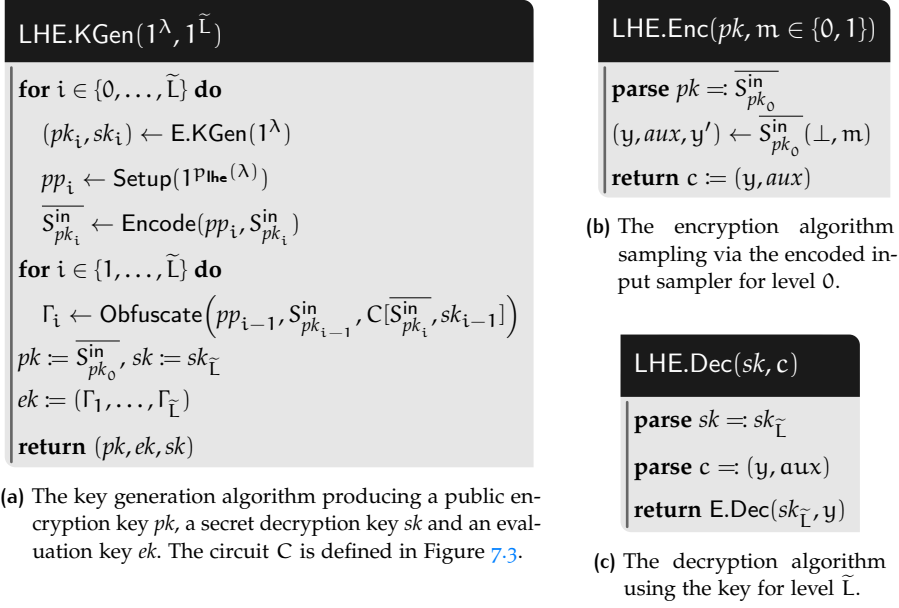
Figure 7.1: Input samplers used for the construction of LHE from dpiO.

contains all input samplers S_{pk}^{in} which on input of a state state and an input $x \in L$, sample a ciphertext y encrypting x via $y := E.\text{Enc}(pk, x)$ and $y' := \perp$ ignoring state , where pk is a public key in the range of $E.\text{KGen}(1^\lambda)$. However, for a more convenient notation, we will slightly abuse notation and treat S_{pk}^{in} as if its input domain L equals $\{0, 1\}$.

Further, let \mathcal{C} be the class of polynomial-size randomized circuits and let $\mathcal{C}^{\text{dyn-ind}}$ be the class of dynamic-input indistinguishable samplers over \mathcal{C} .

THEOREM 7.1. *Let $(\text{Setup}, \text{Encode}, \text{Obfuscate})$ be a 2-source dpiO scheme for $(\text{stf}, \mathcal{S}^{\text{in}}, \mathcal{C}, \mathcal{C}^{\text{dyn-ind}})$ and let E be an IND-CPA secure public-key encryption scheme. Then, LHE as defined in Figure 7.2 is an IND-CPA secure LHE scheme.*

The proof strategy is similar as in [CLTV15]. On a high level, we aim to reduce the IND-CPA security of LHE to the IND-CPA security of the underlying encryption scheme E . However, the evaluation key ek contains information (even though obfuscated) about the secret decryption keys of each level. For the purpose of invoking the security of E on the challenge ciphertext, we need to remove this dependency on sk_0 . Therefore, we gradually (starting from level \tilde{L}) replace the obfuscations of the circuits C with obfuscations of *trapdoor* circuits tC which simply output samples produced by the encoded sampler $\overline{S_{pk}^{\text{in}}}(0)$ on the fixed input 0 (hence, not needing any information on decryption keys). Let $p_{\text{lhe}}(\lambda)$ denote an upper bound on the size of C and tC including all hard-coded values during the proof. These two circuits sample




- (d) The evaluation algorithm using the obfuscated circuits Γ_i to homomorphically evaluate circuit P on the ciphertexts c_1, \dots, c_k .

Figure 7.2: Instantiation of an LHE scheme LHE from some PKE scheme E using a 2-source dpiO scheme dpiO.

from the same encoded sampler $\overline{S_{pk}^{in}}$ but behave differently in which inputs they supply $\overline{S_{pk}^{in}}$ with. Due to the simulatability of encodings property of dpiO and the IND-CPA security of E , these two circuits are dynamic-input indistinguishable. Hence, by the indistinguishability property of dpiO for $\mathcal{C}^{\text{dyn-ind}}$, an honest evaluation key and an evaluation key consisting only of trapdoor circuits are indistinguishable.

Given these modifications, the challenge ciphertext c^* consists of an encryption of a bit b under pk_0 accompanied by some auxiliary information produced by the corresponding encoded sampler. This auxiliary information might leak information on the bit b and thereby prevents to directly employ the IND-CPA security of E . However, as dpiO satisfies simulatability of encodings, this auxiliary information can be simulated without knowledge of b and, hence, contains no information about b . Therefore, by the IND-CPA security of E , LHE is IND-CPA secure.



(a) Definition of the probabilistic circuit C which implements homomorphic evaluation of a single $\overline{\wedge}$ -gate. (b) Definition of the trapdoor analogue tC of the circuit C which is unknowing of any secret decryption keys.

Figure 7.3: Definition of the circuits C and tC .

Proof of Theorem 7.1. We prove that LHE satisfies correctness and IND-CPA security. Let P be a polynomial sized circuit of depth \tilde{L} which consists only of $\overline{\wedge}$ -gates³¹ with input space $\{0, 1\}^k$.

CORRECTNESS. We prove that on every level i , every gate g in P is evaluated exactly as the original gate given the decrypted inputs.

Let $i \in [\tilde{L}]$ be a level and let g be a gate on level i . Further, let $(\chi_\alpha, aux_\alpha, \cdot)$ be in the support of $\overline{S_{pk_{i-1}}^{\text{in}}}(\perp, \alpha)$ and $(\chi_\beta, aux_\beta, \cdot)$ be in the support of $\overline{S_{pk_{i-1}}^{\text{in}}}(\perp, \beta)$. By perfect correctness of E , the circuit $C[\overline{S_{pk_i}^{\text{in}}}, sk_{i-1}]$ on input of $(\chi_\alpha, aux_\alpha)$ and (χ_β, aux_β) always outputs an element from the support of $\overline{S_{pk_i}^{\text{in}}}(\perp, \alpha \overline{\wedge} \beta)$.³² Furthermore, since dpiO is support respecting (Definition 5.4), both the encoded sampler $\overline{S^{\text{in}}}$ and the obfuscated circuit Γ respect the support of the original sampler S^{in} and the original circuit $C[\overline{S_{pk_i}^{\text{in}}}, sk_{i-1}]$, respectively. Hence, the obfuscation Γ of $C[\overline{S_{pk_i}^{\text{in}}}, sk_{i-1}]$ (with respect to $S_{pk_{i-1}}^{\text{in}}$) on input of $(\chi_\alpha, aux_\alpha)$ and (χ_β, aux_β) always outputs an element from the support of $\overline{S_{pk_i}^{\text{in}}}(\perp, \alpha \overline{\wedge} \beta)$. Hence, the probability that gate g is not evaluated correctly is zero.

SECURITY. To prove that LHE is IND-CPA secure, we proceed over a series of hybrids. Let \mathcal{A} be a PPT adversary. Further, let $p_{\text{Ihe}}(\lambda)$ denote an upper bound on the size of C and tC including all hard-coded values during the proof.

GAME G_0^b . This game is exactly the IND-CPA game for LHE, where the challenge ciphertext c^* contains the bit b .

GAME G_1^b . This game is the same as G_0^b with the difference that the evaluation key consists of obfuscations $t\Gamma_i$ of the circuit $tC[\overline{S_{pk_i}^{\text{in}}}]$.

CLAIM 7.1. For every PPT adversary \mathcal{A} , we have that

$$\left| \Pr[out_1^b = 1] - \Pr[out_0^b = 1] \right| \leq \text{negl}(\lambda)$$

for some negligible function negl .

Proof. We define a series of $\tilde{L} + 1$ hybrid games $G_{0,i}^b$ for $i \in \{0, \dots, \tilde{L}\}$ as follows.

GAME $G_{0,i}^b$. $G_{0,i}^b$ is defined in Figure 7.4. Hence, $G_{0,i}^b$ is identical to $G_{0,i-1}^b$

³¹ This is without loss of generality since $\overline{\wedge}$ -gates are a functionally complete set of Boolean operators.

³² We implicitly use that the supports of $\overline{S^{\text{in}}}(\perp, 0)$ and $\overline{S^{\text{in}}}(\perp, 1)$ are disjoint.

```

G0,ib
for  $i \in \{0, \dots, \tilde{L}\}$  do
     $(pk_i, sk_i) \leftarrow \text{E.KGen}(1^\lambda)$ 
     $pp_i \leftarrow \text{Setup}(1^\lambda)$ 
     $\overline{S_{pk_i}^{\text{in}}} \leftarrow \text{Encode}(pp_i, S_{pk_i}^{\text{in}})$ 
    for  $j \in \{1, \dots, \tilde{L} - i\}$  do
         $\Gamma_j \leftarrow \text{Obfuscate}(pp_{j-1}, S_{pk_{j-1}}^{\text{in}}, C[\overline{S_{pk_j}^{\text{in}}}, sk_{j-1}])$ 
    for  $j \in \{\tilde{L} - i + 1, \dots, \tilde{L}\}$  do
         $t\Gamma_j \leftarrow \text{Obfuscate}(pp_{j-1}, S_{pk_{j-1}}^{\text{in}}, tC[\overline{S_{pk_j}^{\text{in}}}] )$ 
     $pk := \overline{S_{pk_0}^{\text{in}}}, sk := sk_{\tilde{L}}$ 
     $ek := (\Gamma_1, \dots, \Gamma_{\tilde{L}-i}, t\Gamma_{\tilde{L}-i+1}, \dots, t\Gamma_{\tilde{L}})$ 
    return  $(pk, ek, sk)$ 
    
```

Figure 7.4: Definition of $\mathbf{G}_{0,i}^b$.

with the difference that the $\tilde{L} - i + 1$ -th obfuscated circuit in the evaluation key ek is produced as

$$\Gamma_{\tilde{L}-i+1} \leftarrow \text{Obfuscate}(pp_{\tilde{L}-i}, S_{pk_{\tilde{L}-i}}^{\text{in}}, C[\overline{S_{pk_{\tilde{L}-i+1}}^{\text{in}}}, sk_{\tilde{L}-i}])$$

in $\mathbf{G}_{0,i-1}^b$, and as

$$t\Gamma_{\tilde{L}-i+1} \leftarrow \text{Obfuscate}(pp_{\tilde{L}-i}, S_{pk_{\tilde{L}-i}}^{\text{in}}, tC[\overline{S_{pk_{\tilde{L}-i+1}}^{\text{in}}}])$$

in $\mathbf{G}_{0,i}^b$. By definition, we have

$$\Pr[out_{0,0}^b = 1] = \Pr[out_0^b = 1] \quad \text{and} \quad (7.1)$$

$$\Pr[out_{0,\tilde{L}}^b = 1] = \Pr[out_1^b = 1]. \quad (7.2)$$

CLAIM 7.2. For every PPT adversary \mathcal{A} , there exists a dynamic-input indistinguishable sampler $D_{0,i+1}^b$ and a PPT adversary $\mathcal{B}_{0,i+1}^b$, such that

$$\left| \Pr[out_{0,i}^b = 1] - \Pr[out_{0,i+1}^b = 1] \right| \leq \text{Adv}_{\text{dpiO}, D_{0,i+1}^b, \mathcal{B}_{0,i+1}^b}^{\text{ind}}(\text{plhe}(\lambda)).$$

By a standard hybrid argument and Equations (7.1) and (7.2), we get $|\Pr[out_0^b = 1] - \Pr[out_1^b = 1]| \leq \tilde{L} \cdot \text{Adv}_{\text{dpiO}, D, \mathcal{B}}^{\text{ind}}(\text{plhe}(\lambda))$ for some dynamic-input indistinguishable sampler D and some PPT adversary \mathcal{B} . By the security of dpiO, this quantity is negligible.

Proof of Claim 7.2. We first show that there exists a PPT adversary $\mathcal{B}_{0,i+1}^b$ and the circuit sampler $D_{0,i+1}^b$ defined in Figure 7.5, such that the distinguishing gap between $\mathbf{G}_{0,i}^b$ and $\mathbf{G}_{0,i+1}^b$ is upper bounded by the advantage of $\mathcal{B}_{0,i+1}^b$ in the dpiO security game. Then we show that $D_{0,i+1}^b$ is a dynamic-input indistinguishable sampler.

We construct a PPT adversary $\mathcal{B}_{0,i+1}^b$ for the security game of dpiO as follows. $\mathcal{B}_{0,i+1}^b$ receives as input public parameters pp' , the circuits C_0 and C_1

```

D0,i+1b(1λ)
pp ← Setup(1λ)
(pk, sk) ← E.KGen(1λ)
(pk', sk') ← E.KGen(1λ)
Spkin ← Encode(pp, Spkin)
C0 := C[Spkin, sk']
C1 := tC[Spkin]
z := (pk, pk', sk', pp)
return (C0, C1, z)

```

Figure 7.5: Circuit sampler used in the proof of **IND-CPA** security of LHE.

together with auxiliary information z produced by $D_{0,i+1}^b$, and an obfuscation $\bar{\Gamma} \leftarrow \text{Obfuscate}(pp', S_{pk'}^{\text{in}}, C_{\tilde{b}})$ for some $\tilde{b} \in \{0, 1\}$. Initially, $\mathcal{B}_{0,i+1}^b$ samples parameters as in $\mathbf{G}_{0,i}^b$ embedding his input as $pp_{\tilde{\Gamma}-i-1} := pp'$, $pk_{\tilde{\Gamma}-i-1} := pk'$, and $pk_{\tilde{\Gamma}-i} := pk$, and defines the public key to be $pk := S_{pk_0}^{\text{in}}$ for $S_{pk_0}^{\text{in}} \leftarrow \text{Encode}(pp_0, S_{pk}^{\text{in}})$. $\mathcal{B}_{0,i+1}^b$ produces the obfuscations Γ_j for $j \in \{1, \dots, \tilde{\Gamma} - i - 1\}$ and $t\Gamma_j$ for $j \in \{\tilde{\Gamma} - i + 1, \dots, \tilde{\Gamma}\}$ as in $\mathbf{G}_{0,i}^b$ and defines the evaluation key as $ek := (\Gamma_1, \dots, \Gamma_{\tilde{\Gamma}-i-1}, \bar{\Gamma}, t\Gamma_{\tilde{\Gamma}-i+1}, \dots, t\Gamma_{\tilde{\Gamma}})$.

$\mathcal{B}_{0,i+1}^b$ produces the challenge ciphertext c^* for the bit b as $c^* \leftarrow \overline{S_{pk_0}^{\text{in}}}(b)$. Finally, $\mathcal{B}_{0,i+1}^b$ calls \mathcal{A} on input (pk, ek, c^*) and outputs \mathcal{A} 's output. If $\tilde{b} = 0$, $\mathcal{B}_{0,i+1}^b$ perfectly simulates $\mathbf{G}_{0,i}^b$ for \mathcal{A} , else $\mathcal{B}_{0,i+1}^b$ perfectly simulates $\mathbf{G}_{0,i+1}^b$ for \mathcal{A} . Hence,

$$\left| \Pr \left[\text{out}_{0,i+1}^b = 1 \right] - \Pr \left[\text{out}_{0,i}^b = 1 \right] \right| \leq \text{Adv}_{\text{dpiO}, D_{0,i+1}^b, \mathcal{B}_{0,i+1}^b}^{\text{ind}}(\text{plhe}(\lambda)).$$

What remains to prove is that $D_{0,i+1}^b$ is a dynamic-input indistinguishable sampler.

CLAIM 7.3. The circuit sampler $D_{0,i+1}^b$ as defined in Figure 7.5 is a dynamic-input indistinguishable sampler.

Proof of Claim 7.3. We prove this using a sequence of two hybrids. Let \mathcal{A} be a **PPT** adversary in the dynamic-input indistinguishability game for the circuit sampler $D_{0,i+1}^b$.

GAME \mathbf{H}_0 . This is the dynamic-input indistinguishability game as defined in Figure 2.2a. In detail, \mathbf{H}_0 samples (C_0, C_1, z) according to $D_{0,i+1}^b$, gives C_0, C_1 together with the state z to \mathcal{A} who adaptively chooses an input x . \mathbf{H}_0 evaluates the circuit $C_{b'}$ (for a random bit b') at input of x (using fresh random coins) and sends the output to \mathcal{A} .

GAME \mathbf{H}_1 . By the simulatability of encodings property of dpiO , there exists a simulator $\text{Sim} = (\text{Sim}_0, \text{Sim}_1)$, such that outputs of the encoded sampler (including the auxiliary information aux) and outputs of the original sampler together with simulated auxiliary information are indistinguishable. The game \mathbf{H}_1 is the same as \mathbf{H}_0 with the difference that the public parameters are simulated as $(pp, td) \leftarrow \text{Sim}_0(1^\lambda)$. Given the adversarially chosen input

$x = (x_\alpha, x_\beta)$, \mathbf{H}_1 computes $(m_0, m_1) := (\alpha \bar{\wedge} \beta, 0)$ using the secret key sk' (which is part of z). Furthermore, \mathbf{H}_1 produces $(y, \cdot) \leftarrow S_{pk}^{\text{in}}(\perp, m_{b'})$ and $aux \leftarrow \text{Sim}_1(td, S_{pk'}^{\text{in}}, y, \perp)$.

In order to upper bound the distinguishing gap between \mathbf{H}_1 and \mathbf{H}_0 , we construct an adversary $\mathcal{B}_{\mathbf{H}.1}$ for the simulatability of encodings property of dpiO . On input of public parameters pp , $\mathcal{B}_{\mathbf{H}.1}$ uses these as public parameters and simulates \mathbf{H}_1 for \mathcal{A} . However, instead of sampling y and aux with $S_{pk}^{\text{in}}(\perp, m_{b'})$ and $\text{Sim}_1(td, S_{pk'}^{\text{in}}, y, \perp)$, respectively, $\mathcal{B}_{\mathbf{H}.1}$ calls the encoding oracle provided by the simulatability game on input of $(S_{pk'}^{\text{in}}, \perp, m_{b'})$ to obtain y and aux . If $\mathcal{B}_{\mathbf{H}.1}$ interacts with $\text{Exp}_{\text{dpiO}, \mathcal{B}_{\mathbf{H}.1}}^{(0)\text{-enc}}(\text{plhe}(\lambda))$, $\mathcal{B}_{\mathbf{H}.1}$ perfectly simulates \mathbf{H}_0 for \mathcal{A} . Otherwise, if $\mathcal{B}_{\mathbf{H}.1}$ interacts with $\text{Exp}_{\text{dpiO}, \mathcal{B}_{\mathbf{H}.1}}^{(1)\text{-enc}}(\text{plhe}(\lambda))$, $\mathcal{B}_{\mathbf{H}.1}$ perfectly simulates \mathbf{H}_1 for \mathcal{A} . Hence, we have

$$\left| \Pr[out_{\mathbf{H}_1} = 1] - \Pr[out_{\mathbf{H}_0} = 1] \right| \leq \text{Adv}_{\text{dpiO}, \text{Sim}, \mathcal{B}_{\mathbf{H}.1}}^{\text{enc}}(\text{plhe}(\lambda))$$

which is negligible by the simulatability of encodings property of dpiO .

The advantage of \mathcal{A} in game \mathbf{H}_1 is negligible. To realize that, we construct an adversary $\mathcal{D}_{\mathbf{H}.1}$ for the **IND-CPA** security of E with respect to the public key pk . On input of pk , $\mathcal{D}_{\mathbf{H}.1}$ samples (C_0, C_1, z) in the same way as in \mathbf{H}_1 (embedding pk) and calls \mathcal{A} on input of (C_0, C_1, z) to obtain (x, st) . $\mathcal{D}_{\mathbf{H}.1}$ decrypts $x = (x_\alpha, x_\beta)$ to obtain α and β using sk' , and outputs $(m_0, m_1) := (\alpha \bar{\wedge} \beta, 0)$ to the **IND-CPA** game. In return, $\mathcal{D}_{\mathbf{H}.1}$ receives a ciphertext c^* , simulates the corresponding auxiliary information $aux \leftarrow \text{Sim}_1(pp, td, S_{pk'}^{\text{in}}, c^*, \perp)$, and invokes \mathcal{A} on input of his previous state st and (c^*, aux) . If c^* encrypts $m_{\tilde{b}}$, y is distributed exactly as in \mathbf{H}_1 conditioned on $b' = \tilde{b}$. Finally, $\mathcal{D}_{\mathbf{H}.1}$ outputs \mathcal{A} 's output. Hence, by the **IND-CPA** security of E ,

$$\left| \Pr[out_{\mathbf{H}_1} = 1] - \frac{1}{2} \right| \leq \frac{1}{2} \cdot \text{Adv}_{E, \mathcal{D}_{\mathbf{H}.1}}^{\text{ind-cpa}}(\lambda)$$

which is negligible. This concludes the proof of Claim 7.3. □

Together with Claim 7.3, this concludes the proof of Claim 7.2. □

Hence, we have that $|\Pr[out_1^b = 1] - \Pr[out_0^b = 1]|$ is negligible in λ , what was to be demonstrated. □

Towards being able to conduct a reduction to the **IND-CPA** security of E for the challenge ciphertext, we need to ensure that the auxiliary information which accompanies the ciphertext from $S_{pk_0}^{\text{in}}(b)$ does not leak any information about the plaintext b . This is formalized via another game transition.

GAME \mathbf{G}_2^b . This game is identical to \mathbf{G}_1^b except for the fact that the public parameters pp_0 are simulated by Sim_0 (additionally yielding a trapdoor td_0). Furthermore, the challenge ciphertext c^* is produced by drawing a sample $(y; y')$ from $S_{pk_0}^{\text{in}}(b)$ and simulating the auxiliary information aux using $\text{Sim}_1(pp_0, td_0, S_{pk_0}^{\text{in}}, y, \perp)$. We recall that for every **PPT** adversary \mathcal{A} such a simulator $\text{Sim} = (\text{Sim}_0, \text{Sim}_1)$ exists as dpiO satisfies simulatability of encodings. Hence, there exists a **PPT** adversary \mathcal{B}_2^b , such that the distinguishing gap

$$\left| \Pr[out_2^b = 1] - \Pr[out_1^b = 1] \right| \leq \text{Adv}_{\text{dpiO}, \text{Sim}, \mathcal{B}_2^b}^{\text{enc}}(\text{plhe}(\lambda))$$

which, by simulatability of encodings, is negligible in λ .

It remains to argue, that the games \mathbf{G}_2^0 and \mathbf{G}_2^1 are indistinguishable. This can be reduced to the **IND-CPA** security of E with respect to pk_0 as in \mathbf{G}_2^b neither the corresponding secret key sk_0 nor the plaintext b are necessary for simulation.

CLAIM 7.4. For every **PPT** adversary \mathcal{A} , the distinguishing gap $|\Pr[out_2^0 = 1] - \Pr[out_2^1 = 1]|$ is negligible.

Proof. We construct a **PPT** adversary \mathcal{B}'_2 against the **IND-CPA** security of E . On input of pk , \mathcal{B}'_2 defines $pk_0 := pk$, samples key pairs $(pk_j, sk_j) \leftarrow E.KGen(1^\lambda)$ for $j \in \{1, \dots, \tilde{L}\}$ and produces ek as in \mathbf{G}_2 . \mathcal{B}'_2 outputs the messages 0 and 1 to the **IND-CPA** game and in turn receives a ciphertext c'^* . Furthermore, \mathcal{B}'_2 simulates the auxiliary information by $aux^* \leftarrow \text{Sim}_1(pp_0, td_0, S_{pk_0}^{\text{in}}, c'^*, \perp)$. Finally, \mathcal{B}'_2 invokes \mathcal{A} on input of $(pk_0, ek, (c'^*, aux^*))$ and outputs \mathcal{A} 's output. If c'^* encrypts 0, \mathcal{B}'_2 perfectly simulates \mathbf{G}_2^0 for \mathcal{A} , and if c'^* encrypts 1, \mathcal{B}'_2 perfectly simulates \mathbf{G}_2^1 for \mathcal{A} . Hence,

$$\left| \Pr[out_2^0 = 1] - \Pr[out_2^1 = 1] \right| \leq \text{Adv}_{E, \mathcal{B}'_2}^{\text{ind-cpa}}(\lambda),$$

which is negligible by the **IND-CPA** security of E . □

□

Together with Theorem 6.1 from Chapter 6, we obtain the following corollary.

COROLLARY 7.1. *Assuming polynomially secure indistinguishability obfuscation and extremely lossy functions, then there exists a leveled homomorphic encryption scheme.*

Note that **IND-CPA** secure schemes, as required in our construction, can be constructed from (polynomially secure) **IO** and one-way functions (the latter being implied by **ELFs**). Previously, constructions of **LHE** were only known from the learning with error assumption, or from *subexponentially secure* indistinguishability obfuscation (together with lossy encryption, which can be based, for instance, on **DDH**, see also Section 9.5).

FULLY HOMOMORPHIC ENCRYPTION. [CLTV15] provide a generic **LHE-to-FHE** transformation. We briefly sketch the strategy. If an **LHE** scheme can handle a (slightly) super-polynomial number of levels \tilde{L} , then it is fully homomorphic for all polynomial-size circuits. However, this would require generating and storing a super-polynomial number of evaluation keys $(ek_1, \dots, ek_{\tilde{L}})$. To overcome this issue, [CLTV15] (inspired by [BGL+15]) use an obfuscated master key generation program **MProg**, which generates the evaluation keys *on the fly*: on input of i , **MProg** returns ek_i . It was observed in [CLTV15] that **pIO** cannot be directly used to obfuscate this program, because the coins used by **MProg** must be correlated between two consecutive levels (for example, in our **LHE** construction, ek_i is an obfuscated program with a secret key sk_i and a public key pk_{i+1} hard-coded; hence, ek_{i+1} must contain the matching secret key for pk_{i+1}). Nevertheless, “de-randomizing” this program, using the standard technique of generating the coins using a pseudorandom function F (the coins for sk_i are generated as $F(K, i)$, and the

coins for pk_{i+1} are generated as $F(K, i + 1)$, which guarantees the appropriate correlation between the coins used in ek_i and ek_{i+1} , and obfuscating with standard **IO**, allows to prove security of this approach.

Applying this generic transformation, we obtain the following corollary.

COROLLARY 7.2. *Assuming slightly super-polynomially secure indistinguishability obfuscation and extremely lossy functions, there exists a fully homomorphic encryption scheme.*

Previously, constructions of **FHE** were only known from circular security assumptions over lattice-based cryptosystems, or subexponentially secure indistinguishability obfuscation and lossy encryption.

Applying the **LHE-toFHE** transformation from [CLTV15] requires to use super-polynomially secure **IO**. However, using the same techniques as we used to build **dpIO**, we can actually base this transformation on polynomially secure **IO** and **ELFs**.

In the original publication [ACH20], we provide an **LHE-to-FHE** transformation using the techniques developed in this part which only requires polynomially secure **IO** and **ELFs**. This is not part of this thesis.

Part II

INSTANTIATING THE ALGEBRAIC GROUP
MODEL FROM OBFUSCATION

II INSTANTIATING THE ALGEBRAIC GROUP MODEL FROM OBFUSCATION	
8	INTRODUCTION 87
8.1	Technical Overview 90
8.2	Related Work 92
9	PRELIMINARIES 93
9.1	Notations and Cryptographic Groups 93
9.2	The Algebraic Group Model 93
9.3	Subset Membership Problem 94
9.4	Dual-mode NIWI Proof System 95
9.5	Re-Randomizable Public-Key Encryption 97
10	STATISTICALLY CORRECT PIO 99
10.1	Statistically Correct pIO 100
10.2	Puncturable Pseudorandom Functions 102
10.3	Construction 104
11	HOW TO SIMULATE EXTRACTION 109
11.1	Group Schemes and Wrappers 110
11.2	An Algebraic Wrapper 111
12	CONSTRUCTION 115
12.1	Main Theorem and Security Analysis 117
13	SIGNED ELGAMAL 129

8

INTRODUCTION

In this part, we present the results of [AHK20]. This part is taken verbatim from [AHK20] with minor changes.

THE GENERIC GROUP MODEL. In order to analyze the plausibility and relative strength of computational assumptions in cyclic groups, Shoup [Sho97] and Maurer [Mau05] have proposed the *generic group model (GGM)*. In the GGM, any adversary can only interact with the modeled group through an oracle. In particular, all computations in that group must be explicitly expressed in terms of the group operation. To prevent an adversary from locally performing computations, that adversary gets to see only truly random strings (in [Sho97]) or independent handles (in [Mau05]) as representations of group elements.³³

The discrete logarithm and even many Diffie-Hellman-style problems are hard generically (i. e., when restricting group operations in the above way) [Sho97; MW98]. Hence, the only way to break such a generically hard assumption in a concrete group is to use the underlying group representation in a nontrivial way. In that sense, the GGM can be very useful as a sanity check for the validity of a given assumption, or even the security of a given cryptographic scheme. However, generic groups cannot be implemented: there exist cryptographic schemes that are secure in the GGM, but insecure when instantiated with *any* concrete group [Den02].

THE ALGEBRAIC GROUP MODEL. The algebraic group model (AGM), introduced in [FKL18], is a relaxation of the GGM that tries to avoid impossibilities as in [Den02] while preserving the GGM's usefulness. Specifically, the AGM only considers *algebraic* (rather than generic) adversaries. An algebraic adversary \mathcal{A} can make arbitrary use of the representation of group elements, but must supply an explicit decomposition for any of its output group elements in terms of input group elements. In other words, \mathcal{A} must also output an explanation of how any group element in its output was computed from its input using the group operation.

Now [FKL18] show that many GGM proofs only use this type of algebraicity of an adversary, and carry over to the AGM. At the same time, GGM impossibilities like [Den02] do not apply to the AGM, since algebraic adversaries are able to work with the actual group (and not only with random or abstract representations of group elements).

THE AGM AND KNOWLEDGE-TYPE ASSUMPTIONS. The AGM is closely related to the notions of knowledge-type assumptions and extractability. To illustrate, assume that for any (possibly non-algebraic) adversary \mathcal{A} , we can find an extractor \mathcal{E} that manages to extract from \mathcal{A} a decomposition of \mathcal{A} 's output in terms of \mathcal{A} 's input. Then, composing \mathcal{E} and \mathcal{A} yields an algebraic adversary \mathcal{A}_{alg} . In this situation, we can then say that without loss of generality, any adversary can be assumed to be algebraic.³⁴ Conversely,

³³ Other black-box abstractions of groups with similar ramifications exist [Nec94; BL96].

³⁴ This observation about algebraic adversaries has already been made in [BV98; PV05]. Also, similar but more specific knowledge-type assumptions have been used to prove concrete cryptographic constructions secure, e. g., [Dam92; HT98; BPO4; Den06].

any algebraic adversary by definition yields the results of such an extraction in its output.

This observation also provides a blueprint to *instantiating* the **AGM**: simply prove that any adversary \mathcal{A} can be replaced by an algebraic adversary \mathcal{A}_{alg} , possibly using an extraction process as above. If this extraction requires \mathcal{A} 's code and randomness but no other trapdoor, we obtain an **AGM** instantiation based on a knowledge-type assumption such as the knowledge-of-exponent assumption [Dam92]. Indeed, this was recently done by [KP19] under a very strong generalized version of the knowledge-of-exponent assumption [WS07]. Unfortunately, such knowledge-type assumptions are not falsifiable in the sense of Naor [Nao03] making them (arguably) undesirable for use in cryptography. Furthermore, it is not entirely clear how to assess the plausibility of such a universal and strong knowledge-type assumption. Naturally, the question arises whether an **AGM** implementation inherently requires such strong and non-falsifiable assumptions. Or, more generally:

*Can we achieve knowledge-type properties
from falsifiable assumptions?*

Note that in the **AGM**, the discrete logarithm assumption implies the existence of extractable one-way functions (**EOWFs**) with unbounded auxiliary input. The existence of such **EOWFs**, however, conflicts with the existence of indistinguishability obfuscation [BCPR16]. Due to this barrier, we can only hope for an instantiation of some suitably relaxed variant of the **AGM** from falsifiable assumptions.

OUR STRATEGY: PRIVATE EXTRACTION. Since extraction as defined in the **AGM** inherently requires knowledge-type assumptions, we need to conceive a different extraction mechanism. Namely, we show that it is possible to extract a decomposition of \mathcal{A} 's outputs *from these outputs* if a suitable (secret) extraction trapdoor is known. In other words, our idea is to avoid non-falsifiable knowledge-type assumptions by assuming that extraction requires a special trapdoor which can be generated alongside the public parameters of the group. This entails a number of technical difficulties (see below), but allows us to rely entirely on falsifiable assumptions.

Specifically, our main result is an *algebraic wrapper* that transforms a given cyclic group into a new one which allows for an extraction of representations. More specifically, an element of the new group carries an encrypted representation of this group element relative to a fixed set of group elements (henceforth denoted “basis”). The group operation guarantees that newly produced group elements carry decompositions which are consistent with the decompositions contained in the input group elements. Further, the corresponding decryption key (which can be generated alongside the public parameters) serves as a trapdoor which allows to extract such decompositions.

OUR RESULTS. Our algebraic wrapper allows us to retrieve several **AGM** results (from [FKL18; FPS20]) in the standard model, in the sense that the group can be concretely implemented from falsifiable assumptions.³⁵ In particular, we show that in the algebraic wrapper,

³⁵ Note that by “standard model”, we mean that the group itself is formulated without idealizations and can be concretely implemented. While our construction itself does not rely on the **ROM**, we still can transfer some **ROM** proofs in the **AGM** to **ROM** proofs using our concrete group instantiation. We stress that a standard model instantiation of the (full-fledged) **AGM** from very strong non-falsifiable assumptions is already known due to [KP19].

- (i) the discrete logarithm assumption, the computational Diffie-Hellman assumption, the square Diffie-Hellman assumption, and the linear-combination Diffie-Hellman assumption (see [FKL18]) are all equivalent,
- (ii) the security of the Schnorr signature scheme [Sch91] can be *tightly* reduced to the discrete logarithm assumption escaping impossibility results due to [FJS19],
- (iii) the IND-CCA2 security of Schnorr-signed ElGamal [Jak98; TY98] can be *tightly* reduced to the discrete logarithm assumption.³⁶

We stress that Items **i** and **ii** are not the contribution of this author and are hence not included in this thesis. Item **iii** is detailed in Chapter 13 and the conception and instantiation of the algebraic wrapper are detailed in Chapters 11 and 12, respectively. We refer the reader to the original publication [AHK20] for more details on Items **i** and **ii**.

While, on a technical level, the AGM proofs from [FKL18; FPS20] need to be adapted by adding some preparatory steps, the general AGM proof strategies can be replicated.

LIMITATIONS. We note that not all known AGM proofs can be transported to the standard model using the above strategy. For instance, [FKL18] also prove the Boneh-Lynn-Shacham (BLS) [BLS04] signature scheme tightly secure in the AGM. Their reduction relies on the fact that the view of a signature forger is statistically independent of how simulated signatures are prepared by the reduction. However, with our algebraic wrapper, group elements (and thus BLS signatures) always carry an encrypted representation. This representation allows to infer information on how these group elements were generated. This additional information is problematic in the AGM-based proof of BLS signatures from [FKL18]. We believe it is an interesting open problem to obtain a tight security proof for the BLS scheme with our group.³⁷

Furthermore, the amount of information which can be carried by a group element is limited by the size of that group element. In particular, in settings in which no a-priori bound on the size of a desired algebraic representation is known, our techniques do not apply. This can be problematic, e. g., for constructions that depend on q-type assumptions.

OUR ASSUMPTIONS. Our algebraic wrapper relies on a strong but falsifiable computational assumption: the existence of subexponentially strong indistinguishability obfuscation.³⁸ Additionally, we assume a re-randomizable encryption scheme. Together with subexponential IO, this implies a number of other strong primitives that we use: a variant of statistically correct probabilistic IO (see [CLTV15] and Chapter 10), fully homomorphic encryption (see [CLTV15]), and dual-mode non-interactive zero-knowledge (see [HU19]).

INTERPRETATION. Due to their inefficiency, we view algebraic wrappers not as a tool to obtain practical cryptographic primitives. Rather, we believe that algebraic wrappers show that the AGM is a useful abstraction and proof strategies developed for the AGM can in fact be replicated in the standard model, and even without resorting to knowledge-type assumptions.

³⁶ Tight security reductions provide a tight relation between the security of cryptographic schemes and the hardness of computational problems. Apart from their theoretical importance, tight reductions are also beneficial for practice, since they allow smaller keylength recommendations.

³⁷ Note that impossibility results for tight reductions of schemes like BLS (e. g., [Coroo]) do not apply in our case, as the representation of our group elements is not unique.

³⁸ We note that IO and certain knowledge-type assumptions (which require universal extraction) contradict each other [BCPR16]. However, we stress that the notion of private extractability we obtain does not contradict IO.

ON IMPLEMENTING IDEALIZED MODELS. Replacing idealized (heuristic) models with concrete standard-model implementations is a widely studied intriguing problem. A well-known example for this is the line of work on programmable hash functions. A programmable hash function due to [HK12] is a cryptographic primitive which can be used to replace random oracles in several cryptographic schemes. Following their introduction, a line of work [FHPS13; HSW13; HSW14] leveraged multilinear maps or indistinguishability obfuscation to transport proofs from the random oracle model to the standard model. Our results can be interpreted as following this endeavor by leveraging indistinguishability obfuscation to replace the [AGM](#) with a standard model implementation (from falsifiable assumptions). From this angle, our algebraic wrapper relates to the [AGM](#) as programmable hash functions relate to the [ROM](#).

8.1 TECHNICAL OVERVIEW

ALGEBRAIC WRAPPERS. We employ the notion of *group schemes* [AFH+16] (also called encoding schemes in [GGH13]) as a generalization of groups with potentially non-unique encodings of group elements. In such a group scheme, a dedicated algorithm is required to determine if two given group element encodings are equal. That is, the group is defined as the quotient set of all well-formed encodings modulo the equivalence relation induced by the equality test. Our algebraic wrapping process takes group parameters of G (which we refer to as “base group”) as input, and outputs group parameters of a new group scheme \mathbb{H} which allows for an efficient extraction process. To be specific, every \mathbb{H} -element \hat{h} can be viewed as a G -element $h \in G$, plus auxiliary information *aux*.

The core advantage of group schemes with non-unique encodings lies in the ability to let group elements carry such auxiliary information *aux*. In our algebraic wrapper, *aux* contains encrypted information expressing h as a linear combination of fixed base group elements $b_1, \dots, b_n \in G$. The corresponding decryption key (which can be generated alongside the group parameters) allows to extract this information, and essentially yields similar information any algebraic adversary (in the sense of the [AGM](#)) would have to provide for any output group element. However, we are facing a number of technical problems:

- (a) Upon creation of each new group element encodings, a valid representation with respect to the basis elements b_1, \dots, b_n needs to be embedded.
- (b) It should be possible to unnoticeably switch the basis elements b_1, \dots, b_n to an application-dependent basis. (For instance, in reductions to the discrete logarithm problem $dlog_g(h)$, one would desire to set the basis to g, h .)
- (c) To preserve tightness of reductions from the [AGM](#) (which is necessary for some of our applications), it should be possible to re-randomize group element encodings *statistically*.

Our solution largely follows the group scheme from [AFH+16]. In particular, we address Item [a](#) by requiring each group element encoding $\hat{h} \in \mathbb{H}$ to carry

an encrypted linear combination z_1, \dots, z_n of $\mathbb{G} \ni h = \sum_i b_i^{z_i}$ relative to the basis and force validity of this linear combination using a consistency proof. Consequently, sampling of new group elements requires explicit knowledge of such a linear combination. This is also true for the group operation. More precisely, the group operation may produce the ciphertext for the new group element encoding via homomorphic evaluation of the input ciphertexts but is required to extract the contained linear combination to produce the consistency proof. Note that the adversary must be able to evaluate the group operation and needs to produce corresponding consistency proofs. This process, however, particularly the involved random coins must not be under adversarial control. Thus, the group operation is obfuscated using probabilistic IO.

For Item [b](#), we adapt a “switching” lemma from [\[AFH+16\]](#). In [\[AFH+16\]](#), that lemma allows to switch between two different representations of the same group element, but under a fixed basis. In our case, we show that similar techniques allow to additionally switch the group elements that form this basis. This issue is dealt with by additional linear randomization of the base group elements. Furthermore, since the basis is computationally hidden, sampling group element encodings which are tailored to contain a certain linear combination requires knowledge of a secret trapdoor.

Note that the switching property already implies a notion of computational re-randomizability. For Item [c](#), we introduce a re-randomization lemma using techniques from Item [b](#) in conjunction with a novel notion for *statistically correct* probabilistic IO.

At this point, one main conceptual difference to the line of work [\[AFH+16; AH18; FHHL18; AMP20\]](#) is that the basis elements b_1, \dots, b_n appear as part of the functionality of the new group scheme \mathbb{H} , not only in a proof.

Another main conceptual difference to [\[AFH+16; AH18; FHHL18; AMP20\]](#) is the notion of statistical re-randomizability of group elements. The group schemes from [\[AFH+16; AH18; FHHL18; AMP20\]](#) do not satisfy this property. This will be resolved by developing a stronger notion of *statistically correct* probabilistic IO which may be of independent interest.

APPLICATIONS. The applications of our algebraic wrapper have already been considered for the [AGM](#) in [\[FKL18; FPS20\]](#). Hence, in this description, we focus on the technical differences that our extraction approach entails for these proofs. Further, we focus on the applications from [\[AHK20\]](#) which are the contributions of this author.

Recall that in the [AGM](#) by [\[FKL18\]](#), an adversary outputs an algebraic representation of each output group element to the basis of its input group elements. Therefore, the choice of this basis depends on the application.

The authors of [\[FPS20\]](#) prove the tight security of Schnorr-signed ElGamal based on the discrete logarithm assumption (in the [AGM](#) and [ROM](#)). Schnorr-signed ElGamal uses a public key of the form $(g, Y := g^y)$, where y is the corresponding secret key. Encryption chooses a random exponent x and uses a (random oracle) hash $H'(Y^x)$ to blind the message. Further, a Schnorr signature relative to the public key (g, g^x) is used to certify g^x . This procedure ensures plaintext-awareness [\[BR95\]](#). Hence, a ciphertext consists of a key $H'(Y^x)$ and the group element g^x and a Schnorr signature $(s := r + H(g^r, g^x) \cdot x \bmod p, g^r)$ for a random exponent r .

The reduction from [FPS20] ultimately needs to solve a discrete logarithm challenge $d\log_g(Y)$. We choose the basis for the algebraic wrapper as $\{g, Y\}$ in order to enable extraction relative to both g and Y and to allow for (private) sampling of wrapper group elements which depend on Y (whose discrete logarithm must not be known during a reduction to the discrete logarithm problem). For the purpose of reproducing the proof of [FPS20], we first make a conceptual change as in [FPS20] to choose x depending on y (note that at this point of time, y is still known to the experiment). Namely, we use a publicly sampled group element encoding for g^r to query the random oracle $H(g^r, g^x)$ but do not include this encoding in the signature. Instead, we sample a fresh encoding for the group element $g^{s-H(g^r, g^x)\cdot x}$ and use this in place of the encoding of g^r . Since this encoding represents the same group element, the view of the adversary remains unchanged. Then, we apply switching to use the basis $\{g, Y\}$ and sample all encodings using the second decomposition entry instead of explicitly using y in the first decomposition entry. These preparations pave the way to conclude the proof following the lines of [FPS20].

Another technical challenge is to information-theoretically hide whether group elements were sampled directly or derived using the group operation. We achieve this by exploiting the statistical re-randomizability of group elements. This technique is described in this thesis and is required for the tight security proof of Schnorr signatures which is not part of this thesis.

8.2 RELATED WORK

This part builds upon the line of work [AFH+16; AH18; FHHL18; AMP20] who build group schemes from IO. Albrecht et al. [AFH+16] lay the conceptual foundations for the construction of group schemes with non-unique encodings from IO and use this framework to equip groups with multilinear maps. [FHHL18] extend this approach by allowing partial evaluations of the multilinear map yielding a graded encoding scheme. [AMP20] extend [AFH+16; FHHL18] such that computational assumptions in the base group hold, thereby proving equivalence between IO and multilinear maps (modulo subexponential reductions and relatively mild assumptions). In contrast to [AFH+16; FHHL18], the authors of [AH18] do not extend the functionality of an underlying group, but build a group scheme with reduced functionality (group elements lack a unique representation). The resulting group scheme allows to mimic commonly used proof techniques from the generic group model. This is demonstrated by proving the validity of an adaptive variant of the Uber assumption family [Boyo8] in the constructed group scheme. Our results can hence be viewed as an extension of [AH18].

[KP19] make a first step towards instantiating the AGM. The authors identify an equivalence between the AGM and a very strong generalized version of the knowledge-of-exponent assumption [Dam92], thus giving rise to the first instantiation of the AGM.

9

PRELIMINARIES

In this chapter, we introduce the necessary preliminaries for this part. Since this part is dedicated to the study of cryptographic groups, we start by introducing additional notations for this purpose (cf. Section 9.1). In Section 9.2 we define the algebraic group model. Subsequently, we introduce useful cryptographic primitives such as subset membership problems in Section 9.3 and dual-mode non-interactive witness indistinguishable proofs in Section 9.4 which are necessary for our construction.

9.1 NOTATIONS AND CRYPTOGRAPHIC GROUPS

Let \mathbb{K} denote a field and let V be a vector space over \mathbb{K} of finite dimension n . For $i \in [n]$, e_i denotes the unit vector, i. e., the vector which carries 1 in its i -th entry and 0 in all other entries. Unless stated otherwise, we work with column vectors.

Let G denote a finite cyclic group with generator g and order p . For $x \in \mathbb{Z}_p$, the notation $[x]_G$ denotes the group element g^x . Note that using this notation does not imply knowledge of x . The group G is associated with a group parameter generation algorithm $G\text{Gen}$. On input of a security parameter (in unary), $G\text{Gen}$ samples the parameters pp_G fully describing the group, such that the group operation can be computed in time polynomial in the security parameter. We define hardness assumptions always relative to a group parameter generation algorithm.

DEFINITION 9.1 (Discrete logarithm assumption). The *discrete logarithm (DL) assumption* holds relative to $G\text{Gen}$ if for all **PPT** adversaries \mathcal{A} ,

$$\text{Adv}_{G\text{Gen}, \mathcal{A}}^{\text{dl}}(\lambda) := \Pr \left[\begin{array}{l} pp_G \leftarrow G\text{Gen}(1^\lambda) \\ x \leftarrow \mathbb{Z}_p \end{array} : \mathcal{A}(pp_G, g^x) = x \right]$$

is negligible in λ .

9.2 THE ALGEBRAIC GROUP MODEL

The algebraic group model was introduced in [FKL18] as a model that lies between the generic group model (**GGM**) [Sho97; Mau05] and the standard model. The algebraic group model is weaker than the generic group model in the sense that in the **AGM**, an adversary has explicit access to (encodings of) group elements and may very well derive information from them. On the other hand, the algebraic group model still requires that new group elements can only be produced by generic group operations. This is formalized by the ability of the adversary to present a decomposition of all newly produced group elements relative to given group elements.

Formally, [FKL18] only quantify over the set of *algebraic adversaries* (which can explain how they produced new group elements).

DEFINITION 9.2 (Algebraic algorithm (informal), [FKL18]). An algorithm \mathcal{A} is *algebraic* if for all group elements h which \mathcal{A} outputs, it additionally provides a linear decomposition of h relative to all input group elements (only using generic group operations).

[KP19] introduce an equivalent notion of *algebraic groups* where every adversary is bound to be algebraic. By this view of the algebraic group model boils down to an extremely strong knowledge-type assumption since every adversary who outputs a group element *must know* how he computed it. More precisely, [KP19] prove that the **AGM** is equivalent to an extremely strong generalized version of the knowledge-of-exponent assumption. That is, the **AGM** (formulated as an assumption³⁹) does not enjoy efficient falsifiability due to Naor [Na03].

³⁹ “There exists a group, such that all adversaries in this group are algebraic.”

Further, the **AGM** as in [FKL18] allows algebraic adversaries to receive “non-group-element inputs” (which are required to be independent of group elements). However, these “non-group-element inputs” to the adversary may contain an obfuscation of a program which on input of a group element $[1]_G$ outputs $[1]_G^{F(K, [1]_G)} = [F(K, [1]_G)]_G$ (where F is a pseudorandom function and K is a random **PRF** key). Clearly, this program description is independent of group elements (only the input to the program depend on group elements). An adversary who evaluates this obfuscated program on, e. g., the generator of the group and outputs the resulting group element is unable to explain how he computed this group element. In other words, the **AGM** (defined via adversarial knowledge as in [KP19]) allowing arbitrarily long unstructured “non-group-element inputs” implies extractable one-way functions with unbounded auxiliary information (cf. Definition 2.21) which are known to conflict with indistinguishability obfuscation [BCPR16].

9.3 SUBSET MEMBERSHIP PROBLEM

DEFINITION 9.3 (Subset membership problem, [CS02]). A hard *subset membership problem* (**SMP**) **SMP** is defined by an **SMP** instance generator Gens_{SMP} .

- Gens_{SMP} , on input of 1^λ , samples a description of a language and a universe $L \subseteq X$ together with a relation R , such that for all $x \in X$,

$$x \in L \Leftrightarrow \exists! w \in \{0, 1\}^{\text{poly}(\lambda)} : (x, w) \in R.$$

We require that uniform sampling from L (together with the unique witness⁴⁰) and uniform sampling from $X \setminus L$ is efficiently possible. Further, we require that for every **PPT** adversary \mathcal{A} , the advantage

⁴⁰ Note that unique witness relations are related to the complexity class unambiguous polynomial-time (**UP**) [Val76].

$$\text{Adv}_{\text{Gens}_{\text{SMP}}, \mathcal{A}}^{\text{smp}}(\lambda) := \left| \Pr \left[\begin{array}{l} (L, X, R) \leftarrow \text{Gens}_{\text{SMP}}(1^\lambda) \\ x \leftarrow L \end{array} : \mathcal{A}(1^\lambda, x) = 1 \right] - \Pr \left[\begin{array}{l} (L, X, R) \leftarrow \text{Gens}_{\text{SMP}}(1^\lambda) \\ x \leftarrow X \setminus L \end{array} : \mathcal{A}(1^\lambda, x) = 1 \right] \right|$$

is negligible in λ .

Let G be a cyclic group (of prime order) with group parameter generation algorithm $G\text{Gen}$ in which **DDH** is assumed to hold relative to $G\text{Gen}$. The Diffie-Hellman language is a hard subset membership problem. The corresponding **SMP** instance generator Gen_{SMP} samples $pp_G \leftarrow G\text{Gen}$ and a uniformly random exponent $x \leftarrow \mathbb{Z}_p$, and outputs the language and the universe

$$L_{pp_G, [(1, x)]_G} := \{ [(y, xy)]_G \mid y \in \mathbb{Z}_p \} \subseteq G \times G = X_{pp_G}$$

and the relation

$$R_{pp_G, x} := \{ [(a, b)]_G \mid b = x \cdot a \pmod p \}.$$

Another instantiation of Definition 9.3 is the language containing all commitments to a fixed value using a perfectly binding non-interactive commitment scheme with unique opening. The universe X is then defined as the set of all valid commitments.

9.4 DUAL-MODE NIWI PROOF SYSTEM

A dual-mode non-interactive witness indistinguishable (**NIWI**) proof system is a variant of general **NIWI** proofs [FS90] additionally offering two computationally indistinguishable modes to setup the common reference string (**CRS**). A **CRS** generated in binding mode provides perfect soundness guarantees whereas a **CRS** generated in hiding mode provides perfect witness indistinguishability guarantees.

Note that **NIWI** proof systems do not inherently require a trusted setup such as a common reference string, e. g., [BP15a]. These instantiations, however, do not satisfy perfect soundness and perfect witness-indistinguishability simultaneously. For our purpose, we require a notion of **NIWI** proof system which offers two indistinguishable modes of operation: one satisfying perfect soundness and the other satisfying perfect witness-indistinguishability.

DEFINITION 9.4 (Dual-mode **NIWI** proof system, [GS08; AFH+16]). A *dual mode non-interactive witness-indistinguishable (**NIWI**) proof system* for a relation \mathcal{R} is a tuple of **PPT** algorithms $\text{NIWI} = (\text{Setup}, \text{HSetup}, \text{Prove}, \text{Verify}, \text{Ext})$.

- $\text{Setup}(1^\lambda)$ outputs a *perfectly binding* common reference string σ and a corresponding extraction trapdoor td_{ext} .
- $\text{HSetup}(1^\lambda)$ outputs a *perfectly hiding* common reference string σ .
- Prove , on input of a **CRS** σ , a statement x and a corresponding witness w , produces a proof π .
- Verify , on input of a **CRS** σ , a statement x and a proof π , deterministically outputs 1 if the proof is valid and 0 otherwise.
- Ext , on input of the extraction trapdoor td_{ext} , a statement x and a proof π , outputs a witness w .

We require **NIWI** to satisfy the following properties.

CRS INDISTINGUISHABILITY. For all **PPT** adversaries \mathcal{A} ,

$$\text{Adv}_{\text{NIWI}, \mathcal{A}}^{\text{crs-ind}}(\lambda) := \left| \Pr \left[(\sigma, td_{\text{ext}}) \leftarrow \text{Setup}(1^\lambda): \mathcal{A}(1^\lambda, \sigma) = 1 \right] - \Pr \left[\sigma \leftarrow \text{HSetup}(1^\lambda): \mathcal{A}(1^\lambda, \sigma) = 1 \right] \right|$$

is negligible.

PERFECT COMPLETENESS. For all $\lambda \in \mathbb{N}$, all **PPT** adversaries \mathcal{A} , we have that

$$\Pr \left[\begin{array}{l} (\sigma, \cdot) \leftarrow \text{Setup}(1^\lambda) \\ (x, w) \leftarrow \mathcal{A}(\sigma) \\ \pi \leftarrow \text{Prove}(\sigma, x, w) \end{array} : \begin{array}{l} (x, w) \notin \mathcal{R} \\ \vee \\ \text{Verify}(\sigma, x, \pi) = 1 \end{array} \right] = 1,$$

where the probability is over the random coins of Setup, Prove and \mathcal{A} . The same holds for HSetup. In other words, every proof π produced via Prove for a valid statement-witness-pair will be accepted by Verify.

PERFECT SOUNDNESS IN BINDING MODE. For all $\lambda \in \mathbb{N}$ and all **PPT** adversaries \mathcal{A} , we have that

$$\Pr \left[\begin{array}{l} (\sigma, \cdot) \leftarrow \text{Setup}(1^\lambda) \\ (x, \pi) \leftarrow \mathcal{A}(\sigma) \end{array} : x \in L \wedge \text{Verify}(\sigma, x, \pi) = 1 \right] = 0,$$

where the probability is over the random coins of Setup and \mathcal{A} . In other words, all statements $x \notin L$, there exists no proof such that Verify accepts.

PERFECT EXTRACTABILITY IN BINDING MODE. For all $\lambda \in \mathbb{N}$ and all **PPT** adversaries \mathcal{A} , we have

$$\Pr \left[\begin{array}{l} (\sigma, td_{ext}) \leftarrow \text{Setup}(1^\lambda) \\ (x, \pi) \leftarrow \mathcal{A}(\sigma) \end{array} : \begin{array}{l} \text{Verify}(\sigma, x, \pi) = 0 \\ \vee \\ (x, \text{Ext}(td_{ext}, x, \pi)) \in \mathcal{R} \end{array} \right] = 1,$$

where the probability is over the random coins of Setup, \mathcal{A} and Ext. In other words, from every accepting statement-proof-pair (x, π) , the extraction algorithm Ext extracts a witness w such that $(x, w) \in \mathcal{R}$.

PERFECT WITNESS-INDISTINGUISHABILITY IN HIDING MODE. For every $\lambda \in \mathbb{N}$, all $\sigma \in \text{supp}(\text{HSetup}(1^\lambda))$, all $(x, w_0), (x, w_1) \in \mathcal{R}$ (of arbitrary but polynomial length in λ), the output distributions induced by

$$\text{Prove}(\sigma, x, w_0) \quad \text{and} \quad \text{Prove}(\sigma, x, w_1)$$

are identical.

There are several instantiations of dual-mode **NIWI** proof systems satisfying the above definition. The construction [GS08] relies on pairing-friendly groups where either the pairing is asymmetric and the **SXDH** assumption holds, or the pairing is symmetric and the **DLIN** assumption holds. [PS19] build a (statistically secure, statistically extractable and statistically zero-knowledge) dual-mode **NIZK** from the learning with errors assumption. The recent paper [HU19] proposes an instantiation of (a statistically secure, statistically extractable and statistically witness-indistinguishable) dual-mode **NIWI** proof system based on strong but mostly unstructured assumptions, namely, subexponentially secure indistinguishability obfuscation, subexponentially secure one-way functions and a secure lossy encryption scheme. This instantiation is sufficient for our purpose, but, for simplicity, we work with dual-mode **NIWI** proof systems providing perfect guarantees in our proofs in Chapter 12.

9.5 RE-RANDOMIZABLE PUBLIC-KEY ENCRYPTION

A re-randomizable public-key encryption scheme allows to perfectly re-randomize any ciphertext. This can be a useful property in scenarios, where it is necessary to *information-theoretically* hide the randomness which was used for ciphertext generation.

DEFINITION 9.5 (Perfectly re-randomizable public-key encryption, [PR07]). A perfectly re-randomizable public-key encryption scheme with message space $\{0, 1\}^*$ is a tuple of PPT algorithms $E = (\text{KGen}, \text{Enc}, \text{Dec}, \text{Rerand})$ such that $(\text{KGen}, \text{Enc}, \text{Dec})$ is a perfectly correct PKE scheme such that the following additional properties are met.

- For all $(pk, \cdot) \in \text{supp}(\text{KGen}(1^\lambda))$, all messages $m \in \{0, 1\}^*$, all ciphertexts $C \in \text{supp}(\text{Enc}(pk, m))$, $\text{Rerand}(C)$ is distributed identically to $\text{Enc}(pk, m)$.
- For all $(pk, \cdot) \in \text{supp}(\text{KGen}(1^\lambda))$, all (maliciously chosen) ciphertexts C , and all $C' \in \text{supp}(\text{Rerand}(C))$, $\text{Dec}(sk, C') = \text{Dec}(sk, C)$.

For our purposes it suffices to let $E.\text{Rerand}$ receive the public key as input. Furthermore, in contrast to [PR07], we do not require that with overwhelming probability over the choice of $(pk', \cdot) \leftarrow E.\text{KGen}(1^\lambda)$, the ciphertext spaces under pk and pk' are disjoint, i. e., $\text{supp}(E.\text{Enc}(pk, \cdot)) \cap \text{supp}(E.\text{Enc}(pk', \cdot)) = \emptyset$. The ElGamal encryption scheme [EIG85], the Goldwasser-Micali encryption scheme [GM84], the Paillier encryption scheme [Pai99] and the Damgård-Jurik encryption scheme [DJ01] are perfectly re-randomizable public-key encryption schemes, [PR07; HLOV11].

It is currently unclear, whether lossy PKE schemes (and hence re-randomizable PKE schemes, [HLOV11]⁴¹) can be built from IO in conjunction with unstructured generic assumptions like OWFs.

Note that the FHE construction due to Canetti et al. [CLTV15] as well as our FHE construction from Part I preserve re-randomizability of the underlying PKE scheme. That is, any perfectly re-randomizable PKE scheme can be converted into a perfectly re-randomizable FHE scheme. Such a FHE scheme will be necessary for our construction in Chapter 12.

⁴¹ [HLOV11] build a lossy encryption scheme from a re-randomizable encryption scheme, where an “injective” public key contains a ciphertext C_0 of 0 and a ciphertext C_1 of 1 and a “lossy” public key contains two ciphertexts C_0, C_1 of 0. Encryption of a bit b simply re-randomizes the ciphertext C_b and decryption works as in the underlying scheme.

10

STATISTICALLY CORRECT PIO

Probabilistic IO allows to obfuscate a randomized computation such that the user of the obfuscated code is unable to control the used randomness. Moreover, the randomness used by the obfuscated circuit is *pseudorandom* from the view of the circuit user – even though he possesses the code. Hence, probabilistic IO can be used to obfuscate a randomized circuit which produces a ciphertext of some plaintext which should remain hidden from the circuit user. In other words, pIO enables to protect the randomness of some probabilistic process from its user. Consequently, pIO fills an important gap which is left open by (deterministic) indistinguishability obfuscation.

On the other hand, it may be undesirable to leak information on *how* some output was computed. In particular, one may be interested to hide whether the obfuscated circuit or the original randomized circuit (with fresh randomness) was used. Probabilistic IO hides this information against computationally bounded observers. However, since the obfuscated program is deterministic, it is information-theoretically evident whether the original randomized circuit or the obfuscated (deterministic) circuit was used. This is an inherent limitation of probabilistic IO. However, there are scenarios, where the origin of some output is required to be *information-theoretically* hidden.⁴² A natural question that arises is whether this is feasible.

In the following, we define a notion of *statistically correct* probabilistic IO and provide an instantiation of this notion. More precisely, we require statistical closeness between evaluations of the original (probabilistic) circuit and the obfuscated (deterministic) circuit. By definition of probabilistic IO, this is impossible since the obfuscated circuit is deterministic and, hence, has no source of entropy other than its input. However, we observe that as long as a portion of the circuit’s input is guaranteed to be outside the view of the adversary (and has sufficiently high min-entropy), the output of the obfuscated circuit and the actual probabilistic circuit can be statistically close.

For this purpose, we extend the definition of probabilistic IO as follows. Before obfuscating a probabilistic circuit, we compile this circuit such that it receives an “auxiliary input” *aux* but simply ignores this input in its computation. Even though the obfuscated circuit is deterministic, this auxiliary input can be used as a source of *actual entropy*.

FIRST TRY. We recall that the pIO construction from [CLTV15] obfuscates a probabilistic circuit C by using IO to obfuscate the deterministic circuit $\bar{C}(x) := C(x; F_K(x))$. A natural idea to achieve statistical correctness is to modify this construction such that the auxiliary input *aux* is directly XORed on the random tape which is derived using F , i.e. to obfuscate the circuit $\bar{C}(x, aux; F_K(x) \oplus aux)$. For uniform auxiliary input *aux*, statistical correctness follows immediately. However, security breaks down. Consider two circuits C_1 and C_2 such that C_1 outputs the first bit on its random tape and C_2 outputs the second bit on its random tape. Since C_1 and C_2 produce identical output distributions, it is desirable that a probabilistic indistinguishability

⁴² One such scenario is described and studied later in Chapter 12.

obfuscator conceals which of the two circuits was obfuscated. However, this construction admits a successful attack. An adversary can evaluate the obfuscated circuit Λ on inputs (x, aux) and $(x, aux \oplus 1)$. If both evaluations yield identical outputs, C_2 was obfuscated, otherwise C_1 was obfuscated.

USING AN EXTRACTING PRF. Our construction of statistically correct **PIO** applies an extracting puncturable **PRF** on the entire input of the circuit to derive the random tape for the probabilistic circuit. An extracting **PRF** guarantees that **PRF** outputs are uniformly distributed (even given the **PRF** key) as long as the input has high min-entropy. This is achieved using a universal hash function and the leftover hash lemma.

OUTLINE. In Section 10.1, we formally define statistically correct **PIO**. In Section 10.2 we introduce preliminaries which are necessary for our instantiation. Finally, in Section 10.3, we provide an instantiation of statistically correct **PIO**.

10.1 STATISTICALLY CORRECT PIO

In this section, we formally define our novel notion of statistically correct input-expanding **PIO**. We start by defining expanding compilers and expanding circuit samplers.

DEFINITION 10.1 (ℓ -expanding compiler). An ℓ -expanding compiler \mathcal{E}_ℓ takes as input a probabilistic circuit C of size $p'(\lambda)$, expecting inputs $x \in \{0, 1\}^{n'(\lambda)}$ and randomness $r \in \{0, 1\}^{m(\lambda)}$, and outputs a circuit \widehat{C} of polynomial size $p(\lambda) = p'(\lambda) + \ell(\lambda)$ expecting inputs $(x, aux) \in \{0, 1\}^{n'(\lambda)} \times \{0, 1\}^{\ell(\lambda)}$, randomness $r \in \{0, 1\}^{m(\lambda)}$ such that for all $x \in \{0, 1\}^{n'(\lambda)}$, all $aux \in \{0, 1\}^{\ell(\lambda)}$ and all $r \in \{0, 1\}^{m(\lambda)}$,

$$C(x; r) = \widehat{C}(x, aux; r).$$

The compiler \mathcal{E}_ℓ which simply appends $\ell(\lambda)$ input gates (without any additional edges) to the original circuit satisfies the above definition.

Our construction of **PIO** will first expand the given circuit as above and then use a slightly modified version of the **PIO** scheme of [CLTV15] to obfuscate the expanded circuit. In the following, we formally define the properties of this expanding probabilistic indistinguishability obfuscator. Let $\mathcal{C} = (\mathcal{C}_\lambda)_{\lambda \in \mathbb{N}}$ be an ensemble of sets \mathcal{C}_λ of probabilistic circuits. The set \mathcal{C}_λ contains circuits of size λ with input length $n'(\lambda)$ expecting (at most) $m(\lambda)$ random bits. Let $\widehat{\mathcal{C}} = \{\widehat{\mathcal{C}}_\lambda\}_{\lambda \in \mathbb{N}}$ be the circuit class which results from ℓ -expanding the circuit class \mathcal{C} . That is, $\widehat{\mathcal{C}}_{\lambda + \ell(\lambda)} = \mathcal{E}_\ell(\mathcal{C}_\lambda)$.⁴³

Let D be a circuit sampler for the circuit class \mathcal{C} . Let $\widehat{D} := \{\widehat{D}_\lambda\}_{\lambda \in \mathbb{N}}$ be the corresponding ℓ -expanding circuit sampler. More formally, $\widehat{D}_{\lambda + \ell(\lambda)}$ samples (C_0, C_1, z) from the distribution D_λ and outputs the circuits $\widehat{C}_0 := \mathcal{E}_\ell(C_0)$, $\widehat{C}_1 := \mathcal{E}_\ell(C_1)$ and auxiliary information $\widehat{z} := (C_0, C_1, z)$.⁴⁴ The circuit sampler \widehat{D} is a circuit sampler for the circuit class $\widehat{\mathcal{C}}$.

By the security of **PIO** for the class of $\mathcal{C}^{X\text{-ind}}$ due to [CLTV15], our input-expanding **PIO** scheme is secure for the following class of input-expanding X -ind samplers.

⁴³ Note that this is well-defined since we implicitly fix the function $\ell(\lambda)$ and require that $\ell(\lambda)$ is strictly increasing. This is without loss of generality.

⁴⁴ Note that since the compilation process of \mathcal{E}_ℓ is reversible, we could actually omit C_0, C_1 from \widehat{z} .

DEFINITION 10.2 (ℓ -expanding X -ind sampler). An ℓ -expanding X -ind sampler D for the circuit class \mathcal{C} is a circuit sampler for \mathcal{C} such that its corresponding ℓ -expanding circuit sampler \widehat{D} is an X -ind sampler according to Definition 2.18. We denote the class of all ℓ -expanding X -ind samplers $\mathfrak{C}_\ell^{X-(*)\text{ind}}$.

Unfortunately, not all X -ind samplers D (for \mathcal{C}) induce a circuit sampler \widehat{D} (for $\widehat{\mathcal{C}}$) which is an \widehat{X} -ind sampler. In order for \widehat{D} to satisfy the \widehat{X} -differing inputs property, we need to set the function \widehat{X} such that $\widehat{X}(\lambda + \ell(\lambda)) = X(\lambda) \cdot 2^{\ell(\lambda)} \leq 2^{\lambda + \ell(\lambda)}$, where λ is the security parameter used for D and $\lambda + \ell(\lambda)$ is the security parameter for \widehat{D} .⁴⁵ \widehat{X} -indistinguishability, however, requires that

$$\widehat{X}(\lambda + \ell(\lambda)) \cdot \text{Adv}_{\widehat{D}, \mathcal{A}}^{\text{sel-ind}}(\lambda + \ell(\lambda)) = X(\lambda) \cdot 2^{\ell(\lambda)} \cdot \text{Adv}_{D, \mathcal{A}}^{\text{sel-ind}}(\lambda + \ell(\lambda))$$

is negligible (for all PPT adversaries $\widehat{\mathcal{A}}$). Due to X -indistinguishability of D , we are only guaranteed that $X(\lambda) \cdot \text{Adv}_{D, \mathcal{A}}^{\text{sel-ind}}(\lambda)$ is negligible (for all PPT adversaries \mathcal{A}), which does not suffice to prove \widehat{X} -indistinguishability of \widehat{D} .

Nevertheless, ℓ -expanding X -ind samplers cover a wide class of circuit samplers. In the following we present a lemma which facilitates the usage of ℓ -expanding X -ind samplers.

LEMMA 10.1. Let D be a circuit sampler for \mathcal{C} which outputs two circuits C_0 and C_1 such that for all inputs x , $C_0(x)$ and $C_1(x)$ produce the exact same distribution. Then, $D \in \mathfrak{C}_\ell^{X-(*)\text{ind}}$.

Proof. Let \widehat{D} be the ℓ -expanding circuit sampler corresponding to D . Let $\widehat{X}(\lambda + \ell(\lambda)) := 2^{\lambda + \ell(\lambda)}$ and $\widehat{\mathcal{X}} := \{0, 1\}^{\lambda + \ell(\lambda)}$. Since there is no input for the circuits which is outside of $\widehat{\mathcal{X}}$, the \widehat{X} -differing inputs property is trivially satisfied. Furthermore, as for all inputs x the distributions $C_0(x)$ and $C_1(x)$ are identical, the same holds for the distributions $\mathcal{E}_\ell(C_0)(x, \text{aux})$ and $\mathcal{E}_\ell(C_1)(x, \text{aux})$ for all inputs x and all $\text{aux} \in \{0, 1\}^{\ell(\lambda)}$. Thus, for all adversaries \mathcal{A} , $\text{Exp}_{\widehat{D}, \mathcal{A}}^{\text{sel-ind}}(\lambda + \ell(\lambda)) = \frac{1}{2}$. Hence, $\widehat{D} \in \mathfrak{C}^{X\text{-ind}}$ and, therefore, $D \in \mathfrak{C}_\ell^{X-(*)\text{ind}}$. \square

Lemma 10.1 directly yields the following corollary.

COROLLARY 10.1. Let D be a circuit sampler for \mathcal{C} which outputs two circuits C_0 and C_1 which always behave exactly identically on identical inputs and random tapes. Then, $D \in \mathfrak{C}_\ell^{X-(*)\text{ind}}$.

Our expanding pIO scheme satisfies similar correctness and security properties as defined in [CLTV15] but additionally guarantees statistical correctness.

Let $\mathcal{C} = (\mathcal{C}_\lambda)_{\lambda \in \mathbb{N}}$ be an ensemble of sets \mathcal{C}_λ of probabilistic circuits. The set \mathcal{C}_λ contains circuits of size λ with input length $n'(\lambda)$ expecting (at most) $m(\lambda)$ random bits. Further, let \mathfrak{C} be a class of circuit samplers over \mathcal{C} .

DEFINITION 10.3 (ℓ -expanding pIO for the class of samplers \mathfrak{C}). An ℓ -expanding probabilistic indistinguishability obfuscator for the class of samplers \mathfrak{C} over \mathcal{C} is a uniform PPT algorithm piO_ℓ^* , satisfying the following properties.

INPUT EXPANDING CORRECTNESS. For all PPT adversaries \mathcal{A} and all circuits $C \in \mathcal{C}$,

$$\left| \Pr \left[\mathcal{A}^{\text{OC}(\cdot, \cdot)}(1^\lambda, C) = 1 \right] - \Pr \left[\Lambda \leftarrow \text{piO}_\ell^*(1^\lambda, C): \mathcal{A}^{\text{O}\Lambda(\cdot, \cdot)}(1^\lambda, C) = 1 \right] \right|$$

⁴⁵ Note that \widehat{D} is called with security parameter $\lambda + \ell(\lambda)$ to compensate for the expanded circuit size.

is negligible as a function in λ , where the oracles \mathcal{O}_C and \mathcal{O}_Λ are defined in Figure 10.1 and must not be called twice on the same input (x, aux) .



Figure 10.1: The oracles used in the definition of correctness of ℓ -expanding pIO.

SECURITY WITH RESPECT TO \mathcal{C} . For all circuit samplers $D \in \mathcal{C}$, for all PPT adversaries \mathcal{A} , the advantage

$$\text{Adv}_{\text{piO}_\ell^*, D, \mathcal{A}}^{\text{pio-ind}(\star)}(\lambda) := \left| \Pr \left[(C_0, C_1, z) \leftarrow D(1^\lambda): \mathcal{A}(C_0, C_1, z, \text{piO}_\ell^*(1^\lambda, C_0)) = 1 \right] - \Pr \left[(C_0, C_1, z) \leftarrow D(1^\lambda): \mathcal{A}(C_0, C_1, z, \text{piO}_\ell^*(1^\lambda, C_1)) = 1 \right] \right|$$

is negligible in λ .

SUPPORT RESPECTING. Let $n'(\lambda)$ denote the input length of circuits in \mathcal{C}_λ . For all circuits $C \in \mathcal{C}_\lambda$, all inputs $x \in \{0, 1\}^{n'(\lambda)}$, all $aux \in \{0, 1\}^{\ell(\lambda)}$, all $\Lambda \in \text{supp}(\text{piO}_\ell^*(1^\lambda, C))$, we have that

$$\Lambda(x, aux) \in \text{supp}(C(x)).$$

STATISTICAL CORRECTNESS WITH ERROR $2^{-e(\lambda)}$. For all circuits $C \in \mathcal{C}_\lambda$ and all joint distributions (Z_1, Z_2) over $\{0, 1\}^{n'(\lambda)} \times \{0, 1\}^{\ell(\lambda)}$ with average min-entropy $\ell(\lambda) \geq \tilde{H}_\infty(Z_2 | Z_1) > m(\lambda) + 2e(\lambda) + 2$, the statistical distance

$$\Delta \left(\left\{ \Lambda \leftarrow \text{piO}_\ell^*(1^\lambda, C): (\Lambda, \Lambda(Z_1, Z_2)) \right\}, \left\{ \Lambda \leftarrow \text{piO}_\ell^*(1^\lambda, C): (\Lambda, C(Z_1; \mathbf{U}_{m(\lambda)})) \right\} \right)$$

is at most $2^{-e(\lambda)}$.

Setting $\ell := 0$ recovers the original definition of pIO due to [CLTV15], see Definition 2.16. Furthermore, instantiating an easy modification of the construction of piO for $e^{\text{X-ind}}$ due to [CLTV15] with a suitably extracting PRF family satisfies our definition of ℓ -expanding pIO.

10.2 PUNCTURABLE PSEUDORANDOM FUNCTIONS

As preparation for our construction, we introduce several variants of puncturable pseudorandom functions (pPRFs). Puncturable PRFs allow to puncture a key at a certain amount of inputs, such that the punctured key works on all remaining inputs as before. Furthermore, evaluations at punctured points are pseudorandom even given the punctured key. We refer the reader to Definition 2.11 in Section 2.4 for more details.

DEFINITION 10.4 (Statistically injective puncturable PRF, [SW14]). A *statistically injective* puncturable PRF family with error $\epsilon(\cdot)$ is a family of puncturable PRFs F mapping $n(\lambda)$ bits to $m(\lambda)$ bits such that $F(K, \cdot)$ is injective with probability at least $1 - \epsilon(\lambda)$.

Such PRFs are known to exist from one-way functions [SW14].

LEMMA 10.2 ([SW14]). *If one-way functions exist, then for all efficiently computable functions $n(\cdot), m(\cdot), e(\cdot)$ such that $m(\lambda) \geq 2n(\lambda) + e(\lambda)$, there exists a statistically injective puncturable PRF family with error $2^{-e(\lambda)}$ mapping $n(\lambda)$ bits to $m(\lambda)$ bits.*

As already indicated, we instantiate the construction of [CLTV15] with a pPRF which allows to use the entropy provided via the auxiliary input aux such that the obfuscated circuit is evaluated with statistically uniform random coins. The following definition which extends the definition of extracting PRFs (cf. Definition 2.12 in Section 2.4) captures this property.

DEFINITION 10.5 (Special extracting puncturable PRF, [SW14]). A *special extracting* puncturable PRF family with error $\epsilon(\cdot)$ for average min-entropy $k(\cdot)$ is a family of puncturable PRFs F mapping $n(\lambda) = n'(\lambda) + n''(\lambda)$ bits to $m(\lambda)$ bits such that for all λ and all joint distributions (Z_1, Z_2) over $\{0, 1\}^{n'(\lambda)} \times \{0, 1\}^{n''(\lambda)}$ with average min-entropy $\tilde{H}_\infty(Z_2 | Z_1) > k(\lambda)$, the statistical distance

$$\Delta\left(\left\{K \leftarrow \text{KGen}(1^\lambda): (K, Z_1, F(K, Z))\right\}, \left\{K \leftarrow \text{KGen}(1^\lambda): (K, Z_1, U_{m(\lambda)})\right\}\right) \leq \epsilon(\lambda),$$

where the distribution Z is defined as $Z := (Z_1, Z_2)$.

Note that extracting PRFs (cf. Definition 2.12 in Section 2.4) do not suffice to deal with joint distributions (Z_1, Z_2) , where only a guarantee on the average min-entropy $\tilde{H}_\infty(Z_2 | Z_1)$ is known. However, this is required by Definition 10.3.

The leftover hash lemma states that universal hash functions are good randomness extractors, cf. Definition 2.4 in Section 2.1. In other words, if the input of a universal hash function has sufficiently high (average) min-entropy, the output of that hash function is statistically close to uniform even given the function description.

LEMMA 10.3 (Leftover Hash Lemma for average min-entropy, [HILL99]). *Let \mathcal{H} be a 2-universal hash function family mapping n to m bits. If $\tilde{H}_\infty(Z | E) \geq k$ and $m = k - 2 \log(\frac{1}{2\epsilon})$, then*

$$\Delta\left(\left\{H \leftarrow \mathcal{H}: (H, E, H(Z))\right\}, \left\{H \leftarrow \mathcal{H}: (H, E, U_m)\right\}\right) \leq \epsilon.$$

For our purpose, we require that the statistical distance between

$$\left\{H \leftarrow \mathcal{H}: (H, E, H(Z, E))\right\} \quad \text{and} \quad \left\{H \leftarrow \mathcal{H}: (H, E, U_m)\right\}$$

is at most ϵ . Since $\tilde{H}_\infty(Z | E) = \tilde{H}_\infty((Z, E) | E)$, this is implied by Lemma 10.3.

Now we are prepared to extend Theorem 3 in the full version of [SW14] to support joint distributions.

THEOREM 10.1. *If one-way functions exist, then for all efficiently computable functions $n'(\cdot), n''(\cdot), m(\cdot), k(\cdot)$ and $e(\cdot)$ such that $n(\lambda) = n'(\lambda) + n''(\lambda) \geq k(\lambda) \geq m(\lambda) + 2e(\lambda) + 2$, there exists a special extracting puncturable PRF family mapping $n(\lambda)$ bits to $m(\lambda)$ bits with error $2^{-e(\lambda)}$ for average min-entropy $k(\lambda)$.*

Proof. The proof is very similar to the proof of Theorem 3 in the full version of [SW14]. Let F be a family of statistically injective puncturable PRFs with error $2^{-(e(\lambda)+1)}$ mapping $n(\lambda)$ bits to $2n(\lambda) + e(\lambda) + 1$ bits. By Lemma 10.2, such a PRF family exists from one-way functions. Let \mathcal{H} be a family of 2-universal hash functions mapping $2n(\lambda) + e(\lambda) + 1$ bits to $m(\lambda)$ bits. In Figure 10.2, we define a family $F' = (\text{KGen}', \text{Eval}', \text{Punct}')$ of puncturable PRFs.

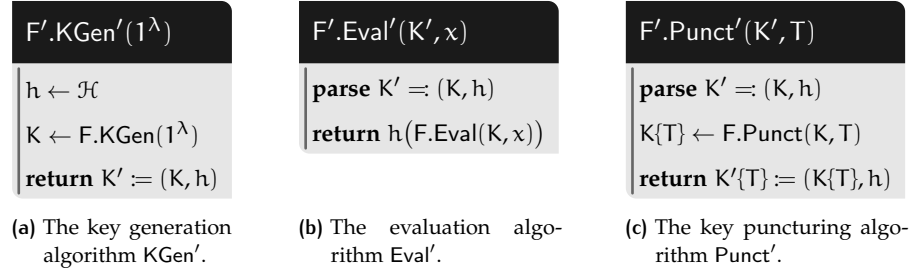


Figure 10.2: The definition of our special extracting puncturable PRF F' .

Due to [SW14], F' is a family of puncturable PRFs mapping $n(\lambda)$ bits to $m(\lambda)$ bits. Let (Z_1, Z_2) be a joint distribution over $\{0, 1\}^{n(\lambda)} = \{0, 1\}^{n'(\lambda)} \times \{0, 1\}^{n''(\lambda)}$ such that $\tilde{H}_\infty(Z_2 | Z_1) \geq m(\lambda) + 2e(\lambda) + 2$. Fix a key K such that $F(K, \cdot)$ is injective. Hence, $\tilde{H}_\infty(Z_2 | Z_1) = \tilde{H}_\infty(F(K, (Z_1, Z_2)) | Z_1)$. Lemma 10.3 implies that the statistical distance

$$\Delta\left(\left\{h \leftarrow \mathcal{H}: (h, Z_1, h(F(K, (Z_1, Z_2))))\right\}, \left\{h \leftarrow \mathcal{H}: (h, Z_1, U_{m(\lambda)})\right\}\right)$$

is at most $2^{-(e(\lambda)+1)}$. The probability that a randomly sampled key K yields a non-injective PRF is at most $2^{-(e(\lambda)+1)}$. Therefore, the statistical distance

$$\Delta\left(\left\{K' \leftarrow F'.\text{KGen}(1^\lambda): (K', Z_1, F'(K', (Z_1, Z_2)))\right\}, \left\{K' \leftarrow F'.\text{KGen}(1^\lambda): (K', Z_1, U_{m(\lambda)})\right\}\right) \leq 2^{-e(\lambda)}$$

concluding the proof. □

10.3 CONSTRUCTION

Figure 10.3 defines our construction of an ℓ -expanding pIO scheme piO_ℓ^* . We use the pIO scheme for X-ind samplers piO from [CLTV15] instantiated with a special extracting pPRF as a subroutine.

THEOREM 10.2. *Let $e(\lambda)$ be an efficiently computable function. Let F be a subexponentially secure special extracting PRF family with distinguishing advantage $2^{-\lambda^\epsilon}$ (for some constant ϵ) and error $2^{-e(\lambda)}$ mapping $n(\lambda) := n'(\lambda) + \ell(\lambda)$ bits to $m(\lambda)$ bits which is extracting if the input average min-entropy is greater than*

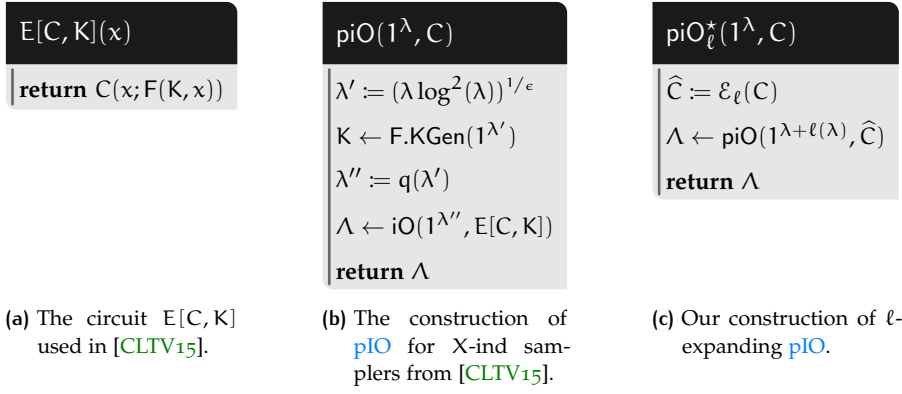


Figure 10.3: Construction of piO for X-ind samplers from [CLTV15] (Figures 10.3a and 10.3b) and of ℓ-expanding piO (Figure 10.3c). Both definitions are for circuits of size λ. In Figure 10.3b, the polynomial q(λ′) denotes an upper bound on the size of the circuit E[C, K] and F denotes a special extracting subexponentially secure pPRF with distinguishing gap 2^{−λ^ε} (for a constant ε).

$m(\lambda) + 2e(\lambda) + 2$. Let piO denote the construction of piO from [CLTV15] instantiated with F (cf. Figure 10.3b). Then, piO_ℓ^{*} as defined in Figure 10.3c is a statistically correct input-expanding piO for the class of samplers $\mathcal{C}_\ell^{X-(*)\text{ind}}$.

Proof. We recall that due to [CLTV15], piO satisfies correctness and security with respect to X-ind samplers. We now prove that piO_ℓ^{*} defined in Figure 10.3 is a statistically correct input expanding piO scheme for the class of samplers $\mathcal{C}_\ell^{X-(*)\text{ind}}$.

INPUT EXPANDING CORRECTNESS. Note that input expanding correctness does not follow in a black-box manner from the correctness of piO since input expanding correctness allows an adversary to query the randomized circuit multiple times on the same input x using different auxiliary inputs aux. This cannot be simulated by an adversary in the correctness game of piO, cf. Definition 2.16.

Input expanding correctness follows from pseudorandomness of the PRF used in the construction of piO and correctness of iO. Let C be a circuit in \mathcal{C}_λ . Let \mathcal{A} be a PPT adversary. Let \mathbf{G}_0 denote the game in which the adversary gets C as input and has oracle access to $\mathcal{O}_C(\cdot, \cdot)$ and let \mathbf{G}_2 denote the game in which the adversary gets C as input and has oracle access to $\mathcal{O}_\Lambda(\cdot, \cdot)$ for $\Lambda \leftarrow \text{piO}_\ell^*(1^\lambda, C)$. Both oracles must not be called twice on the same input. Consider an intermediate game \mathbf{G}_1 , where the oracle $\mathcal{O}_\Lambda(x, aux)$ evaluates the circuit $\widehat{C}(x, aux; \mathcal{R}(x, aux))$ for a truly random function $\mathcal{R} \in \text{maps}(\{0, 1\}^{n'(\lambda)+\ell(\lambda)}, \{0, 1\}^{m(\lambda)})$. By the security of the PRF F, oracle access to F(K, ·) (for a random key K) and a truly random function are indistinguishable. Hence, by perfect correctness of iO and the security of F, \mathbf{G}_1 and \mathbf{G}_2 are indistinguishable. Furthermore, since $\widehat{C}(x, aux; \mathcal{R}(x, aux)) = C(x; \mathcal{R}(x, aux))$ and by perfect correctness of iO, \mathbf{G}_0 and \mathbf{G}_1 are identically distributed.

SECURITY WITH RESPECT TO $\mathcal{C}_\ell^{X\text{-ind}}$. Since piO is secure with respect to X-ind samplers (and $\mathcal{E}_\ell(C) \in \widehat{\mathcal{C}}$), piO_ℓ^{*} is secure with respect to ℓ-expanding X-ind samplers.

More formally, let $D \in \mathfrak{C}_\ell^{X^{-(*)}\text{-ind}}$ be a circuit sampler and let \mathcal{A} be a **PPT** distinguisher for piO_ℓ^* -obfuscations of C_0 and C_1 as sampled from D , cf. Definition 10.3. We construct a circuit sampler \widehat{D} (as described above) and an adversary \mathcal{B} for the security property of piO . Let \widehat{D} be the ℓ -expanding circuit sampler corresponding to D . That is, $\widehat{D}_{\lambda+\ell(\lambda)}$, samples from D_λ to obtain (C_0, C_1, z) and outputs $(\mathcal{E}_\ell(C_0), \mathcal{E}_\ell(C_1), \widehat{z} := (z, C_0, C_1))$. On input of $(\mathcal{E}_\ell(C_0), \mathcal{E}_\ell(C_1), \widehat{z}, \Lambda)$, \mathcal{B} calls \mathcal{A} on input of (C_0, C_1, z, Λ) and outputs \mathcal{A} 's output. Hence, for $b \in \{0, 1\}$, we have

$$\begin{aligned} & \Pr \left[(C_0, C_1, z) \leftarrow D(1^\lambda) : \mathcal{A}(C_0, C_1, z, \text{piO}_\ell^*(1^\lambda, C_b)) = 1 \right] \\ &= \Pr \left[(\mathcal{E}_\ell(C_0), \mathcal{E}_\ell(C_1), (C_0, C_1, z)) \leftarrow \widehat{D}(1^{\lambda+\ell(\lambda)}) : \right. \\ & \quad \left. \mathcal{B}(\mathcal{E}_\ell(C_0), \mathcal{E}_\ell(C_1), (C_0, C_1, z), \text{piO}(1^{\lambda+\ell(\lambda)}, \mathcal{E}_\ell(C_b))) = 1 \right]. \end{aligned}$$

Therefore, for all $D \in \mathfrak{C}_\ell^{X^{-(*)}\text{-ind}}$, all **PPT** adversaries \mathcal{A} , there exists a sampler $\widehat{D} \in \mathfrak{C}^{X\text{-ind}}$ and a **PPT** adversary \mathcal{B} , such that

$$\text{Adv}_{\text{piO}_\ell^*, D, \mathcal{A}}^{\text{pio-ind}^{(*)}}(\lambda) = \text{Adv}_{\text{piO}, \widehat{D}, \mathcal{B}}^{\text{pio-ind}}(\lambda + \ell(\lambda))$$

which is negligible by the security of piO for X -ind samplers and since $\widehat{D} \in \mathfrak{C}^{X\text{-ind}}$.

SUPPORT RESPECTING. Let $\lambda \in \mathbb{N}$, $C \in \mathcal{C}_\lambda$ and $\Lambda \in \text{supp}(\text{piO}_\ell^*(1^\lambda, C))$. By construction of piO_ℓ^* and piO and perfect correctness of iO , for every input $(x, \text{aux}) \in \{0, 1\}^{n'(\lambda)} \times \{0, 1\}^{\ell(\lambda)}$, we have that

$$\begin{aligned} \Lambda(x, \text{aux}) &= \mathcal{E}_\ell(C)(x, \text{aux}; F(K, (x, \text{aux}))) \\ &= C(x; F(K, (x, \text{aux}))) \in \text{supp}(C(x)). \end{aligned}$$

STATISTICAL CORRECTNESS WITH ERROR $2^{-e(\lambda)}$. Let $e(\lambda)$ be an efficiently computable function. Let $\lambda \in \mathbb{N}$ and $C \in \mathcal{C}_\lambda$ be a circuit expecting inputs $x \in \{0, 1\}^{n'(\lambda)}$ and randomness $r \in \{0, 1\}^{m(\lambda)}$, and let $\widehat{C} := \mathcal{E}_\ell(C)$ be its corresponding ℓ -expanded circuit. Let $Z := (Z_1, Z_2)$ be a joint distribution over $\{0, 1\}^{n'(\lambda)} \times \{0, 1\}^{\ell(\lambda)}$ with average min-entropy $\ell(\lambda) \geq \widetilde{H}_\infty(Z_2 | Z_1) > m(\lambda) + 2e(\lambda) + 2$. Then, by Theorem 10.1, the statistical distance

$$\Delta \left(\left\{ K \leftarrow \text{KGen}(1^\lambda) : (K, Z_1, F(K, Z)) \right\}, \left\{ K \leftarrow \text{KGen}(1^\lambda) : (K, Z_1, U_{m(\lambda)}) \right\} \right)$$

is at most $2^{-e(\lambda)}$. Since, for all distributions A, B and for every randomized function f (see, e. g., [MG02]), the statistical distance between $f(A)$ and $f(B)$ is upper bounded by the statistical distance between A and B , we have that the statistical distance

$$\Delta \left(\left\{ K \leftarrow \text{KGen}(1^\lambda) : (K, C(Z_1; F(K, Z))) \right\}, \left\{ K \leftarrow \text{KGen}(1^\lambda) : (K, C(Z_1; U_{m(\lambda)})) \right\} \right)$$

and the statistical distance

$$\Delta \left(\left\{ \Lambda \leftarrow \text{piO}_\ell^*(C) : (\Lambda, \overbrace{C(Z_1; F(K, Z))}^{=\Lambda(Z_1, Z_2)}) \right\}, \left\{ \Lambda \leftarrow \text{piO}_\ell^*(C) : (\Lambda, C(Z_1; U_{m(\lambda)})) \right\} \right)$$

are both upper bounded by $2^{-e(\lambda)}$. □

REMARK 10.1. Our construction also achieves security with respect to X -ind samplers. This, however, comes at the cost of a non-black-box proof reproducing the hybrid argument of [CLTV15].

11

HOW TO SIMULATE
EXTRACTION

In order to instantiate the algebraic group model, we first require an abstract definitional framework of what it means to be a group in a cryptographic sense which allows to simulate crucial properties of the [AGM](#). The notion of a *group scheme* or *encoding scheme* [\[GGH13\]](#) abstracts the properties of a cryptographic group. In a nutshell, a group scheme provides an interface of algorithms abstracting the handling of a cryptographic group. As we want to prove hardness of certain problems based on hardness assumptions in an already existing *base group*, we incorporate this existing group into our group scheme. Each element of our group scheme corresponds to exactly one element from a base group in a structure-preserving way. In other words, every element group of our group schemes “wraps” an element of the base group. We call the resulting notion a *group wrapper*.

For the purpose of simulating extraction properties, we introduce the notion of an *algebraic wrapper*. An algebraic wrapper is a group wrapper which provides the additional functionality that group element encodings allow for efficient extraction of representations which – similar to the algebraic group model – can be used in a security reduction. A similar approach has been taken by [\[KP19\]](#). [\[KP19\]](#) define their group scheme as a linear subspace of $G \times G$ for an existing group G in such a way that the Generalized Knowledge of Exponent Assumption ([GKEA](#)) [\[WSo7\]](#) can be used to extract a representation (membership can for instance be tested via a symmetric pairing). However, [\[KP19\]](#) relies on [GKEA](#) in the base group which more or less directly yields an equivalence between algebraic groups and [GKEA](#). Moreover, the existence of algebraic groups implies the existence of extractable one-way functions with unbounded auxiliary input as per [Definition 2.21](#) (since the [AGM](#) allows an additional unstructured input from $\{0, 1\}^*$). This notion of extractable one-way functions conflicts with the existence of indistinguishability obfuscation [\[BCPR16\]](#). More precisely, consider the additional unstructured input from $\{0, 1\}^*$ encodes an obfuscated circuit which on input of a group element $[a]_G$ computes $[a]_G^{F_K([a]_G)}$ using some hard-coded [PRF](#) key K . An adversary who simply evaluates this circuit on the group generator and outputs the result is clearly unable to explain how this output was derived algebraically. Due to this contradiction and the difficulty to assess the plausibility of knowledge-type assumptions, we strive for a weaker model which can purely be based on falsifiable assumptions.

EXTRACTION TRAPDOORS. In [\[KP19\]](#), extraction is possible as long as the code and the randomness which were used to produce a group element are known. Since we seek to avoid non-falsifiable knowledge-type assumptions, we need to find a different mechanism of what enables extraction. We observe that in order to reproduce proof strategies from the algebraic group model, extraction is only necessary during security reductions. Since the reduction to some assumption in the base group is in control of the parameter generation

of the wrapper, the reduction may use corresponding trapdoor information which we define to enable extraction. We call this notion *private extractability*.

In the following we provide formal definitions of group schemes and our notion of algebraic wrappers.

11.1 GROUP SCHEMES AND WRAPPERS

A group scheme or encoding scheme [GGH13] abstracts the properties of mathematical groups used in cryptography. Group schemes have recently been studied in [AFH+16; AH18; FHHL18; KP19; AMP20]. In contrast to traditional groups, group elements are not bound to be represented by some unique bitstring – henceforth referred to as *encoding*. This allows to encode auxiliary information inside group elements.

Formally, a group scheme \mathbb{H} consists of the algorithms $(\text{GGen}_{\mathbb{H}}, \text{Sam}_{\mathbb{H}}, \text{Val}_{\mathbb{H}}, \text{Add}_{\mathbb{H}}, \text{Eq}_{\mathbb{H}}, \text{GetID}_{\mathbb{H}})$.

- $\text{GGen}_{\mathbb{H}}$ is a group generation algorithm which, given 1^λ , samples group parameters $pp_{\mathbb{H}}$.
- $\text{Sam}_{\mathbb{H}}$ is a sampling algorithm which, on input of the group parameters $pp_{\mathbb{H}}$ and an exponent a , produces an encoding corresponding to the exponent a .⁴⁶
- $\text{Val}_{\mathbb{H}}$ is a validation algorithm which, given the group parameters $pp_{\mathbb{H}}$ and a bitstring, decides whether the given bitstring is a valid encoding.
- $\text{Add}_{\mathbb{H}}$ implements the group operation, i. e., $\text{Add}_{\mathbb{H}}$ expects the group parameters $pp_{\mathbb{H}}$ and two encodings as input, and produces an encoding of the group element which results from applying the group operation on its input encodings.
- $\text{Eq}_{\mathbb{H}}$ is an equality testing algorithm. Since group elements do not necessarily possess unique encodings, $\text{Eq}_{\mathbb{H}}$ enables to test whether two given encodings correspond to the same group element (with respect to the given group parameters). Note that $\text{Eq}_{\mathbb{H}}(pp_{\mathbb{H}}, \cdot)$ defines an equivalence relation on the set of valid bitstrings.
- $\text{GetID}_{\mathbb{H}}$ is a “get-identifier” algorithm which, given the group parameters $pp_{\mathbb{H}}$ and an encoding of a group element, produces a bitstring which is unique for all encodings of the same group element.⁴⁷ This, again, compensates for the non-uniqueness of encodings.

⁴⁶ For instance, $\text{Sam}_{\mathbb{H}}$ can be implemented to simply raise a fixed group generator to the power of a .

⁴⁷ Previous work refers to this algorithm as extraction algorithm. However, to avoid overloading the word “extraction”, we rename this algorithm in this part.

Note that $\text{Eq}_{\mathbb{H}}(pp_{\mathbb{H}}, a, b)$ can be implemented using $\text{GetID}_{\mathbb{H}}$ by simply comparing $\text{GetID}_{\mathbb{H}}(pp_{\mathbb{H}}, a)$ and $\text{GetID}_{\mathbb{H}}(pp_{\mathbb{H}}, b)$ as bitstrings. The “get-identifier” algorithm allows to convert group elements into unique bitstrings which can, for instance, be used as keys. In some settings, group schemes are additionally equipped with a multilinear map [AFH+16; FHHL18; AMP20]. Hence, a group schemes supports all features cryptography demands from ordinary cyclic groups.

For a group scheme it is required that the quotient set

$$\left\{ a \in \{0,1\}^* \mid \text{Val}_{\mathbb{H}}(pp_{\mathbb{H}}, a) = 1 \right\} / \text{Eq}_{\mathbb{H}}(pp_{\mathbb{H}}, \cdot)$$

equipped with the operation defined via $\text{Add}_{\mathbb{H}}(pp_{\mathbb{H}}, \cdot, \cdot)$ defines a mathematical group (with overwhelming probability over the choice of $pp_{\mathbb{H}} \leftarrow \text{GGen}_{\mathbb{H}}(1^\lambda)$). We say that a bitstring a is (an encoding of) a group element (relative to $pp_{\mathbb{H}}$), written as $a \in \mathbb{H}$, if and only if $\text{Val}_{\mathbb{H}}(pp_{\mathbb{H}}, a) = 1$.

A group scheme requires that encodings corresponding to the same group element are computationally indistinguishable as formalized by the “Switching Lemma(s)” in [AFH+16; AH18; FHHL18].

We henceforth use the notation \hat{h} to mark encodings of a group element.

11.2 AN ALGEBRAIC WRAPPER

Given a cyclic group, an algebraic wrapper is a group wrapper which equips a given group G with a notion of extractability while preserving its group structure and computational hardness guarantees. In particular, we achieve a property which we refer to as “private extractability” with respect to a given set of group elements in the base group. More precisely, our group generation algorithm expects group parameters pp_G of the base group together with a set of group elements $[b]_G \in G^n$ from that base group – henceforth referred to as *basis* – and produces group parameters $pp_{\mathbb{H}}$ of the wrapper group together with a corresponding trapdoor $\tau_{\mathbb{H}}$. We exploit the non-uniqueness of encodings to let all encodings carry some representation relative to this basis $[b]_G$. The trapdoor $\tau_{\mathbb{H}}$ enables to extract this representation from every encoding. A “representation” of a group element relative to the basis $[b]_G$ is a decomposition of the group element as linear combination of the group elements in $[b]_G$. Hence, our notion of extractability is strictly weaker than the notion provided by the algebraic group model (AGM) [FKL18], but also circumvents the implausibility results entailed by the AGM. Nonetheless, this notion suffices to simulate proof strategies developed for the algebraic group model in the very absence of knowledge-type assumptions.

More precisely, encodings which wrap some base group element $h \in G$ carry (computationally hidden) representation vectors $z \in \mathbb{Z}_p^n$, such that z contains the coefficients of a linear combination of h relative to $[b]_G$, i.e., $z^T \cdot [b]_G = h$. Private extraction allows to recover this representation vector. Given the trapdoor, we require that it is possible to “privately” sample encodings which carry some specific dictated representation vector. We require that publicly sampled encodings and privately sampled encodings are computationally indistinguishable. We refer to this property as “switching”. In order to preserve tightness of security reductions when implementing AGM proofs with our algebraic wrapper, we require a statistical re-randomization property. Furthermore, we require that representation vectors compose additively (in \mathbb{Z}_p^n) with the group operation and do not change when encodings are re-randomized.

Let

$$\mathcal{B}_{pp_G}^n := \left\{ ([1]_G, [x_2]_G, \dots, [x_n]_G)^T \in G^n \mid x_2, \dots, x_n \in \mathbb{Z}_p^\times \right\}$$

be the set of what we call “legitimate basis vectors”. Note that we require the first group element to be the generator of the base group $[1]_G$ to allow for public sampling (using $\text{Sam}_{\mathbb{H}}$). This is necessary because $\text{Sam}_{\mathbb{H}}$ must be able to encode a matching representation vector for every produced group element encoding.

DEFINITION 11.1 (Algebraic wrapper for \mathbb{G}). An algebraic wrapper \mathbb{H} for \mathbb{G} is a tuple of PPT algorithms $(\text{GGen}_{\mathbb{H}}, \text{Sam}_{\mathbb{H}}, \text{Val}_{\mathbb{H}}, \text{Add}_{\mathbb{H}}, \text{Eq}_{\mathbb{H}}, \text{GetID}_{\mathbb{H}}, \text{Rerand}_{\mathbb{H}}, \text{PrivSam}_{\mathbb{H}}, \text{PrivExt}_{\mathbb{H}}, \text{Unwrap}_{\mathbb{H}})$ such that the tuple $(\text{GGen}_{\mathbb{H}}, \text{Sam}_{\mathbb{H}}, \text{Val}_{\mathbb{H}}, \text{Add}_{\mathbb{H}}, \text{Eq}_{\mathbb{H}}, \text{GetID}_{\mathbb{H}})$ constitutes a group scheme and the following properties are satisfied.

G-WRAPPING. The algorithm $\text{Unwrap}_{\mathbb{H}}(pp_{\mathbb{H}}, \cdot)$ is deterministic and for all $pp_{\mathbb{G}} \in \text{supp}(\text{GGen}_{\mathbb{G}}(1^\lambda))$, all legitimate basis vectors $[\mathbf{b}]_{\mathbb{G}} \in \mathcal{B}_{pp_{\mathbb{G}}}^n$, all $(pp_{\mathbb{H}}, \tau_{\mathbb{H}}) \in \text{supp}(\text{GGen}_{\mathbb{H}}(pp_{\mathbb{G}}, [\mathbf{b}]_{\mathbb{G}}))$,

$$\text{Unwrap}_{\mathbb{H}}(pp_{\mathbb{H}}, \cdot): \mathbb{H} \rightarrow \mathbb{G}$$

defines a group isomorphism.

EXTRACTABILITY. The algorithm $\text{PrivExt}_{\mathbb{H}}$ is deterministic. Furthermore, for all $pp_{\mathbb{G}} \in \text{supp}(\text{GGen}_{\mathbb{G}}(1^\lambda))$, all legitimate basis vectors $[\mathbf{b}]_{\mathbb{G}} \in \mathcal{B}_{pp_{\mathbb{G}}}^n$, all $(pp_{\mathbb{H}}, \tau_{\mathbb{H}}) \in \text{supp}(\text{GGen}_{\mathbb{H}}(pp_{\mathbb{G}}, [\mathbf{b}]_{\mathbb{G}}))$, all $\widehat{h} \in \mathbb{H}$, we require that $\text{PrivExt}_{\mathbb{H}}$ always extracts a representation of $[x]_{\mathbb{G}} := \text{Unwrap}_{\mathbb{H}}(pp_{\mathbb{H}}, \widehat{h})$ with respect to $[\mathbf{b}]_{\mathbb{G}}$, i. e., for $\mathbf{z} := \text{PrivExt}_{\mathbb{H}}(\tau_{\mathbb{H}}, \widehat{h})$,

$$[\mathbf{z}^\top \cdot \mathbf{b}]_{\mathbb{G}} = \text{Unwrap}_{\mathbb{H}}(pp_{\mathbb{H}}, \widehat{h}).$$

CORRECTNESS OF EXTRACTION. For all $pp_{\mathbb{G}} \in \text{supp}(\text{GGen}_{\mathbb{G}}(1^\lambda))$, all $[\mathbf{b}]_{\mathbb{G}} \in \mathcal{B}_{pp_{\mathbb{G}}}^n$, all $(pp_{\mathbb{H}}, \tau_{\mathbb{H}}) \in \text{supp}(\text{GGen}_{\mathbb{H}}(pp_{\mathbb{G}}, [\mathbf{b}]_{\mathbb{G}}))$, all $\widehat{h}_0, \widehat{h}_1 \in \mathbb{H}$, we require that private extraction respects the group operation in the sense that for all $\widehat{h}_2 \in \text{supp}(\text{Add}_{\mathbb{H}}(pp_{\mathbb{H}}, \widehat{h}_0, \widehat{h}_1))$, $\mathbf{z}^{(i)} := \text{PrivExt}_{\mathbb{H}}(\tau_{\mathbb{H}}, \widehat{h}_i)$ for $i \in \{0, 1, 2\}$, satisfy

$$\mathbf{z}^{(2)} = \mathbf{z}^{(0)} + \mathbf{z}^{(1)}.$$

Additionally, for all $pp_{\mathbb{G}} \in \text{supp}(\text{GGen}_{\mathbb{G}}(1^\lambda))$, all $[\mathbf{b}]_{\mathbb{G}} \in \mathcal{B}_{pp_{\mathbb{G}}}^n$, all $(pp_{\mathbb{H}}, \tau_{\mathbb{H}}) \in \text{supp}(\text{GGen}_{\mathbb{H}}(pp_{\mathbb{G}}, [\mathbf{b}]_{\mathbb{G}}))$, all $\widehat{h} \in \mathbb{H}$, we require that re-randomization does not interfere with private extraction in the sense that for all $\widehat{h}' \in \text{supp}(\text{Rerand}_{\mathbb{H}}(pp_{\mathbb{H}}, \widehat{h}))$,

$$\text{PrivExt}_{\mathbb{H}}(\tau_{\mathbb{H}}, \widehat{h}) = \text{PrivExt}_{\mathbb{H}}(\tau_{\mathbb{H}}, \widehat{h}').$$

CORRECTNESS OF SAMPLING. For all $pp_{\mathbb{G}} \in \text{supp}(\text{GGen}_{\mathbb{G}}(1^\lambda))$, all $[\mathbf{b}]_{\mathbb{G}} \in \mathcal{B}_{pp_{\mathbb{G}}}^n$, all $(pp_{\mathbb{H}}, \tau_{\mathbb{H}}) \in \text{supp}(\text{GGen}_{\mathbb{H}}(pp_{\mathbb{G}}, [\mathbf{b}]_{\mathbb{G}}))$, we require that

- for all $\mathbf{v} \in \mathbb{Z}_p^n$, $\Pr[\text{PrivExt}_{\mathbb{H}}(\tau_{\mathbb{H}}, \text{PrivSam}_{\mathbb{H}}(\tau_{\mathbb{H}}, \mathbf{v})) = \mathbf{v}] = 1$,
- for all $x \in \mathbb{Z}_p$, $\Pr[\text{PrivExt}_{\mathbb{H}}(\tau_{\mathbb{H}}, \text{Sam}_{\mathbb{H}}(pp_{\mathbb{H}}, x \cdot \mathbf{e}_1)) = x \cdot \mathbf{e}_1] = 1$.

k-SWITCHING. We say a PPT adversary \mathcal{A} is a legitimate k -switching adversary if on input of base group parameters $pp_{\mathbb{G}}$, \mathcal{A} outputs two bases $([\mathbf{b}^{(j)}]_{\mathbb{G}})_{j \in \{0,1\}}$ and two lists comprising k representation vectors $(\mathbf{v}^{(j),(i)})_{i \in [k], j \in \{0,1\}}$ (and an internal state st) such that $[\mathbf{b}^{(0)}]_{\mathbb{G}}, [\mathbf{b}^{(1)}]_{\mathbb{G}} \in \mathcal{B}_{pp_{\mathbb{G}}}^n$ and $\mathbf{v}^{(0),(i)}, \mathbf{v}^{(1),(i)} \in \mathbb{Z}_p^n$ for some $n \in \mathbb{N}$ and all $i \in [k]$ and

$$[(\mathbf{v}^{(0),(i)})^\top \cdot \mathbf{b}^{(0)}]_{\mathbb{G}} = [(\mathbf{v}^{(1),(i)})^\top \cdot \mathbf{b}^{(1)}]_{\mathbb{G}}$$

for all $i \in [k]$.

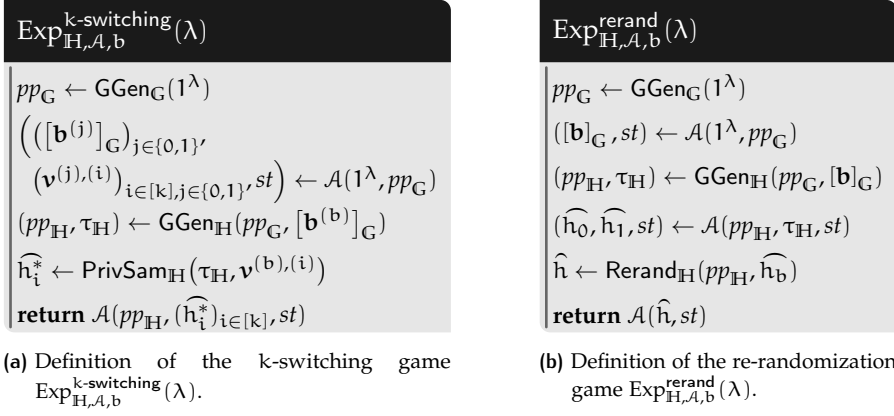


Figure 11.1: The re-randomization and k-switching games.

For all legitimate k-switching PPT adversaries \mathcal{A} ,

$$\text{Adv}_{\mathbb{H},\mathcal{A}}^{\text{k-switching}}(\lambda) := \left| \Pr \left[\text{Exp}_{\mathbb{H},\mathcal{A},0}^{\text{k-switching}}(\lambda) = 1 \right] - \Pr \left[\text{Exp}_{\mathbb{H},\mathcal{A},1}^{\text{k-switching}}(\lambda) = 1 \right] \right|$$

is negligible, where $\text{Exp}_{\mathbb{H},\mathcal{A},b}^{\text{k-switching}}(\lambda)$ (for $b \in \{0,1\}$) is defined in Figure 11.1a.

STATISTICALLY RE-RANDOMIZABLE. We say an unbounded adversary \mathcal{A} is a legitimate re-randomization adversary if, on input of base group parameters pp_G , \mathcal{A} outputs $[\mathbf{b}]_G$ and a state st such that $[\mathbf{b}]_G \in \mathcal{B}_{pp_G}^n$ and, in a second phase, \mathcal{A} on input of (pp_H, τ_H, st) outputs two valid encodings $\widehat{h}_0, \widehat{h}_1$ (and a state st) such that

$$\text{PrivExt}_H(\tau_H, \widehat{h}_0) = \text{PrivExt}_H(\tau_H, \widehat{h}_1).$$

For all unbounded legitimate re-randomization adversaries \mathcal{A} ,

$$\text{Adv}_{\mathbb{H},\mathcal{A}}^{\text{rerand}}(\lambda) := \left| \Pr \left[\text{Exp}_{\mathbb{H},\mathcal{A},0}^{\text{rerand}}(\lambda) = 1 \right] - \Pr \left[\text{Exp}_{\mathbb{H},\mathcal{A},1}^{\text{rerand}}(\lambda) = 1 \right] \right| \leq \frac{1}{2^\lambda},$$

where $\text{Exp}_{\mathbb{H},\mathcal{A},b}^{\text{rerand}}(\lambda)$ (for $b \in \{0,1\}$) is defined in Figure 11.1b.

For simplicity, we require that encodings are always in $\{0,1\}^{\text{penc}(\lambda)}$ for a fixed polynomial $\text{penc}(\lambda)$.

The k-switching property allows to simultaneously switch the representation vectors of multiple group element encodings. It is necessary to switch all encodings simultaneously since private sampling can only be simulated knowing the trapdoor τ_H which is not the case in $\text{Exp}_{\mathbb{H},\mathcal{A},b}^{\text{k-switching}}(\lambda)$. More precisely, suppose we want to apply 1-switching twice. After switching the first representation, simulation of the resulting hybrid game requires to use private sampling and, hence, possession of the trapdoor τ_H . Consequently, a further reduction to 1-switching is not possible, since an adversary against 1-switching does not know τ_H .

Further, we note that group element encodings should always be re-randomized for the following reason. While publicly sampled encodings and

privately sampled encodings are indistinguishable due to k -switching (provided that they encode the same group elements), Definition 11.1 requires no such indistinguishability guarantee for, e. g., a privately sampled encoding and an encoding which is computed using the group operation. However, if these two encodings carry identical representation vectors, statistical re-randomizability guarantees indistinguishability if these encodings are re-randomized.

12 | CONSTRUCTION

Our construction follows the ideas from [AFH+16; AH18; FHHL18]. Let $\text{GGen}_{\mathbb{G}}$ be a group generator for a cyclic group \mathbb{G} . Let Gen_{SMP} be a hard subset membership problem. Let $\text{FHE} = (\text{KGen}, \text{Enc}, \text{Dec}, \text{Eval}, \text{Rerand})$ be a perfectly correct and perfectly re-randomizable fully homomorphic public-key encryption scheme. Let $pp_{\mathbb{G}}$ be group parameters for \mathbb{G} and $[\mathbf{\Omega}]_{\mathbb{G}} \in \mathbb{G}^n$ for some $n \in \mathbb{N}$. Let $L \subseteq X$ together with the relation R denote a subset membership problem instance sampled from Gen_{SMP} and let $y \leftarrow X \setminus L$. Further, let pk be a public key for FHE. For ease of notation, we define $pars := (pp_{\mathbb{G}}, (L, X, R), y, pk, [\mathbf{\Omega}]_{\mathbb{G}})$. Let $\text{NIWI} := (\text{Setup}, \text{Prove}, \text{Verify}, \text{HSetup}, \text{Ext})$ be a perfectly complete, perfectly extractable and perfectly witness-indistinguishable dual-mode [NIWI](#) proof system for the language

$$\mathcal{L} := \{y := (pars, [x]_{\mathbb{G}}, C) \mid \exists w: (y, w) \in \mathcal{R} := \mathcal{R}_1 \vee \mathcal{R}_2 \vee \mathcal{R}_3\}.$$

The relations $\mathcal{R}_1, \mathcal{R}_2, \mathcal{R}_3$ are defined as follows.

$$\begin{aligned} \mathcal{R}_1 &= \left\{ \left((pars, [x]_{\mathbb{G}}, C), (sk, \mathbf{v}) \right) \mid \begin{array}{l} \text{KGen}(1^\lambda; sk) = (pk, sk) \\ \text{Dec}(sk, C) = \mathbf{v} \\ [\mathbf{\Omega}^T \cdot \mathbf{v}]_{\mathbb{G}} = [x]_{\mathbb{G}} \end{array} \right\} \\ \mathcal{R}_2 &= \left\{ \left((pars, [x]_{\mathbb{G}}, C), (r, \mathbf{v}) \right) \mid \begin{array}{l} \text{Enc}(pk, \mathbf{v}; r) = C \\ [\mathbf{\Omega}^T \cdot \mathbf{v}]_{\mathbb{G}} = [x]_{\mathbb{G}} \end{array} \right\} \\ \mathcal{R}_3 &= \left\{ \left((pars, [x]_{\mathbb{G}}, C), (w_y) \right) \mid (y, w_y) \in \mathcal{R} \right\} \end{aligned}$$

Encodings \hat{h} of group elements in \mathbb{H} are of the form

$$\hat{h} := ([x]_{\mathbb{G}}, C, \pi),$$

where $[x]_{\mathbb{G}}$ is the “wrapped” element from the base group \mathbb{G} , C is a ciphertext and π is a consistency proof π produced for \mathcal{L} . During “honest” use of the algebraic wrapper, encodings carry proofs produced for relation \mathcal{R}_1 or relation \mathcal{R}_2 . Both relations bind a “perturbed” representation vector \mathbf{v} to the base group element $[x]_{\mathbb{G}}$. These two relations only differ in the necessary witnesses. Relation \mathcal{R}_1 requires the secret key which is associated with the public key in the public parameters $pp_{\mathbb{H}}$ as witness whereas relation \mathcal{R}_2 requires the randomness used for producing the ciphertext C . Hence, relation \mathcal{R}_1 can be used in settings, where the ciphertext is not produced freshly but, for instance, via homomorphic evaluation, but the secret key can be used. This will be necessary in certain obfuscated circuits. Relation \mathcal{R}_2 , on the other hand, requires only information which is accessible for an honest user and, hence, allows public sampling of group elements. Relation \mathcal{R}_3 is a trapdoor branch which enables simulation. Note that during “honest” use of the algebraic wrapper $y \notin L$ and, hence, due to perfect soundness of [NIWI](#), there exists no proof for relation \mathcal{R}_3 .

With $m'(\lambda)$ we denote a polynomial upper bound on the number of random bits $\text{FHE.Rerand}(1^\lambda, \cdot, \cdot)$ expects and with $m''(\lambda)$ we denote a polynomial upper bound on the number of random bits $\text{NIWI.Prove}(1^\lambda, \cdot, \cdot, \cdot)$ expects. Let $\ell(\lambda) := m'(\lambda) + m''(\lambda) + 2(\lambda + 1) + 3$. Let $\mathcal{C}_{\text{add}} := \{\mathcal{C}_{\text{add}}, \mathcal{C}_{\text{add}}^{(1)}, \mathcal{C}_{\text{add}}^{(2)}, \mathcal{C}_{\text{add}}^{(3)}\}$ and let $\mathcal{C}_{\text{rerand}} := \{\mathcal{C}_{\text{rerand}}, \mathcal{C}_{\text{rerand}}^{(1)}, \mathcal{C}_{\text{rerand}}^{(2)}, \mathcal{C}_{\text{rerand}}^{(3)}\}$, see Figures 12.4, 12.5 and 12.6. Let piO be a **piO** scheme for the class of samplers $\mathcal{C}^{\text{X-ind}}$ over \mathcal{C}_{add} and let piO_ℓ^* be an ℓ -expanding **piO** scheme for the class of samplers $\mathcal{C}_\ell^{\text{X-}(\cdot)\text{-ind}}$ over $\mathcal{C}_{\text{rerand}}$. Further, let $p_{\text{add}}(\lambda)$ denote a polynomial upper bound on the size of addition circuits in \mathcal{C}_{add} and $p_{\text{rerand}}(\lambda)$ denote a polynomial upper bound on the size of re-randomization circuits in \mathcal{C}_{add} .

Our algebraic wrapper \mathbb{H} is composed of the **PPT** algorithms $(\text{GGen}_{\mathbb{H}}, \text{Sam}_{\mathbb{H}}, \text{Val}_{\mathbb{H}}, \text{Add}_{\mathbb{H}}, \text{Eq}_{\mathbb{H}}, \text{Rerand}_{\mathbb{H}}, \text{PrivExt}_{\mathbb{H}}, \text{PrivSam}_{\mathbb{H}}, \text{GetID}_{\mathbb{H}}, \text{Unwrap}_{\mathbb{H}})$ defined in Figures 12.1 and 12.2. We note that the algorithm $\text{Val}_{\mathbb{H}}$ which is evaluated inside \mathcal{C}_{add} and $\mathcal{C}_{\text{rerand}}$ only requires a certain part of the public parameters as input. In particular, $\text{Val}_{\mathbb{H}}$ does not depend on Λ_{add} and Λ_{rerand} .

DIFFERENCES TO [AFH+16; AH18; FHHL18]. [AFH+16; FHHL18] introduce similar constructions of a group scheme featuring a multilinear map and of a graded encoding scheme, respectively. More precisely, [AFH+16; FHHL18] equip a base group with encodings carrying auxiliary information which can be used (in an obfuscated circuit) to “multiply in the exponent”. We observe that these constructions already *wrap* a given base group in the sense that “unwrapping” encodings yields a group isomorphism to the base group.

Our construction builds upon these group schemes. In order to enable extractability with respect to a dynamically chosen basis⁴⁸, our group parameters must be generated depending on that basis.

This modification, however, comes at the cost of the multilinear map functionality. This is because any implementation of a multilinear map requires knowledge of discrete logarithms of each group element encoding to a fixed generator. This is undesirable for our purposes, since we want to be able to use sets of group elements as basis which we do not know discrete logarithms of (for instance group elements provided by a reduction). Thus, we have to give up the multiplication functionality.

Furthermore, looking ahead, we crucially require that the basis can be altered via computational game hops during proofs. We solve this problem by linearly perturbing the given basis $[\mathbf{b}]_{\mathbb{G}}$ (except for its first entry to enable meaningful public sampling). We refer to this perturbed basis as $[\mathbf{\Omega}]_{\mathbb{G}}$. Our group element encodings are defined to carry representation vectors with respect to $[\mathbf{\Omega}]_{\mathbb{G}}$. By construction of \mathcal{C}_{add} , these representation vectors are treated homomorphically by the group operation.

To preserve tightness of security reductions, we additionally introduce a statistical re-randomization mechanism.

As opposed to [AFH+16; FHHL18], [AH18] uses a quite different approach. In [AH18], the group scheme is constructed from scratch, meaning there is no necessity for an underlying group. The consequences are twofold. On one hand, very strong decisional assumptions can be proven to hold in the resulting group scheme. On the other hand, however, the group scheme from [AH18] lacks a $\text{GetID}_{\mathbb{H}}$ algorithm limiting its applicability.

⁴⁸ With basis we mean a set of group elements in the base group.

REMARK 12.1. We note that for our construction, we implicitly build a controlled malleable [NIZK](#) as per [\[CKLM12\]](#).

12.1 MAIN THEOREM AND SECURITY ANALYSIS

In the following, we formally prove that \mathbb{H} as defined in Figures 12.1 and 12.2 is an algebraic wrapper.

THEOREM 12.1. *Let (i) $\text{GGen}_{\mathbb{G}}$ be a group parameter generator for a cyclic group \mathbb{G} , (ii) Gen_{SMP} be a hard subset membership problem, (iii) $\text{FHE} = (\text{KGen}, \text{Enc}, \text{Dec}, \text{Eval}, \text{Rerand})$ be a perfectly correct and perfectly re-randomizable fully homomorphic public-key encryption scheme, (iv) $\text{NIWI} := (\text{Setup}, \text{Prove}, \text{Verify}, \text{HSetup}, \text{Ext})$ be a perfectly complete, perfectly sound, perfectly extractable and perfectly witness-indistinguishable dual-mode [NIWI](#) proof system for the language \mathcal{L} , (v) piO be a [pIO](#) scheme for the class of samplers $\mathcal{C}^{\text{X-ind}}$ and (vi) piO_{ℓ}^* be an ℓ -expanding [pIO](#) scheme for the class of samplers $\mathcal{C}_{\ell}^{\text{X-}(\ast)\text{-ind}}$. Then, \mathbb{H} defined in Figures 12.1 and 12.2 is an algebraic wrapper.*

Proof. Since the algorithms defined in Figure 12.1 equip the base group \mathbb{G} with non-unique encodings but respect its group structure, the tuple $(\text{GGen}_{\mathbb{H}}, \text{Sam}_{\mathbb{H}}, \text{Val}_{\mathbb{H}}, \text{Eq}_{\mathbb{H}}, \text{Add}_{\mathbb{H}}, \text{GetID}_{\mathbb{H}})$ forms a group scheme such that $\text{Unwrap}_{\mathbb{H}}(\text{pp}_{\mathbb{H}}, \cdot)$ defines a group isomorphism from \mathbb{H} to \mathbb{G} . Hence, by the correctness of FHE , piO , piO_{ℓ}^* and the completeness of NIWI , \mathbb{H} satisfies \mathbb{G} -wrapping.

LEMMA 12.1. *The group scheme \mathbb{H} defined in Figures 12.1 and 12.2 satisfies extractability.*

Proof of Lemma 12.1. The algorithm $\text{PrivExt}_{\mathbb{H}}$ is deterministic. Let $\text{pp}_{\mathbb{G}} \in \text{supp}(\text{GGen}_{\mathbb{G}}(1^{\lambda}))$, $[\mathbf{b}]_{\mathbb{G}} \in \mathcal{B}_{\text{pp}_{\mathbb{G}}}^n$ and $(\text{pp}_{\mathbb{H}}, \tau_{\mathbb{H}}) \in \text{supp}(\text{GGen}_{\mathbb{H}}(\text{pp}_{\mathbb{G}}, [\mathbf{b}]_{\mathbb{G}}))$. Since $y \notin \mathcal{L}$ and because NIWI is perfectly sound, every valid encoding $\hat{h} = ([x]_{\mathbb{G}}, C, \pi)_{\mathbb{H}}$ must satisfy either relation \mathcal{R}_1 or \mathcal{R}_2 with respect to pars . Hence, decryption of C yields a vector \mathbf{v} such that $[\mathbf{\Omega}^T \cdot \mathbf{v}]_{\mathbb{G}} = [x]_{\mathbb{G}}$ or C was produced as an encryption of a vector \mathbf{v} such that the above is true. Due to perfect correctness of FHE , $\text{FHE.Dec}(sk, C)$ recovers this \mathbf{v} in both cases. Therefore, the output \mathbf{z} produced by $\text{PrivExt}_{\mathbb{H}}(\tau_{\mathbb{H}}, \hat{h})$ satisfies

$$\underbrace{(v_1 \cdot \alpha_1, \dots, v_n \cdot \alpha_n)}_{=: \mathbf{z}} \cdot [\mathbf{b}]_{\mathbb{G}} = (v_1, \dots, v_n) \cdot \underbrace{\begin{bmatrix} b_1 \cdot \alpha_1 \\ \vdots \\ b_n \cdot \alpha_n \end{bmatrix}}_{=: [\mathbf{\Omega}]_{\mathbb{G}}} = [x]_{\mathbb{G}}.$$

Since $[x]_{\mathbb{G}} = \text{Unwrap}_{\mathbb{H}}(\text{pp}_{\mathbb{H}}, \hat{h})$, extractability follows. \square

LEMMA 12.2. *The group scheme \mathbb{H} defined in Figures 12.1 and 12.2 satisfies correctness of extraction.*

GGen_H(pp_G, [b]_G)

```

parse [b]G := [(b1, ..., bn)T]G
α1 := 1, α2, ..., αn ← Zp×
[Ω]G := ([b1]Gα1, ..., [bn]Gαn)T
(pk, sk) ← FHE.KGen(1λ)
σ ← NIZK.Setup(1λ)
(L, X, R) ← GenSMP, y ← X \ L
Λadd ← piO(1Padd(λ), Cadd)
Λrerand ← piOℓ*(1Prerand(λ), Crerand)
pars := (ppG, (L, X, R), y, pk, [Ω]G)
ppH := (σ, pars, Λadd, Λrerand)
τH := (ppH, sk, α1, ..., αn, [b]G)
return (ppH, τH)

```

(a) The parameter generation algorithm.

Sam_H(pp_H, v ∈ Z_pⁿ)

```

C = FHE.Enc(pk, v; τ)
[x]G := [Ω]T · vG
π = Prove(σ, (pars, [x]G, C), (τ, v))
return ĥ := ([x]G, C, π)H

```

(b) The (public) sampling algorithm.

GetID_H(pp_H, ĥ)

```

if ¬ValH((σ, pars), ĥ) then
  return ⊥
parse ĥ := ([x]G, C, π)H
return [x]G

```

(c) The “extraction” algorithm producing a unique identifier.

Val_H(pp_H, ĥ)

```

parse ĥ := ([x]G, C, π)H
return Verify(σ, (pars, [x]G, C), π)

```

(d) The validation algorithm.

Eq_H(pp_H, ĥ₁, ĥ₂)

```

if ∃j ∈ [2]: ¬ValH((σ, pars), ĥj) then
  return ⊥
parse ĥi := ([xi]G, Ci, πi)H
return [x1]G = [x2]G

```

(e) The equality testing algorithm.

Add_H(pp_H, ĥ₁, ĥ₂)

```

return Λadd(ĥ1, ĥ2)

```

(f) The algorithm for computing the group operation.

C_{add}[pars, σ, sk](ĥ₁, ĥ₂; τ)

```

if ∃j ∈ [2]: ¬ValH((σ, pars), ĥj) then
  return ⊥
parse ĥi := ([xi]G, Ci, πi)H
[xout]G := [x1]G · [x2]G
Cout ← FHE.Eval(pk, P(+)[Zpn], C1, C2)
  // P(+)[Zpn] computes addition in Zpn
vi ← FHE.Dec(sk, Ci)
vout := v1 + v2
πout ← Prove(σ,
  (pars, [xout]G, Cout), (sk, vout))
return ĥout := ([xout]G, Cout, πout)

```

(g) The group operation circuit which is obfuscated and part of the group parameters. The algorithm Add_H simply evaluates this obfuscated circuit.

Figure 12.1: Definition of the algorithms GGen_H, Sam_H, Val_H, Eq_H, GetID_H, Add_H, Unwrap_H and the circuit C_{add}.

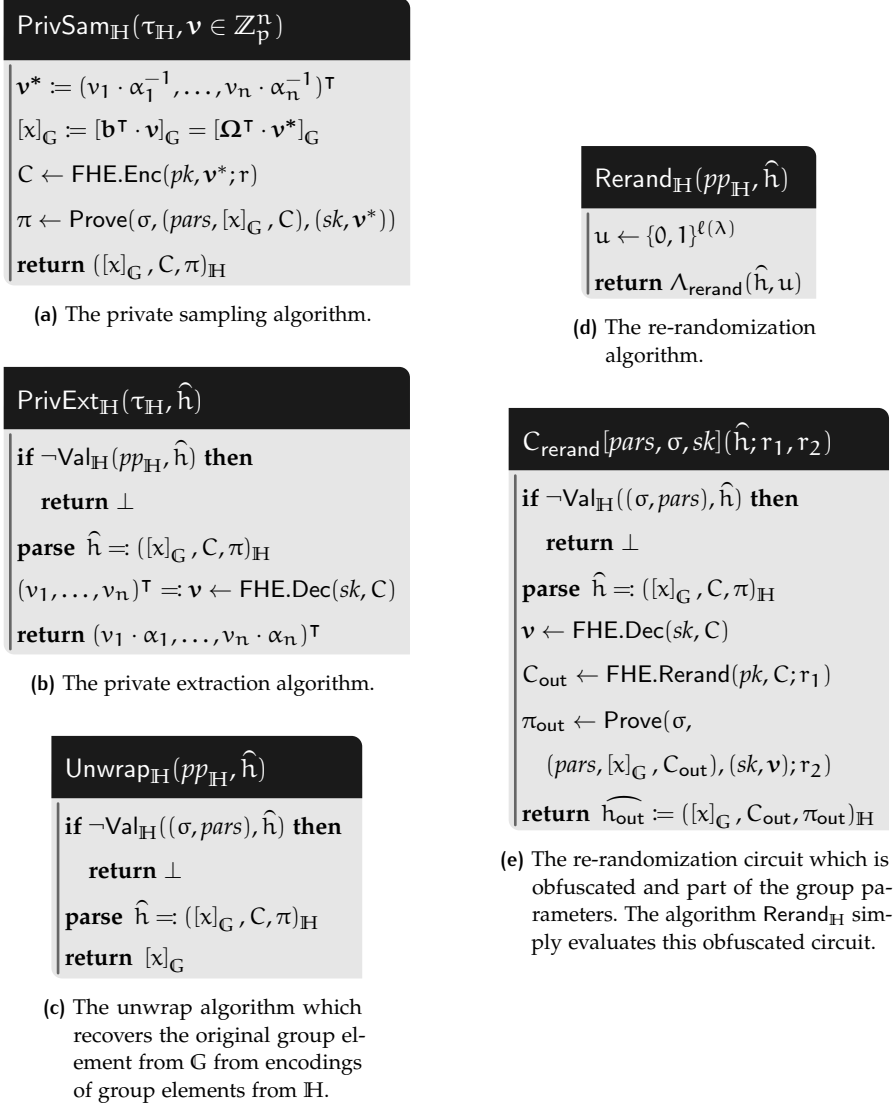


Figure 12.2: Definition of the algorithms PrivSam_H, PrivExt_H, Rerand_H, Unwrap_H and the circuit C_{rerand} .

Proof of Lemma 12.2. We first prove, that Add_H respects private extraction in \mathbb{Z}_p^n . Let $pp_G \in \text{supp}(\text{GGen}_G(1^\lambda))$, let $[\mathbf{b}]_G \in \mathcal{B}_{pp_G}^n$ and let $(pp_H, \tau_H) \in \text{supp}(\text{GGen}_H(pp_G, [\mathbf{b}]_G))$. Let

$$\begin{aligned} \widehat{h}_0 &:= ([x_0]_G, C_0, \pi_0)_H \in H \quad \text{with} \quad \mathbf{v}^{(0)} := \text{Dec}(sk, C_0) \\ \widehat{h}_1 &:= ([x_1]_G, C_1, \pi_1)_H \in H \quad \text{with} \quad \mathbf{v}^{(1)} := \text{Dec}(sk, C_1). \end{aligned}$$

Further, let

$$([x_2]_G, C_2, \pi_2)_H := \widehat{h}_2 := \text{Add}_H(pp_H, \widehat{h}_0, \widehat{h}_1) \quad \text{with} \quad \mathbf{v}^{(2)} := \text{Dec}(sk, C_2).$$

Since FHE is perfectly correct and piO is support respecting, we have $\mathbf{v}^{(2)} = \mathbf{v}^{(0)} + \mathbf{v}^{(1)}$ (in \mathbb{Z}_p^n). Therefore, we have

$$\begin{aligned} \mathbf{v}^{(2)} \circ (\alpha_1, \dots, \alpha_n)^\top &= \underbrace{\mathbf{v}^{(0)} \circ (\alpha_1, \dots, \alpha_n)^\top}_{= \text{PrivExt}_H(\tau_H, \widehat{h}_2)} + \underbrace{\mathbf{v}^{(1)} \circ (\alpha_1, \dots, \alpha_n)^\top}_{= \text{PrivExt}_H(\tau_H, \widehat{h}_0)} + \underbrace{\mathbf{v}^{(1)} \circ (\alpha_1, \dots, \alpha_n)^\top}_{= \text{PrivExt}_H(\tau_H, \widehat{h}_1)} \in \mathbb{Z}_p^n. \end{aligned}$$

Note that the operator “ \circ ” denotes the Hadamard product, i. e., $(\mathbf{a} \circ \mathbf{b})_i = (\mathbf{a})_i \cdot (\mathbf{b})_i$.

Furthermore, $\text{Rerand}_{\mathbb{H}}$ does not interfere with private extraction. Let $pp_{\mathbb{G}} \in \text{supp}(\text{GGen}_{\mathbb{G}}(1^\lambda))$, $[\mathbf{b}]_{\mathbb{G}} \in \mathcal{B}_{pp_{\mathbb{G}}}^n$ and $(pp_{\mathbb{H}}, \tau_{\mathbb{H}}) \in \text{supp}(\text{GGen}_{\mathbb{H}}(pp_{\mathbb{G}}, [\mathbf{b}]_{\mathbb{G}}))$. Let $\hat{\mathbf{h}} := ([x]_{\mathbb{G}}, C, \pi)_{\mathbb{H}} \in \mathbb{H}$ with $\mathbf{v} := \text{Dec}(sk, C)$. Further, let

$$\begin{aligned} ([x']_{\mathbb{G}}, C', \pi')_{\mathbb{H}} &:= \hat{\mathbf{h}} \in \text{supp}(\text{Rerand}_{\mathbb{H}}(pp_{\mathbb{H}}, \hat{\mathbf{h}})) \quad \text{with} \\ \mathbf{v}' &:= \text{Dec}(sk, C'). \end{aligned}$$

Since FHE is perfectly correct and piO_{ℓ}^* is support respecting, $\mathbf{v} = \mathbf{v}'$ and thus

$$\text{PrivExt}_{\mathbb{H}}(\tau_{\mathbb{H}}, \hat{\mathbf{h}}) = \mathbf{v} \circ (\alpha_1, \dots, \alpha_n)^{\top} = \text{PrivExt}_{\mathbb{H}}(\tau_{\mathbb{H}}, \hat{\mathbf{h}}').$$

□

LEMMA 12.3. *The group scheme \mathbb{H} defined in Figures 12.1 and 12.2 satisfies correctness of sampling.*

Proof of Lemma 12.3. Let $pp_{\mathbb{G}} \in \text{supp}(\text{GGen}_{\mathbb{G}}(1^\lambda))$, bases $[\mathbf{b}]_{\mathbb{G}} \in \mathcal{B}_{pp_{\mathbb{G}}}^n$ and $(pp_{\mathbb{H}}, \tau_{\mathbb{H}}) \in \text{supp}(\text{GGen}_{\mathbb{H}}(pp_{\mathbb{G}}, [\mathbf{b}]_{\mathbb{G}}))$.

We prove that $\text{PrivExt}_{\mathbb{H}}(\tau_{\mathbb{H}}, \cdot)$ and $\text{PrivSam}_{\mathbb{H}}(\tau_{\mathbb{H}}, \cdot)$ act inversely to each other. Let $\mathbf{v} \in \mathbb{Z}_{\mathbb{p}}^n$ and let $([x]_{\mathbb{G}}, C, \pi)_{\mathbb{H}} := \hat{\mathbf{h}} \in \text{supp}(\text{PrivSam}_{\mathbb{H}}(\tau_{\mathbb{H}}, \mathbf{v}))$. By correctness of FHE, $\text{FHE.Dec}(sk, C) = \mathbf{v} \circ (\alpha_1^{-1}, \dots, \alpha_n^{-1})^{\top}$. Hence, the algorithm $\text{PrivExt}_{\mathbb{H}}(\tau_{\mathbb{H}}, \hat{\mathbf{h}})$ outputs $((v_1 \cdot \alpha_1^{-1}) \cdot \alpha_1, \dots, (v_n \cdot \alpha_n^{-1}) \cdot \alpha_n)^{\top} = \mathbf{v}$.

Next, we prove that $\text{PrivExt}_{\mathbb{H}}(\tau_{\mathbb{H}}, \cdot)$ and $\text{Sam}_{\mathbb{H}}(pp_{\mathbb{H}}, \cdot)$ act inversely to each other for multiples of \mathbf{e}_1 . Let $x \in \mathbb{Z}_{\mathbb{p}}$ and let $([x]_{\mathbb{G}}, C, \pi)_{\mathbb{H}} := \hat{\mathbf{h}} \in \text{supp}(\text{Sam}_{\mathbb{H}}(pp_{\mathbb{H}}, x \cdot \mathbf{e}_1))$. By correctness of FHE, $\text{FHE.Dec}(sk, C) = x \cdot \mathbf{e}_1 \circ (\alpha_1^{-1}, \dots, \alpha_n^{-1})^{\top} = x \cdot \mathbf{e}_1$ since $\alpha_1 = 1$. Hence, $\text{PrivExt}_{\mathbb{H}}(\tau_{\mathbb{H}}, \hat{\mathbf{h}})$ outputs $(x \cdot \alpha_1, 0, \dots, 0)^{\top} = x \cdot \mathbf{e}_1$ since $\alpha_1 = 1$. □

LEMMA 12.4 (Removing information from $pp_{\mathbb{H}}$). *Let $\text{GGen}'_{\mathbb{H}}$ denote the distribution of public parameters sampled as in $\text{GGen}_{\mathbb{H}}$ with the difference that $\mathbf{y} \leftarrow L$ is a yes-instance of the subset membership problem, σ is sampled according to HSetup and $\Lambda_{\text{add}}, \Lambda_{\text{rerand}}$ are produced for the circuits $C_{\text{add}}^{(3)}$ and $C_{\text{rerand}}^{(3)}$, see Figures 12.4, 12.5 and 12.6. Then, for all legitimate PPT adversaries \mathcal{A} ,*

$$\text{Adv}_{\mathbb{H}, \mathcal{A}}^{\text{swap}}(\lambda) := \left| \Pr \left[\text{Exp}_{\mathbb{H}, \mathcal{A}, 0}^{\text{swap}}(\lambda) = 1 \right] - \Pr \left[\text{Exp}_{\mathbb{H}, \mathcal{A}, 1}^{\text{swap}}(\lambda) = 1 \right] \right|$$

is negligible, where $\text{Exp}_{\mathbb{H}, \mathcal{A}, b}^{\text{swap}}(\lambda)$ is defined in Figure 12.3 and where legitimate means that \mathcal{A} on input of $pp_{\mathbb{G}}$ guarantees $[\mathbf{b}]_{\mathbb{G}} \in \mathcal{B}_{pp_{\mathbb{G}}}^n$ for some $n \in \mathbb{N}$.

Proof of Lemma 12.4. The proof strategy is closely related to the proof of the “swap lemma” from [AFH+16; AH18]. After some preparations, we replace the no-instance of the subset membership problem with a yes-instance $\mathbf{y} \in L$ enabling relation \mathcal{R}_3 to be satisfied. The extraction trapdoor td_{ext} for σ used inside the obfuscated circuits can then be replaced by the unique witness $w_{\mathbf{y}}$ for $(\mathbf{y}, w_{\mathbf{y}}) \in R$. This enables to switch σ from binding to hiding mode. Due to perfect witness-indistinguishability, the obfuscated circuits can be replaced by variants which *always* use $w_{\mathbf{y}}$ to simulate proofs. We refer the reader to Table 12.1 for an overview of the hybrid games. We assume that the circuits $C_{\text{add}}, C_{\text{add}}^{(1)}, C_{\text{add}}^{(2)}, C_{\text{add}}^{(3)}$ are padded to size $p_{\text{add}}(\lambda)$ and the circuits $C_{\text{rerand}},$

$\text{Exp}_{\mathbb{H},\mathcal{A},0}^{\text{swap}}(\lambda)$	$\text{Exp}_{\mathbb{H},\mathcal{A},1}^{\text{swap}}(\lambda)$
$pp_{\mathbb{G}} \leftarrow \text{GGen}_{\mathbb{G}}(1^\lambda)$	$pp_{\mathbb{G}} \leftarrow \text{GGen}_{\mathbb{G}}(1^\lambda)$
$([\mathbf{b}]_{\mathbb{G}}, st) \leftarrow \mathcal{A}(1^\lambda, pp_{\mathbb{G}})$	$([\mathbf{b}]_{\mathbb{G}}, st) \leftarrow \mathcal{A}(1^\lambda, pp_{\mathbb{G}})$
$(pp_{\mathbb{H}}, \tau_{\mathbb{H}}) \leftarrow \text{GGen}_{\mathbb{H}}(pp_{\mathbb{G}}, [\mathbf{b}]_{\mathbb{G}})$	$(pp_{\mathbb{H}}, \tau_{\mathbb{H}}) \leftarrow \text{GGen}'_{\mathbb{H}}(pp_{\mathbb{G}}, [\mathbf{b}]_{\mathbb{G}})$
return $\mathcal{A}(pp_{\mathbb{H}}, \tau_{\mathbb{H}}, st)$	return $\mathcal{A}(pp_{\mathbb{H}}, \tau_{\mathbb{H}}, st)$

- (a) In game $\text{Exp}_{\mathbb{H},\mathcal{A},0}^{\text{swap}}(\lambda)$, \mathcal{A} has access to parameters distributed according to $\text{GGen}_{\mathbb{H}}$. (b) In game $\text{Exp}_{\mathbb{H},\mathcal{A},1}^{\text{swap}}(\lambda)$, \mathcal{A} has access to parameters distributed according to $\text{GGen}'_{\mathbb{H}}$.

Figure 12.3: Indistinguishability between $\text{GGen}_{\mathbb{H}}$ and $\text{GGen}'_{\mathbb{H}}$.

Table 12.1: Overview of proof steps of Lemma 12.4. Changes to previous games are highlighted in boxes.

	$y \in L$	Λ_{add}	Λ_{rerand}	σ	Remark
\mathbf{G}_0	NO	C_{add}	C_{rerand}	Setup	
\mathbf{G}_2	NO	$C_{\text{add}}^{(1)}$	$C_{\text{rerand}}^{(1)}$	Setup	piO and piO^*_ℓ
\mathbf{G}_3	YES	$C_{\text{add}}^{(1)}$	$C_{\text{rerand}}^{(1)}$	Setup	subset membership problem
\mathbf{G}_5	YES	$C_{\text{add}}^{(2)}$	$C_{\text{rerand}}^{(2)}$	Setup	piO and piO^*_ℓ
\mathbf{G}_6	YES	$C_{\text{add}}^{(2)}$	$C_{\text{rerand}}^{(2)}$	HSetup	CRS indistinguishability
\mathbf{G}_8	YES	$C_{\text{add}}^{(3)}$	$C_{\text{rerand}}^{(3)}$	HSetup	piO and piO^*_ℓ

$C_{\text{rerand}}^{(1)}$, $C_{\text{rerand}}^{(2)}$, $C_{\text{rerand}}^{(3)}$ are padded to size $p_{\text{rerand}}(\lambda)$, see Figures 12.4, 12.5 and 12.6.

GAME \mathbf{G}_0 . This game is identical to $\text{Exp}_{\mathbb{H},\mathcal{A},0}^{\text{swap}}(\lambda)$.

GAME \mathbf{G}_1 . \mathbf{G}_1 is defined as \mathbf{G}_0 but Λ_{add} is produced via $\text{piO}(1^{\text{Padd}(\lambda)}, C_{\text{add}}^{(1)})$, see Figure 12.4.

Indistinguishability between \mathbf{G}_0 and \mathbf{G}_1 follows from the security of piO with respect to $\mathcal{E}^{\text{X-ind}}$. More formally, let \mathcal{A} be a legitimate PPT adversary distinguishing \mathbf{G}_0 and \mathbf{G}_1 .

Let D_1 be the circuit sampler which on input of $1^{\text{Padd}(\lambda)}$ samples $pp_{\mathbb{G}} \leftarrow \text{GGen}_{\mathbb{G}}(1^\lambda)$, calls $\mathcal{A}(pp_{\mathbb{G}})$ to obtain $([\mathbf{b}]_{\mathbb{G}}, st)$, produces parameters exactly as $\text{GGen}_{\mathbb{H}}(pp_{\mathbb{G}}, [\mathbf{b}]_{\mathbb{G}})$ (except for the obfuscated circuit Λ_{add}), and outputs $C_0 := C_{\text{add}}$, $C_1 := C_{\text{add}}^{(1)}$ and auxiliary information $z := (\sigma, \text{pars} = (pp_{\mathbb{G}}, (L, X, R), y, pk, [\mathbf{\Omega}]_{\mathbb{G}}), \Lambda_{\text{rerand}}, (sk, \alpha_1, \dots, \alpha_n, [\mathbf{b}]_{\mathbb{G}}), st)$. Since $y \notin L$ and NIWI is perfectly sound, the circuits C_{add} and $C_{\text{add}}^{(1)}$ behave exactly identically on identical inputs and random tapes. Hence, for the function $X(\lambda) := 0$ and a differing domain $\mathcal{X} := \emptyset$, the circuit sampler D_1 satisfies X -differing inputs and X -indistinguishability and, hence, is an X -ind sampler.

We construct an adversary \mathcal{B}_1 on the security of piO . On input of $(1^{\text{Padd}(\lambda)}, C_0, C_1, z, \Lambda)$, \mathcal{B}_1 defines $pp_{\mathbb{H}} := (\sigma, \text{pars}, \Lambda, \Lambda_{\text{rerand}})$ and $\tau_{\mathbb{H}} := (pp_{\mathbb{H}}, sk, \alpha_1, \dots, \alpha_n, [\mathbf{b}]_{\mathbb{G}})$, calls $\mathcal{A}(pp_{\mathbb{H}}, \tau_{\mathbb{H}}, st)$ and outputs \mathcal{A} 's output. If Λ is produced via $\text{piO}(1^{\text{Padd}(\lambda)}, C_{\text{add}})$, D_1 together with \mathcal{B}_1 perfectly simulate \mathbf{G}_0 for \mathcal{A} .

Otherwise, if Λ is produced via $\text{piO}(1^{\text{Padd}(\lambda)}, C_{\text{add}}^{(1)})$, D_1 together with \mathcal{B}_1 perfectly simulate \mathbf{G}_1 for \mathcal{A} . Hence,

$$|\Pr[out_1 = 1] - \Pr[out_0 = 1]| \leq \text{Adv}_{\text{piO}, D_1, \mathcal{B}_1}^{\text{pio-ind}}(\text{padd}(\lambda)).$$

GAME \mathbf{G}_2 . The game \mathbf{G}_2 is identical to \mathbf{G}_1 except that Λ_{rerand} is produced via $\text{piO}_\ell^*(1^{\text{Prerand}(\lambda)}, C_{\text{rerand}}^{(1)})$, see Figure 12.4. Indistinguishability between \mathbf{G}_1 and \mathbf{G}_2 follows from the security of piO_ℓ^* with respect to $\mathfrak{C}_\ell^{\text{X-}(\star)\text{ind}}$. The proof is very similar to the previous game hop. The main difference is that the used circuit sampler needs to be in $\mathfrak{C}_\ell^{\text{X-}(\star)\text{ind}}$.

Let \mathcal{A} be a legitimate PPT adversary distinguishing \mathbf{G}_1 and \mathbf{G}_2 . Let D_2 be the circuit sampler which on input of $1^{\text{Prerand}(\lambda)}$ samples $pp_G \leftarrow \text{GGen}_G(1^\lambda)$, calls $\mathcal{A}(pp_G)$ to obtain $([\mathbf{b}]_G, st)$ and produces parameters as $\text{GGen}_H(pp_G, [\mathbf{b}]_G)$ (except for the obfuscated circuit Λ_{rerand} and such that Λ_{add} is produced for $C_{\text{add}}^{(1)}$), and outputs $C_0 := C_{\text{rerand}}$, $C_1 := C_{\text{rerand}}^{(1)}$ and auxiliary information $z := (\sigma, \text{pars} = (pp_G, (L, X, R), y, pk, [\mathbf{\Omega}]_G), \Lambda_{\text{add}}, (sk, \alpha_1, \dots, \alpha_n, [\mathbf{b}]_G), st)$. Since $y \notin L$ and NIWI is perfectly sound, the circuits C_{rerand} and $C_{\text{rerand}}^{(1)}$ behave exactly identically on identical inputs and random tapes. Therefore, by Corollary 10.1, $D_2 \in \mathfrak{C}_\ell^{\text{X-}(\star)\text{ind}}$. Hence, there exists a PPT adversary \mathcal{B}_2 such that

$$|\Pr[out_2 = 1] - \Pr[out_1 = 1]| \leq \text{Adv}_{\text{piO}_\ell^*, D_2, \mathcal{B}_2}^{\text{pio-ind}(\star)}(\text{Prerand}(\lambda)).$$

GAME \mathbf{G}_3 . The game \mathbf{G}_3 is defined as game \mathbf{G}_2 except that $y \leftarrow L$ is a yes-instance of the subset membership problem. This game hop is justified by the hardness of the subset membership problem. More precisely, for every PPT distinguisher \mathcal{A} between \mathbf{G}_2 and \mathbf{G}_3 , there exists an adversary \mathcal{B}_3 such that

$$|\Pr[out_3 = 1] - \Pr[out_2 = 1]| \leq \text{Adv}_{\text{Gen}_{\text{SMP}}, \mathcal{B}_3}^{\text{smP}}(\lambda).$$

GAME \mathbf{G}_4 . \mathbf{G}_4 is defined as \mathbf{G}_3 but Λ_{add} is produced via $\text{piO}(1^{\text{Padd}(\lambda)}, C_{\text{add}}^{(2)})$, see Figure 12.5. Due to the perfect extractability of NIWI and the fact that for every $y \in L$ there exists exactly one witness w_y for the statement $y \in L$, the two circuits $C_{\text{add}}^{(1)}$ and $C_{\text{add}}^{(2)}$ behave exactly identically on identical inputs and random tapes. Hence, for every PPT adversary \mathcal{A} , there exists a circuit sampler $D_4 \in \mathfrak{C}^{\text{X-ind}}$ and a PPT adversary \mathcal{B}_4 such that

$$|\Pr[out_4 = 1] - \Pr[out_3 = 1]| \leq \text{Adv}_{\text{piO}, D_4, \mathcal{B}_4}^{\text{pio-ind}}(\text{padd}(\lambda)).$$

GAME \mathbf{G}_5 . The game \mathbf{G}_5 is identical to \mathbf{G}_4 except that Λ_{rerand} is produced via $\text{piO}_\ell^*(1^{\text{Prerand}(\lambda)}, C_{\text{rerand}}^{(2)})$, see Figure 12.5. Indistinguishability between \mathbf{G}_4 and \mathbf{G}_5 follows from the security of piO_ℓ^* with respect to $\mathfrak{C}_\ell^{\text{X-}(\star)\text{ind}}$. Again, due to the perfect extractability of NIWI and the fact that for every $y \in L$ there exists exactly one witness w_y for the statement $y \in L$, the two circuits $C_{\text{rerand}}^{(1)}$ and $C_{\text{rerand}}^{(2)}$ behave exactly identically on identical inputs and random tapes. Hence, by Corollary 10.1, for every PPT adversary \mathcal{A} , there exists a circuit sampler $D_5 \in \mathfrak{C}_\ell^{\text{X-}(\star)\text{ind}}$ and a PPT adversary \mathcal{B}_5 such that

$$|\Pr[out_5 = 1] - \Pr[out_4 = 1]| \leq \text{Adv}_{\text{piO}_\ell^*, D_5, \mathcal{B}_5}^{\text{pio-ind}(\star)}(\text{Prerand}(\lambda)).$$

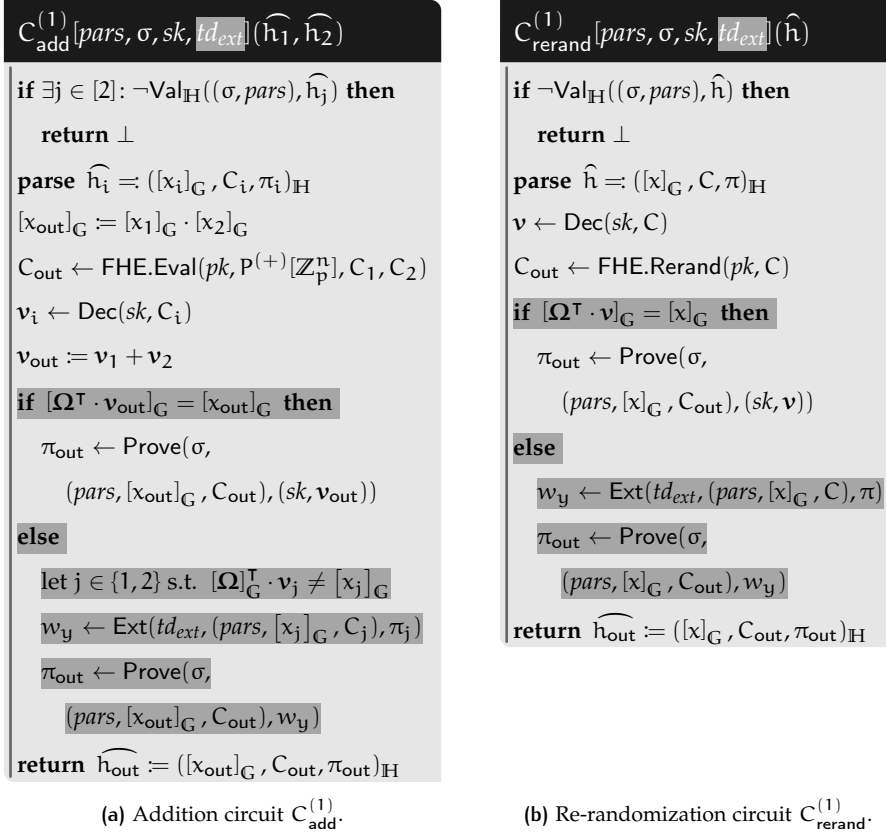


Figure 12.4: Addition and re-randomization circuits $C_{\text{add}}^{(1)}$ and $C_{\text{rerand}}^{(1)}$ which are capable to deal with simulated proofs by extracting a witness for relation \mathcal{R}_3 .

Thus, \mathbf{G}_5 does not make use of the extraction trapdoor td_{ext} corresponding to σ anymore.

GAME \mathbf{G}_6 . The game \mathbf{G}_6 is defined as \mathbf{G}_5 , except that σ is produced via $\text{HSetup}(1^\lambda)$, i. e., in hiding mode. This game hop is justified by CRS indistinguishability of NIWI. Hence, there exists a PPT adversary \mathcal{B}_6 such that

$$|\Pr[out_6 = 1] - \Pr[out_5 = 1]| \leq \text{Adv}_{\text{NIWI}, \mathcal{B}_6}^{\text{crs-ind}}(\lambda).$$

GAME \mathbf{G}_7 . \mathbf{G}_7 is defined as \mathbf{G}_6 but Λ_{add} is produced via $\text{piO}(1^{\text{padd}(\lambda)}, C_{\text{add}}^{(3)})$, see Figure 12.6. Let D_7 be a circuit sampler, which produces public parameters as in \mathbf{G}_6 and outputs the circuits $C_0 := C_{\text{add}}^{(2)}$ and $C_1 := C_{\text{add}}^{(3)}$ (and suitable auxiliary information z). The circuit sampler D_7 is an \mathcal{X} -ind sampler for $\mathcal{X}(\lambda) := 2^{2^{\text{penc}(\lambda)}} \leq 2^{\text{padd}(\lambda)}$ and $\mathcal{X} := \{0, 1\}^{2^{\text{penc}(\lambda)}}$ since the \mathcal{X} -differing inputs property is trivially satisfied and \mathcal{X} -indistinguishability is satisfied since perfect witness indistinguishability of NIWI implies that for every input $x = (\widehat{a}_1, \widehat{a}_2) \in \{0, 1\}^{2^{\text{penc}(\lambda)}}$, the output distributions produced by $C_{\text{add}}^{(2)}(x)$ and $C_{\text{add}}^{(3)}(x)$ are identical. Hence, for every PPT adversary \mathcal{A} , there exists a circuit sampler $D_7 \in \mathcal{C}^{\mathcal{X}\text{-ind}}$ and a PPT adversary \mathcal{B}_7 such that

$$|\Pr[out_7 = 1] - \Pr[out_6 = 1]| \leq \text{Adv}_{\text{piO}, D_7, \mathcal{B}_7}^{\text{pio-ind}}(\text{padd}(\lambda)).$$

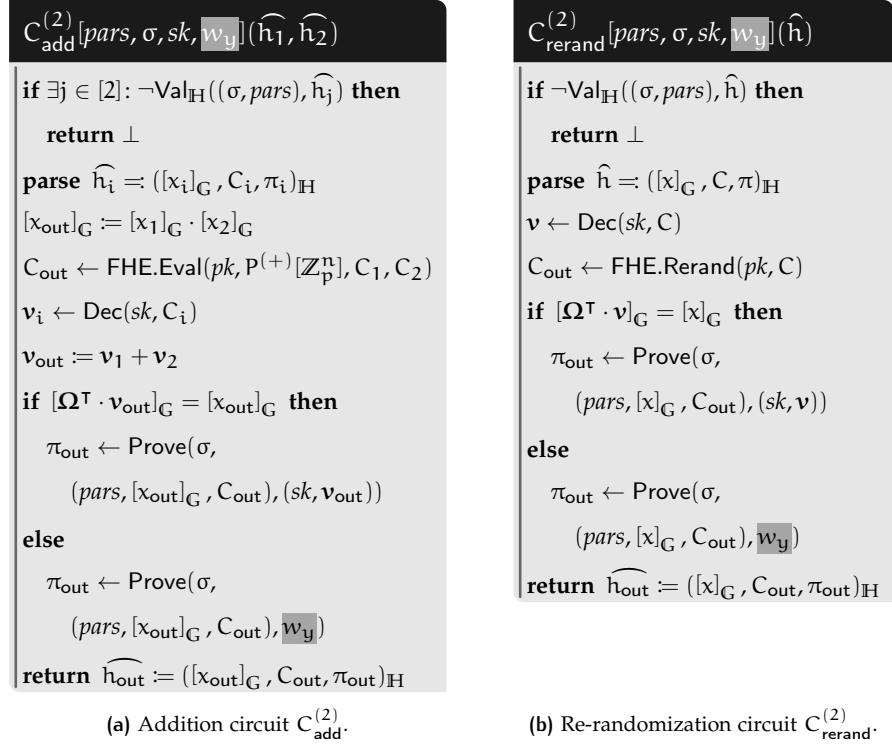


Figure 12.5: Addition and re-randomization circuits $C_{\text{add}}^{(2)}$ and $C_{\text{rerand}}^{(2)}$ which use a hard-coded witness w_y for relation \mathcal{R}_3 .

GAME \mathbf{G}_8 . The game \mathbf{G}_8 is identical to \mathbf{G}_7 except that Λ_{rerand} is produced via $\text{piO}_{\ell}^*(1^{\text{Prerand}(\lambda)}, C_{\text{rerand}}^{(3)})$, see Figure 12.6. Indistinguishability between \mathbf{G}_7 and \mathbf{G}_8 follows from the security of piO_{ℓ}^* with respect to $\mathcal{C}_{\ell}^{X-(*)\text{ind}}$. Again, due to the perfect witness indistinguishability of NIWI we have that for every input $x = \widehat{a} \in \{0, 1\}^{\text{penc}(\lambda)}$, the output distributions produced by $C_{\text{rerand}}^{(2)}(x)$ and $C_{\text{rerand}}^{(3)}(x)$ are identical.

Let D_8 be the circuit sampler which on input of $1^{\text{Prerand}(\lambda)}$ samples $pp_{\mathbb{G}} \leftarrow \text{GGen}_{\mathbb{G}}(1^\lambda)$, calls $\mathcal{A}(pp_{\mathbb{G}})$ to obtain $([b]_{\mathbb{G}}, st)$ and produces parameters as in \mathbf{G}_7 (except for the obfuscated circuit Λ_{rerand}) and outputs $C_0 := C_{\text{rerand}}^{(2)}$, $C_1 := C_{\text{rerand}}^{(3)}$ and auxiliary information $z := (\sigma, pars = (pp_{\mathbb{G}}, (L, X, R), y, pk, [\Omega]_{\mathbb{G}}), \Lambda_{\text{add}}, (sk, \alpha_1, \dots, \alpha_n, [b]_{\mathbb{G}}), st)$. By Lemma 10.1, $D_8 \in \mathcal{C}_{\ell}^{X-(*)\text{ind}}$ and, thus, for every PPT adversary \mathcal{A} , there exists a circuit sampler $D_8 \in \mathcal{C}_{\ell}^{X-(*)\text{ind}}$ and a PPT adversary \mathcal{B}_8 such that

$$|\Pr[out_8 = 1] - \Pr[out_7 = 1]| \leq \text{Adv}_{\text{piO}_{\ell}^*, D_8, \mathcal{B}_8}^{\text{pio-ind}(\ast)}(\text{Prerand}(\lambda))$$

is negligible.

Note that \mathbf{G}_8 is defined as $\text{Exp}_{\mathbb{H}, \mathcal{A}, 1}^{\text{swap}}(\lambda)$. Therefore,

$$\text{Adv}_{\mathbb{H}, \mathcal{A}}^{\text{swap}}(\lambda) \leq |\Pr[out_0 = 1] - \Pr[out_8 = 1]|$$

is negligible in λ . □

LEMMA 12.5. *The group scheme \mathbb{H} defined in Figures 12.1 and 12.2 satisfies k -switching.*

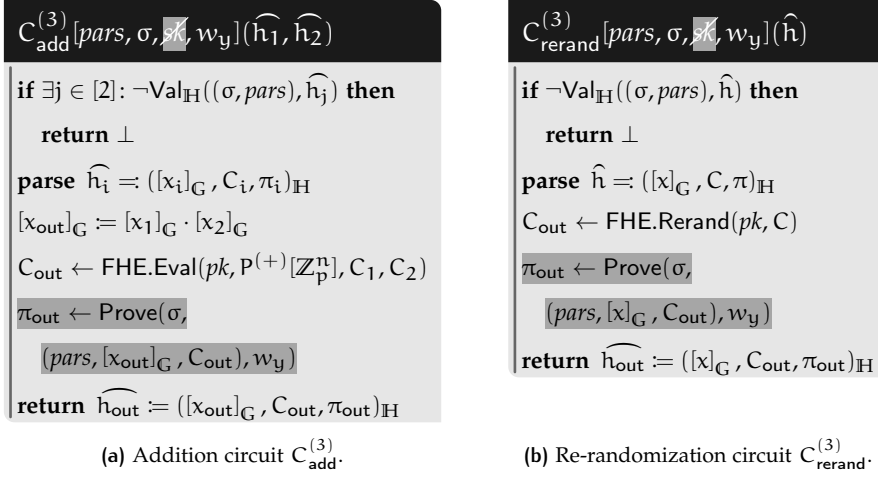


Figure 12.6: Addition and re-randomization circuits $C_{\text{add}}^{(3)}$ and $C_{\text{rerand}}^{(3)}$ which always produce consistency proofs for relation \mathcal{R}_3 .

Proof of Lemma 12.5. We recall that an adversary \mathcal{A} is a legitimate k -switching adversary if \mathcal{A} , on input of base group parameters $pp_{\mathbb{G}}$, always guarantees that for all $i \in [k]$, $[\mathbf{b}^{(0)}]_{\mathbb{G}}, [\mathbf{b}^{(1)}]_{\mathbb{G}} \in \mathcal{B}_{pp_{\mathbb{G}}}^n$ and $\mathbf{v}^{(0),(i)}, \mathbf{v}^{(1),(i)} \in \mathbb{Z}_p^n$ and $[x_i^*]_{\mathbb{G}} := [(\mathbf{v}^{(0),(i)})^{\top} \cdot \mathbf{b}^{(0)}]_{\mathbb{G}} = [(\mathbf{v}^{(1),(i)})^{\top} \cdot \mathbf{b}^{(1)}]_{\mathbb{G}}$. We proceed over a series of hybrids as defined in Figure 12.7.

GAME G_0^b . This game is the original switching game $\text{Exp}_{\mathbb{H}, \mathcal{A}, b}^{\text{k-switching}}(\lambda)$ as defined in Figure 11.1a.

GAME G_1^b . The difference between G_1^b and G_0^b is only conceptual. The only difference is that $[x_i^*]_{\mathbb{G}}$ is computed as $[(\mathbf{b}^{(0)})^{\top} \cdot \mathbf{v}^{(0),(i)}]_{\mathbb{G}}$ instead of $[(\mathbf{b}^{(b)})^{\top} \cdot \mathbf{v}^{(b),(i)}]_{\mathbb{G}}$. Since \mathcal{A} is legitimate, $[x_i^*]_{\mathbb{G}} = [(\mathbf{b}^{(0)})^{\top} \cdot \mathbf{v}^{(0),(i)}]_{\mathbb{G}} = [(\mathbf{b}^{(1)})^{\top} \cdot \mathbf{v}^{(1),(i)}]_{\mathbb{G}}$, $[x_i^*]_{\mathbb{G}}$ is independent of b . Hence, $\Pr[\text{out}_1^b = 1] = \Pr[\text{out}_0^b = 1]$.

GAME G_2^b . The only difference between the games G_1^b and G_2^b is that $pp_{\mathbb{H}}$ is produced via $\text{GGen}_{\mathbb{H}}(pp_{\mathbb{G}}, [\mathbf{b}^{(b)}]_{\mathbb{G}})$ and $\text{GGen}'_{\mathbb{H}}(pp_{\mathbb{G}}, [\mathbf{b}^{(b)}]_{\mathbb{G}})$, respectively. Let \mathcal{A} be a distinguisher between G_1^b and G_2^b . We construct an adversary \mathcal{B}_2^b against the property of Lemma 12.4. On input of $pp_{\mathbb{G}}$, \mathcal{B}_2^b calls \mathcal{A} on input of $pp_{\mathbb{G}}$ to obtain $([\mathbf{b}^{(0)}]_{\mathbb{G}}, [\mathbf{b}^{(1)}]_{\mathbb{G}}, (\mathbf{v}^{(0),(i)})_{i \in [k]}, (\mathbf{v}^{(1),(i)})_{i \in [k]}, st)$. \mathcal{B}_2^b outputs $([\mathbf{b}^{(b)}]_{\mathbb{G}}, st)$ and obtains $(pp_{\mathbb{H}}, \tau_{\mathbb{H}})$ which is either sampled according to $\text{GGen}_{\mathbb{H}}(pp_{\mathbb{G}}, [\mathbf{b}^{(b)}]_{\mathbb{G}})$ or according to $\text{GGen}'_{\mathbb{H}}(pp_{\mathbb{G}}, [\mathbf{b}^{(b)}]_{\mathbb{G}})$. Since, $\tau_{\mathbb{H}}$ contains sk and $(\alpha_1, \dots, \alpha_n)$, \mathcal{B}_2^b is able to simulate the game G_1^b (respectively, G_2^b) for \mathcal{A} . More precisely, \mathcal{B}_2^b samples \widehat{h}_i^* as in G_1^b , calls \mathcal{A} on input of $(pp_{\mathbb{H}}, (\widehat{h}_i^*)_{i \in [k]})$ and outputs \mathcal{A} 's output. If $pp_{\mathbb{H}}$ is sampled using $\text{GGen}_{\mathbb{H}}$, \mathcal{B}_2^b perfectly simulates G_1^b for \mathcal{A} and if $pp_{\mathbb{H}}$ is sampled using $\text{GGen}'_{\mathbb{H}}$, \mathcal{B}_2^b perfectly simulates G_2^b for \mathcal{A} . Hence,

$$\left| \Pr[\text{out}_2^b = 1] - \Pr[\text{out}_1^b = 1] \right| \leq \text{Adv}_{\mathbb{H}, \mathcal{B}_2^b}^{\text{swap}}(\lambda).$$

GAME G_3^b . G_3^b is identical to G_2^b but instead of using the witnesses (sk, \mathbf{v}_i^*) to compute the consistency proofs π_i^* , G_3^b uses w_y as witness for all consistency

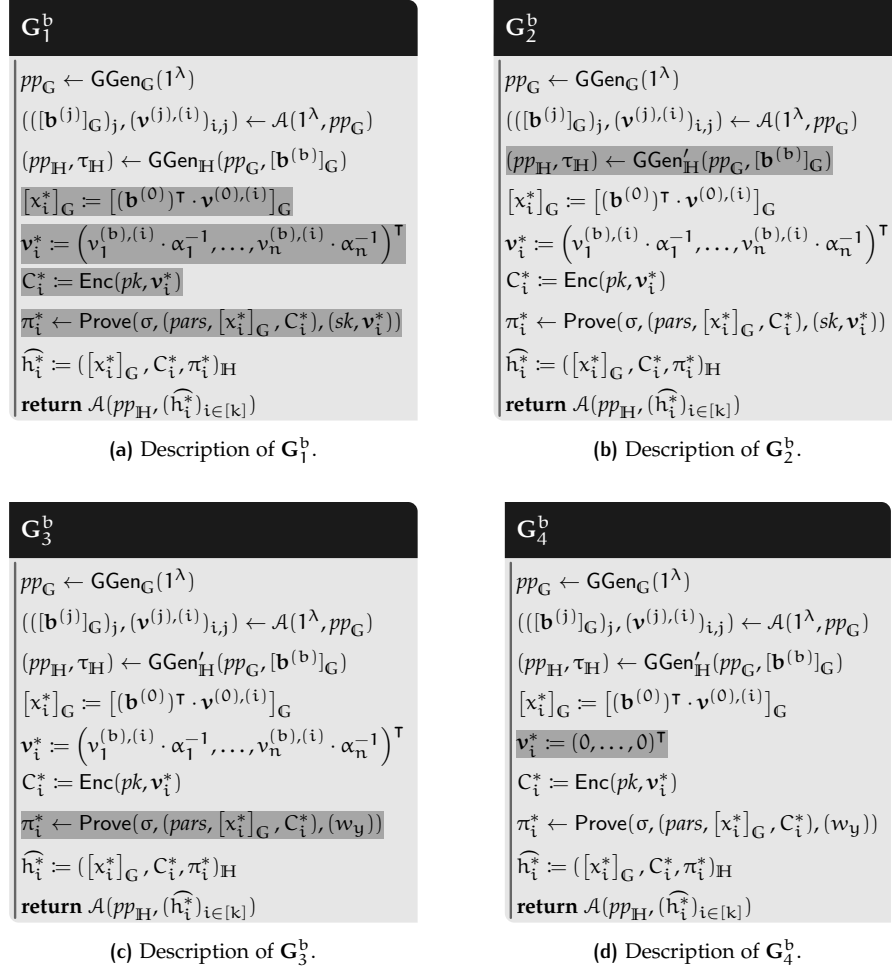


Figure 12.7: Games used in proof of Lemma 12.5. Recall that $\alpha_1 = 1$ and $\alpha_2, \dots, \alpha_n$ are uniformly distributed over \mathbb{Z}_p^\times .

proofs. Since σ is produced via $\text{HSetup}(1^\lambda)$, the proofs produced in G_3^b and G_2^b are *identically* distributed due to the perfect witness indistinguishability of NIWI. Hence, $\Pr[out_2^b = 1] = \Pr[out_3^b = 1]$.

GAME G_4^b . G_4^b is identical to G_3^b except for the fact that G_4^b always uses the zero vector $(0, \dots, 0)^\top$ as \mathbf{v}_i^* instead of $\mathbf{v}_i^* := (v_1^{(b),(i)} \cdot \alpha_1^{-1}, \dots, v_n^{(b),(i)} \cdot \alpha_n^{-1})^\top$. Note that the secret decryption key sk is not necessary to simulate G_3^b and G_4^b . Additionally, \mathbf{v}_i^* is only used to call $\text{Enc}(pk, \mathbf{v}_i^*)$ (in G_3^b). This allows to apply the IND-CPA security of FHE.

CLAIM 12.1. For all legitimate k -switching PPT adversaries \mathcal{A} , there exists a (legitimate) PPT adversary \mathcal{B}_4^b such that

$$\left| \Pr[out_4^b = 1] - \Pr[out_3^b = 1] \right| \leq k \cdot \text{Adv}_{\text{FHE}, \mathcal{B}_4^b}^{\text{ind-cpa}}(\lambda).$$

Proof of Claim 12.1. We use a standard hybrid argument.

GAME $G_{3,1}^b$. The hybrid game $G_{3,1}^b$ is defined as G_3^b except that for $j < i$, C_j^* is produced via $\text{Enc}(pk, \mathbf{0})$ and for $j \geq i$, C_j^* is produced for $\text{Enc}(pk, \mathbf{v}_j^*)$. Hence, $G_{3,1}^b$ is identical to G_3^b and $G_{3,k+1}^b$ is identical to G_4^b . We construct an adversary $\mathcal{B}_{3,1}^b$ for the IND-CPA security of FHE. On input of pk , $\mathcal{B}_{3,1}^b$ guesses

an index $\nu \in [k]$, samples $pp_G \leftarrow \text{GGen}_G(1^\lambda)$, calls $\mathcal{A}(pp_G)$, obtains \mathcal{A} 's output $(([\mathbf{b}^{(j)}]_G)_{j \in \{0,1\}}, (\mathbf{v}^{(j),(i)})_{i \in [k], j \in \{0,1\}})$ and produces pp_H as in \mathbf{G}_3^b . Note that $\mathcal{B}_{3,1}^b$ knows $\alpha_1, \dots, \alpha_n$ and w_y . For $\mu \in \{1, \dots, \nu - 1\}$, $\mathcal{B}_{3,1}^b$ produces \widehat{h}_μ^* as in \mathbf{G}_3^b . For $\mu \in \{\nu + 1, \dots, k\}$, $\mathcal{B}_{3,1}^b$ produces \widehat{h}_μ^* as in \mathbf{G}_4^b . For $\mu = \nu$, $\mathcal{B}_{3,1}^b$ outputs

$$\begin{aligned} M_0 &:= \left(v_1^{(b),(\mu)} \cdot \alpha_1^{-1}, \dots, v_n^{(b),(\mu)} \cdot \alpha_n^{-1} \right)^\top \\ M_1 &:= (0, \dots, 0)^\top \end{aligned}$$

to the IND-CPA game together with his internal state. On input of C^* , $\mathcal{B}_{3,1}^b$ uses C^* as C_μ^* , computes

$$\begin{aligned} [x_\mu^*]_G &:= \left[(\mathbf{b}^{(0)})^\top \cdot \mathbf{v}^{(0),(\mu)} \right]_G, \\ \pi_\mu^* &\leftarrow \text{Prove} \left(\sigma, (pars, [x_\mu^*]_G, C_\mu^*), (w_y) \right). \end{aligned}$$

Finally, $\mathcal{B}_{3,1}^b$ calls \mathcal{A} on input $(pp_H, (\widehat{h}_i^*)_{i \in [k]})$ and outputs \mathcal{A} 's output. If C^* is an encryption of M_0 , $\mathcal{B}_{3,1}^b$ perfectly simulates $\mathbf{G}_{3,\nu}^b$, otherwise $\mathcal{B}_{3,1}^b$ perfectly simulates $\mathbf{G}_{3,\nu+1}^b$. Hence,

$$\left| \Pr[out_4^b = 1] - \Pr[out_3^b = 1] \right| \leq k \cdot \text{Adv}_{\text{FHE}, \mathcal{B}_{3,1}^b}^{\text{ind-cpa}}(\lambda).$$

□

Recall that every legitimate k -switching adversary always guarantees

$$\left[(\mathbf{v}^{(0),(i)})^\top \cdot \mathbf{b}^{(0)} \right]_G = \left[(\mathbf{v}^{(1),(i)})^\top \cdot \mathbf{b}^{(1)} \right]_G$$

and $[\mathbf{b}^{(0)}]_G, [\mathbf{b}^{(1)}]_G \in \mathcal{B}_{pp_G}^n$. Therefore, \mathbf{G}_4^0 and \mathbf{G}_4^1 only differ in the fact that in \mathbf{G}_4^b the vector $[\mathbf{b}^{(b)}]_G$ is used to compute $[\Omega]_G$. That is, $[\Omega]_G = ([1]_G, [b_{2,b}]_G^{\alpha_2}, \dots, [b_{n,b}]_G^{\alpha_n})^\top$ for $\alpha_2, \dots, \alpha_n$ chosen uniformly at random from \mathbb{Z}_p^n . Except for $[\Omega]_G$, the view of \mathcal{A} is independent of $\alpha_2, \dots, \alpha_n$.⁴⁹ Thus, $\Pr[out_4^0 = 1] = \Pr[out_4^1 = 1]$.

Note that since G has unique encodings, \mathcal{A} is unable to extract auxiliary information from the encodings of $[x_i^*]_G$. This is crucial since such auxiliary information may for instance reveal how $[x_i^*]_G$ was computed. □

LEMMA 12.6. *The group scheme \mathbb{H} defined in Figures 12.1 and 12.2 satisfies statistical re-randomizability.*

Proof of Lemma 12.6. The circuit C_{rerand} takes inputs from $\{0, 1\}^{\text{penc}(\lambda)}$ and expects randomness from $\{0, 1\}^{m'(\lambda)} \times \{0, 1\}^{m''(\lambda)}$. We recall that piO_ℓ^* is an ℓ -expanding **PIO** scheme for $\ell(\lambda) = m'(\lambda) + m''(\lambda) + 2(\lambda + 1) + 3$. Since for every distribution Z_1 over $\{0, 1\}^{\text{penc}(\lambda)}$,

$$\widetilde{H}_\infty(\mathcal{U}_{\ell(\lambda)} \mid Z_1) = \ell(\lambda) > m'(\lambda) + m''(\lambda) + 2(\lambda + 1) + 2,$$

the statistical distance

$$\begin{aligned} \Delta \left(\left\{ \Lambda_{\text{rerand}} \leftarrow \text{piO}_\ell^*(C_{\text{rerand}}) : (\Lambda_{\text{rerand}}, \Lambda_{\text{rerand}}(Z_1, Z_2)) \right\}, \right. \\ \left. \left\{ \Lambda_{\text{rerand}} \leftarrow \text{piO}_\ell^*(C_{\text{rerand}}) : (\Lambda_{\text{rerand}}, C_{\text{rerand}}(Z_1; \mathcal{U}_{m'(\lambda)+m''(\lambda)})) \right\} \right) \end{aligned} \quad (12.1)$$

⁴⁹ Note that since $[\mathbf{b}^{(0)}]_G, [\mathbf{b}^{(1)}]_G \in \mathcal{B}_{pp_G}^n$, we have $[b_{i,b}]_G \neq [0]_G$ (for all $i \in [n]$).

is at most $2^{-(\lambda+1)}$.

Let $\widehat{h}_0 = ([x_0]_{\mathbb{G}}, C_0, \pi_0)_{\mathbb{H}}$, $\widehat{h}_1 = ([x_1]_{\mathbb{G}}, C_1, \pi_1)_{\mathbb{H}} \in \mathbb{H}$ be the encodings chosen by the adversary \mathcal{A} . Since \mathcal{A} is a legitimate re-randomization adversary, $\text{PrivExt}_{\mathbb{H}}(\tau_{\mathbb{H}}, \widehat{h}_0) = \text{PrivExt}_{\mathbb{H}}(\tau_{\mathbb{H}}, \widehat{h}_1)$. Due to perfect correctness of FHE and since $\alpha_1, \dots, \alpha_n \in \mathbb{Z}_p^\times$ are invertible, $\text{Dec}(sk, C_0) = \text{Dec}(sk, C_1)$. Due to perfect re-randomizability of FHE, the ciphertexts produced via $C_{\text{rerand}}(\widehat{h}_0)$ and $C_{\text{rerand}}(\widehat{h}_1)$ are identically distributed. Furthermore, since $C_{\text{rerand}}(\widehat{h}_b)$ produces the consistency proof using the witness $(sk, \text{Dec}(sk, C_b))$, the statistical distance

$$\Delta\left(C_{\text{rerand}}\left(\widehat{h}_0; \mathcal{U}_{m'(\lambda)+m''(\lambda)}\right), C_{\text{rerand}}\left(\widehat{h}_1; \mathcal{U}_{m'(\lambda)+m''(\lambda)}\right)\right) = 0. \quad (12.2)$$

Therefore, combining Equations (12.1) and (12.2) yields

$$\text{Adv}_{\mathbb{H}, \mathcal{A}}^{\text{rerand}}(\lambda) \leq 2 \cdot 2^{-(\lambda+1)} = 2^{-\lambda}.$$

Note that since \mathbb{G} has unique encodings, \mathcal{A} is unable to extract auxiliary information from the encodings of $\text{Unwrap}_{\mathbb{H}}(pp_{\mathbb{H}}, \widehat{h})$. This is crucial since such auxiliary information may be used to distinguish whether \widehat{h}_0 or \widehat{h}_1 was used to derive \widehat{h} . \square

Theorem 12.1 then follows by combining Lemmas 12.1, 12.2, 12.3, 12.5 and 12.6. \square

The ElGamal public-key encryption scheme [ELG85] is a classical public-key encryption scheme based on cyclic groups. Given some cyclic group \mathbb{G} of prime order with generator g , a public-key secret-key pair is of the form (g^y, y) for $y \in \mathbb{Z}_p$. Encryption of some message $M \in \mathbb{G}$ chooses a random $x \leftarrow \mathbb{Z}_p$ and outputs the ciphertext $(C_1, C_2) := (g^x, g^x \cdot M)$. Decryption recovers M by computing $C_2 \cdot (C_1)^{-y}$. The ElGamal encryption scheme is IND-CPA secure under the DDH assumption in \mathbb{G} , [TY98]. That is, no adversary can distinguish encryptions of two messages, even if the public key is known. The de-facto security notion for public-key encryption is IND-CCA2 security introduced in [RS92]. IND-CCA2 security requires that no adversary is able to distinguish encryptions of two messages even when granted the ability to decrypt arbitrary ciphertexts during the entire experiment (even after given the challenge ciphertext), as long as he does not use this ability to decrypt the challenge ciphertext. The ElGamal encryption scheme is homomorphic and can hence not be IND-CCA2 secure. As detailed in [FPS20], the ElGamal encryption scheme can be made IND-CCA2 secure by making it plaintext-aware [BR95] or by considering its hashed variant [ABR01].

One possibility of making ElGamal plaintext-aware is to use a Schnorr signature with g^x as public key. The resulting scheme is called Schnorr-signed ElGamal [Jak98; TY98]. [SJ00] show that Schnorr-signed ElGamal is IND-CCA2 secure in the generic group model and the random oracle model and [TY98] prove Schnorr-signed ElGamal IND-CCA2 secure under some non-standard knowledge-type assumptions on Schnorr signatures.

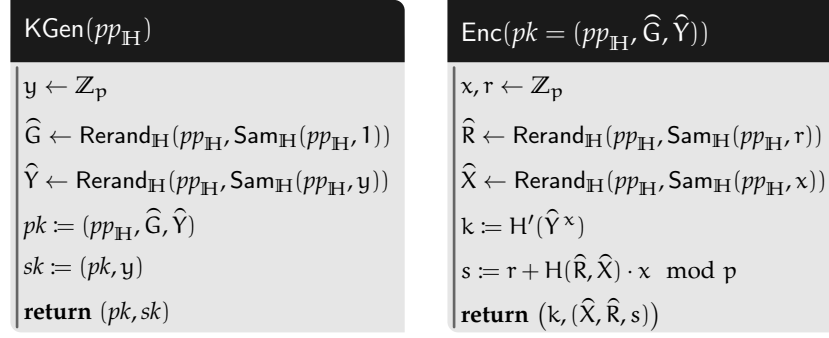
The hashed variant of the ElGamal encryption scheme derives the key via $H((g^y)^x)$ (instead of $(g^y)^x$). More precisely, in the hashed ElGamal key-encapsulation mechanism (KEM), a public key is a group element Y , the corresponding secret key is $y = \text{dlog}_g(Y)$. For encryption, one picks a random exponent $x \leftarrow \mathbb{Z}_p$ to compute a key $H(Y^x)$ accompanied by an encapsulation $X := g^x$. Given the encapsulation and the secret key y , the receiver can recover that key $K = H(X^y)$. In the random oracle model the hashed ElGamal key encapsulation algorithm (KEM) is IND-CCA2 secure under the gap Diffie-Hellman assumption⁵⁰ [CS03].

[FPS20] showed that a combination of Schnorr-signed ElGamal and hashed ElGamal (henceforth simply called Schnorr-signed ElGamal) is tightly⁵¹ IND-CCA2 secure under the DL assumption in the algebraic group model and the random oracle model. Schnorr-signed ElGamal (see Figure 13.1) works similarly as hashed ElGamal but every encapsulation is accompanied by a Schnorr signature for message X under public key X . Decryption works as before with the difference that decryption aborts if the provided Schnorr signature is invalid.

In this chapter, we demonstrate that the algebraic wrapper can be applied to mimic the proof of tight IND-CCA2 security of Schnorr-signed ElGamal E_{SIG} from [FPS20]. Note that in contrast to the tight reduction for Schnorr


⁵⁰ The gap Diffie-Hellman assumption states that the computational Diffie-Hellman problem is hard even in the presence of a decisional Diffie-Hellman oracle.

⁵¹ A cryptographic building block is tightly secure under some assumption A if every adversary breaking the building block can be transformed into a problem solver solving A with roughly the same runtime and success probability as the adversary.



(a) The key generation algorithm.

(b) The encapsulation algorithm.



(c) The decapsulation algorithm.

Figure 13.1: The Schnorr-signed ElGamal encryption scheme $E_{s\text{ElG}}$ instantiated with an algebraic wrapper \mathbb{H} . Note that to compensate for the non-uniqueness of group element encodings, the (random oracle) hash value of a group element encoding is computed for the identifier produced by $\text{GetID}_{\mathbb{H}}(pp_{\mathbb{H}}, \cdot)$ (which is unique for all group element encodings of the same group element). The hash function H maps tuples of group elements from \mathbb{H} to elements in \mathcal{K} and the hash function H' maps group elements from \mathbb{H} to \mathbb{Z}_p elements.

signatures from [AHK20], the tightness for Schnorr-signed ElGamal does not require the “origin element trick” in combination with statistical re-randomizability since it is not necessary to switch the representations of oracle responses. Note that the “origin element trick” is not part of this thesis.

THEOREM 13.1. *Let $\text{GGen}_{\mathbb{G}}$ be a group generator for a cyclic group \mathbb{G} such that **DL** is hard relative to $\text{GGen}_{\mathbb{G}}$ and let \mathbb{H} be an algebraic wrapper for \mathbb{G} . Then, $E_{s\text{ElG}}$ in \mathbb{H} is tightly **IND-CCA2** secure in the random oracle model.*

*More precisely, for all **PPT** adversaries \mathcal{A} , there exists a **PPT** adversary \mathcal{B} and a legitimate switching adversary \mathcal{A}' both running in time $T(\mathcal{B}) \approx T(\mathcal{A}) + (q_d + q_h) \cdot \text{poly}(\lambda)$ and $T(\mathcal{A}') \approx T(\mathcal{A}) + (q_d + q_h) \cdot \text{poly}(\lambda)$ such that*

$$\text{Adv}_{E_{s\text{ElG}}, \mathcal{A}}^{\text{ind-cca2}}(\lambda) \leq \text{Adv}_{\mathcal{B}, \mathbb{G}}^{\text{dl}}(\lambda) + \text{Adv}_{\mathcal{A}', \mathbb{H}}^{2\text{-switching}}(\lambda) + \frac{\mathcal{O}(q_d + q_h)}{2^\lambda},$$

where q_h is a polynomial upper bound on the number of random oracle queries, q_d is a polynomial upper bound on the number of decryption queries and poly is a polynomial independent of q_d and q_h .

Proof. The proof strategy follows (up to some preparations) the outline of [FPS20]. The hybrid \mathbf{G}_0 is identical to $\text{Exp}_{E_{s\text{ElG}}, \mathcal{A}}^{\text{ind-cca2}}(\lambda)$. The initial game transitions until hybrid \mathbf{G}_3 are preparation steps due to the algebraic wrapper. The following hybrids \mathbf{G}_4 , \mathbf{G}_5 and \mathbf{G}_6 correspond exactly to the hybrids \mathbf{G}_1 , \mathbf{G}_2 , \mathbf{G}_3 used in [FPS20], respectively. The preparation steps set up the randomness for the challenge ciphertext as $x^* := z \cdot y$. Further, the randomness for

the signature in the challenge ciphertext is chosen using an x^* -component. Subsequently, re-randomizability and switching are applied such that the public key \hat{Y} uses the representation vector $(0, 1)^\top$ and the randomness for the challenge ciphertext \hat{X}^* uses the representation vector $(0, z)^\top$. The remaining proof proceeds as in [FPS20].

For notational convenience, we introduce the notation $\hat{A} \stackrel{=}{\mathbb{H}} \hat{B}$ for the equality testing algorithm $\text{Eq}_{\mathbb{H}}(pp_{\mathbb{H}}, \hat{A}, \hat{B})$. Note that the hash values for a group element encoding \hat{A} are produced for the bitstring derived via $\text{GetID}_{\mathbb{H}}(pp_{\mathbb{H}}, \hat{A})$ which is guaranteed to be unique for all group element encodings \hat{B} satisfying $\hat{A} \stackrel{=}{\mathbb{H}} \hat{B}$.

We proceed over a series of games starting from the **IND-CCA₂** game in the random oracle model. The hash functions $\mathbb{H}: \mathbb{H} \times \mathbb{H} \rightarrow \mathbb{Z}_p$ and $\tilde{\mathbb{H}}: \mathbb{H} \rightarrow \mathcal{K}$ behave exactly as their counterparts \mathbb{H} and $\tilde{\mathbb{H}}$, respectively, and act solely as helper functions. The adversary \mathcal{A} only has access to the oracles \mathbb{H} and $\tilde{\mathbb{H}}$ (and Dec). Throughout the proof, the behavior of \mathbb{H} and $\tilde{\mathbb{H}}$ will *not* be altered.

GAME \mathbf{G}_0 . \mathbf{G}_0 is identical to the **IND-CCA₂** security game in the random oracle model, see Figure 13.2.

GAME \mathbf{G}_1 . We first change how the signature in the challenge ciphertext is generated. Particularly, the randomness used for the signature is chosen using a y -component. More precisely, \mathbf{G}_1 (defined in Figure 13.3) is identical to \mathbf{G}_0 except for two key differences. First, x^* is not sampled directly from \mathbb{Z}_p but is computed as $z \cdot y$ for some uniformly chosen $z \leftarrow \mathbb{Z}_p$. Second, in \mathbf{G}_0 , the first part of the signature \hat{R}^* is produced via $\text{Rerand}_{\mathbb{H}}(pp_{\mathbb{H}}, \text{Sam}_{\mathbb{H}}(pp_{\mathbb{H}}, x^*))$ and is used to obtain the hash value $c^* := \tilde{\mathbb{H}}(\hat{R}^*, \hat{X}^*)$ and is itself part of the signature. In \mathbf{G}_1 , the group element which is used to obtain the hash value c^* and the group element which is actually part of the signature are decoupled. To obtain c^* , \mathbf{G}_1 uses \hat{R}_1^* which is produced as in \mathbf{G}_0 via $\text{Rerand}_{\mathbb{H}}(pp_{\mathbb{H}}, \text{Sam}_{\mathbb{H}}(pp_{\mathbb{H}}, x^*))$. The group element \hat{R}_2^* which is part of the signature is computed later, depending on c^* , via $\text{Rerand}_{\mathbb{H}}(pp_{\mathbb{H}}, \text{Sam}_{\mathbb{H}}(pp_{\mathbb{H}}, s^* - c^* \cdot x^*))$.

Since x^* is in both games uniformly distributed and $r^* = s^* - c^* \cdot x^* \pmod p$ and thus

$$\text{GetID}_{\mathbb{H}}(pp_{\mathbb{H}}, \hat{R}_1^*) = \text{GetID}_{\mathbb{H}}(pp_{\mathbb{H}}, \hat{R}_2^*),$$

the view of \mathcal{A} in \mathbf{G}_0 and \mathbf{G}_1 is identical and $\Pr[out_0 = 1] = \Pr[out_1 = 1]$.

GAME \mathbf{G}_2 . \mathbf{G}_2 (see Figure 13.4) is identical to \mathbf{G}_1 except that the encodings \hat{Y} , \hat{X}^* and \hat{R}_2^* are produced via private sampling or via the group operation instead of public sampling. Since these encodings are re-randomized subsequently, this game hop is justified by the re-randomizability of the algebraic wrapper \mathbb{H} . More precisely, we successively replace

$$\begin{aligned} & \text{Rerand}_{\mathbb{H}}(pp_{\mathbb{H}}, \text{Sam}_{\mathbb{H}}(pp_{\mathbb{H}}, y)) \text{ by } \text{Rerand}_{\mathbb{H}}(pp_{\mathbb{H}}, \text{PrivSam}_{\mathbb{H}}(\tau_{\mathbb{H}}, y)), \\ & \text{Rerand}_{\mathbb{H}}(pp_{\mathbb{H}}, \text{Sam}_{\mathbb{H}}(pp_{\mathbb{H}}, x^*)) \text{ by } \text{Rerand}_{\mathbb{H}}(pp_{\mathbb{H}}, \text{PrivSam}_{\mathbb{H}}(\tau_{\mathbb{H}}, x^*)), \\ & \text{Rerand}_{\mathbb{H}}(pp_{\mathbb{H}}, \text{Sam}_{\mathbb{H}}(pp_{\mathbb{H}}, s^* - c^* \cdot x^*)) \text{ by} \\ & \quad \text{Rerand}_{\mathbb{H}}(pp_{\mathbb{H}}, \hat{G}^{s^*} \cdot (\hat{X}^*)^{-c^*}). \end{aligned}$$

Due to correctness of sampling, we have

$$\begin{aligned} \text{PrivExt}_{\mathbb{H}}(\tau_{\mathbb{H}}, \text{Sam}_{\mathbb{H}}(pp_{\mathbb{H}}, y)) &= y \cdot \mathbf{e}_1 \\ &= \text{PrivExt}_{\mathbb{H}}(\tau_{\mathbb{H}}, \text{PrivSam}_{\mathbb{H}}(\tau_{\mathbb{H}}, y)), \\ \text{PrivExt}_{\mathbb{H}}(\tau_{\mathbb{H}}, \text{Sam}_{\mathbb{H}}(pp_{\mathbb{H}}, x^*)) &= x^* \cdot \mathbf{e}_1 \\ &= \text{PrivExt}_{\mathbb{H}}(\tau_{\mathbb{H}}, \text{PrivSam}_{\mathbb{H}}(\tau_{\mathbb{H}}, x^*)). \end{aligned}$$

Further, due to correctness of sampling and correctness of extraction,

$$\begin{aligned} &\text{PrivExt}_{\mathbb{H}}(\tau_{\mathbb{H}}, \text{Sam}_{\mathbb{H}}(pp_{\mathbb{H}}, s^* - c^* \cdot x^*)) \\ &= (s^* - c^* \cdot x^*) \cdot \mathbf{e}_1 \\ &= \text{PrivExt}_{\mathbb{H}}(\tau_{\mathbb{H}}, \widehat{G}^{s^*} \cdot (\widehat{X}^*)^{-c^*}). \end{aligned}$$

Hence, for each of these changes we can apply statistical re-randomizability of \mathbb{H} ,

$$|\Pr[out_2 = 1] - \Pr[out_1 = 1]| \leq \frac{3}{2\lambda}.$$

GAME \mathbf{G}_3 . Towards removing the necessity to know y for the simulation, we change the basis to be $[\mathbf{b}]_{\mathbb{G}} := ([1]_{\mathbb{G}}, [y]_{\mathbb{G}})^{\top}$ and adapt the representation vectors used for private sampling of \widehat{Y} and \widehat{X}^* accordingly. More precisely, \mathbf{G}_3 is identical to \mathbf{G}_2 except for the fact that instead of the basis $([1]_{\mathbb{G}}, [\beta_2]_{\mathbb{G}})^{\top}$ for some arbitrary $\beta_2 \in \mathbb{Z}_p$, we use the basis $([1]_{\mathbb{G}}, [y]_{\mathbb{G}})^{\top}$. Additionally, the encodings \widehat{Y} and \widehat{X}^* are privately sampled such that they carry the representation vectors $(0, 1)^{\top}$ and $(0, z)^{\top}$, respectively. Figure 13.5 provides a formal definition of \mathbf{G}_3 .

This game hop is justified by 2-switching of \mathbb{H} . We construct an adversary \mathcal{B}_3 on 2-switching as follows. Initially, on input of $pp_{\mathbb{G}}$, \mathcal{B}_3 outputs two basis vectors

$$[\mathbf{b}^{(\mathbf{G}_2)}]_{\mathbb{G}} := ([1]_{\mathbb{G}}, [\beta_2]_{\mathbb{G}})^{\top} \text{ and } [\mathbf{b}^{(\mathbf{G}_3)}]_{\mathbb{G}} := ([1]_{\mathbb{G}}, [y]_{\mathbb{G}})^{\top}$$

and representation vectors

$$\begin{aligned} \mathbf{v}^{(1),(\mathbf{G}_2)} &:= (y, 0)^{\top}, & \mathbf{v}^{(2),(\mathbf{G}_2)} &:= (z \cdot y, 0)^{\top} \\ \mathbf{v}^{(1),(\mathbf{G}_3)} &:= (0, 1)^{\top}, & \mathbf{v}^{(2),(\mathbf{G}_3)} &:= (0, z)^{\top}. \end{aligned}$$

In return, \mathcal{B}_3 receives public parameters $pp_{\mathbb{H}}$ and two encodings $\widehat{C}^{(1)}$ and $\widehat{C}^{(2)}$. \mathcal{B}_3 computes $\widehat{Y} \leftarrow \text{Rerand}_{\mathbb{H}}(pp_{\mathbb{H}}, \widehat{C}^{(1)})$ and $\widehat{X}^* \leftarrow \text{Rerand}_{\mathbb{H}}(pp_{\mathbb{H}}, \widehat{C}^{(2)})$ and simulates the remaining game as in \mathbf{G}_2 . Note that this is possible since $\tau_{\mathbb{H}}$ is not necessary. \mathcal{B}_3 simulates either \mathbf{G}_2 or \mathbf{G}_3 for \mathcal{A} depending on the challenge provided by $\text{Exp}_{\mathbb{H}, \mathcal{A}, \mathbf{b}}^{\text{2-switching}}(\lambda)$. Hence,

$$|\Pr[out_3 = 1] - \Pr[out_2 = 1]| \leq \text{Adv}_{\mathbb{H}, \mathcal{B}_3}^{\text{2-switching}}(\lambda).$$

Since \mathcal{B}_3 is a legitimate adversary, $\text{Adv}_{\mathbb{H}, \mathcal{B}_3}^{\text{2-switching}}(\lambda)$ is negligible.

GAME \mathbf{G}_4 . From this point on, we are able to closely follow the lines of [FPS20]. \mathbf{G}_4 (see Figure 13.6) is identical to \mathbf{G}_3 except for the following modifications. In \mathbf{G}_4 , the oracle \mathbb{H} stores the private extractions of the encodings used to

call H in a list U . Furthermore, the decryption oracle obtains representation vectors corresponding to the supplied encodings \widehat{R} and \widehat{X} by first looking for a matching entry in U and, if no such entry is present, by applying private extraction. Let (ν, μ) and (ν', μ') be the thus obtained representation vectors of \widehat{R} and \widehat{X} , respectively. Game G_4 additionally introduces an abort condition. If $\mu + \mu' \cdot c = 0$ and $\mu' \neq 0$, G_4 aborts and outputs a random bit. The games G_3 and G_4 only differ if G_4 aborts.

Note that all values in the table T are set in an adversarial call to either H or Dec , except for $c^* = T[(\text{GetID}_{\mathbb{H}}(pp_{\mathbb{H}}, \widehat{R}_2^*), \text{GetID}_{\mathbb{H}}(pp_{\mathbb{H}}, \widehat{X}^*))]$ which is set using \widetilde{H} within the game.⁵² If $\text{Dec}(\widehat{R}, \widehat{X}, s) \neq \perp$, then $(\widehat{R}, \widehat{X}) \neq (\widehat{R}_2^*, \widehat{X}^*)$ since otherwise $s = s^*$. Hence, the value $c = T[(\text{GetID}_{\mathbb{H}}(pp_{\mathbb{H}}, \widehat{R}), \text{GetID}_{\mathbb{H}}(pp_{\mathbb{H}}, \widehat{X}))]$ is independent of (ν, μ, ν', μ') . Therefore, the probability that G_4 aborts is upper bounded by the probability that c is chosen as $c = -\frac{\mu}{\mu'} \bmod p$ which can be upper bounded by $\frac{1}{p} \leq 2^{-\lambda}$. By a union bound, the probability that G_4 aborts is upper bounded by $\frac{q_d}{2^\lambda}$. Since G_3 and G_4 behave identical unless G_4 aborts, we have

$$|\Pr[out_4 = 1] - \Pr[out_3 = 1]| \leq \frac{q_d}{2^\lambda}.$$

GAME G_5 . G_5 , defined in Figure 13.7, is identical to G_4 except for the following changes. Instead of sampling r^* and querying \widetilde{H} for c^* to obtain $s^* = r^* + c^* \cdot x^*$, G_5 samples s^* and c^* independently and computes $\widehat{R}_2^* = \text{Rerand}_{\mathbb{H}}(pp_{\mathbb{H}}, \widehat{G}^{s^*} \cdot (\widehat{X}^*)^{-c^*})$ as in G_4 . This behavior is identical to G_4 except for the event that the tuple $(\widehat{R}_2^*, \widehat{X}^*)$ already has an entry in T . If this event occurs, G_5 aborts. Since \widehat{R}_2^* and \widehat{X}^* are uniformly random and T contains at most $q_d + q_h$ many entries after at most q_d Dec-queries and q_h H-queries, the probability that G_5 aborts, but G_4 does not, can be upper bounded by $\frac{q_d + q_h}{2^{2\lambda}}$. Hence,

$$|\Pr[out_5 = 1] - \Pr[out_4 = 1]| \leq \frac{q_d + q_h}{2^{2\lambda}}.$$

GAME G_6 . G_6 (see Figure 13.8) introduces two further abort conditions (\star) and $(\star\star)$. G_6 aborts as per (\star) if H' is called on \widehat{K} such that $\widehat{K} = \widehat{X}^{*y}$ and aborts as per $(\star\star)$ if $\mu + \mu' \cdot \widetilde{H}(\widehat{R}, \widehat{X}) \neq 0 \bmod p$, where μ and μ' are the $[y]_{\mathbb{G}}$ -components in the decomposition extracted from \widehat{R} and \widehat{X} , respectively.

As in [FPS20], we show that if G_6 differs from G_5 , then we can solve discrete logarithms.

We construct an adversary \mathcal{B}_6 on the discrete logarithm problem. Given $(pp_{\mathbb{G}}, [1]_{\mathbb{G}}, [y]_{\mathbb{G}})$, \mathcal{B}_6 produces $(pp_{\mathbb{H}}, \tau_{\mathbb{H}}) \leftarrow \text{GGen}_{\mathbb{H}}(pp_{\mathbb{G}}, ([1]_{\mathbb{G}}, [y]_{\mathbb{G}})^{\top})$ and simulates G_6 for \mathcal{A} . Note that \widehat{Y} and \widehat{X}^* can be sampled without knowing y (and x^*).

- \mathcal{B}_6 simulates queries to H' as follows. When \mathcal{A} queries H' for \widehat{K} , \mathcal{B}_6 computes $(\nu'', \mu'') \leftarrow \text{PrivExt}_{\mathbb{H}}(\tau_{\mathbb{H}}, \widehat{K})$. Hence,

$$\text{Unwrap}_{\mathbb{H}}(pp_{\mathbb{H}}, \widehat{K}) = (\nu'', \mu'') \cdot ([1]_{\mathbb{G}}, [y]_{\mathbb{G}})^{\top}.$$

To test whether $\widehat{K} = (\widehat{X}^*)^y$ which in turn (implicitly) equals $\widehat{G}^{z \cdot y^2}$, \mathcal{B}_6 solves the equation

$$z \cdot y^2 - \mu'' \cdot y - \nu'' = 0 \bmod p$$

⁵² Recall that T keeps track of previous oracle responses of H .

for y . If one solution is the discrete logarithm of the given DL challenge game \mathbf{G}_6 aborts and \mathcal{B}_6 outputs y . (Note that due to (\star) , if the game does not abort, \mathcal{A} 's view is independent if it receives k_b or k_1 .)

- \mathcal{B}_6 simulates queries to Dec as follows. As argued above, if $\text{Dec}(\widehat{R}, \widehat{X}, s)$ does not return \perp , then $(\widehat{R}, \widehat{X}) \neq (\widehat{R}^*, \widehat{X}^*)$. We have that

$$\begin{aligned}\text{Unwrap}_{\mathbb{H}}(pp_{\mathbb{H}}, \widehat{R}) &= (\nu, \mu) \cdot ([1]_{\mathbb{G}}, [y]_{\mathbb{G}})^{\top} \bmod p, \\ \text{Unwrap}_{\mathbb{H}}(pp_{\mathbb{H}}, \widehat{X}) &= (\nu', \mu') \cdot ([1]_{\mathbb{G}}, [y]_{\mathbb{G}})^{\top} \bmod p.\end{aligned}$$

If Dec does not return \perp , we have $s = r + c \cdot x \bmod p$ and hence

$$y \cdot (\mu + \mu' \cdot c) = s - \nu - \nu' \cdot c \bmod p. \quad (13.1)$$

If $\mu + \mu' \cdot c \neq 0 \bmod p$, \mathbf{G}_6 aborts and \mathcal{B}_6 solves Equation (13.1) for y . If $\mu + \mu' \cdot c = 0 \bmod p$ and $\mu' \neq 0$ then both \mathbf{G}_5 and \mathbf{G}_6 abort. If $\mu + \mu' \cdot c = 0 \bmod p$ and $\mu' = 0$ then $\mu = 0$ and $d\log_{[1]_{\mathbb{G}}}(\text{Unwrap}_{\mathbb{H}}(pp_{\mathbb{H}}, \widehat{X})) = x = \nu'$ allowing the reduction to simulate Dec response as $k := \widetilde{\mathbb{H}}'(\widehat{Y}^x)$.

Therefore,

$$|\Pr[out_6 = 1] - \Pr[out_5 = 1]| \leq \text{Adv}_{\mathbf{G}, \mathcal{B}_6}^{\text{dl}}(\lambda).$$

□


```

G0
ppG ← GGenG(1λ)
(ppH, τH) ← GGenH(ppG, ([1]G, [β2]G)T)
y ← Zp
Ĝ ← RerandH(ppH, SamH(ppH, 1))
Ŷ ← RerandH(ppH, SamH(ppH, y))
pk := (ppH, Ĝ, Ŷ)
T, T' := []
x*, r* ← Zp
X̂* ← RerandH(ppH, SamH(ppH, x*))
R̂* ← RerandH(ppH, SamH(ppH, r*))
c* := H̃(R̂*, X̂*)
s* := r* + c* · x* mod p
k0 := H̃'(Ŷx*), k1 ← K
b' ← AH, H', Dec(pk, kb, (R̂*, X̂*, s*))
return b = b'

```

(a) Description of G_0 .

```

Dec(R̂, X̂, s)
if R̂ =H R̂* ∧ X̂ =H X̂* ∧ s = s* then
  return ⊥
c := H̃(R̂, X̂)
if [s]H ≠H R̂ · X̂c then
  return ⊥
k := H̃'(X̂y)
return k

```

(b) Decryption oracle associated with G_0 .

```

H(R̂, X̂)
if T[(GetIDH(ppH, R̂), GetIDH(ppH, X̂))] = ⊥ then
  T[(GetIDH(ppH, R̂), GetIDH(ppH, X̂))] ← Zp
return T[(GetIDH(ppH, R̂), GetIDH(ppH, X̂))]

```



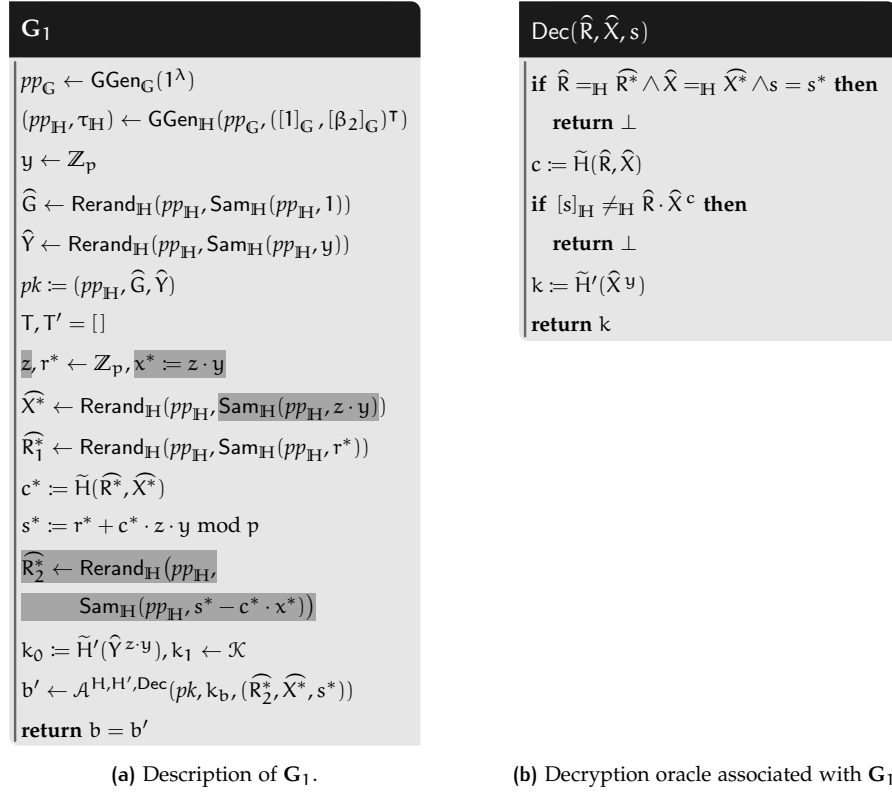
(c) Random oracle H associated with G_0 .

```

H'(R̂)
if T'[GetIDH(ppH, R̂)] = ⊥ then
  T'[GetIDH(ppH, R̂)] ← K
return T'[GetIDH(ppH, R̂)]

```

(d) Random oracle H' associated with G_0 .Figure 13.2: The description of G_0 . G_0 is identical to $\text{Exp}_{\mathcal{E}, \text{ElG}, \mathcal{A}}^{\text{ind-cca}2}(\lambda)$.

Figure 13.3: The description of G₁.

```

G2
ppG ← GGenG(1λ)
(ppH, τH) ← GGenH(ppG, ([1]G, [β2]G)τ)
y ← Zp
Ĝ ← RerandH(ppH, SamH(ppH, 1))
Ŷ ← RerandH(ppH, PrivSamH(τH, y))
pk := (ppH, Ĝ, Ŷ)
T, T' = []
z, r* ← Zp, x* := z · y
X̂* ← RerandH(ppH, PrivSamH(τH, z · y))
R̂*1 ← RerandH(ppH, SamH(ppH, r*))
c* := H̃(R̂*, X̂*)
s* := r* + c* · z · y mod p
R̂*2 ← RerandH(ppH, Gs* · (X̂*)-c*)
k0 := H'(Ŷz·y), k1 ← K
b' ← AH, H', Dec(pk, kb, (R̂*2, X̂*, s*))
return b = b'

```

(a) Description of G_2 .

```

Dec(R̂, X̂, s)
if R̂ =H R̂* ∧ X̂ =H X̂* ∧ s = s* then
  return ⊥
c := H̃(R̂, X̂)
if [s]H ≠H R̂ · X̂c then
  return ⊥
k := H'(X̂y)
return k

```

(b) Decryption oracle associated with G_2 .

```

H(R̂, X̂)
if T[(GetIDH(ppH, R̂), GetIDH(ppH, X̂))] = ⊥ then
  T[(GetIDH(ppH, R̂), GetIDH(ppH, X̂))] ← Zp
return T[(GetIDH(ppH, R̂), GetIDH(ppH, X̂))]

```



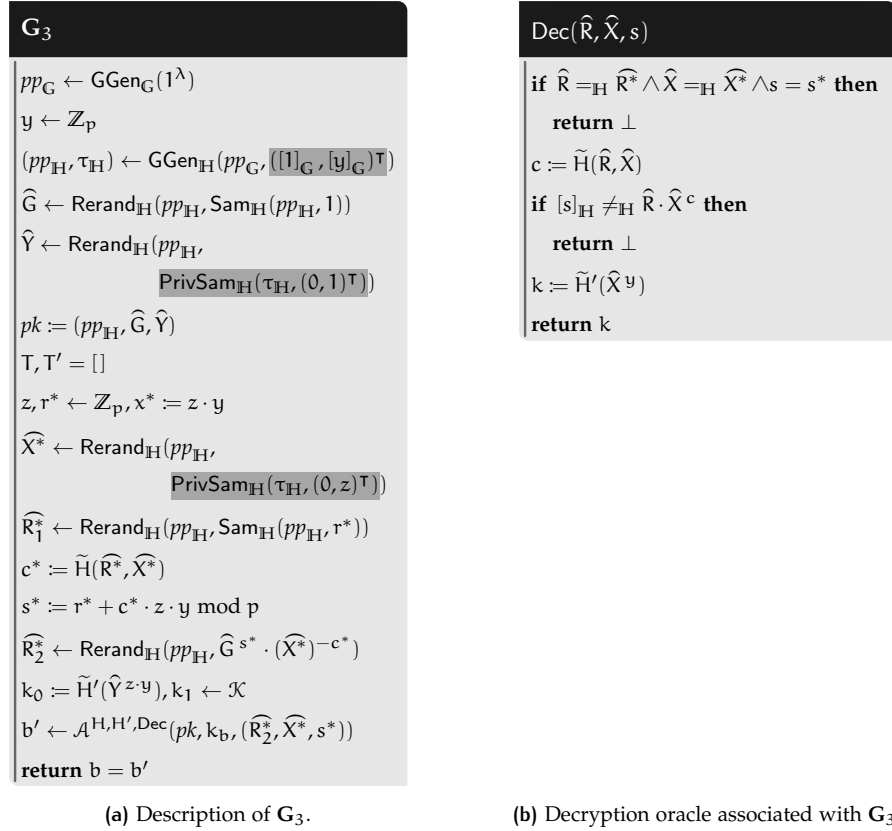
(c) Random oracle H associated with G_2 .

```

H'(R̂)
if T'[GetIDH(ppH, R̂)] = ⊥ then
  T'[GetIDH(ppH, R̂)] ← K
return T'[GetIDH(ppH, R̂)]

```

(d) Random oracle H' associated with G_2 .Figure 13.4: The description of G_2 .

Figure 13.5: The description of G₃.

```

G4
ppG ← GGenG(1λ)
y ← Zp
(ppH, τH) ← GGenH(ppG, ([1]G, [y]G)T)
Ĝ ← RerandH(ppH, SamH(ppH, 1))
Ŷ ← RerandH(ppH,
  PrivSamH(τH, (0, 1)T))
pk := (ppH, Ĝ, Ŷ)
T, T' = [], U := []
z, r* ← Zp, x* := z · y
X̂* ← RerandH(ppH,
  PrivSamH(τH, (0, z)T))
R̂* ← RerandH(ppH, SamH(ppH, r*))
c* := H̃(R̂*, X̂*)
s* := r* + c* · z · y mod p
R̂*2 ← RerandH(ppH, Ĝs* · (X̂*)-c*)
k0 := H̃'(Ŷz·y), k1 ← K
b' ← AH, H', Dec(pk, kb, (R̂*2, X̂*, s*))
return b = b'

```

(a) Description of G_4 .

```

Dec(R̂, X̂, s)
if R̂ =H R̂*2 ∧ X̂ =H X̂* ∧ s = s* then
  return ⊥
c := H̃(R̂, X̂)
if [s]H ≠H R̂ · X̂c then
  return ⊥
(v, μ) ← PrivExtH(τH, R̂)
(v', μ') ← PrivExtH(τH, X̂)
if U[(GetIDH(ppH, R̂),
  GetIDH(ppH, X̂))] ≠ ⊥ then
  (v, μ, v', μ') := U[(GetIDH(ppH, R̂),
  GetIDH(ppH, X̂))]
if (μ + μ' · c = 0) ∧ (μ' ≠ 0) then
  abort and output random bit
k := H̃'(X̂y)
return k

```

(b) Decryption oracle associated with G_4 .

```

H(R̂, X̂)
if T[(GetIDH(ppH, R̂), GetIDH(ppH, X̂))] = ⊥ then
  T[(GetIDH(ppH, R̂), GetIDH(ppH, X̂))] ← Zp
  U[(GetIDH(ppH, R̂), GetIDH(ppH, X̂))] :=
  (PrivExtH(τH, R̂), PrivExtH(τH, X̂))
return T[(GetIDH(ppH, R̂), GetIDH(ppH, X̂))]

```

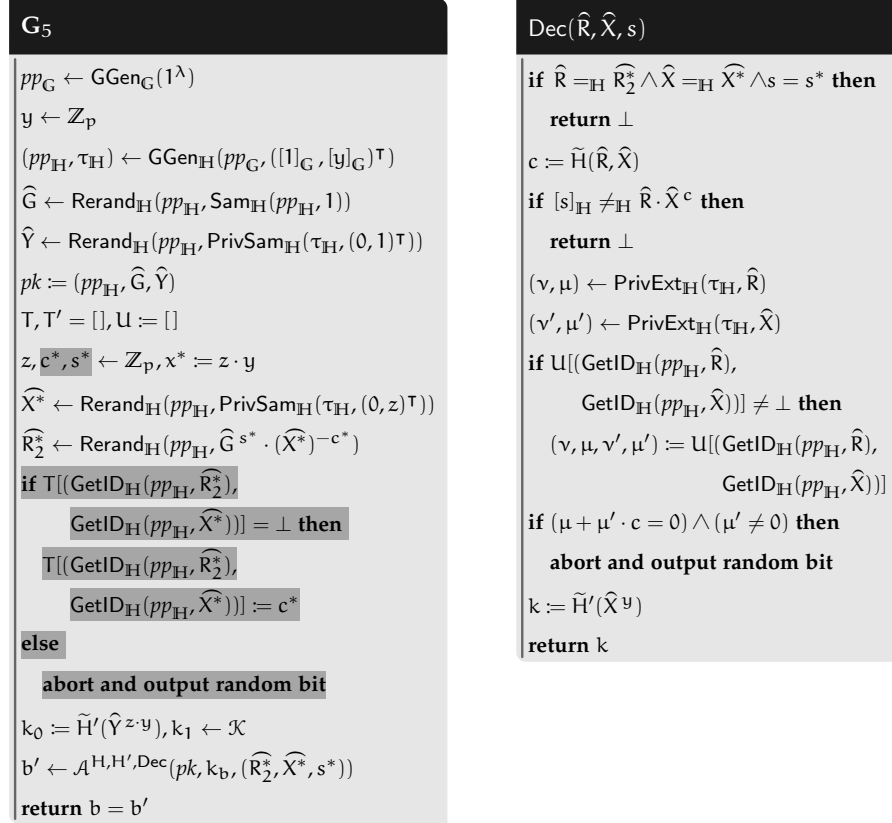
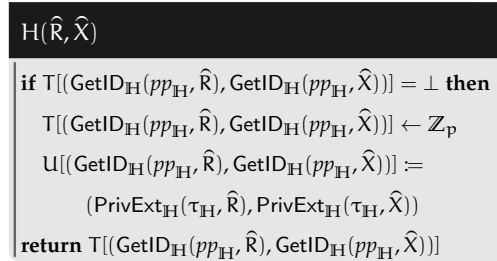

(c) Random oracle H associated with G_4 .

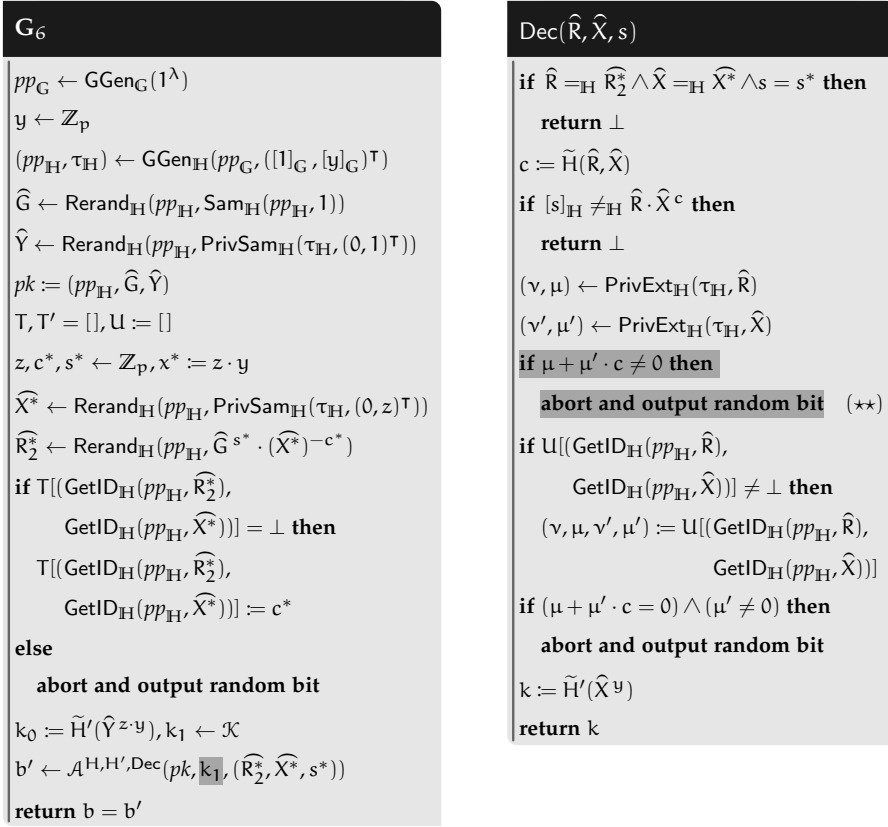
```

H'(R̂)
if T'[(GetIDH(ppH, R̂)] = ⊥ then
  T'[(GetIDH(ppH, R̂)] ← K
return T'[(GetIDH(ppH, R̂)]

```

(d) Random oracle H' associated with G_4 .Figure 13.6: The description of G_4 . This game corresponds to G_1 from [FPS20].

(a) Description of G₅.(b) Decryption oracle associated with G₅.(c) Random oracle H associated with G₅.(d) Random oracle H' associated with G₅.Figure 13.7: The description of G₅. This game corresponds to G₂ from [FPS20].



(a) Description of G₆.

(b) Decryption oracle associated with G₆.

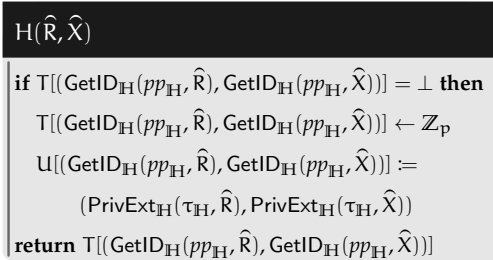
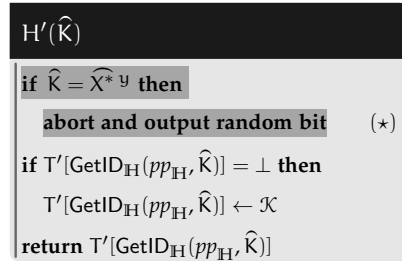
(c) Random oracle H associated with G₆.(d) Random oracle H' associated with G₆.

Figure 13.8: The description of G₆. This game corresponds to G₃ from [FPS20].

Part III

ON PSEUDORANDOM ENCODINGS

III ON PSEUDORANDOM ENCODINGS

14	INTRODUCTION	147
14.1	Flavors of Pseudorandom Encodings	149
14.1.1	Randomized, Computational Pseudorandom Encodings	151
14.2	Implications of our Results	154
14.2.1	New Results for Adaptively Secure Computation . . .	154
14.2.2	Steganography and Covert Multi-Party Computation .	154
14.2.3	Other Results	156
14.3	Negative Results	156
14.3.1	Deterministic, Statistical Pseudorandom Encodings . .	156
14.3.2	Deterministic, Computational Pseudorandom Encodings	157
14.3.3	Randomized, Statistical Pseudorandom Encodings . .	157
14.4	Open questions and subsequent work	158
14.5	Overview of Techniques	159
15	THE PSEUDORANDOM ENCODING HYPOTHESIS	165
15.1	The Pseudorandom Encoding Hypothesis with Setup	169
15.2	Static-to-Adaptive Transformation	171
16	PSEUDORANDOM ENCODINGS AND INVERTIBLE SAMPLING	177
16.1	The Invertible Sampling Hypothesis	177
16.2	The Invertible Sampling Hypothesis with Setup	178
16.3	Equivalence between PREH and ISH	181
16.3.1	Every Inverse Samplable Distribution can be Pseudo- randomly Encoded	181
16.3.2	Every Pseudorandomly Encodable Distribution can be Inverse Sampled	184
17	A TAXONOMY OF PSEUDORANDOM ENCODINGS	189
17.1	Deterministic Encoding Algorithm	189
17.1.1	Information-theoretic Guarantees and Compression .	190
17.1.2	Computational Guarantees and Pseudoentropy	195
17.2	Randomized Encoding Algorithm	200
17.2.1	(Generalized) Extractable One-way Functions	200
17.2.2	Information-theoretic Guarantees and Privately Verifi- able GEOWFs	201
17.2.3	Computational Guarantees and EOWFs with Common Auxiliary Information	204
17.3	Static Pseudorandom Encodings from IO	212

17.4 Bootstrapping Pseudorandom Encodings with URC	216
18 APPLICATIONS OF PSEUDORANDOM ENCODINGS	221
18.1 Adaptively Secure Multi-Party Computation	221
18.1.1 Adaptive MPC	221
18.1.2 Pseudorandom Encodings Imply Adaptive MPC	224
18.1.3 Adaptive MPC Implies Pseudorandom Encodings	226
18.2 Honey Encryption	229
18.3 Keyless Steganography	231
18.4 Deniable Encryption	233
18.5 Non-Committing Encryption	237
18.6 Super-Polynomial Encoding	238

In this part, we present the results of [ACI+20]. This part is taken verbatim from [ACI+20] with minor changes.

The problem of *compression* has been extensively studied in the field of information theory and, more recently, in computational complexity and cryptography [GS85; Wee04; TVZ05; HLR07]. Informally, given a distribution X , compression aims to efficiently encode samples from X as short strings while at the same time being able to efficiently recover these samples. While the typical information-theoretic study of compression considers the case of compressing multiple independent samples from the same source X , its study in computer science, and in particular in this part, considers the “single-shot” case. Compression in this setting is closely related to *randomness condensers* [RR99; TV00; TUZ01; DRV12] and *resource-bounded Kolmogorov complexity* [Sol64; Kol68; Cha69; LV90; LV19] – two well-studied problems in computational complexity. Randomness condensers, which relax randomness extractors, are functions that efficiently map an input distribution into an output distribution with a higher entropy rate. A randomness condenser can be viewed as an efficient compression algorithm, without a corresponding efficient decompression algorithm. The resource-bounded Kolmogorov complexity of a string is the smallest description length of an efficient program that outputs this string. This program description can be viewed as a compressed string, such that decoding is efficiently possible, while finding the compressed string may be inefficient.

An important property of efficient compression algorithms, which combines the efficiency features of randomness condensers and resource-bounded Kolmogorov complexity, is their ability to efficiently produce “random-looking” outputs while allowing the original input to be efficiently recovered. Despite the large body of work on compression and its computational variants, this fundamental property has, to our knowledge, never been the subject of a dedicated study. In this part, we fill this gap by initiating such a study. Before formalizing the problem, we give a simple motivating example.

Consider the goal of encrypting a sample x from a distribution X (say, a random 5-letter English word from the Merriam-Webster Dictionary) using a low-entropy secret key k . Applying a standard symmetric-key encryption scheme with a key derived from k gives rise to the following brute-force attack: Try to decrypt with different keys until obtaining x' in the support of X . In the typical case that wrong keys always lead to x' outside the support of X , this attack successfully recovers x . Variants of this attack arise in different scenarios, including password-authenticated key exchange [BM92], honey encryption [JR14], subliminal communication and steganography [vHL05; CGOS07; HPRV19], and more. A natural solution is to use perfect compression: if x can be compressed to a *uniformly random* string $\hat{x} \in \{0, 1\}^n$ before being encrypted, it cannot be distinguished from another random string $\hat{x}' \in \{0, 1\}^n$ obtained by trying the wrong key. Note, however, that compression may be an overkill for this application. Instead, it suffices to efficiently

encode x into a (possibly longer) *pseudorandom* string from which x can be efficiently decoded. This more general solution motivates the question we consider in this part.

ENCODING INTO THE UNIFORM DISTRIBUTION. We initiate the study of encoding distributions into a random-looking distribution. Informally, we say that a distribution ensemble X_λ admits a *pseudorandom encoding* if there exist efficient encoding and decoding algorithms (E_X, D_X) , where D_X is deterministic, such that

$$\Pr [y \leftarrow X_\lambda : D_X(E_X(y)) = y] \text{ is overwhelming and} \quad (14.1)$$

$$\{y \leftarrow X_\lambda : E_X(y)\} \approx U_{n(\lambda)}. \quad (14.2)$$

Here, “ \approx ” denotes some notion of indistinguishability (we will consider both computational and statistical indistinguishability), and the probability is over the randomness of both E_X and X_λ . The polynomial $n(\lambda)$ denotes the output length of the encoding algorithm E_X . We refer to Equation (14.1) as *correctness* and to Equation (14.2) as *pseudorandomness*. It will also be useful to consider distribution ensembles parameterized by an input m from a language L . We say that such a distribution ensemble $(X_m)_{m \in L}$ admits a pseudorandom encoding if there exist efficient algorithms (E_X, D_X) as above satisfying correctness and pseudorandomness for all $m \in L$, where E_X and D_X both additionally receive m as input. Note that we insist on the decoding algorithm being efficient. This is required for our motivating applications.⁵³ Note also that encoding and decoding above are *keyless*; that is, we want encoded samples to be close to uniform *even though anyone can decode them*. This is a crucial distinction from, for instance, encryption schemes with pseudorandom ciphertexts, which look uniformly distributed to everyone except the owner of the decryption key, and cannot be efficiently decrypted without knowing the decryption key. Here, we seek to simultaneously achieve pseudorandomness and correctness for all parties.

Our motivation for studying pseudorandom encodings stems from the fact that this very natural problem appears in a wide variety of – sometimes seemingly unrelated – problems in cryptography. We already mentioned steganography, honey encryption, and password-authenticated key exchange; we will cover more such connections in this part. Yet, this notion of encoding has to our knowledge never been studied systematically. In this part, we study several natural flavors of this notion, obtain positive and negative results about realizing them, and map their connections with other problems in cryptography.

The main focus of this part is on the hypothesis that *all* efficiently samplable distributions admit a pseudorandom encoding. Henceforth, we refer to this hypothesis as the *pseudorandom encoding hypothesis (PREH)*.

For describing our results, it will be convenient to use the following general notion of efficiently samplable distributions. A distribution family ensemble $(X_m)_{m \in L}$ (for some language $L \subseteq \{0, 1\}^*$) is efficiently samplable if there exists a probabilistic polynomial time (PPT) algorithm S such that $S(m)$ is distributed according to X_m for every $m \in L$. In case the distribution does not depend on additional inputs, L can be considered equal to \mathbb{N} .

OVERVIEW OF CONTRIBUTIONS. Following is a brief summary of our main contributions along with pointers to the relevant technical chapters. We

⁵³ Without this requirement, the problem can be solved using non-interactive commitment schemes with the additional property that commitments are pseudorandom (which exist under standard cryptographic assumptions).

will give an expanded overview of the contributions and the underlying techniques in the rest of this chapter.

- We provide a unified study of different flavors of pseudorandom encodings (PRE) and identify computational, randomized PRE in the CRS model as a useful and achievable notion (Chapter 17 and Theorem 17.13).
- We establish a two-way relation between PRE and the previously studied notion of invertible sampling (Theorem 16.1). This reveals unexpected connections between seemingly unrelated problems in cryptography (e.g., between adaptively secure computation for general functionalities and “honey encryption”).
- We bootstrap “adaptive PRE” from “static PRE” (Theorem 15.1). As a consequence, one can base succinct adaptively secure computation on standard IO as opposed to subexponential IO [CsW19] (Corollary 16.3).
- In [ACI+20], we use PRE to obtain a compiler from standard secure multiparty computation (MPC) protocols to *covert* MPC protocols. This result is not part of this thesis.

14.1 FLAVORS OF PSEUDORANDOM ENCODINGS

The notion of pseudorandom encoding has several natural flavors, depending on whether the encoding algorithm is allowed to use randomness or not, and whether the pseudorandomness property satisfies a computational or information-theoretic notion of indistinguishability. We denote the corresponding hypotheses that every efficiently samplable distribution can be pseudorandomly encoded according to the above variants as $\text{PREH}_{\approx_c}^{\text{rand}}$, $\text{PREH}_{\equiv_s}^{\text{rand}}$, $\text{PREH}_{\approx_c}^{\text{det}}$ and $\text{PREH}_{\equiv_s}^{\text{det}}$.⁵⁴

We note that not all efficiently samplable distributions can be pseudorandomly encoded with a deterministic encoding algorithm. For instance, a distribution which has one very likely event and many less likely ones requires one specific encoding to appear with high probability. Thus, we formally restrict the deterministic variants of the pseudorandom encoding hypothesis to only hold for “compatible” samplers, which still results in interesting connections. In this overview, however, we ignore this restriction.

Further, we explore relaxations which rely on a trusted setup assumption: we consider the pseudorandom encoding hypothesis in the *common reference string model*, in which a common string sampled in a trusted way from some distribution is made available to the parties. This is the most common setup assumption in cryptography and it is standard to consider the feasibility of cryptographic primitives in this model to overcome limitations in the plain model. That is, we ask whether for every efficiently samplable distribution X , there exists an efficiently samplable CRS distribution and efficient encoding and decoding algorithms (E_X, D_X) as above, such that correctness and pseudorandomness hold, where the encoding and decoding algorithm as well as the distinguisher receive the CRS as input, and the distributions in Equations (14.1) and (14.2) are additionally over the choice of the CRS.

Considering distributions which may depend on an input $m \in L$ further entails two different flavors. On the one hand, we consider the notion where

⁵⁴ The subscripts \equiv_s and \approx_c denote statistical and computational guarantees, respectively. The superscripts *det* and *rand* indicate whether the encoding algorithm is deterministic or randomized.

inputs m are chosen adversarially but *statically* (that is, independent of the **CRS**) and, on the other hand, we consider the stronger notion where inputs m are chosen adversarially and *adaptively* depending on the **CRS**. We henceforth denote these variants by the prefix “c” and “ac”, respectively.⁵⁵

⁵⁵ The prefix “c” denotes the static variants with a **CRS** and the prefix “ac” denotes the adaptive variants with a **CRS**.

STATIC-TO-ADAPTIVE TRANSFORMATION. The adaptive notion, where inputs may be chosen depending on the **CRS**, is clearly stronger than the static notion. However, surprisingly, the very nature of pseudorandom encodings allows one to apply an indirection argument similar to the one used in [HJK+16; CPR17; CPV17], which yields a static-to-adaptive transformation.

THEOREM (informal). *If all efficiently samplable distributions can be pseudorandomly encoded in the **CRS** model with a static choice of inputs, then all efficiently samplable distributions can be pseudorandomly encoded in the **CRS** model with an adaptive choice of inputs.*

Static-to-adaptive transformations in cryptography are generally nontrivial, and often come at a big cost in security when they rely on a “complexity leveraging” technique. This connection and its application we will discuss below are a good demonstration of the usefulness of the notion of pseudorandom encodings.

RELAXING COMPRESSION. The notion of statistical deterministic pseudorandom encodings recovers the notion of perfect compression. Hence, this conflicts with the existence of one-way functions.⁵⁶ In our systematic study of pseudorandom encodings, we gradually relax perfect compression in several dimensions, while maintaining one crucial property – the indistinguishability of the encoded distribution from true randomness.

⁵⁶ If perfect compression exists, pseudorandom generators cannot exist (observation attributed to Levin in [GS85]).

EXAMPLE. To illustrate the importance of this property, we elaborate on the example we outline at the beginning of this chapter, focusing more specifically on password-authenticated key exchange (**PAKE**). A **PAKE** protocol allows two parties holding a (low entropy) common password to jointly and confidentially generate a (high entropy) secret key, such that the protocol is resilient against offline dictionary attacks, and no adversary can establish a shared key with a party if he does not know the matching password. A widely used **PAKE** protocol due to Bellare and Merritt [BM92] has a very simple structure: the parties use their low-entropy password to encrypt the flows of a key-exchange protocol using a block cipher. When the block cipher is modeled as a random cipher, it has the property that decrypting a ciphertext (of an arbitrary plaintext) under an incorrect secret key yields a fresh random plaintext. Thus, Bellare and Merritt point out that the security of their **PAKE** protocol requires that “the message to be encrypted by the password must be indistinguishable from a random number.” This is easy to achieve for Diffie-Hellman key exchange over the multiplicative group of integers modulo a prime p . However, for elliptic curve groups this is no longer the case, and one needs to resort to alternative techniques including nontrivial *point compression* algorithms that compress the representation of a random group element into a nearly uniform bitstring [BMN01].

Clearly, our relaxation of compression does not preserve the useful property of obtaining outputs that are *shorter* than the inputs. However, the

remaining pseudorandomness property is good enough for many applications.

In the following, we elaborate on our weakest notion of pseudorandom encodings, that is, pseudorandom encodings allowing the encoding algorithm to be randomized and providing a computational pseudorandomness guarantee. We defer the discussion on the stronger statistical or deterministic variants to Section 14.3, where we derive negative results for most of these stronger notions, which leaves computational randomized pseudorandom encodings as the “best possible” notion that can be realized for general distributions.

14.1.1 Randomized, Computational Pseudorandom Encodings

Computational randomized pseudorandom encodings allow the encoding algorithm to be randomized and require only computational pseudorandomness.

RELATION TO INVERTIBLE SAMPLING. We show a simple but unexpected connection with the notion of *invertible sampling* [CFGN96; DN00; GKM+00]. Informally, invertible sampling refers to the task of finding, given samples from a distribution, random coins that *explain* the sample. Invertible sampling allows to *obliviously* sample from distributions, that is, sampling from distributions *without knowing the corresponding secrets*. This can be useful for, e. g., sampling common reference strings without knowing the random coins or public keys without knowing the corresponding secret keys. A natural relaxation of this notion was systematically studied by Ishai, Kumarasubramanian, Orlandi, and Sahai [IKOS10]. Concretely, a PPT sampler S is *inverse samplable* if there exists an alternative PPT sampler \bar{S} and a PPT inverse sampler \bar{S}^{-1} such that

$$\begin{aligned} \{y \leftarrow S(1^\lambda): y\} &\approx_c \{y \leftarrow \bar{S}(1^\lambda): y\}, \\ \{y \leftarrow \bar{S}(1^\lambda; r): (r, y)\} &\approx_c \{y \leftarrow \bar{S}(1^\lambda): (\bar{S}^{-1}(1^\lambda, y), y)\}. \end{aligned}$$

Note that the inverse sampling algorithm is only required to efficiently inverse-sample from another distribution \bar{S} , but this distribution must be computationally close to the distribution induced by S . The main question studied in [IKOS10] is whether *every* efficient sampler admits such an invertible sampler. They refer to this hypothesis as the *invertible sampling hypothesis (ISH)*, and show that ISH is equivalent to adaptive MPC for general randomized functionalities that may hide their internal randomness. In this part, we show the following two-way relation with pseudorandom encoding.

THEOREM (informal). *A distribution admits a pseudorandom encoding if and only if it admits invertible sampling.*

Intuitively, the efficient encoding algorithm corresponds to the inverse sampling algorithm, and decoding an encoded string corresponds to sampling with the de-randomized alternative sampler \bar{S} . This equivalence immediately extends to all variants of pseudorandom encodings and corresponding variants of invertible sampling we introduce in this part. Invertible sampling is itself connected to other useful cryptographic notions, such as oblivious sampling, trusted common reference string generations, and adaptively secure computation (which we will elaborate upon below).

Building on this connection, the impossibility result of [IKOS10] translates to our setting. On a high level, extractable one-way functions (EOWFs) conflict with invertible sampling because they allow to extract a “secret” (in this case a pre-image) from an image, independently of how it was computed. This conflicts with invertible sampling because invertible sampling is about sampling without knowing the secrets.

THEOREM (informal, [IKOS10]). *Assuming the existence of extractable one-way functions (EOWF) and a non-interactive zero-knowledge proof system, $\text{PREH}_{\approx_c}^{\text{rand}}$ does not hold.*

This suggests that towards a realizable notion of pseudorandom encodings, a further relaxation is due. Thus, we ask whether the above impossibility result extends to the CRS model. In the CRS model, the above intuition why ISH conflicts with EOWFs fails, because the CRS can contain an obfuscated program that samples an image using some secret, but does not output this secret.

Dachman-Soled, Katz, and Rao [DKR15] (building on the universal deniable encryption construction of Sahai and Waters [SW14]) construct a so-called “explainability compiler” that implies $\text{cSH}_{\approx_c}^{\text{rand}}$ based on indistinguishability obfuscation (IO). By our equivalence theorem above, this implies pseudorandom encodings for all efficiently samplable distributions in the CRS model, with static choice of inputs, from IO. Invoking the static-to-adaptive transformation detailed above, this also applies to the adaptive variant.

THEOREM (informal). *Assuming the existence of (polynomially secure) indistinguishability obfuscation and one-way functions, $\text{acPREH}_{\approx_c}^{\text{rand}}$ holds.*

Note that [IKOS10] claim that their impossibility result extends to the CRS model, whereas the above theorem seems to suggest the opposite. We show that technically the result of [IKOS10] does extend to the CRS model at the cost of assuming *unbounded auxiliary-input* extractable one-way functions, a strong flavor of EOWFs that seems very unlikely to exist but cannot be unconditionally ruled out.

THEOREM (informal). *Assuming the existence of extractable one-way functions with unbounded common auxiliary input and a non-interactive zero-knowledge proof system, $\text{cPREH}_{\approx_c}^{\text{rand}}$ does not hold.*

In fact, this apparent contradiction has been the source of some confusion in previous works: the work of [IKOS10] makes an informal claim that their impossibility result for ISH extends to the CRS model. However, due to the connection between ISH and adaptively secure MPC (which we will discuss in more details later on), this claim was challenged in [DKR15]: the authors achieve a construction of adaptively secure MPC for all functionalities assuming IO, which seemingly contradicts the claim of [IKOS10]. The authors of [DKR15] therefore stated that the “impossibility result of Ishai et al. [...] does not hold in the CRS model.” Our extension clarifies that the distinction is in fact more subtle: the result of [IKOS10] *does* extend to the CRS model, but at the cost of assuming EOWF with *unbounded auxiliary inputs*. This does not contradict the constructions based on IO, because IO and EOWF with unbounded auxiliary inputs are known to be contradictory [BCPR16].

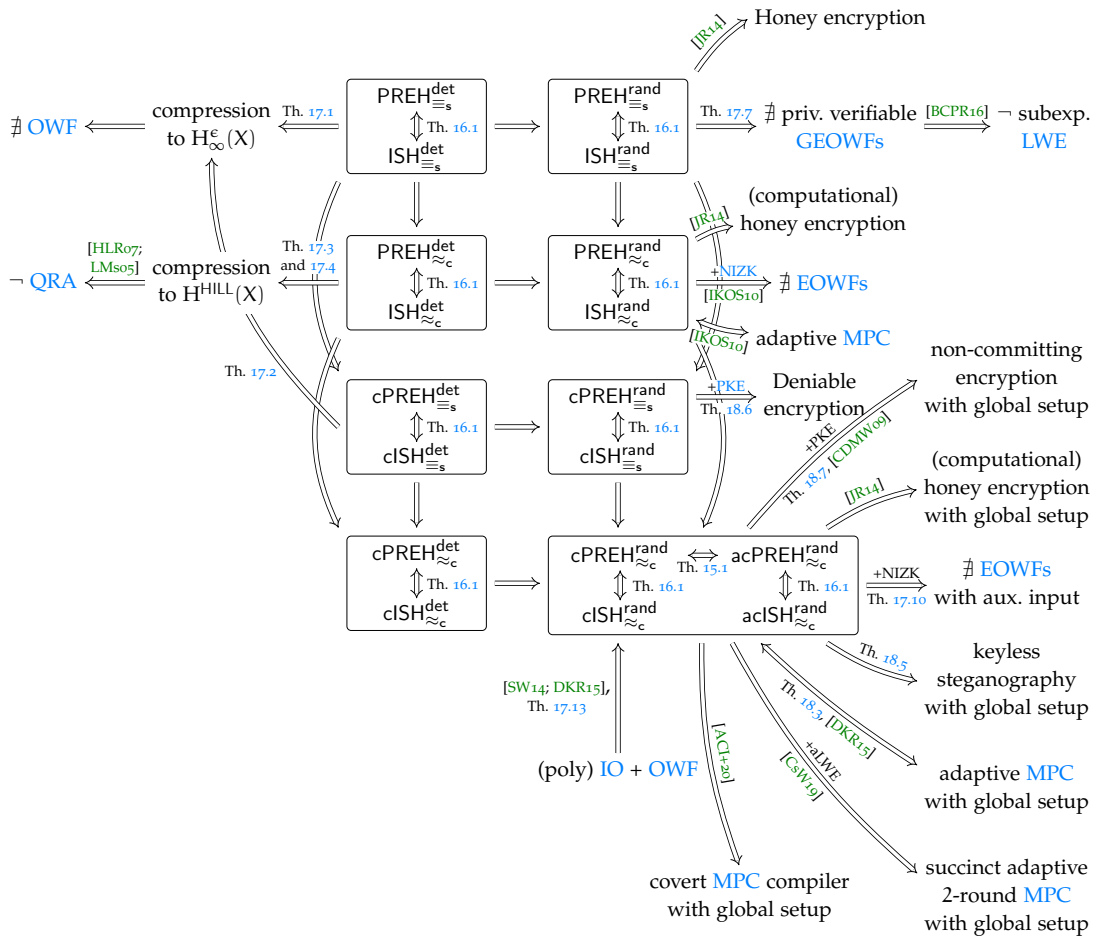


Figure 14.1: An overview of the connections of the pseudorandom encoding hypothesis and other fields of cryptography and computational complexity theory. For simplicity, our static-to-adaptive transformation only appears in the computational, randomized setting in this overview, but also applies to the other settings. (Since the deterministic variants of the pseudorandom encoding hypothesis are impossible for some pathologic samplers, the arrows between deterministic and randomized variants of the pseudorandom encoding hypothesis are to be read as if the deterministic variant is true for some sampler, then the corresponding randomized variant is true for that sampler.)

OVERVIEW. In Figure 14.1, we provide a general summary of the many flavors of the pseudorandom encoding hypothesis, and how they relate to a wide variety of other primitives.

FURTHER RELAXATION. We further study an additional relaxation of pseudorandom encodings, where we allow the encoding algorithm to run in super-polynomial time. We show that this relaxed variant can be achieved from cryptographic primitives similar to *extremely lossy functions* [Zha16], which can be based on the exponential hardness of the decisional Diffie-Hellman problem – a strong assumption, but (still) more standard than indistinguishability obfuscation. However, the applicability of the resulting notion turns out to be rather restricted.

14.2 IMPLICATIONS OF OUR RESULTS

In this section, we elaborate on the implications of the techniques we develop and the results we obtain for a variety of other cryptographic primitives.

14.2.1 New Results for Adaptively Secure Computation

As mentioned above, a sampler admits invertible sampling if and only if it can be pseudorandomly encoded. Ishai et al. [IKOS10] establish a two-way connection between invertible sampling and *adaptively* secure MPC for general randomized functionalities. An MPC protocol allows two or more parties to jointly evaluate a (possibly randomized) functionality \mathcal{F} on their inputs without revealing anything to each other except what follows from their inputs and outputs. This should hold even in the presence of an adversary who can corrupt *any* number of parties in an adaptive (sequential) fashion. When we write “adaptive MPC”, we mean adaptive MPC for *all* randomized functionalities. This should be contrasted with weaker notions of adaptive MPC for strict subsets of corrupted parties [BH93; CFGN96; GS12] or for adaptively well-formed functionalities⁵⁷ [CLOS02] which can both be achieved from mild assumptions. The connection from [IKOS10] shows that adaptive MPC for all randomized functions is possible if and only if every PPT sampler admits invertible sampling, i. e., the invertible sampling hypothesis is true.

⁵⁷ Adaptively well-formed functionalities do not hide internal randomness.

We show that this result generalizes to the global CRS model. More precisely, we prove the adaptive variant of the pseudorandom encoding hypothesis in the CRS model $\text{acPREH}_{\approx_c}^{\text{rand}}$ to be equivalent to adaptive MPC in the global CRS model.⁵⁸

⁵⁸ Together with the conflict between $\text{cPREH}_{\approx_c}^{\text{rand}}$ and EOWFs with unbounded auxiliary input, this corrects a claim made in [DKR15] that the impossibility result of adaptive MPC from [IKOS10] would not extend to the CRS model.

As detailed above, the static variant of the pseudorandom encoding hypothesis $\text{cPREH}_{\approx_c}^{\text{rand}}$ in the CRS model follows from IO (and one-way functions). Applying our static-to-adaptive transformation, the same holds for the adaptive variant. Thus, we obtain the first instantiation of an adaptive explainability compiler [DKR15] without complexity leveraging and, hence, based only on polynomial hardness assumptions. The recent work of Cohen, shelat, and Wichs [CsW19] uses such an adaptive explainability compiler to obtain succinct adaptive MPC, where “succinct” means that the communication complexity is sublinear in the complexity of the evaluated function. Due to our instantiation of $\text{acPREH}_{\approx_c}^{\text{rand}}$ from polynomial IO, we improve the results of [CsW19] by relaxing the requirement for subexponentially secure IO to polynomially secure IO in a black-box way.

COROLLARY (informal). *Assuming the existence of polynomially secure indistinguishability obfuscation and the adaptive hardness of the learning with errors problem, then malicious, two-round, UC-secure adaptive MPC and sublinear communication complexity is possible (in the local CRS model, for all deterministic functionalities).*

14.2.2 Steganography and Covert Multi-Party Computation

We explore the connection of the pseudorandom encoding hypothesis to various flavors of steganography. The goal of steganography, informally, is to embed secret messages in distributions of natural-looking messages, in

order to hide them from external observers. While the standard setting for steganography relies on shared secret keys to encode the messages, we show that pseudorandom encodings naturally give rise to a strong form of *keyless steganography*. Namely, one can rely on pseudorandom encodings to encode any message into an innocent-looking distribution, without truly hiding the message (since anyone can decode the stream), but providing *plausible deniability*, in the sense that, even with the decoded message, it is impossible to tell apart whether this message was indeed encoded by the sender, or whether it is simply the result of decoding the innocent distribution.

COROLLARY (informal). *Assuming pseudorandom encodings, then there exists a keyless steganographic protocol which provides plausible deniability.*

Plausible deniability is an important security notion; in particular, an important cryptographic primitive in this area is the notion of (sender-)deniable encryption [CDNO97], which is known to exist assuming indistinguishability obfuscation [SW14]. Deniable encryption enables to “explain” ciphertexts produced for some message to any arbitrary other message by providing corresponding random coins for a faked encryption process. We view it as an interesting open problem to build deniable encryption under the pseudorandom encoding hypothesis together with more standard cryptographic primitives; we make a first step in this direction and show the following: the *statistical* variant of pseudorandom encodings, together with the existence of public-key encryption, implies deniable encryption. Interestingly, we also show that the computational randomized pseudorandom encoding hypothesis suffices to imply non-committing encryption, a weaker form of deniable encryption allowing to explain only *simulated* ciphertexts to arbitrary messages [CFGN96].

COVERT SECURE COMPUTATION. Covert MPC [vHL05; CGOS07] is an intriguing flavor of MPC that aims at achieving the following strong security guarantee: if the output of the protocol is not “favorable”, the transcript of the interaction should not leak any information to the parties, including whether any given party was *actually taking part in the protocol*. This strong form of MPC aims at providing security guarantees when the very act of starting a computation with other parties should remain hidden. As an example, [vHL05] suppose that a CIA agent who infiltrated a terrorist group wants to make a handshake with another individual to find out whether she is also a CIA agent. Here, we show that pseudorandom encodings give rise to a general compiler transforming a standard MPC protocol into a covert one, in a round-preserving way. The idea is to encode each round of the protocol such that encoded messages look random. Together with the equivalence between adaptively secure MPC and pseudorandom encodings, this gives a connection between two seemingly unrelated notions of secure computation.

COROLLARY (informal). *Assuming adaptively secure MPC for all functionalities, there exists a round-preserving compiler that transforms a large class of “natural” MPC protocols into covert MPC protocols (in the static, semi-honest setting).*

We emphasize that the result on covert secure computation is *not* part of this thesis. We refer the reader to the original publication [ACI+20] for more details on this result.

14.2.3 Other Results

Due to our infeasibility results of $\text{PREH}_{\approx_s}^{\text{rand}}$, distribution transforming encoders (DTEs) for all efficiently samplable distributions are infeasible. Even the computational relaxation of DTEs is infeasible assuming extractable one-way functions. Since all currently known constructions of honey encryption rely on DTEs, we conditionally refute the existence of honey encryption based on the DTE-then-encrypt framework from [JR14]. On the positive side, due to our feasibility result of $\text{acPREH}_{\approx_c}^{\text{rand}}$, computational honey encryption is feasible in the CRS model.

THEOREM (informal). *Assuming $\text{acPREH}_{\approx_c}^{\text{rand}}$ and a suitable symmetric-key encryption scheme (modeled as a random cipher), computational honey encryption for all efficiently samplable distributions exists in the CRS model.*

14.3 NEGATIVE RESULTS FOR STRONGER NOTIONS

Below we describe how we gradually relax perfect compression via different notions of pseudorandom encodings and derive infeasibility results for all variants of pseudorandom encodings which restrict the encoding algorithm to be deterministic or require an information-theoretic pseudorandomness guarantee. This leaves computational randomized pseudorandom encodings as the best possible achievable notion.

14.3.1 Deterministic, Statistical Pseudorandom Encodings

The notion of pseudorandom encodings with a deterministic encoding algorithm and information-theoretic indistinguishability is perhaps the simplest notion one can consider. As we will prove in this paper, this notion recovers the notion of perfect compression: since the encoding algorithm for some source X is deterministic, it can be seen with an entropy argument that the output size of E_X must be at most $H_\infty(X)$, the min-entropy of X ; otherwise, the distribution $\{E_X(X)\}$ can necessarily be distinguished from random with some statistically non-negligible advantage. Therefore, E_X is a perfect and efficient compression algorithm for X , with decompression algorithm D_X ; this is true even for the relaxation in the CRS model. The existence of efficient compression algorithms for various categories of samplers was thoroughly studied [TVZ05]. In particular, the existence of compression algorithms for all efficiently samplable sources implies the inexistence of one-way functions (this is an observation attributed to Levin in [GS85]) since compressing the output of a pseudorandom generator to its entropy would distinguish it from a random string, and the existence of one-way functions implies the existence of pseudorandom generators [HILL99].

THEOREM (informal). *Assuming the existence of one-way functions, then neither $\text{PREH}_{\approx_s}^{\text{det}}$ nor $\text{cPREH}_{\approx_s}^{\text{det}}$ hold.*

This is a strong impossibility result, as one-way functions dwell among the weakest assumptions in cryptography, [Imp95]. One can circumvent this impossibility by studying whether compression can be achieved for more restricted classes of distributions, as was done for instance in [TVZ05].

Our work can be seen as pursuing an orthogonal direction. We seek to determine whether a relaxed notion of compression can be achieved for *all* efficiently samplable distributions. The relaxations we consider comprise the possibility to use randomness in the encoding algorithm, and relaxing the requirement on the encoded distribution to be only computationally indistinguishable from random. Clearly, these relaxations remove one of the most important features of compression algorithms, which is that their outputs are smaller than their inputs (i. e., they compress). Nevertheless, the indistinguishability of the encoded distribution from the uniform distribution is another crucial feature of perfect compression algorithms, which has independent applications.

14.3.2 Deterministic, Computational Pseudorandom Encodings

We now turn towards a relaxation where the encoded distribution is only required to be computationally indistinguishable from random, but the encoding algorithm is still required to be deterministic. This flavor is strongly connected to an important problem in cryptography: the problem of separating HILL entropy [HILL99] from Yao entropy [Yao82]. HILL and Yao entropy are different approaches of formalizing computational entropy, i. e., the amount of entropy a distribution appears to have from the viewpoint of a computationally bounded entity. Informally, a distribution has high HILL entropy if it is computationally close to a distribution with high min-entropy; a distribution has high Yao entropy if it cannot be compressed efficiently. Finding a distribution which, under standard cryptographic assumptions, has high Yao entropy, but low HILL entropy constitutes a long standing open problem in cryptography. Currently, only an oracle separation [Weeo4] and a separation for conditional distributions [HLRo7] are known. To establish the connection between $\text{PREH}_{\approx_c}^{\text{det}}$ and this problem, we proceed as follows: informally, a deterministic pseudorandom encoding must necessarily *compress its input to the HILL entropy of the distribution*. That is, the output size of the encoding cannot be much larger than the HILL entropy of the distribution. This, in turn, implies that a distribution which admits such a pseudorandom encoding cannot have high Yao entropy.

In this work, we formalize the above argument, and show that the *conditional* separation of HILL and Yao entropy from [HLRo7] suffices to refute $\text{PREH}_{\approx_c}^{\text{det}}$. This separation holds under the assumption that non-interactive zero-knowledge proofs with some appropriate structural properties exist (which in turn can be based on standard assumptions such as the quadratic residuosity assumption). Thus, we obtain the following infeasibility result:

THEOREM (informal). *If the quadratic residuosity assumption (QRA) holds, then $\text{PREH}_{\approx_c}^{\text{det}}$ does not hold.*

Hence, we may conclude that towards a feasible variant of pseudorandom encodings for all efficiently samplable distributions, requiring the encoding algorithm to be deterministic poses a strong restriction.

14.3.3 Randomized, Statistical Pseudorandom Encodings

We now consider the relaxation of perfect compression by allowing the encoding algorithm to be randomized while still requiring information-

theoretic indistinguishability from randomness. This flavor of pseudorandom encoding was used in the context of *honey encryption* [JR14]. Honey encryption is a cryptographic primitive which has been introduced to mitigate attacks on encryption schemes resulting from the use of low-entropy passwords. Honey encryption has the property that decrypting a ciphertext with an incorrect key always yields a valid-looking plaintext which seems to come from the expected distribution, thereby mitigating brute-force attacks. This is the same property that was useful in the previous PAKE example.

The study of honey encryption was initiated in [JR14], where it was shown that honey encryption can naturally be constructed by composing a block cipher (modeled as a random cipher) with a *distribution transforming encoder* (DTE), a notion which is equivalent to our notion of pseudorandom encoding with randomized encoding and statistical pseudorandomness. The focus of [JR14] was on obtaining such DTEs for simple and useful distributions. In contrast, we seek to understand the feasibility of this notion for *arbitrary* distributions. Intuitively, it is not straightforward to encode any efficient distribution into the uniform distribution; consider for example the distribution over RSA moduli, i. e., products of two random n -bit primes. Since no efficient algorithm is known to test membership in the support of this distribution, natural approaches seem to break down. In fact, we show in this work that this difficulty is inherent: building on techniques from [BCPR16; IKOS10], we demonstrate the impossibility of (randomized, statistical) pseudorandom encodings for all efficiently samplable distributions, under a relatively standard cryptographic assumption.

THEOREM (informal). *Assuming the subexponential hardness of the learning with errors (LWE) problem, $\text{PREH}_{\equiv_s}^{\text{rand}}$ does not hold.*

This result directly implies that under the same assumption, there exist efficiently samplable distributions (with input) for which no distribution transforming encoder exists. We view it as an interesting open problem whether this result can be extended to rule out the existence of honey encryption for arbitrary distributions under the same assumption.

14.4 OPEN QUESTIONS AND SUBSEQUENT WORK

The most intriguing question left open by our work is whether the weakest variant of the pseudorandom encoding hypothesis $\text{cPREH}_{\approx_c}^{\text{rand}}$, which is implied by IO, also *implies* IO. Very recently, this question was settled in the affirmative by Wee and Wichs [WW20] under the LWE assumption. More concretely, by modifying a heuristic IO construction of Brakerski et al. [BDGM20a], they show that IO is implied by LWE if one is additionally given an *oblivious LWE-sampler* in the CRS model. Such a sampler, given a matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$, generates outputs that are indistinguishable from LWE samples $\mathbf{A} \cdot \mathbf{s} + \mathbf{e}$ without knowing the secrets \mathbf{s} or the noise \mathbf{e} . The existence of an oblivious LWE sampler is nontrivial even under the LWE assumption, because \mathbf{A} can be such that $\mathbf{A} \cdot \mathbf{s} + \mathbf{e}$ is not pseudorandom; however, such a sampler still follows from the invertible sampling hypothesis [IKOS10], which we show to be equivalent to the pseudorandom encoding hypothesis. By proposing an explicit heuristic construction of (a relaxed flavor of) an

oblivious **LWE** sampler, the end result of [WW20] is a construction of **IO** from a new “falsifiable” assumption.

Whether $\text{cPREH}_{\approx_c}^{\text{rand}}$ implies **IO** under weaker or different assumptions than **LWE** remains open. A potentially easier goal is using $\text{cPREH}_{\approx_c}^{\text{rand}}$ to construct public-key encryption from one-way functions. This is related to the possibility of constructing oblivious transfer from any public-key encryption in which public keys and ciphertexts are obviously samplable [EGL85; GKM+00], which is implied by public-key encryption and $\text{cPREH}_{\approx_c}^{\text{rand}}$. Here $\text{cPREH}_{\approx_c}^{\text{rand}}$ is used to bypass the black-box separation between public-key encryption and oblivious transfer [GKM+00].

Finally, there is a lot of room for relaxing the intractability assumptions we use to rule out the statistical ($\text{cPREH}_{\approx_s}^{\text{rand}}$) and deterministic ($\text{cPREH}_{\approx_c}^{\text{det}}$) flavors of pseudorandom encodings.

14.5 OVERVIEW OF TECHNIQUES

In this section, we elaborate on some of our technical results in more detail. In the following, we identify a **PPT** sampler S with the distribution (family) ensemble it induces.

THE RELATION TO INVERTIBLE SAMPLING. In this paragraph, we elaborate on the equivalence between pseudorandom encodings and invertible sampling. Recall that a **PPT** sampler S is *inverse samplable* [CFGN96; DN00; GKM+00], if there exists an alternative sampler \bar{S} inducing a distribution which is computationally indistinguishable to the distribution induced by S such that the computations of \bar{S} can be efficiently inverted. Efficiently inverting the computation of \bar{S} means that there exists an efficient inverse sampler \bar{S}^{-1} which, given an output of \bar{S} , recovers a well-distributed random tape for \bar{S} to compute the given output in the following sense. The inverse sampled random tape is required to be computationally indistinguishable from the actually used random tape. More formally, a **PPT** sampler S is inverse samplable if there exists an efficient alternative sampler \bar{S} and an efficient inverse sampler \bar{S}^{-1} such that

$$\{y \leftarrow S(1^\lambda): y\} \approx_c \{y \leftarrow \bar{S}(1^\lambda): y\}, \quad (14.3)$$

$$\{y \leftarrow \bar{S}(1^\lambda; r): (r, y)\} \approx_c \{y \leftarrow \bar{S}(1^\lambda): (\bar{S}^{-1}(1^\lambda, y), y)\}. \quad (14.4)$$

We refer to Equation (14.3) as *closeness* and to Equation (14.4) as *invertibility*. If the sampler S admits an input m , the above is required to hold for all inputs m in the input space L , where \bar{S} and \bar{S}^{-1} both additionally receive m as input. In accordance with [IKOS10], we refer to the hypothesis that all **PPT** algorithms with input are inverse samplable as the *invertible sampling hypothesis*. Restricting the invertible sampling hypothesis to algorithms which do not admit inputs is denoted the *weak invertible sampling hypothesis*.

The concepts of inverse samplability and pseudorandom encodings are tightly connected. Suppose a **PPT** algorithm S is inverse samplable. Then, there exists an alternative and an inverse sampler (\bar{S}, \bar{S}^{-1}) satisfying closeness and invertibility. Invertibility guarantees that the inverse sampler \bar{S}^{-1} on input of a sample y from $\bar{S}(1^\lambda)$, outputs a computationally well-distributed

random tape r . Hence, with overwhelming probability over the choice of $y \leftarrow \bar{S}(1^\lambda)$ and $r \leftarrow \bar{S}^{-1}(y)$, the alternative sampler on input of r , recovers y . In other words, the inverse sampler \bar{S}^{-1} can be seen as encoding a given sample y , whereas the *de-randomized* alternative sampler \bar{S} given this encoding as random tape, is able to recover y . Looking through the lens of pseudorandom encoding, this almost proves correctness except that y is sampled according to $\bar{S}(1^\lambda)$ instead of $S(1^\lambda)$. This difference can be bridged due to closeness. We now turn towards showing pseudorandomness of the encoded distribution. Due to closeness, the distributions $\{y \leftarrow \bar{S}(1^\lambda): (\bar{S}^{-1}(1^\lambda, y), y)\}$ and $\{y \leftarrow S(1^\lambda): (\bar{S}^{-1}(1^\lambda, y), y)\}$ are computationally indistinguishable. Invertibility guarantees that, given a sample y from $\bar{S}(1^\lambda)$, an encoding of y is indistinguishable to uniformly chosen randomness conditioned on the fact that decoding yields y . Removing y from this distribution almost corresponds to pseudorandomness, except that y is sampled according to $\bar{S}(1^\lambda)$ instead of $S(1^\lambda)$. Again, we are able to bridge this gap due to closeness. Note that we crucially use the fact that the initial randomness used by \bar{S} resides outside of the view of an adversary. Summing up, if a PPT sampler S is inverse samplable, then it can be pseudorandomly encoded.

Interestingly, this connection turns out to be bidirectional. Suppose a PPT algorithm S can be pseudorandomly encoded. Then, there exists an efficient encoding algorithm E_S and an efficient deterministic decoding algorithm D_S satisfying correctness and pseudorandomness. Looking through the lens of invertible sampling, we identify the decoding algorithm to correspond to the alternative sampler (viewing the random tape of the alternative sampler as explicit input to D_S) and the encoding algorithm to correspond to the inverse sampler. Pseudorandomness guarantees that $E_S(S(1^\lambda))$ is indistinguishable from uniform randomness. Hence, applying the decode algorithm D_S on uniform randomness is indistinguishable from applying D_S to outputs of $E_S(S(1^\lambda))$. Correctness guarantees that $D_S(E_S(y))$ for y sampled according to $S(1^\lambda)$ recovers y with overwhelming probability. Thus, the distribution induced by applying D_S on uniform randomness is computationally close to the distribution induced by $S(1^\lambda)$. This shows closeness. For the purpose of arguing about invertibility, consider the distribution $A := \{y \leftarrow D_S(r): (r, y)\}$. Due to pseudorandomness r can be considered an encoded sample from $S(1^\lambda)$. Hence, A is indistinguishable to the distribution, where r is produced by $E_S(y')$ for some *independent* $y' \leftarrow S(1^\lambda)$, i. e.,

$$\{y \leftarrow D_S(r): (r, y)\} \approx_c \{y' \leftarrow S(1^\lambda), r \leftarrow E_S(y'), y \leftarrow D_S(r): (r, y)\}.$$

Note that by correctness, y and y' are identical with overwhelming probability. Therefore, A is indistinguishable to $\{y' \leftarrow S(1^\lambda), r \leftarrow E_S(y'): (r, y')\}$. Since sampling y' via D_S applied on uniform randomness is computationally close to the above distribution due to closeness, invertibility follows. Summing up, a sampler S can be pseudorandomly encoded *if and only if* it is inverse samplable.

Likewise to the variations and relaxations described for pseudorandom encodings, we vary and relax the notion of invertible sampling. The inverse sampler can be required to be deterministic or allowed to be randomized. Further, closeness and invertibility can be required to hold information-theoretically or computationally. We denote these variants as $\text{ISH}_{\equiv_s}^{\text{rand}}$, $\text{ISH}_{\approx_c}^{\text{rand}}$, $\text{ISH}_{\approx_c}^{\text{det}}$ and $\text{ISH}_{\equiv_s}^{\text{det}}$. To circumvent impossibilities in the plain model, we also define

the relaxations in the common reference string model in static and adaptive flavors, denoted by the prefix “c” and “ac”, respectively. The above equivalence extends to all introduced variations of the pseudorandom encoding and invertible sampling hypotheses.

THE STATIC-TO-ADAPTIVE TRANSFORMATION. The static variant of pseudorandom encodings in the CRS model only guarantees correctness and pseudorandomness as long as the input m for the sampler S is chosen independently of the CRS. The adaptive variant, on the other hand, provides correctness and pseudorandomness even for adaptive choices of inputs. Adaptive notions always imply their static analogues. Interestingly, for pseudorandom encodings, the opposite direction is true as well. The core idea is to use an *indirection* argument (similar to [HJK+16; CPR17; CPV17]) to delay CRS generation until during the actual encoding process. Thus, the advantage stemming from adaptively choosing the input is eliminated.

Suppose that the static variant of the pseudorandom encoding hypothesis in the CRS model is true and let S be some PPT sampler. Since S can be pseudorandomly encoded in the CRS model with static choice of inputs, there exist algorithms (Setup', E', D') such that static correctness and pseudorandomness hold. Further, the algorithm Setup' can also be pseudorandomly encoded as above. Let $(\text{Setup}'', E'', D'')$ be the corresponding algorithms such that static correctness and pseudorandomness hold. Note that since the sampler Setup' does not expect an input, static and adaptive guarantees are equivalent.

Then, the sampler S can be pseudorandomly encoded in the CRS model with *adaptive* choice of inputs as follows. Initially, we sample a common reference string crs'' via $\text{Setup}''(1^\lambda)$ and make it available to the parties. Given crs'' and a sample y from $S(m)$, adaptive encoding works in two phases. First, a fresh CRS crs' is sampled via $\text{Setup}'(1^\lambda)$ and pseudorandomly encoded via $r_1 \leftarrow E''(\text{crs}'', \text{crs}')$. Second, the given sample y is pseudorandomly encoded via $r_2 \leftarrow E'(\text{crs}', m, y)$. The encoding of y then consists of (r_1, r_2) . To decode, the CRS crs' is restored via $D''(\text{crs}'', r_1)$. Then, using crs' , the original sample y is recovered via $D'(\text{crs}', m, r_2)$.

Since crs' is chosen freshly during the encoding process, the input m which may depend on crs'' , cannot depend on crs' . Further, the distribution Setup'' does not expect an input. Hence, static guarantees suffice.

To realize that adaptive pseudorandomness holds, consider the encoding of $S(m)$ for some adaptively chosen message m . Since the view of \mathcal{A} when choosing the message m is independent of crs' , static pseudorandomness can be applied to replace the distribution $E'(\text{crs}', m, S(m))$ with uniform randomness. Further, since the sampler Setup' does not expect any input, static pseudorandomness suffices to replace the distribution $E''(\text{crs}'', \text{Setup}'(1^\lambda))$ with uniform randomness. This proves adaptive pseudorandomness.

The adaptive variant of correctness follows similarly from the static variant of correctness. Consider the distribution of decoding an encoded sample of $S(m)$, where m is adaptively chosen. Since the sampler Setup' does not expect an input, static correctness can be applied to replace decoding $D''(\text{crs}'', r_1)$ with the crs' sampled during encoding. Again, since crs' does not lie in the view of the adversary when choosing the message m , static correctness guarantees that decoding succeeds with overwhelming probability. This proves adaptive correctness.

DETERMINISTIC PSEUDORANDOM ENCODING AND COMPRESSION. The notion of pseudorandom encoding is inspired by the notion of compression. A tuple of deterministic functions (E_X, D_X) is said to compress a source X_λ to length $n(\lambda)$ with decoding error $\epsilon(\lambda)$, if (i) $\Pr[D_X(E_X(X_\lambda)) \neq X_\lambda] \leq \epsilon(\lambda)$ and (ii) $\mathbb{E}[|E_X(X_\lambda)|] \leq n(\lambda)$, see [Wee04; TVZ05]. Pseudorandom encoding partially recovers the notion of compression if we require the encoding algorithm to be deterministic. If a source X_λ can be pseudorandomly encoded with a deterministic encoding algorithm having output length $n(\lambda)$, then X_λ is compressible to length $n(\lambda)$. Note, however, that the converse direction is not true. Compression and decompression algorithms for a compressible source do not necessarily encode that source pseudorandomly. The output of a compression algorithm is not required to look pseudorandom and, in some cases, admits a specific structure which makes it easily distinguishable from uniform randomness, e. g., instances using Levin search, [TVZ05].

Clearly, the requirement for correctness poses a lower bound on the encoding length $n(\lambda)$ [Sha48]. Conversely, requiring the encoding algorithm E_X to be deterministic means that the only source of entropy in the distribution $E_X(X_\lambda)$ originates from the source X_λ itself. Hence, for the distributions $E_X(X_\lambda)$ and the uniform distribution over $\{0, 1\}^{n(\lambda)}$ to be indistinguishable, the encoding length $n(\lambda)$ must be “sufficiently small”. We observe that correctness together with the fact that E_X is deterministic implies that the event $E_X(D_X(E_X(X_\lambda))) = E_X(X_\lambda)$ occurs with overwhelming probability. Applying pseudorandomness yields that $E_X(D_X(U_{n(\lambda)})) = U_{n(\lambda)}$ holds with overwhelming probability, wherefore we can conclude that D_X operates almost injectively on the set $\{0, 1\}^{n(\lambda)}$. Hence, the (smooth) min-entropy of $D_X(U_{n(\lambda)})$ is at least $n(\lambda)$.

Considering information-theoretical pseudorandomness, the distributions $D_X(U_{n(\lambda)})$ and X_λ are statistically close. Hence, by the reasoning above, the encoding length $n(\lambda)$ is upper bounded by the (smooth) min-entropy of the source X_λ . In conclusion, if a distribution can be pseudorandomly encoded such that the encoding algorithm is deterministic satisfying statistical pseudorandomness, then this distribution is compressible to its (smooth) min-entropy. Using a technical “Splitting Lemma”, this extends to the relaxed variant of the pseudorandom encoding hypothesis in the CRS model.

Considering computational pseudorandomness, by a similar argument as above, we obtain that X_λ is computationally close to a distribution with min-entropy $n(\lambda)$. This does not yield a relation between the encoding length and the min-entropy of the source. However, we do obtain relations to computational analogues of entropy. Computational entropy is the amount of entropy a distribution appears to have from the perspective of a computationally bounded entity. The notion of HILL entropy [HILL99] is defined via the computational indistinguishability from a truly random distribution. More formally, a distribution X_λ has HILL entropy at least k , if there exists a distribution with min-entropy k which is computationally indistinguishable from X_λ . Hence, the encoding length $n(\lambda)$ is upper bounded by the HILL entropy of the source X_λ . Another important notion of computational entropy is the notion of Yao entropy [Yao82]. Yao entropy is defined via the incompressibility of a distribution. More precisely, a distribution X_λ has Yao entropy at least k if X_λ cannot be efficiently compressed to length less than k (and successfully decompressed). If a distribution can be pseudorandomly encoded with deterministic encoding, then it can be compressed to the en-

coding length $n(\lambda)$. This poses an upper bound on the Yao entropy of the source. In summary, this yields

$$n(\lambda) \leq H^{\text{HILL}}(X_\lambda) \quad \text{and} \quad H^{\text{Yao}}(X_\lambda) \leq n(\lambda). \quad (14.5)$$

However, due to [HLR07; LMs05], if the Quadratic Residuosity Assumption (QRA) is true, then there exist distributions which have low *conditional* HILL entropy while being *conditionally* incompressible, i. e., have high conditional Yao entropy.⁵⁹ The above observations, particularly Equation (14.5), can be extended to conditional HILL and conditional Yao entropy, by considering $\text{PREH}_{\approx_c}^{\text{det}}$ for PPT algorithms with input. Therefore, if the Quadratic Residuosity Assumption is true, $\text{PREH}_{\approx_c}^{\text{det}}$ cannot be true for those distributions.

Unfortunately, we do not know whether this extends to the relaxed variants of the pseudorandom encoding hypothesis admitting access to a CRS. On a high level, the problem is that the HILL entropy, in contrast to the min-entropy, does not remain untouched when additionally conditioning on some common reference string distribution, even though the initial distribution is independent of the CRS. Hence, the splitting technique cannot be applied here.

⁵⁹ Let (X, Z) be a joint distribution. The conditional computational entropy is the entropy X appears to have to a bounded adversary when additionally given Z .

15

THE PSEUDORANDOM ENCODING HYPOTHESIS

A source X is a probability distribution on strings. A family of sources is a probability ensemble $(X_\lambda)_{\lambda \in \mathbb{N}}$, where X_λ is distributed over $\{0, 1\}^{k(\lambda)}$ for some polynomial p . A family of sources $(X_m)_{m \in L}$ can also be indexed by strings from some language $L \subseteq \{0, 1\}^+$, where X_m is distributed over $\{0, 1\}^{k(|m|)}$ for some polynomial p . We say that a source X_λ is efficiently samplable if there is a PPT algorithm S such that $S(1^\lambda)$ is distributed according to X_λ for all $\lambda \in \mathbb{N}$ (using random coins from $\{0, 1\}^{p(\lambda)}$). We further say that a source X_m indexed by strings is efficiently samplable if there exists a PPT algorithm S such that $S(m)$ is distributed according to X_m for all $m \in L$ (using random coins from $\{0, 1\}^{p(|m|)}$).

We initiate the study of the ability to encode efficiently samplable distributions into the uniform distribution. In the following, an efficiently samplable distribution will be defined by the corresponding sampler S with input space L . A distribution defined via S can be pseudorandomly encoded if there exists an efficient potentially randomized encoding algorithm E_S and an efficient deterministic decoding algorithm D_S such that for all $m \in L$, the probability for the event $D_S(E_S(S(m))) = S(m)$ is overwhelming and the distribution $E_S(S(m))$ is indistinguishable from the uniform distribution $U_{n(\lambda)}$. We work with the hypothesis that every efficiently samplable distribution can be pseudorandomly encoded. In this chapter, we formally define the pseudorandom encoding hypothesis and its variants.

DEFINITION 15.1 (Pseudorandom encoding hypothesis, $\text{PREH}_{\approx_c}^{\text{rand}}$). For every PPT algorithm S , there exist efficient algorithms E_S (the encoding algorithm) with output length $n(\lambda)$ and D_S (the decoding algorithm), where D_S is deterministic and E_S is randomized satisfying the following two properties.

CORRECTNESS. For all inputs $m \in L$,

$$\epsilon_{(E_S, D_S), m}^{\text{dec-error}}(\lambda) := \Pr [y \leftarrow S(m) : D_S(m, E_S(m, y)) \neq y]$$

is negligible.


PSEUDORANDOMNESS. For all PPT adversaries \mathcal{A} and all inputs $m \in L$,

$$\text{Adv}_{(E_S, D_S), \mathcal{A}, m}^{\text{pre}}(\lambda) := \left| \Pr \left[\text{Exp}_{(E_S, D_S), \mathcal{A}, m}^{(0)\text{-pre}}(\lambda) = 1 \right] - \Pr \left[\text{Exp}_{(E_S, D_S), \mathcal{A}, m}^{(1)\text{-pre}}(\lambda) = 1 \right] \right|$$

is negligible, where $\text{Exp}_{\mathcal{A}, m, 0}^{\text{pre}}$ and $\text{Exp}_{\mathcal{A}, m, 1}^{\text{pre}}$ are defined in Figure 15.1.

REMARK 15.1. Definition 15.1 formulated for PPT algorithms S which do not admit an input m is called the *weak* $\text{PREH}_{\approx_c}^{\text{rand}}$.

DEFINITION 15.2 ($\text{PREH}_{\approx_c}^{\text{det}}$, $\text{PREH}_{\equiv_s}^{\text{rand}}$, $\text{PREH}_{\equiv_s}^{\text{det}}$). Definition 15.1 can be tuned in two dimensions: the encode algorithm can be required to be deterministic or allowed to be randomized, and the pseudorandomness property



(a) Pseudorandomness game using encoded samples. (b) Pseudorandomness game using true randomness.

Figure 15.1: The pseudorandomness games.

can be required to hold statistically or computationally. We denote these variants as $\text{PREH}_{\alpha}^{\beta}$, where $\alpha \in \{\approx_c, \equiv_s\}$ and $\beta \in \{\text{rand}, \text{det}\}$.

REMARK 15.2. Definition 15.1 demands indistinguishability between encoded samples and the uniform distribution over all bitstrings of some length $n(\lambda)$. This requirement can be relaxed to indistinguishability from the uniform distribution over some efficiently samplable and efficiently recognizable set \mathcal{R} of size N , where elements in \mathcal{R} can be represented with $\mathcal{O}(\log N)$ bits.

DETERMINISTIC ENCODING AND COMPATIBLE SAMPLERS. Requiring the encoding algorithm E_S to be deterministic entails the existence of what we call “incompatible samplers” for which $\text{PREH}_{\equiv_s}^{\text{det}}$ and even $\text{PREH}_{\approx_c}^{\text{det}}$ are unconditionally false. For instance, consider the sampler S^* which on input of 1^λ uniformly chooses an element from the set $\{00, 01, 10\} \subset \{0, 1\}^2$. Assume $\text{PREH}_{\approx_c}^{\text{det}}$ is true for S^* . Then, correctness requires that with overwhelming probability over the sampling process $y \leftarrow S^*(1^\lambda)$, $D_{S^*}(E_{S^*}(y)) = y$. Hence, E_{S^*} must map into the set $\{0, 1\}^k$ for $k \geq 2$ (otherwise there is an correctness error of at least $\frac{1}{3}$). Pseudorandomness, on the other hand, requires that $E_{S^*}(S^*(1^\lambda))$ is computationally indistinguishable from the uniform distribution over $\{0, 1\}^k$. However, since E_{S^*} is a deterministic algorithm, $|\text{supp}(E_{S^*}(S^*(1^\lambda)))| = 3$. Therefore, there exists at least one element in $\{0, 1\}^k \setminus \text{supp}(E_{S^*}(S^*(1^\lambda)))$. An adversary can easily determine this element by evaluating E_{S^*} on each element in the support of S^* (if the support of the sampler was super-polynomial, this would not be possible for a PPT adversary, but very well possible for an unbounded one).

Another example of such an incompatible sampler is a sampler with large support but low min-entropy (i. e., a sampler that has (at least one) very likely output and many much less likely outputs). For instance, consider the sampler S' with probability distribution

$$\Pr[S'(1^\lambda) = 0^\lambda] = \frac{1}{2},$$

$$\Pr[S'(1^\lambda) = 1 \parallel x] = \frac{1}{2^\lambda} \text{ for each } x \in \{0, 1\}^{\lambda-1}.$$

Assume $\text{PREH}_{\approx_c}^{\text{det}}$ is true for S' . Then, correctness requires that an overwhelming fraction of the elements of the form $1 \parallel x$ for $x \in \{0, 1\}^{\lambda-1}$ are correctly decodable. Furthermore, since the element 0^λ has a non-negligible probability to occur, it needs to be correctly decodable. Let the support of $E_{S'}(S'(1^\lambda))$ be (a subset of) $\{0, 1\}^{\lambda-1}$. Pseudorandomness requires that $E_{S'}(S'(1^\lambda))$ is indistinguishable from the uniform distribution over $\{0, 1\}^{\lambda-1}$. However, since

$E_{S'}(0^\lambda)$ is deterministic, the value $E_{S'}(0^\lambda)$ occurs with probability at least $\frac{1}{2}$ and, due to correctness, all other values $E_{S'}(1 \parallel x)$ occur with far lower probability.

In order to obtain a meaningful definition of $\text{PREH}_{\approx c}^{\text{det}}$ and $\text{PREH}_{\equiv s}^{\text{det}}$, we restrict these hypotheses to only hold for specific classes which we refer to as “compatible samplers” $\mathcal{S}^{\approx c}$ and $\mathcal{S}^{\equiv s}$, respectively.

DEFINITION 15.3 (Compatibility with deterministic encodings). A sampler S is *statistically compatible with deterministic encodings* if there exists a set A whose cardinality is a power of 2, such that $\Pr[S(1^\lambda) \in A]$ is overwhelming, and S is ϵ -flat on A , i. e., for all $a \in A$ we have $|\Pr[S(1^\lambda) = a] - \frac{1}{|A|}| \leq \epsilon(\lambda)$, for some negligible function ϵ . The class $\mathcal{S}^{\equiv s}$ contains all samplers which are statistically compatible with deterministic encodings.

A sampler S is *computationally compatible with deterministic encodings* if $S \in \mathcal{S}^{\equiv s}$ or if $|\text{supp}(S(1^\lambda))|$ is super-polynomial and the min-entropy $H_\infty(S(1^\lambda)) \in \omega(\log(\lambda))$ (i. e., the most likely event occurs with negligible probability). The class $\mathcal{S}^{\approx c}$ contains all samplers which are computationally compatible with deterministic encodings.

If S admits an input $z \in L$, S is statistically or computationally compatible with deterministic encodings if the corresponding criterion is met for all $z \in L$.

The above criterion for $\mathcal{S}^{\equiv s}$ may seem unnatural. We note that by relaxing Definition 15.1 as noted in Remark 15.2, requiring high min-entropy suffices for a sampler to be statistically compatible with deterministic encodings.

Note that ϵ -flatness is a weaker criterion than statistical closeness to the uniform distribution over A . ϵ -flatness on A corresponds to closeness to the uniform distribution over A with respect to the infinity norm. Statistical closeness, however, is formalized with respect to the Manhattan norm. The deterministic pseudorandom encodings restricted to compatible samplers still has interesting connections.

Let iPRG be an injective PRG with stretch ℓ . Let S be the sampler which on input of 1^λ draws a uniform seed $s \in \{0, 1\}^\lambda$ and outputs $\text{iPRG}(s) \in \{0, 1\}^{\ell(\lambda)}$. Clearly, S is statistically compatible with deterministic encodings.

LEMMA 15.1. *Let PRG be a PRG with stretch ℓ and let S be the sampler which on input of 1^λ produces the distribution $\text{PRG}(U_\lambda)$. Then, $S \in \mathcal{S}^{\equiv s}$.*

Proof. Let $A' := \text{PRG}(\{0, 1\}^\lambda)$ and let $A'' \subseteq \{0, 1\}^{\ell(\lambda)} \setminus A'$ such that for $A := A' \cup A''$ we have $|A| = 2^\lambda$. We have $\Pr[S(1^\lambda) \in A] = 1$.

For flatness on A , we upper bound the probability of the most likely event in A . Due to pseudorandomness of PRG, all (non-uniform) PPT adversaries distinguishing the distributions $S(1^\lambda)$ and $U_{\ell(\lambda)}$ have a negligible advantage. Assume there exists an $a \in A'$, such that $\Pr[S(1^\lambda) = a] \geq \delta(\lambda)$, for a non-negligible function δ . Then, the most likely value in A' could be used as polynomial advice string allowing a non-uniform adversary to recognize the distribution $S(1^\lambda)$ with non-negligible probability $\delta(\lambda)$. (A uniform adversary can sample a random seed and evaluate the PRG. The thereby obtained output equals the most likely event a with probability $\delta(\lambda)$. Hence, a uniform adversary has advantage at least $\delta(\lambda)^2$ in distinguishing.) Therefore, for all $a \in A'$, $\Pr[S(1^\lambda) = a] \leq \epsilon(\lambda)$ for some negligible function ϵ . Hence, for all $a \in A$, we have $|\Pr[S(1^\lambda) = a] - \frac{1}{|A|}| \leq \epsilon(\lambda) + 2^{-\lambda}$ which is negligible. \square

LEMMA 15.2. *Let PRG be a PRG with stretch ℓ . Let S be the sampler which on input of $x \in L$ produces the distribution $\{y_1 \leftarrow \text{PRG}(U_{|x|}), y_2 \leftarrow D_{x,y_1} : (y_1, y_2)\}$ for any distribution D . Then, $S \in \mathcal{S}^{\approx c}$.*

Proof. Let $\lambda := |x|$. Using the argument of Lemma 15.1, all bitstrings in $\text{PRG}(\{0, 1\}^\lambda)$ have only a negligible probability to occur. Therefore, for all $x \in L$ and all $(y_1, y_2) \in \text{supp}(S(x))$, we have $\Pr[S(1^\lambda, x) = (y_1, y_2)]$ is negligible. Therefore, $S \in \mathcal{S}^{\approx c}$. \square

IMPOSSIBILITY OF UNIVERSAL ENCODING. It is essential that the decoding algorithm D_S depends on the sampler S , since due to pseudorandomness, D_S on input of a random string needs to produce a sample that is in some sense close to the distribution produced by S . This argument does not hold necessarily for the encode algorithm. For instance, in the context of compression, the encoding algorithm does not necessarily depend on the distribution to be compressed. [TVZ05] study the notion of universal compression. This translates to the following definition of PREH with universal encoding.

DEFINITION 15.4 (PREH $_{\alpha}^{\beta}$ with universal encoding). Let \mathcal{S} be a class of sampling algorithms. We say PREH $_{\alpha}^{\beta}$ with universal sampling is true for the class \mathcal{S} , if there exists a universal encoding algorithm E , such that for every PPT algorithm $S \in \mathcal{S}$, there exists an efficient deterministic decoding algorithm D_S , such that correctness and pseudorandomness are satisfied.

In contrast to universal compression, pseudorandom encoding with universal encoding is impossible.

LEMMA 15.3. *PREH with universal encoding is false.*

Proof. Consider the class of samplers $\mathcal{S} = \{S_1, S_2, S_3\}$ over $\{0, 1\}^{k(\lambda)}$ with support $Y_{i,\lambda} := \text{supp}(S_i(1^\lambda))$. We require that for $i \in \{1, 2\}$, $|Y_{i,\lambda}| = 2^{k+1}$, $|Y_{1,\lambda} \cap Y_{2,\lambda}| = 2^k$, $Y_{1,\lambda} \cap Y_{2,\lambda} = Y_{3,\lambda}$ for $k \in \mathcal{O}(\log \lambda)$, and that S_1, S_2 and S_3 produce uniform samples over their support (i.e., correspond to flat distributions). We note that $S_1, S_2, S_3 \in \mathcal{S}^{\equiv s}$. For notational convenience we omit the dependency on λ in the following.

Let $\alpha \in \{\approx c, \equiv s\}$ and $\beta \in \{\text{rand}, \text{det}\}$. Assume toward a contradiction, that PREH $_{\alpha}^{\beta}$ with universal encoding is true for the class \mathcal{S} of sampling algorithms above. Let $\{0, 1\}^{n(\lambda)}$ be the range of E . Due to correctness, for $i \in \{1, 2\}$, there is a negligible function ϵ , such that

$$\begin{aligned} \Pr_{y \leftarrow Y_i} [D_{S_i}(E(y)) \neq y] &= \sum_{y' \in Y_i} \Pr_{y \leftarrow Y_i} [D_{S_i}(E(y)) \neq y \wedge y = y'] \\ &= \frac{1}{|Y_i|} \sum_{y' \in Y_i} \Pr_{y \leftarrow Y_i} [D_{S_i}(E(y)) \neq y \mid y = y'] \\ &\leq \epsilon(\lambda). \end{aligned}$$

Since $|Y_i|$ is polynomial, for each $y \in Y_i$, $\Pr[D_{S_i}(E(y)) = y] \geq 1 - \epsilon'(\lambda)$ for some negligible function ϵ' , where the probability is solely over the randomness of E . Hence, for all $y \in Y_3 = Y_1 \cap Y_2$,

$$\Pr[u \leftarrow E(y) : y = D_{S_1}(u) = D_{S_2}(u) = D_{S_3}(u)] \geq 1 - \epsilon''(\lambda)$$

for some negligible function ϵ'' .

Due to Theorem 16.1 (see Section 16.3), for $i \in \{1, 2, 3\}$, the distribution $\{u \leftarrow U_{n(\lambda)} : D_{S_i}(u)\}$ is computationally (statistically) close to the distribution produced by $S_i(1^\lambda)$. Hence, (up to some negligible fraction ϵ''') the algorithms D_{S_1} and D_{S_2} map at most half of the strings in $\{0, 1\}^{n(\lambda)}$ into the same set $Y_1 \cap Y_2 = Y_3$.

We build an adversary \mathcal{A} on pseudorandomness with respect to sampler S_3 as follows. On input of u , \mathcal{A} outputs 1 if and only if $D_{S_1}(u) = D_{S_2}(u)$.

$$\begin{aligned} \Pr \left[\text{Exp}_{(E_S, D_S), \mathcal{A}}^{(0)\text{-pre}}(\lambda) = 1 \right] &\geq 1 - \epsilon''(\lambda) \\ \Pr \left[\text{Exp}_{(E_S, D_S), \mathcal{A}}^{(1)\text{-pre}}(\lambda) = 1 \right] &\leq \frac{1}{2} - \epsilon'''(\lambda) \end{aligned}$$

Therefore, \mathcal{A} has a non-negligible advantage $\text{Adv}_{(E, D_{S_3}), \mathcal{A}}^{\text{pre}}(\lambda)$. \square

15.1 THE PSEUDORANDOM ENCODING HYPOTHESIS WITH SETUP

We obtain a natural relaxation of the pseudorandom encoding hypothesis by introducing public parameters. We refer to these public parameters as *global* or *non-programmable* common reference string. That is, a distribution defined via S can be pseudorandomly encoded in this relaxed sense, if there exists a probabilistic setup algorithm Setup_S and encode and decode algorithms as before such that for all $m \in L$, the event $D_S(\text{crs}, E_S(\text{crs}, S(m))) = S(m)$ is overwhelming, where the probability is also over the choice of crs , and the distribution $(\text{Setup}_S(1^\lambda), E_S(\text{Setup}_S(1^\lambda), S(m)))$ is indistinguishable from the distribution $(\text{Setup}_S(1^\lambda), U_{n(\lambda)})$.

DEFINITION 15.5 ((Static) pseudorandom encoding hypothesis with setup, $\text{cPREH}_{\approx_c}^{\text{rand}}$). For every PPT algorithm S , there exist a PPT algorithm Setup_S and efficient algorithms (E_S, D_S) , where D_S is deterministic and E_S is randomized (with output length $n(\lambda)$) satisfying the following two requirements.

CORRECTNESS. For all PPT adversaries \mathcal{A} ,

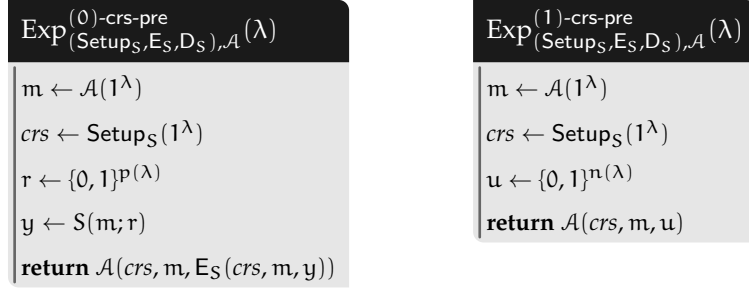
$$\begin{aligned} e_{(\text{Setup}_S, E_S, D_S), \mathcal{A}}^{\text{c-dec-error}}(\lambda) := \\ \Pr \left[\begin{array}{l} m \leftarrow \mathcal{A}(1^\lambda) \\ \text{crs} \leftarrow \text{Setup}_S(1^\lambda) : D_S(\text{crs}, m, E_S(\text{crs}, m, y)) \neq y \\ y \leftarrow S(m) \end{array} \right] \end{aligned}$$

is negligible.

PSEUDORANDOMNESS. For all PPT adversaries \mathcal{A} ,

$$\text{Adv}_{(\text{Setup}_S, E_S, D_S), \mathcal{A}}^{\text{crs-pre}}(\lambda) := \left| \Pr \left[\text{Exp}_{(\text{Setup}_S, E_S, D_S), \mathcal{A}}^{(0)\text{-crs-pre}}(\lambda) = 1 \right] - \Pr \left[\text{Exp}_{(\text{Setup}_S, E_S, D_S), \mathcal{A}}^{(1)\text{-crs-pre}}(\lambda) = 1 \right] \right|$$

is negligible, where $\text{Exp}_{(\text{Setup}_S, E_S, D_S), \mathcal{A}}^{(0)\text{-crs-pre}}(\lambda)$ and $\text{Exp}_{(\text{Setup}_S, E_S, D_S), \mathcal{A}}^{(1)\text{-crs-pre}}(\lambda)$ are defined in Figure 15.2.



(a) Pseudorandomness game using encoded samples. (b) Pseudorandomness game using true randomness.

Figure 15.2: The static games experiments with setup.

We note that assuming non-uniform adversaries, correctness and pseudorandomness defined in Definition 15.5 can be equivalently defined by quantifying over all messages $m \in L$. Definition 15.5 is static in the sense that the inputs $m \in L$ are required to be chosen statically, i. e., independently of crs . In the following, we define the corresponding adaptive variant.

DEFINITION 15.6 (Adaptive pseudorandom encoding hypothesis with setup, $\text{acPREH}_{\approx_c}^{\text{rand}}$). For every PPT algorithm S , there exist a PPT algorithm Setup_S and efficient algorithms (E_S, D_S) , where D_S is deterministic and E_S is randomized (with output length $n(\lambda)$) satisfying the following two requirements.

CORRECTNESS. For all PPT adversaries \mathcal{A} ,

$$\epsilon_{(\text{Setup}_S, E_S, D_S), \mathcal{A}}^{\text{ac-dec-error}}(\lambda) := \Pr \left[\begin{array}{l} \text{crs} \leftarrow \text{Setup}_S(1^\lambda) \\ m \leftarrow \mathcal{A}(\text{crs}) \\ y \leftarrow S(m) \end{array} : D_S(\text{crs}, m, E_S(\text{crs}, m, y)) \neq y \right]$$

is negligible.

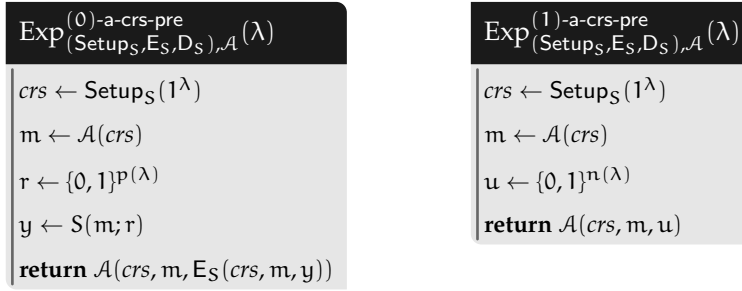
PSEUDORANDOMNESS. For all PPT adversaries \mathcal{A} ,

$$\text{Adv}_{(\text{Setup}_S, E_S, D_S), \mathcal{A}}^{\text{a-crs-pre}}(\lambda) := \left| \Pr \left[\text{Exp}_{(\text{Setup}_S, E_S, D_S), \mathcal{A}}^{(0)\text{-a-crs-pre}}(\lambda) = 1 \right] - \Pr \left[\text{Exp}_{(\text{Setup}_S, E_S, D_S), \mathcal{A}}^{(1)\text{-a-crs-pre}}(\lambda) = 1 \right] \right|$$

is negligible, where $\text{Exp}_{(\text{Setup}_S, E_S, D_S), \mathcal{A}}^{(0)\text{-a-crs-pre}}(\lambda)$ and $\text{Exp}_{(\text{Setup}_S, E_S, D_S), \mathcal{A}}^{(1)\text{-a-crs-pre}}(\lambda)$ are defined in Figure 15.3.

DEFINITION 15.7 ($\text{cPREH}_{\approx_c}^{\text{det}}$, $\text{cPREH}_{\equiv_s}^{\text{rand}}$, $\text{cPREH}_{\equiv_s}^{\text{det}}$, $\text{acPREH}_{\approx_c}^{\text{det}}$, $\text{acPREH}_{\equiv_s}^{\text{rand}}$, $\text{acPREH}_{\equiv_s}^{\text{det}}$). Definitions 15.5 and 15.6 can be tuned in two dimensions: the encoding algorithm can be required to be deterministic or allowed to be randomized, and the correctness and pseudorandomness properties can be required to hold statistically or computationally. We denote these variants as $\text{cPREH}_{\alpha}^{\beta}$ and $\text{acPREH}_{\alpha}^{\beta}$, respectively, where $\alpha \in \{\approx_c, \equiv_s\}$ and $\beta \in \{\text{rand}, \text{det}\}$.

REMARK 15.3 (Universal setup). In Definitions 15.5 and 15.6, we allow the algorithm Setup_S to depend on the sampler S . A natural (but somewhat incomparable) variant of this definition is to require the existence of a (bounded-



(a) Pseudorandomness game using encoded samples. (b) Pseudorandomness game using true randomness.

Figure 15.3: The adaptive pseudorandomness games with setup.

circuit-size) *universal setup algorithm* $\text{Setup}(1^\lambda, B)$ for $B \in \mathbb{N}$ which provides the above guarantees for all samplers which can be represented with B bits. We refer to this variant as $\text{universal cPREH}_\alpha^\beta$ and $\text{universal acPREH}_\alpha^\beta$, respectively.

REMARK 15.4 (Common random string). We will denote the strengthening of the pseudorandom encoding hypothesis as in Definitions 15.5 and 15.6, where the setup algorithm Setup_S is required to sample uniform random strings, as the pseudorandom encoding hypothesis with a *common random string* or *uniformly random CRS (URC)*.

Note that in Definitions 15.5 and 15.6, we implicitly only consider *legitimate* adversaries which guarantee that $m \in L$. For the sake of avoiding notational overhead, we do not make this explicit.

15.2 STATIC-TO-ADAPTIVE TRANSFORMATION

Clearly, Definition 15.6 implies Definition 15.5. Interestingly, the opposite direction is true as well by an “indirection” argument similar to the one from the work on universal samplers [HJK+16]. A similar technique was also used in the context of non-committing encryption [CPR17] and adaptively secure MPC [CPV17].

THEOREM 15.1. *Let $\alpha \in \{\approx_c, \equiv_s\}$ and $\beta \in \{\text{rand}, \text{det}\}$. If $\text{cPREH}_\alpha^\beta$ is true, then $\text{acPREH}_\alpha^\beta$ is true.*

Proof. We prove the statement for the computational randomized case. The proof directly extends to the remaining cases.

Let S be a PPT sampler with input space L . Since $\text{cPREH}_{\approx_c}^{\text{rand}}$ is true, for the PPT sampler S , there exist $(\text{Setup}'_S, E'_S, D'_S)$ with output length $n'(\lambda)$ such that correctness and pseudorandomness hold (statically) as in Definition 15.5. Again, since $\text{cPREH}_{\approx_c}^{\text{rand}}$ is true, for the PPT sampler Setup'_S , there exist $(\text{Setup}''_S, E''_S, D''_S)$ with output length $n''(\lambda)$ such that correctness and pseudorandomness hold (statically).⁶⁰ Note that Setup'_S does not expect an input.

In Figure 15.4, we define algorithms $(\text{Setup}_S, E_S, D_S)$ satisfying *adaptive* correctness and pseudorandomness, as in Definition 15.6.

⁶⁰ For notational convenience, we do not write the sampler Setup'_S as index.

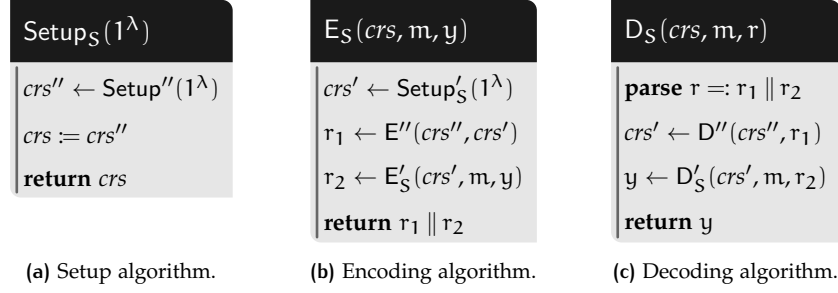


Figure 15.4: Adaptive pseudorandom encodings.

Intuitively, since crs' is chosen freshly and independently after the adversary fixes the message m , selective security suffices. Furthermore, since the distribution of crs' has no input, selective security is sufficient.

ADAPTIVE CORRECTNESS. We define a series of hybrid games to prove correctness, see Figure 15.5.

GAME G₀. This game corresponds to encoding and decoding a sample from $S(m)$ for some adversarially and adaptively chosen message m . G_0 outputs 1 if decoding succeeds.

GAME G₁. G_1 is identical to G_0 except for some reordering. That is, the game hop from G_0 to G_1 is only conceptual and $\Pr[out_0 = 1] = \Pr[out_1 = 1]$.

GAME G₂. G_2 is identical to G_1 except that for decoding y_D , G_2 uses the originally sampled **CRS** crs' instead of the decoded **CRS** crs'_D .

CLAIM 15.1. For all **PPT** adversaries \mathcal{A} , there exists a **PPT** adversary \mathcal{B}_2 , such that $|\Pr[out_2 = 1] - \Pr[out_1 = 1]| \leq \epsilon_{(\text{Setup}'', E'', D''), \mathcal{B}_2}^{\text{c-dec-error}}(\lambda)$.

Proof. The games G_1 and G_2 proceed exactly identically if $crs'_D = crs'$. Let E be the event that $crs' \neq crs'_D$. We have that $out_1 = 1 \wedge \neg E \Leftrightarrow out_2 = 1 \wedge \neg E$. Due to correctness of $(\text{Setup}'', E'', D'')$,

$$\Pr[E] = \Pr \left[\begin{array}{l} crs'' \leftarrow \text{Setup}(1^\lambda) \\ crs' \leftarrow \text{Setup}'_S(1^\lambda) \\ r_1 \leftarrow E''(crs'', crs') \\ crs'_D \leftarrow D''(crs'', r_1) \end{array} : crs'_D \neq crs' \right]$$

is negligible. Hence, the Difference Lemma (due to Shoup [Shoo04]) upper bounds

$$|\Pr[out_2 = 1] - \Pr[out_1 = 1]| \leq \Pr[E].$$

□

GAME G₃. G_3 is identical to G_2 except that r_1 and crs'_D are not produced and that crs' is sampled *after* the adversary chooses the message m (depending on crs''). Hence, the game hop from G_2 to G_3 is only conceptual and $\Pr[out_2 = 1] = \Pr[out_3 = 1]$.

CLAIM 15.2. For all **PPT** adversaries \mathcal{A} , there exists a **PPT** adversary \mathcal{B}_3 , such that $\Pr[out_3 = 1] \geq 1 - \epsilon_{(\text{Setup}'_S, E'_S, D'_S), \mathcal{B}_3}^{\text{c-dec-error}}(\lambda)$.

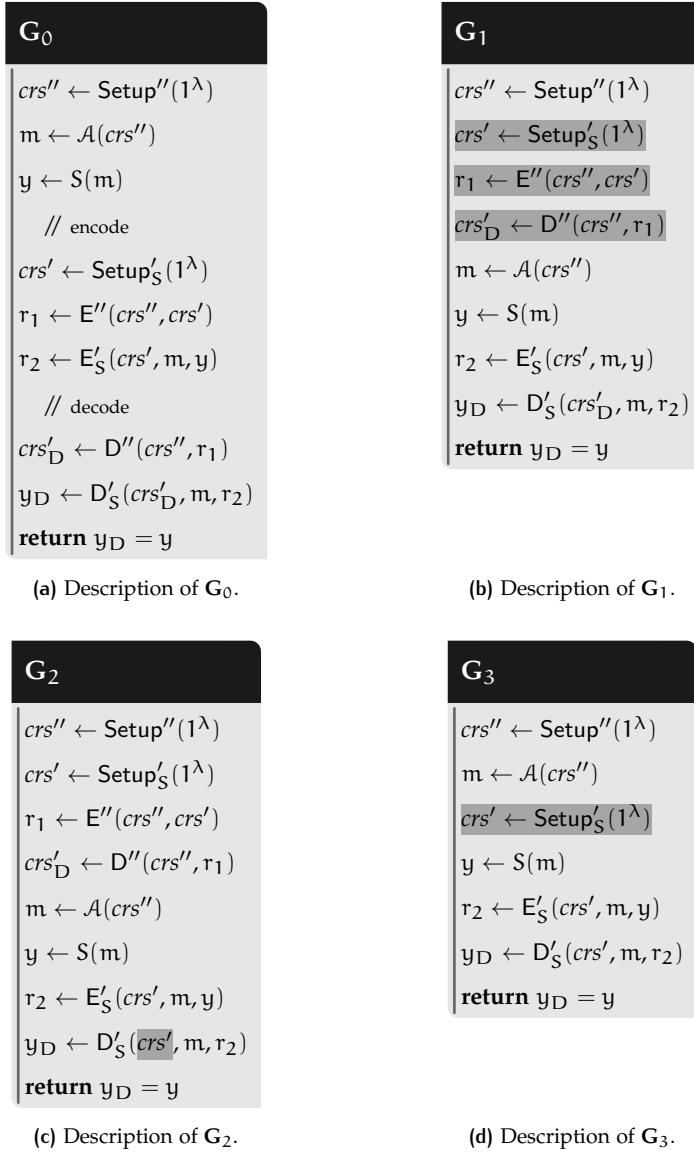


Figure 15.5: Hybrid games for the proof of adaptive correctness.

Proof. Due to correctness of $(\text{Setup}'_S, E'_S, D'_S)$, we have that for all PPT adversaries $\overline{\mathcal{B}}$,

$$\Pr \left[\begin{array}{l} m \leftarrow \overline{\mathcal{B}}(1^\lambda) \\ crs' \leftarrow \text{Setup}'_S(1^\lambda) \\ y \leftarrow S(m) \\ r \leftarrow E'_S(crs', m, y) \\ y_D \leftarrow D'_S(crs', m, r) \end{array} : y = y_D \right]$$

is overwhelming. Consider the PPT adversary \mathcal{B}_3 who proceeds as in Figure 15.6.

Therefore, for all PPT adversaries \mathcal{A} , $\Pr[out_3 = 1]$ is overwhelming. \square

Thus, we have

$$\epsilon_{(\text{Setup}_S, E_S, D_S), \mathcal{A}}^{\text{ac-dec-error}}(\lambda) \leq \epsilon_{(\text{Setup}'', E'', D''), \overline{\mathcal{B}}}^{\text{c-dec-error}}(\lambda) + \epsilon_{(\text{Setup}'_S, E'_S, D'_S), \overline{\mathcal{B}}'}^{\text{c-dec-error}}(\lambda)$$

```

B3(1λ)
crs'' ← Setup''(1λ)
m ← A(crs'')
return m

```

Figure 15.6: Adversary \mathcal{B}_3 used for the game hop from G_2 to G_3 .

for some PPT adversaries $\overline{\mathcal{B}}$ and $\overline{\mathcal{B}'}$.

ADAPTIVE PSEUDORANDOMNESS. We define a series of hybrid games to prove pseudorandomness, see Figure 15.7.

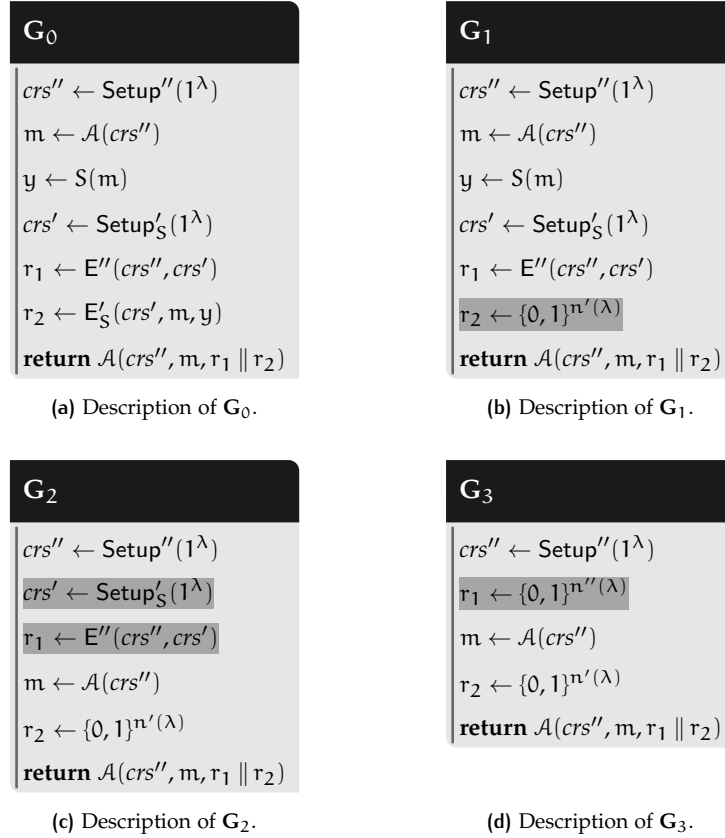


Figure 15.7: Hybrid games for the proof of adaptive pseudorandomness.

GAME G_0 . G_0 is identical to $\text{Exp}_{(\text{Setup}_S, E_S, D_S), \mathcal{A}}^{(0)\text{-a-crs-pre}}(\lambda)$.

GAME G_1 . G_1 is identical to G_0 except for the fact that in G_0 , r_2 is an encoding of y , whereas in G_1 , r_2 is a uniformly random bitstring.

CLAIM 15.3. For all PPT adversaries \mathcal{A} , there exists a PPT adversary \mathcal{B}_1 , such that $|\Pr[\text{out}_1 = 1] - \Pr[\text{out}_0 = 1]| \leq \text{Adv}_{(\text{Setup}'_S, E'_S, D'_S), \mathcal{B}_1}^{\text{crs-pre}}(\lambda)$.

Proof. Construct an adversary \mathcal{B}_1 on static pseudorandomness relative to $(\text{Setup}'_S, E'_S, D'_S)$ as follows. On input of 1^λ , \mathcal{B}_1 samples $\text{crs}'' \leftarrow \text{Setup}''(1^\lambda)$, calls \mathcal{A} on input of crs'' , and outputs the message m produced by \mathcal{A} . In return, \mathcal{B}_1 receives $\text{crs}' \leftarrow \text{Setup}'_S(1^\lambda)$ and either $u := E'_S(\text{crs}', m, S(m))$ or

a uniform random string $u \leftarrow \{0, 1\}^{n'(\lambda)}$ from $\text{Exp}_{(\text{Setup}'_S, E'_S, D'_S), \mathcal{A}}^{(b)\text{-crs-pre}}(\lambda)$. \mathcal{A}_1 computes $r_1 \leftarrow E''(\text{crs}'', \text{crs}')$, calls \mathcal{A} on input of $(\text{crs}'', m, r_1 \parallel u)$ and returns \mathcal{A} 's output. See Figure 15.8 for a description of \mathcal{B}_1 .

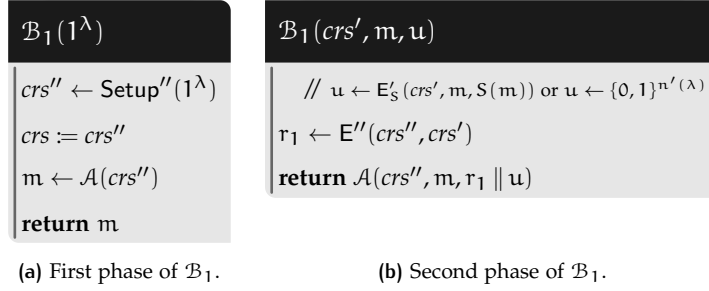


Figure 15.8: Adversary \mathcal{B}_1 used for the game hop from \mathbf{G}_0 to \mathbf{G}_1 .

If \mathcal{B}_1 plays $\text{Exp}_{(\text{Setup}'_S, E'_S, D'_S), \mathcal{B}_1}^{(0)\text{-crs-pre}}(\lambda)$, then it perfectly simulates \mathbf{G}_0 . On the other hand, if \mathcal{B}_1 plays $\text{Exp}_{(\text{Setup}'_S, E'_S, D'_S), \mathcal{B}_1}^{(1)\text{-crs-pre}}(\lambda)$, then it perfectly simulates \mathbf{G}_1 . Hence, we have

$$|\Pr[out_1 = 1] - \Pr[out_0 = 1]| \leq \text{Adv}_{(\text{Setup}'_S, E'_S, D'_S), \mathcal{B}_1}^{\text{crs-pre}}(\lambda).$$

□

GAME \mathbf{G}_2 . \mathbf{G}_2 is identical to \mathbf{G}_1 except for some reordering. More precisely, crs' and r_1 are sampled before \mathcal{A} chooses the message m . The difference between \mathbf{G}_1 and \mathbf{G}_2 is only conceptual and $\Pr[out_2 = 1] = \Pr[out_1 = 1]$.

GAME \mathbf{G}_3 . \mathbf{G}_3 is identical to \mathbf{G}_2 except that in \mathbf{G}_2 , r_1 is produced as an encoding of crs' whereas in \mathbf{G}_3 , r_1 is sampled uniformly at random. Hence, \mathbf{G}_3 is distributed identically to $\text{Exp}_{(\text{Setup}_S, E_S, D_S), \mathcal{A}}^{(1)\text{-a-crs-pre}}(\lambda)$.


CLAIM 15.4. For all PPT adversaries \mathcal{A} , there exists a PPT adversary \mathcal{B}_3 , such that $|\Pr[out_3 = 1] - \Pr[out_2 = 1]| \leq \text{Adv}_{(\text{Setup}'', E'', D''), \mathcal{B}_3}^{\text{crs-pre}}(\lambda)$.

Proof. Construct an adversary \mathcal{B}_3 on static pseudorandomness relative to $(\text{Setup}'', E'', D'')$ as follows. On input of 1^λ , \mathcal{B}_3 returns \perp since the input space L of the sampler $\text{Setup}'_S(1^\lambda)$ is empty. In return, \mathcal{B}_3 receives crs'' sampled via $\text{Setup}''(1^\lambda)$ and u which is either produced via $E''(\text{crs}'', \text{Setup}'(1^\lambda))$ or via uniform sampling from $\{0, 1\}^{n''(\lambda)}$. \mathcal{B}_3 calls \mathcal{A} on input of crs'' and receives a message m from \mathcal{A} . Finally, \mathcal{B}_3 samples $r_2 \leftarrow \{0, 1\}^{n'(\lambda)}$, calls \mathcal{A} on input of $(\text{crs}'', m, u \parallel r_2)$ and outputs his output. See Figure 15.9 for a description of \mathcal{B}_3 .

If \mathcal{B}_3 plays $\text{Exp}_{(\text{Setup}'', E'', D''), \mathcal{B}_3}^{(0)\text{-crs-pre}}(\lambda)$, then it perfectly simulates \mathbf{G}_2 . On the other hand, if \mathcal{B}_3 plays $\text{Exp}_{(\text{Setup}'', E'', D''), \mathcal{B}_3}^{(1)\text{-crs-pre}}(\lambda)$, then it perfectly simulates \mathbf{G}_3 . Hence, we have

$$|\Pr[out_3 = 1] - \Pr[out_2 = 1]| \leq \text{Adv}_{(\text{Setup}'', E'', D''), \mathcal{B}_3}^{\text{crs-pre}}(\lambda).$$

□

(a) First phase of \mathcal{B}_3 .(b) Second phase of \mathcal{B}_3 .Figure 15.9: Adversary \mathcal{B}_3 used for the game hop from \mathbf{G}_2 to \mathbf{G}_3 .

Therefore, we have

$$\begin{aligned}
 \text{Adv}_{(\text{Setup}_S, E_S, D_S), \mathcal{A}}^{\text{a-crs-pre}}(\lambda) &= |\Pr[out_3 = 1] - \Pr[out_0 = 1]| \\
 &\leq \text{Adv}_{(\text{Setup}'', E'', D''), \overline{\mathcal{B}}}^{\text{crs-pre}}(\lambda) \\
 &\quad + \text{Adv}_{(\text{Setup}'_S, E'_S, D'_S), \overline{\mathcal{B}}'}^{\text{crs-pre}}(\lambda)
 \end{aligned}$$

for some PPT adversaries $\overline{\mathcal{B}}$ and $\overline{\mathcal{B}}'$. □

16

PSEUDORANDOM ENCODINGS AND INVERTIBLE SAMPLING

In this chapter, we explain the relation between pseudorandom encodings and invertible sampling, [DN00]. In Sections 16.1 and 16.2, we restate the invertible sampling hypothesis of [IKOS10] and define several variants thereof. In Section 16.3, we prove that a distribution can be pseudorandomly encoded if and only if it is inverse samplable. This extends to all of the introduced variations of the pseudorandom encoding hypothesis and the invertible sampling hypothesis.

16.1 THE INVERTIBLE SAMPLING HYPOTHESIS

A PPT algorithm S is inverse samplable according to [DN00; IKOS10] if there exists an alternative PPT algorithm \bar{S} and a corresponding inverse sampler \bar{S}^{-1} such that \bar{S} (on every input) induces a distribution which is computationally indistinguishable from the distribution induced by S (on identical inputs) and \bar{S}^{-1} inverts the computation of \bar{S} . That is, \bar{S}^{-1} on input of an output produced by \bar{S} produces computationally well-distributed random coins for \bar{S} to produce the given output.

DEFINITION 16.1 (Invertible sampling hypothesis, $\text{ISH}_{\approx_c}^{\text{rand}}$, [IKOS10]). For every PPT algorithm S , there exists a PPT algorithm \bar{S} (the alternate sampler) with randomness space $\{0, 1\}^{n(\lambda)}$ and an efficient randomized algorithm \bar{S}^{-1} (the inverse sampler), satisfying the following two properties.

CLOSENESS. For all PPT adversaries \mathcal{A} and all inputs $m \in L$,

$$\text{Adv}_{(\bar{S}, \bar{S}^{-1}), \mathcal{A}, m}^{\text{close}}(\lambda) := \left| \Pr \left[\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{A}, m}^{(0)\text{-close}}(\lambda) = 1 \right] - \Pr \left[\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{A}, m}^{(1)\text{-close}}(\lambda) = 1 \right] \right|$$

is negligible, where $\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{A}, m}^{(0)\text{-close}}(\lambda)$ and $\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{A}, m}^{(1)\text{-close}}(\lambda)$ are defined in Figure 16.1.

INVERTIBILITY. For all PPT adversaries \mathcal{A} and all inputs $m \in L$,

$$\text{Adv}_{(\bar{S}, \bar{S}^{-1}), \mathcal{A}, m}^{\text{inv}}(\lambda) := \left| \Pr \left[\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{A}, m}^{(0)\text{-inv}}(\lambda) = 1 \right] - \Pr \left[\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{A}, m}^{(1)\text{-inv}}(\lambda) = 1 \right] \right|$$

is negligible, where $\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{A}, m}^{(0)\text{-inv}}(\lambda)$ and $\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{A}, m}^{(1)\text{-inv}}(\lambda)$ are defined in Figure 16.1.

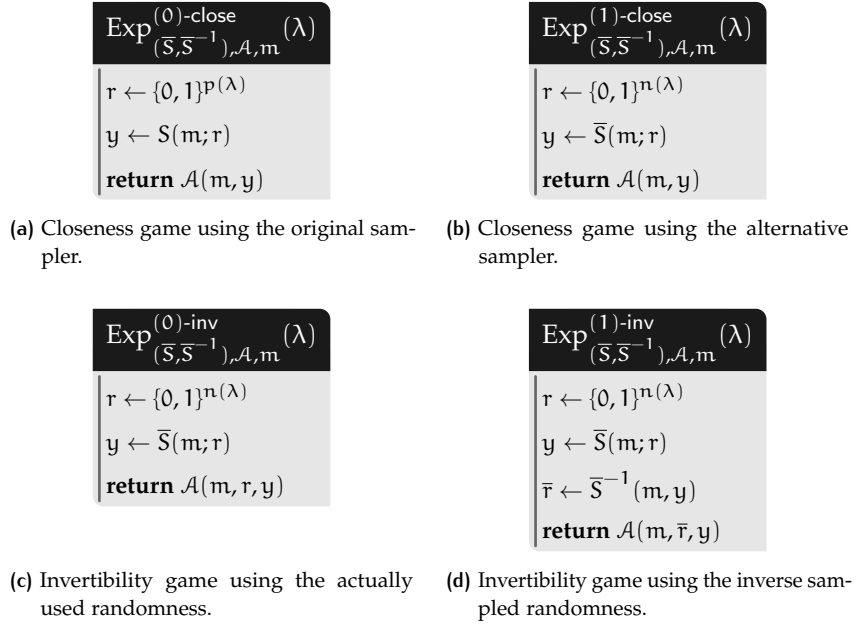


Figure 16.1: The closeness and invertibility experiments.

DEFINITION 16.2 ($\text{ISH}_{\approx_c}^{\text{det}}$, $\text{ISH}_{\equiv_s}^{\text{rand}}$, $\text{ISH}_{\equiv_s}^{\text{det}}$). Definition 16.1 can be tuned in two dimensions: the inverse sampler can be required to be deterministic or allowed to be randomized, and the closeness and invertibility properties can be required to hold statistically or computationally. We denote these variants as $\text{ISH}_{\alpha}^{\beta}$, where $\alpha \in \{\approx_c, \equiv_s\}$ and $\beta \in \{\text{rand}, \text{det}\}$.

16.2 THE INVERTIBLE SAMPLING HYPOTHESIS WITH SETUP

The invertible sampling hypothesis can be naturally relaxed by introducing public parameters (or global **CRS**), henceforth denoted *crs*. This allows the alternative sampler and the inverse sampler to use *crs*. Closeness and invertibility are defined against adversaries knowing the *crs* but choosing the inputs statically.

DEFINITION 16.3 (Invertible sampling hypothesis with setup, $\text{cISH}_{\approx_c}^{\text{rand}}$). For every **PPT** algorithm S there exists a **PPT** algorithm Setup_S , a **PPT** algorithm \bar{S} (with randomness space $\{0, 1\}^{n(\lambda)}$) and an efficient randomized \bar{S}^{-1} satisfying the following two properties.

CLOSENESS. For all **PPT** adversaries \mathcal{A} ,

$$\text{Adv}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{A}}^{\text{crs-close}}(\lambda) := \left| \Pr \left[\text{Exp}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{A}}^{(0)\text{-crs-close}}(\lambda) = 1 \right] - \Pr \left[\text{Exp}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{A}}^{(1)\text{-crs-close}}(\lambda) = 1 \right] \right|$$

is negligible, where $\text{Exp}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{A}}^{(0)\text{-crs-close}}(\lambda)$ and $\text{Exp}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{A}}^{(1)\text{-crs-close}}(\lambda)$ are defined in Figure 16.2.

INVERTIBILITY. For all PPT adversaries \mathcal{A} ,

$$\text{Adv}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{A}}^{\text{crs-inv}}(\lambda) := \left| \Pr \left[\text{Exp}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{A}}^{(0)\text{-crs-inv}}(\lambda) = 1 \right] - \Pr \left[\text{Exp}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{A}}^{(1)\text{-crs-inv}}(\lambda) = 1 \right] \right|$$

is negligible, where $\text{Exp}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{A}}^{(0)\text{-crs-inv}}(\lambda)$ and $\text{Exp}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{A}}^{(1)\text{-crs-inv}}(\lambda)$ are defined in Figure 16.2.

```

Exp(0)-crs-close(SetupS,  $\bar{S}$ ,  $\bar{S}^{-1}$ ),  $\mathcal{A}$ ( $\lambda$ )
m  $\leftarrow$   $\mathcal{A}(1^\lambda)$ 
crs  $\leftarrow$  SetupS( $1^\lambda$ )
r  $\leftarrow$  {0, 1}p( $\lambda$ )
y  $\leftarrow$  S(m; r)
return  $\mathcal{A}(\text{crs}, y)$ 
    
```

(a) Closeness game using the original sampler.

```

Exp(1)-crs-close(SetupS,  $\bar{S}$ ,  $\bar{S}^{-1}$ ),  $\mathcal{A}$ ( $\lambda$ )
m  $\leftarrow$   $\mathcal{A}(1^\lambda)$ 
crs  $\leftarrow$  SetupS( $1^\lambda$ )
r  $\leftarrow$  {0, 1}n( $\lambda$ )
y  $\leftarrow$   $\bar{S}(\text{crs}, m; r)$ 
return  $\mathcal{A}(\text{crs}, y)$ 
    
```

(b) Closeness game using the alternative sampler.

```

Exp(0)-crs-inv(SetupS,  $\bar{S}$ ,  $\bar{S}^{-1}$ ),  $\mathcal{A}$ ( $\lambda$ )
m  $\leftarrow$   $\mathcal{A}(1^\lambda)$ 
crs  $\leftarrow$  SetupS( $1^\lambda$ )
r  $\leftarrow$  {0, 1}n( $\lambda$ )
y  $\leftarrow$   $\bar{S}(\text{crs}, m; r)$ 
return  $\mathcal{A}(\text{crs}, r, y)$ 
    
```

(c) Invertibility game using the actually used randomness.

```

Exp(1)-crs-inv(SetupS,  $\bar{S}$ ,  $\bar{S}^{-1}$ ),  $\mathcal{A}$ ( $\lambda$ )
m  $\leftarrow$   $\mathcal{A}(1^\lambda)$ 
crs  $\leftarrow$  SetupS( $1^\lambda$ )
r  $\leftarrow$  {0, 1}n( $\lambda$ )
y  $\leftarrow$   $\bar{S}(\text{crs}, m; r)$ 
 $\bar{r} \leftarrow \bar{S}^{-1}(\text{crs}, m, y)$ 
return  $\mathcal{A}(\text{crs}, \bar{r}, y)$ 
    
```

(d) Invertibility game using the inverse sampled randomness.

Figure 16.2: The static closeness and invertibility experiments with setup.

Definition 16.3 is static in the sense that closeness and invertibility adversaries are required to statically choose the challenge input $m \in L$. In the following, we define the corresponding adaptive version, where adversaries are allowed to choose the challenge input $m \in L$ depending on crs .

DEFINITION 16.4 (Adaptively secure invertible sampling hypothesis with setup, $\text{aclSH}_{\approx_c}^{\text{rand}}$). For every PPT algorithm S , there exists a PPT algorithm Setup_S , a PPT algorithm \bar{S} (with randomness space $\{0, 1\}^{n(\lambda)}$) and an efficient randomized \bar{S}^{-1} satisfying the following two properties.

(ADAPTIVE) CLOSNESS. For all PPT adversaries \mathcal{A} ,

$$\text{Adv}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{A}}^{\text{a-crs-close}}(\lambda) := \left| \Pr \left[\text{Exp}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{A}}^{(0)\text{-a-crs-close}}(\lambda) = 1 \right] - \Pr \left[\text{Exp}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{A}}^{(1)\text{-a-crs-close}}(\lambda) = 1 \right] \right|$$

is negligible, where $\text{Exp}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{A}}^{(0)\text{-a-crs-close}}(\lambda)$ and $\text{Exp}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{A}}^{(1)\text{-a-crs-close}}(\lambda)$ are defined in Figure 16.3.

(ADAPTIVE) INVERTIBILITY. For all PPT adversaries \mathcal{A} ,

$$\text{Adv}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{A}}^{\text{a-crs-inv}}(\lambda) := \left| \Pr \left[\text{Exp}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{A}}^{(0)\text{-a-crs-inv}}(\lambda) = 1 \right] - \Pr \left[\text{Exp}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{A}}^{(1)\text{-a-crs-inv}}(\lambda) = 1 \right] \right|$$

is negligible, where $\text{Exp}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{A}}^{(0)\text{-a-crs-inv}}(\lambda)$ and $\text{Exp}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{A}}^{(1)\text{-a-crs-inv}}(\lambda)$ are defined in Figure 16.3.

```

Exp(SetupS,  $\bar{S}$ ,  $\bar{S}^{-1}$ ),  $\mathcal{A}$ }(0)-a-crs-close( $\lambda$ )
crs  $\leftarrow$  SetupS(1 $^\lambda$ )
m  $\leftarrow$   $\mathcal{A}$ (crs)
r  $\leftarrow$  {0, 1}p( $\lambda$ )
y  $\leftarrow$  S(m; r)
return  $\mathcal{A}$ (y)
    
```

(a) Closeness game using the original sampler.

```

Exp(SetupS,  $\bar{S}$ ,  $\bar{S}^{-1}$ ),  $\mathcal{A}$ }(1)-a-crs-close( $\lambda$ )
crs  $\leftarrow$  SetupS(1 $^\lambda$ )
m  $\leftarrow$   $\mathcal{A}$ (crs)
r  $\leftarrow$  {0, 1}n( $\lambda$ )
y  $\leftarrow$   $\bar{S}$ (crs, m; r)
return  $\mathcal{A}$ (y)
    
```

(b) Closeness game using the alternative sampler.

```

Exp(SetupS,  $\bar{S}$ ,  $\bar{S}^{-1}$ ),  $\mathcal{A}$ }(0)-a-crs-inv( $\lambda$ )
crs  $\leftarrow$  SetupS(1 $^\lambda$ )
m  $\leftarrow$   $\mathcal{A}$ (crs)
r  $\leftarrow$  {0, 1}n( $\lambda$ )
y  $\leftarrow$   $\bar{S}$ (crs, m; r)
return  $\mathcal{A}$ (r, y)
    
```

(c) Invertibility game using the actually used randomness.

```

Exp(SetupS,  $\bar{S}$ ,  $\bar{S}^{-1}$ ),  $\mathcal{A}$ }(1)-a-crs-inv( $\lambda$ )
crs  $\leftarrow$  SetupS(1 $^\lambda$ )
m  $\leftarrow$   $\mathcal{A}$ (crs)
r  $\leftarrow$  {0, 1}n( $\lambda$ )
y  $\leftarrow$   $\bar{S}$ (crs, m; r)
 $\bar{r} \leftarrow \bar{S}^{-1}$ (crs, m, y)
return  $\mathcal{A}$ ( $\bar{r}$ , y)
    
```

(d) Invertibility game using the inverse sampled randomness.

Figure 16.3: The adaptive closeness and invertibility experiments with setup.

DEFINITION 16.5 ($\text{clSH}_{\approx_c}^{\text{det}}$, $\text{clSH}_{\equiv_s}^{\text{rand}}$, $\text{clSH}_{\equiv_s}^{\text{det}}$, $\text{aclSH}_{\approx_c}^{\text{det}}$, $\text{aclSH}_{\equiv_s}^{\text{rand}}$, $\text{aclSH}_{\equiv_s}^{\text{det}}$). Definitions 16.3 and 16.4 can be tuned in two dimensions: the inverse sampler can be required to be deterministic or allowed to be randomized, and the closeness and invertibility properties can be required to hold statistically or computationally. We denote these variants as $\text{clSH}_{\alpha}^{\beta}$ and $\text{aclSH}_{\alpha}^{\beta}$, respectively, where $\alpha \in \{\approx_c, \equiv_s\}$ and $\beta \in \{\text{rand}, \text{det}\}$.

REMARK 16.1 (Universal setup). In Definitions 16.3 and 16.4, we allow the algorithm Setup_S to depend on the sampler S . A natural (but somewhat incomparable) variant of this definition is to require the existence of a (bounded-circuit-size) *universal setup algorithm* $\text{Setup}(1^\lambda, B)$ for $B \in \mathbb{N}$ which provides the above guarantees for all samplers which can be represented with B bits. We refer to this variant as $\text{universal clSH}_{\alpha}^{\beta}$ and $\text{universal aclSH}_{\alpha}^{\beta}$, respectively.

REMARK 16.2 (Common random string). We will denote the strengthening of the invertible sampling hypothesis as in Definitions 16.3 and 16.4, where the setup algorithm Setup_S is required to sample uniform random strings, as the invertible sampling hypothesis with a *common random string* or *uniformly random CRS* (URC).

Again, in Definitions 16.3 and 16.4, we implicitly only consider *legitimate* adversaries which guarantee that $m \in L$.

16.3 EQUIVALENCE OF PSEUDORANDOM ENCODINGS AND INVERTIBLE SAMPLING

We prove the following theorem.

THEOREM 16.1. *Let $\alpha \in \{\approx_c, \equiv_s\}$ and $\beta \in \{\text{rand}, \text{det}\}$. PREH_α^β is true if and only if ISH_α^β is true.*

It is straight forward to extend Theorem 16.1 to the non-adaptive and adaptive variants with setup.

16.3.1 Every Inverse Samplable Distribution can be Pseudorandomly Encoded

LEMMA 16.1. *Let $\alpha \in \{\approx_c, \equiv_s\}$ and $\beta \in \{\text{rand}, \text{det}\}$. If ISH_α^β holds, then PREH_α^β holds.*

Proof. We prove this for the computational randomized case. The remaining cases are similar.

Assume $\text{ISH}_{\approx_c}^{\text{rand}}$ holds. Let S be a PPT algorithm. $\text{ISH}_{\approx_c}^{\text{rand}}$ implies that there exists an alternative sampler \bar{S} (with randomness space $\{0, 1\}^{n(\lambda)}$) and a corresponding inverse sampler \bar{S}^{-1} satisfying closeness and invertibility.

For $m \in L$, $y \in \{0, 1\}^*$, $r \in \{0, 1\}^{n(\lambda)}$, we define the encode and decode algorithms as follows:

$$\begin{aligned} E_S(m, y) &:= \bar{S}^{-1}(m, y), \\ D_S(m, r) &:= \bar{S}(m; r). \end{aligned}$$

CORRECTNESS. We consider a series of hybrids, see Figure 16.4. In the first game, $D_S(m, r) = \bar{S}(m; r) = y$. Since all games are computationally close, this must hold for the last game as well.

GAME G_0 . G_0 is identical to $\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{A}, m}^{(0)\text{-inv}}(\lambda)$.

GAME G_1 . G_1 is identical to $\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{A}, m}^{(1)\text{-inv}}(\lambda)$. Hence, we have $|\Pr[\text{out}_1 = 1] - \Pr[\text{out}_0 = 1]| \leq \text{Adv}_{(\bar{S}, \bar{S}^{-1}), \mathcal{A}, m}^{\text{inv}}(\lambda)$.

GAME G_2 . G_2 is identical to G_1 except that in G_1 y is sampled using the alternative sampler whereas in G_2 , y is sampled using the original sampler S .

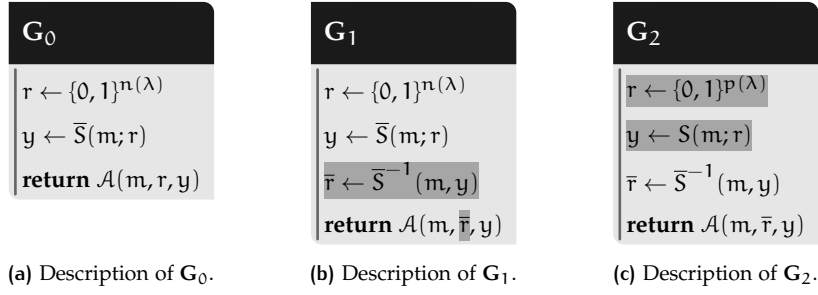


Figure 16.4: Hybrids used in the proof of correctness of Lemma 16.1.

CLAIM 16.1. For all PPT adversaries \mathcal{A} , for all $m \in L$, there exists a PPT adversary \mathcal{B}_2 , such that $|\Pr[out_2 = 1] - \Pr[out_1 = 1]| \leq \text{Adv}_{(\bar{S}, \bar{S}^{-1}), \mathcal{B}_2, m}^{\text{close}}(\lambda)$.

Proof. Construct an adversary \mathcal{B}_2 breaking closeness. On input of (m, y) , \mathcal{B}_2 computes $\bar{r} \leftarrow \bar{S}^{-1}(m, y)$, calls \mathcal{A} on input of (m, \bar{r}, y) and outputs the resulting output. If y is sampled using $\bar{S}(m; r)$ (for $r \leftarrow \{0, 1\}^{n(\lambda)}$), \mathcal{B}_2 perfectly simulates game G_1 for \mathcal{A} . If y is sampled using $S(m; r)$ (for $r \leftarrow \{0, 1\}^{p(\lambda)}$), \mathcal{B}_2 perfectly simulates game G_2 for \mathcal{A} . Therefore,

$$\Pr[out_1 = 1] = \Pr\left[\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{B}_2, m}^{(1)\text{-close}}(\lambda) = 1\right] \text{ and}$$

$$\Pr[out_2 = 1] = \Pr\left[\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{B}_2, m}^{(0)\text{-close}}(\lambda) = 1\right].$$

□

Thus, we have that

$$\begin{aligned} |\Pr[out_2 = 1] - \Pr[out_0 = 1]| &\leq \text{Adv}_{(\bar{S}, \bar{S}^{-1}), \bar{\mathcal{B}}, m}^{\text{close}}(\lambda) \\ &\quad + \text{Adv}_{(\bar{S}, \bar{S}^{-1}), \bar{\mathcal{B}}', m}^{\text{inv}}(\lambda) \end{aligned}$$

for some PPT adversaries $\bar{\mathcal{B}}$ and $\bar{\mathcal{B}}'$. Consider the adversary \mathcal{A} distinguishing between game G_0 and game G_2 who on input of (m, r, y) , outputs 0 if $\bar{S}(m; r) = y$ and outputs 1 otherwise. By definition, \mathcal{A} always outputs 0 in G_0 . Hence,

$$\begin{aligned} e_{(\mathcal{E}_S, \mathcal{D}_S), m}^{\text{dec-error}}(\lambda) &= \Pr\left[y \leftarrow S(m): \bar{S}(m, \bar{S}^{-1}(m, y)) \neq y\right] \\ &= \Pr[out_{2, \mathcal{A}} = 1] \\ &= |\Pr[out_{2, \mathcal{A}} = 1] - \Pr[out_{0, \mathcal{A}} = 1]| \\ &\leq \text{Adv}_{(\bar{S}, \bar{S}^{-1}), \bar{\mathcal{B}}, m}^{\text{close}}(\lambda) + \text{Adv}_{(\bar{S}, \bar{S}^{-1}), \bar{\mathcal{B}}', m}^{\text{inv}}(\lambda). \end{aligned}$$

PSEUDORANDOMNESS. We consider a sequence of hybrids starting from $\text{Exp}_{(\mathcal{E}_S, \mathcal{D}_S), \mathcal{A}, m}^{(0)\text{-pre}}(\lambda)$ and concluding with $\text{Exp}_{(\mathcal{E}_S, \mathcal{D}_S), \mathcal{A}, m}^{(1)\text{-pre}}(\lambda)$, see Figure 16.5.

GAME G_0 . G_0 is identical to $\text{Exp}_{(\mathcal{E}_S, \mathcal{D}_S), \mathcal{A}, m}^{(0)\text{-pre}}(\lambda)$.

GAME G_1 . G_1 is identical to G_0 except that in G_0 , y is sampled using the original sampler S whereas in G_1 , y is sampled using the alternative sampler \bar{S} .

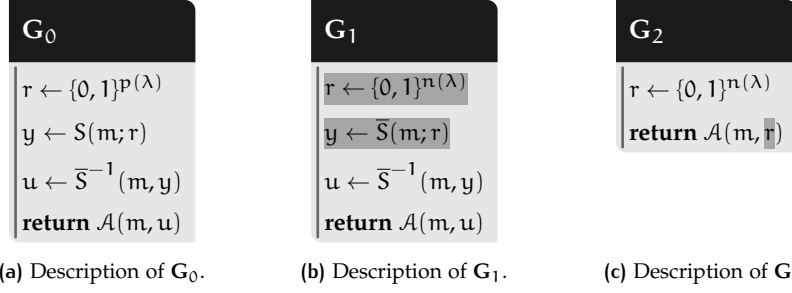


Figure 16.5: Hybrids used in the proof of pseudorandomness of Lemma 16.1.

CLAIM 16.2. For all PPT adversaries \mathcal{A} , for all $m \in L$, there exists a PPT adversary \mathcal{B}_1 , such that $|\Pr[out_1 = 1] - \Pr[out_0 = 1]| \leq \text{Adv}_{(\bar{S}, \bar{S}^{-1}), \mathcal{B}_1, m}^{\text{close}}(\lambda)$.

Proof. Construct a PPT adversary \mathcal{B}_1 on the closeness property as follows. On input of (m, y) , \mathcal{B}_1 calls \mathcal{A} on input of $(m, \bar{S}^{-1}(m, y))$ and outputs the resulting output. If \mathcal{B}_1 plays $\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{B}_1, m}^{(0)\text{-close}}(\lambda)$, he simulates game G_0 . Otherwise, if \mathcal{B}_1 plays $\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{B}_1, m}^{(1)\text{-close}}(\lambda)$, he simulates game G_1 . Hence,

$$\Pr[out_0 = 1] = \Pr\left[\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{B}_1, m}^{(0)\text{-close}}(\lambda) = 1\right] \text{ and}$$

$$\Pr[out_1 = 1] = \Pr\left[\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{B}_1, m}^{(1)\text{-close}}(\lambda) = 1\right].$$

□

GAME G_2 . G_2 is identical to G_1 except for the following difference. In G_1 , u is produced as inverse sampled random coins for y , whereas in G_2 u is a uniformly random bitstring from $\{0, 1\}^{n(\lambda)}$. Hence, G_2 is distributed identically to $\text{Exp}_{(E_S, D_S), \mathcal{A}, m}^{(1)\text{-pre}}(\lambda)$.

CLAIM 16.3. For all PPT adversaries \mathcal{A} , for all $m \in L$, there exists a PPT adversary \mathcal{B}_2 , such that $|\Pr[out_2 = 1] - \Pr[out_1 = 1]| \leq \text{Adv}_{(\bar{S}, \bar{S}^{-1}), \mathcal{B}_2, m}^{\text{inv}}(\lambda)$.

Proof. We construct a PPT adversary \mathcal{B}_2 on the invertibility property. On input of (m, r, y) , \mathcal{B}_2 calls \mathcal{A} on input of (m, r) and outputs \mathcal{A} 's output. If \mathcal{B}_2 plays the game $\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{B}_2, m}^{(0)\text{-inv}}(\lambda)$, he perfectly simulates G_1 . If \mathcal{B}_2 plays $\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{B}_2, m}^{(1)\text{-inv}}(\lambda)$, he perfectly simulates G_2 . Therefore,

$$\Pr[out_1 = 1] = \Pr\left[\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{B}_2, m}^{(0)\text{-inv}}(\lambda) = 1\right] \text{ and}$$

$$\Pr[out_2 = 1] = \Pr\left[\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{B}_2, m}^{(1)\text{-inv}}(\lambda) = 1\right].$$

□

Hence,

$$\begin{aligned} \text{Adv}_{(E_S, D_S), \mathcal{A}, m}^{\text{pre}}(\lambda) &= |\Pr[out_2 = 1] - \Pr[out_0 = 1]| \\ &\leq \text{Adv}_{(\bar{S}, \bar{S}^{-1}), \bar{\mathcal{B}}, m}^{\text{close}}(\lambda) + \text{Adv}_{(\bar{S}, \bar{S}^{-1}), \bar{\mathcal{B}}', m}^{\text{inv}}(\lambda) \end{aligned}$$

for some PPT adversaries $\bar{\mathcal{B}}$ and $\bar{\mathcal{B}}'$. □

16.3.2 Every Pseudorandomly Encodable Distribution can be Inverse Sampled

LEMMA 16.2. *Let $\alpha \in \{\approx_c, \equiv_s\}$ and $\beta \in \{\text{rand}, \text{det}\}$. If PREH_α^β holds, then ISH_α^β holds.*

Proof. We prove the statement for the computational randomized case. The remaining cases are similar.

Assume $\text{PREH}_{\approx_c}^{\text{rand}}$ holds. Let S be a PPT algorithm. $\text{PREH}_{\approx_c}^{\text{rand}}$ implies that for S there exist efficient algorithms E_S (possibly randomized) with output length $n(\lambda)$ and D_S (deterministic) satisfying correctness and pseudorandomness.

For $m \in L$, $r \in \{0, 1\}^{n(\lambda)}$, $y \in \{0, 1\}^*$, we define the alternative sampler and the inverse sampler as follows:

$$\begin{aligned}\bar{S}(m; r) &:= D_S(m, r), \\ \bar{S}^{-1}(m, y) &:= E_S(m, y).\end{aligned}$$

CLOSENESS. Let \mathcal{A} be an adversary breaking closeness. We consider a sequence of games starting from the game $\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{A}, m}^{(0)\text{-close}}(\lambda)$ and concluding with $\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{A}, m}^{(1)\text{-close}}(\lambda)$, see Figure 16.6.

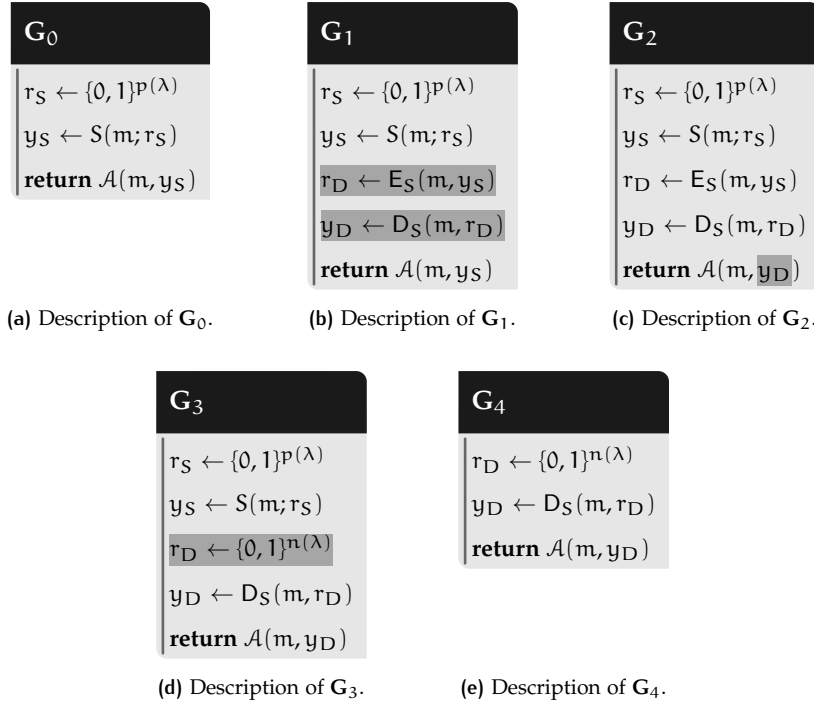


Figure 16.6: Hybrids used in the proof of closeness of Lemma 16.2.

GAME G_0 . G_0 is identical to $\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{A}, m}^{(0)\text{-close}}(\lambda)$.

GAME G_1 . G_1 is the same as G_0 except for the conceptual difference that G_1 additionally samples $r_D \leftarrow E_S(m, y_S)$ and $y_D \leftarrow D_S(m, r_D)$. This difference is only conceptual since the view of the adversary is not affected. Hence, $\Pr[out_0 = 1] = \Pr[out_1 = 1]$.

GAME G_2 . G_2 is identical to G_1 except that G_2 passes the value y_D to \mathcal{A} instead of the value y_S .

Hence, G_1 and G_2 proceed exactly identical if $y_S = y_D$. More formally, let F be the event that $y_S \neq y_D$. Then, $out_1 = 1 \wedge \neg F \Leftrightarrow out_2 = 1 \wedge \neg F$. Hence, the Difference Lemma (due to Shoup [Sho04]) bounds

$$|\Pr[out_2 = 1] - \Pr[out_1 = 1]| \leq \Pr[F].$$

Furthermore, correctness guarantees that for all $m \in L$,

$$\Pr[F] = \Pr[y_S \leftarrow S(m): D_S(m, E_S(m, y_S)) \neq y_S] = \epsilon_{(E_S, D_S), m}^{\text{dec-error}}(\lambda)$$

is negligible.

GAME G_3 . G_3 is identical to G_2 except for one difference. In G_2 , r_D is produced as $E_S(m, y_S)$, whereas in G_3 , r_D is sampled uniformly from $\{0, 1\}^{n(\lambda)}$.

CLAIM 16.4. For all PPT adversaries \mathcal{A} , for all $m \in L$, there exists a PPT adversary \mathcal{B}_3 , such that $|\Pr[out_3 = 1] - \Pr[out_2 = 1]| \leq \text{Adv}_{(E_S, D_S), \mathcal{B}_3, m}^{\text{pre}}(\lambda)$.

Proof. Construct an adversary \mathcal{B}_3 breaking pseudorandomness as follows. On input of $(m, u =: r_D)$, \mathcal{B}_3 calls \mathcal{A} on input $(m, D_S(m, r_D))$ and outputs the resulting output. If u is produced via $E_S(m, y_S)$ for $y_S \leftarrow S(m)$, \mathcal{B}_3 perfectly simulates game G_2 . Otherwise, if u is uniformly random over $\{0, 1\}^{n(\lambda)}$, \mathcal{B}_3 perfectly simulates game G_3 . Hence,

$$\begin{aligned} \Pr[out_3 = 1] &= \Pr\left[\text{Exp}_{(E_S, D_S), \mathcal{B}_3, m}^{(1)\text{-pre}}(\lambda) = 1\right] \text{ and} \\ \Pr[out_2 = 1] &= \Pr\left[\text{Exp}_{(E_S, D_S), \mathcal{B}_3, m}^{(0)\text{-pre}}(\lambda) = 1\right]. \end{aligned}$$

□

GAME G_4 . The game G_4 is identical to G_3 except for the conceptual difference that G_4 does not sample r_D and y_D since these values are never used. Hence, $\Pr[out_3 = 1] = \Pr[out_4 = 1]$. G_4 is identical to $\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{A}, m}^{(1)\text{-close}}(\lambda)$.

Hence,

$$\begin{aligned} \text{Adv}_{(\bar{S}, \bar{S}^{-1}), \mathcal{A}, m}^{\text{close}}(\lambda) &= |\Pr[out_4 = 1] - \Pr[out_0 = 1]| \\ &\leq \text{Adv}_{(E_S, D_S), \bar{\mathcal{B}}, m}^{\text{pre}}(\lambda) + \epsilon_{(E_S, D_S), m}^{\text{dec-error}}(\lambda) \end{aligned}$$

for some PPT adversary $\bar{\mathcal{B}}$.

INVERTIBILITY. We consider a sequence of hybrids starting from the game $\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{A}, m}^{(1)\text{-inv}}(\lambda)$ and concluding with $\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{A}, m}^{(0)\text{-inv}}(\lambda)$, see Figure 16.7.

GAME G_0 . G_0 is identical to $\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{A}, m}^{(1)\text{-inv}}(\lambda)$.

GAME G_1 . G_1 is identical to G_0 up to a small difference. In G_0 , y is sampled using D_S on uniform randomness. In G_1 , y is sampled using the original sampler S .

CLAIM 16.5. For all PPT adversaries \mathcal{A} , for all $m \in L$, there exists a PPT adversary \mathcal{B}_1 , such that $|\Pr[out_1 = 1] - \Pr[out_0 = 1]| \leq \text{Adv}_{(E_S, D_S), \mathcal{B}_1, m}^{\text{pre}}(\lambda) + \epsilon_{(E_S, D_S), m}^{\text{dec-error}}(\lambda)$.

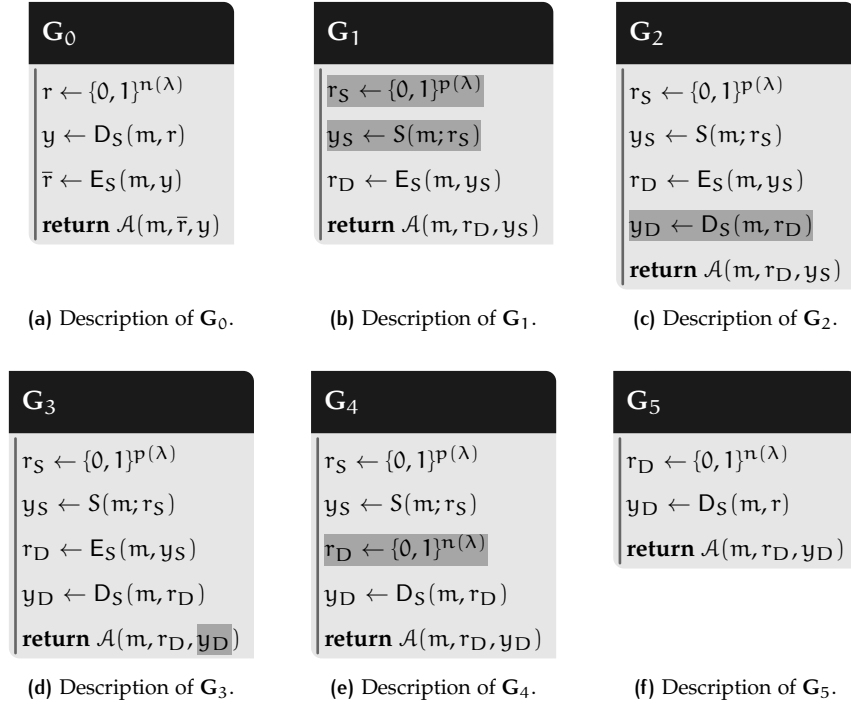


Figure 16.7: Hybrids used in the proof of invertibility of Lemma 16.2.

Proof. Let \mathcal{A} be an adversary distinguishing G_0 and G_1 . Construct an adversary \mathcal{B}_1 on the closeness property. On input of (m, y) , \mathcal{B}_1 computes $\bar{r} \leftarrow E_S(m, y)$ and calls \mathcal{A} on input (m, \bar{r}, y) . If y is sampled via $\bar{S}(m)$, \mathcal{B}_1 simulates game G_0 for \mathcal{A} . Else, if y is sampled via $S(m)$, \mathcal{B}_1 simulates game G_1 for \mathcal{A} . Hence,

$$|\Pr[out_1 = 1] - \Pr[out_0 = 1]| \leq \text{Adv}_{(\bar{S}, \bar{S}^{-1}), \mathcal{B}_1, m}^{\text{close}}(\lambda).$$

□

GAME G_2 . G_2 is identical to G_1 except for the conceptual difference that G_2 additionally computes y_D as $D_S(m, r_D)$ but never uses this value. Hence, $\Pr[out_1 = 1] = \Pr[out_2 = 1]$.

GAME G_3 . G_3 is identical to G_2 except that G_3 passes the value y_D to \mathcal{A} instead of the value y_S . G_2 and G_3 behave identical if $y_D = y_S$. Let F denote the failure event $y_D \neq y_S$. Hence, $out_2 = 1 \wedge \neg F \Leftrightarrow out_3 = 1 \wedge \neg F$. The Difference Lemma (due to Shoup [Sho04]) bounds

$$|\Pr[out_3 = 1] - \Pr[out_2 = 1]| \leq \Pr[F].$$

Due to correctness, for all $m \in L$,

$$\Pr[F] = \Pr[y_S \leftarrow S(m) : D_S(m, E_S(m, y_S)) \neq y_S] = \epsilon_{(E_S, D_S), m}^{\text{dec-error}}(\lambda)$$

is negligible.

GAME G_4 . G_4 is identical to G_3 except that G_3 produces r_D via $E_S(m, y_S)$, whereas G_4 samples r_D uniformly from $\{0, 1\}^{n(\lambda)}$.

CLAIM 16.6. For all PPT adversaries \mathcal{A} , for all $m \in L$, there exists a PPT adversary \mathcal{B}_4 , such that $|\Pr[out_4 = 1] - \Pr[out_3 = 1]| \leq \text{Adv}_{(\mathcal{E}_S, \mathcal{D}_S), \mathcal{B}_4, m}^{\text{pre}}(\lambda)$.

Proof. Construct a PPT adversary \mathcal{B}_4 breaking the pseudorandomness property. On input of (m, u) , \mathcal{B}_4 calls \mathcal{A} on input of $(m, u =: r_D, D_S(m, u) =: y_D)$ and outputs the resulting output. If u is sampled via $\mathcal{E}_S(m, y)$ for $y \leftarrow S(m)$, \mathcal{B}_4 perfectly simulates game \mathbf{G}_3 for \mathcal{A} . Otherwise, if u is uniformly random over $\{0, 1\}^{n(\lambda)}$, \mathcal{B}_4 perfectly simulates game \mathbf{G}_4 for \mathcal{A} . Hence,

$$\begin{aligned} \Pr[out_3 = 1] &= \Pr\left[\text{Exp}_{(\mathcal{E}_S, \mathcal{D}_S), \mathcal{B}_4, m}^{(0)\text{-pre}}(\lambda) = 1\right] \text{ and} \\ \Pr[out_4 = 1] &= \Pr\left[\text{Exp}_{(\mathcal{E}_S, \mathcal{D}_S), \mathcal{B}_4, m}^{(1)\text{-pre}}(\lambda) = 1\right]. \end{aligned}$$

□

GAME \mathbf{G}_5 . The difference between \mathbf{G}_4 and \mathbf{G}_5 is that \mathbf{G}_5 does not produce the obsolete values r_S and y_S . Hence, this game hop is again only conceptual and $\Pr[out_4 = 1] = \Pr[out_5 = 1]$.

Hence,

$$\begin{aligned} |\Pr[out_5 = 1] - \Pr[out_0 = 1]| &\leq 2 \cdot \text{Adv}_{(\mathcal{E}_S, \mathcal{D}_S), \overline{\mathcal{B}}, m}^{\text{pre}}(\lambda) \\ &\quad + 2 \cdot \epsilon_{(\mathcal{E}_S, \mathcal{D}_S), m}^{\text{dec-error}}(\lambda) \end{aligned}$$

for some PPT adversary $\overline{\mathcal{B}}$. □

The above proof directly generalizes to the non-adaptive and adaptive variants of pseudorandom encodings and invertible sampling with public parameters. See Figure 16.8 for a proof sketch for the adaptive version.

Theorem 16.1 together with Theorem 15.1 yields the following corollaries.


COROLLARY 16.1. *Let $\alpha \in \{\approx_c, \equiv_s\}$ and $\beta \in \{\text{rand}, \text{det}\}$. If $\text{cPREH}_\alpha^\beta$ is true, then $\text{aclSH}_\alpha^\beta$ is true.*

COROLLARY 16.2. *Let $\alpha \in \{\approx_c, \equiv_s\}$ and $\beta \in \{\text{rand}, \text{det}\}$. If clSH_α^β is true, then $\text{aclSH}_\alpha^\beta$ is true.*


This is an excellent example of the potential pseudorandom encodings offer. Looking through the lens of invertible sampling, it is rather unclear how the static and adaptive notions relate, whereas pseudorandom encodings directly provide a static-to-adaptive transformation.

Particularly, Corollary 16.2 together with [DKR15] yields the first instantiation of an *adaptive* explainability compiler without complexity leveraging and, hence, based only on polynomial hardness assumptions. The recent paper [CSW19] uses such an adaptive explainability compiler to obtain adaptive MPC with communication complexity which is sublinear in the circuit size. Their construction relies on complexity leveraging which entails a subexponential loss relative to IO and one-way functions. Hence, we obtain the following corollary improving on Theorem 7 in the proceedings version of [CSW19] in a black-box way.


COROLLARY 16.3. *Assuming polynomially secure IO and the adaptive hardness of LWE, then succinct adaptive two-round MPC in the malicious setting is possible (for all deterministic n -party functionalities $f: (\{0, 1\}^{\ell_{in}})^n \rightarrow \{0, 1\}^{\ell_{out}}$ with circuit depth d) in the global CRS model such that the size of the CRS, the communication complexity, and online-computational complexity of the protocol are polynomial in $\lambda, \ell_{in}, \ell_{out}, d$ and n .*



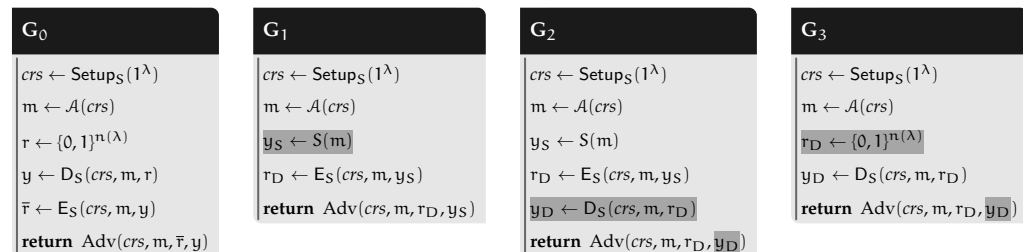
(a) *Adaptive correctness.* In G_0 , for adaptively chosen m , we have $D_S(crs, m, r) = \bar{S}(crs, m; r) = y$. Game hop from G_0 to G_1 follows by adaptive invertibility. Game hop from G_1 to G_2 follows from adaptive closeness. Hence, in G_2 , $D_S(crs, m, r) = \bar{S}(crs, m; r) = y$, where $r = E_S(crs, m, y) = \bar{S}^{-1}(crs, m, y)$.



(b) *Adaptive pseudorandomness.* G_0 corresponds to adaptive pseudorandomness game, where adversary receives encoded randomness, that is $E_S(crs, m, y) = \bar{S}^{-1}(crs, m, y)$. Game hop from G_0 to G_1 is justified by adaptive closeness. Game hop from G_1 to G_2 is justified by adaptive invertibility. G_2 corresponds to adaptive pseudorandomness game, where adversary receives true randomness.



(c) *Adaptive closeness.* G_0 corresponds to the adaptive closeness game, where \mathcal{A} receives a sample from S . If y_S and y_D are identical, G_0 and G_1 are identical. Adaptive correctness guarantees this with overwhelming probability. The game hop from G_1 to G_2 is justified by adaptive pseudorandomness. G_2 corresponds to the adaptive closeness game, where \mathcal{A} receives a sample from \bar{S} .



(d) *Adaptive invertibility.* Game G_0 corresponds to the adaptive invertibility game, where \mathcal{A} receives inverse sampled random coins. The game hop from G_0 to G_1 is justified by adaptive closeness which follows from adaptive correctness and adaptive pseudorandomness. The game hop from G_1 to G_2 is justified by adaptive correctness and the game hop from G_2 to G_3 is justified by adaptive pseudorandomness. G_3 corresponds to the adaptive invertibility game, where \mathcal{A} receives the actual randomness.

Figure 16.8: Proof sketch for the equivalence of the adaptive notions of pseudorandom encodings and invertible sampling.

17

A TAXONOMY OF PSEUDORANDOM ENCODINGS

In this chapter we classify the different variants of the pseudorandom encoding hypothesis. In Section 17.1, we study the pseudorandom encoding hypothesis with deterministic encoding algorithm and identify a relation to compression. In Section 17.2, we study the pseudorandom encoding hypothesis with randomized encoding and its conflicts with extractable one-way functions. In Section 17.3, we describe an instantiation of $\text{cPREH}_{\approx_c}^{\text{rand}}$ (with and without universal setup) based on indistinguishability obfuscation and one-way functions due to [SW14; DKR15]. Finally, in Section 17.4, we bootstrap $\text{cPREH}_{\approx_c}^{\text{rand}}$ with a common *random* string from the construction in Section 17.3 in conjunction with *weak* $\text{cPREH}_{\approx_c}^{\text{rand}}$ with a common *random* string for the setup algorithm from Section 17.3.

17.1 DETERMINISTIC ENCODING ALGORITHM

For the purpose of classifying pseudorandom encodings with a deterministic encoding algorithm, we first introduce some notions of entropy and computational analogues thereof.

Min-entropy captures the ability to guess the value of a distribution in a single attempt, cf. Definition 2.2. However, this is a very pessimistic view. For many purposes, it suffices to work with a distribution which is statistically close (i. e., has negligible statistical distance) to a distribution with high min-entropy.⁶¹

DEFINITION 17.1 (ϵ -smooth min-entropy, [Rey11]). A source X has ϵ -smooth min-entropy at least k , denoted as $H_{\infty}^{\epsilon}(X) \geq k$, if there exists a distribution X' with $\Delta(X, X') \leq \epsilon$ such that $H_{\infty}(X') \geq k$.

In many cases, some information Z that is correlated to the actual source X is known. Since for our purposes, the conditional part Z is not under adversarial control, we use the notion of *average* conditional min-entropy as used in [HLR07; DORS08], cf. Definition 2.3. That is, the average min-entropy of some distribution X conditioned on Z is defined as $-\log(\mathbb{E}_{z \leftarrow Z}[\max_x \Pr[X_z = x]])$.

The notion of average min-entropy can be relaxed in a similar way as Definition 17.1 as follows.

DEFINITION 17.2 (Average ϵ -smooth min-entropy, [DORS08; Rey11]). Let (X, Z) be a joint distribution. The distribution X has *average ϵ -smooth min-entropy* at least k *conditioned on* Z , denoted as $\tilde{H}_{\infty}^{\epsilon}(X | Z) \geq k$, if there exists a joint distribution (X', Z') with $\Delta((X, Z), (X', Z')) \leq \epsilon$ such that $\tilde{H}_{\infty}(X' | Z') \geq k$.

Let X be an efficiently samplable distribution. In the literature, there are several computational notions of entropy. HILL entropy [HILL99] constitutes a natural computational variant of min-entropy. A source has high HILL

⁶¹ We note that if some distribution X is ϵ -close to a distribution Y , $H_{\infty}(X) \geq -\log(2^{-H_{\infty}(Y)} + \epsilon(\lambda)) = H_{\infty}(Y) - \log(1 + \epsilon(\lambda)2^{H_{\infty}(Y)})$.

entropy, if it is computationally indistinguishable from a source that has high min-entropy.

DEFINITION 17.3 (HILL entropy, [HILL99; BSW03]). A distribution X has *HILL entropy* at least k , denoted by $H_{\epsilon,s}^{\text{HILL}}(X) \geq k$, if there exists a distribution Y such that $H_\infty(Y) \geq k$ and $|\Pr[x \leftarrow X: \mathcal{A}(x) = 1] - \Pr[y \leftarrow Y: \mathcal{A}(y) = 1]| \leq \epsilon$ for all \mathcal{A} of size at most s .

Shannon's source coding theorem [Sha48] states that the minimum compression length (over all compression and decompression algorithms) of a distribution equals its average entropy (up to small additive terms). Yao entropy [Yao82] constitutes the corresponding computational counterpart. Intuitively, a source has high Yao entropy, if it cannot be *efficiently* compressed.

DEFINITION 17.4 (Yao entropy, [Yao82; BSW03]). A distribution X has *Yao entropy* at least k , denoted by $H_{\epsilon,s}^{\text{Yao}}(X) \geq k$, if for every pair of circuits (E, D) of total size s with outputs of E having length ℓ , $\Pr_{x \leftarrow X}[D(E(x)) = x] \leq 2^{\ell-k} + \epsilon$.

When we omit the subscripts for H^{HILL} and H^{Yao} , we mean $H_{\epsilon,s}^{\text{HILL}}$ and $H_{\epsilon,s}^{\text{Yao}}$ for any negligible ϵ and polynomial s , respectively. Since compressibility implies distinguishability, HILL entropy implies Yao entropy. The converse, however, is believed to be false, [Wee04; HLR07].

DEFINITION 17.5 (Conditional HILL entropy, [HLR07]). For a distribution (X, Z) , we say that X has HILL entropy at least k conditioned on Z , denoted by $H_{\epsilon,s}^{\text{HILL}}(X | Z) \geq k$, if there exists a collection of distributions Y_z giving rise to a joint distribution (Y, Z) , such that $\tilde{H}_\infty(Y | Z) \geq k$ and $|\Pr[(x, z) \leftarrow (X, Z): \mathcal{A}(x, z) = 1] - \Pr[(y, z) \leftarrow (Y, Z): \mathcal{A}(y, z) = 1]| \leq \epsilon$ for all circuits \mathcal{A} of size at most s .

Conditional Yao entropy is defined by simply giving the compressor and decompressor algorithm the value z as input.

DEFINITION 17.6 (Conditional Yao entropy, [HLR07]). For a distribution (X, Z) , we say that X has Yao entropy at least k conditioned on Z , denoted by $H_{\epsilon,s}^{\text{Yao}}(X | Z) \geq k$, if for every pair of circuits (E, D) of total size at most s with outputs of E having length ℓ , $\Pr_{(x,z) \leftarrow (X,Z)}[D(z, E(z, x)) = x] \leq 2^{\ell-k} + \epsilon$.

17.1.1 Information-theoretic Guarantees and Compression

Traditionally, the theory of compression mostly considers families of sources that are not indexed by strings. Let Σ be some alphabet. In this part, we always consider $\Sigma := \{0, 1\}$. Let X be a source.

DEFINITION 17.7 ([Wee04; TVZ05]). Let $E_X: \Sigma^* \rightarrow \Sigma^*$ and $D_X: \Sigma^* \rightarrow \Sigma^*$ be functions. We say (E_X, D_X) *compresses source X to length m with decoding error ϵ* if

- $\Pr[x \leftarrow X: D_X(E_X(x)) \neq x] \leq \epsilon$, and
- $\mathbb{E}[|E_X(X)|] \leq m$.

DEFINITION 17.8 ([Wee04; TVZ05]). We say source X is *compressible to length (exactly) m* if there exist functions E_X and D_X such that (E_X, D_X) compresses X to length (exactly) m .

The decoding error can be eliminated at the cost of longer codewords.

LEMMA 17.1 ([TVZ05]). *Let X_λ be a source on $\{0, 1\}^\lambda$ which is compressible to length m with decoding error ϵ by algorithms (E_X, D_X) . Further, let $m_0 \in \mathbb{N}$ be a lower bound on the output length of E_X , i.e., such that for all $x \in \text{supp}(X_\lambda)$, $|E_X(x)| \geq m_0$. Then, X_λ is compressible to length $m + \epsilon(\lambda - m_0) + 1$ with decoding error 0.*

Proof sketch. To show this, [TVZ05] construct an encoding algorithm E'_X which on input of x tests if $D_X(E_X(x)) = x$. If this is the case, E'_X outputs $0 \parallel E_X(x)$. Else, E'_X outputs $1 \parallel x$. \square

The statistical deterministic variant of the pseudorandom encoding hypothesis is strongly related to compression.

THEOREM 17.1. *If (weak) $\text{PREH}_{\equiv_s}^{\text{det}}$ is true for X , i.e., there exist deterministic algorithms (E_X, D_X) with E_X having output length n satisfying correctness and pseudorandomness. Then, the ϵ -smooth min-entropy $H_\infty^\epsilon(X) \geq n$ for some negligible function ϵ .*

Proof. Consider the distribution $Y' := D_X(U_n)$. As already seen in the proof of Theorem 16.1, Y' and X are statistically indistinguishable due to correctness and pseudorandomness. Hence, the statistical distance $\Delta(X, Y') \leq \delta(\lambda)$ for some negligible δ .

Due to correctness and since E_X is deterministic,

$$\begin{aligned} \Pr[x \leftarrow X: D_X(E_X(x)) = x] &= \Pr[x \leftarrow X: E_X(D_X(E_X(x))) = E_X(x)] \\ &\geq 1 - \nu(\lambda) \end{aligned}$$

for some negligible function ν . Applying pseudorandomness, we get that the probability

$$\Pr[u \leftarrow U_n: E_X(D_X(u)) = u] \geq 1 - \nu'(\lambda)$$

for some negligible function ν' , where the probability is only over the choice of u . Therefore, D_X operates almost injectively on the set $\{0, 1\}^n$. More formally, let $V_0 \subset \{0, 1\}^n$ denote the set of all u such that $E_X(D_X(u)) = u$ and let $V_1 := \{0, 1\}^n \setminus V_0$. We have that, $|V_1|/|\{0, 1\}^n| \leq \nu'(\lambda)$. Let $\overline{V_1}$ be some arbitrary subset of $\{0, 1\}^{k(\lambda)} \setminus D_X(V_0)$ such that $|\overline{V_1}| = |V_1|$ (for some polynomial k). Let Y'' be the uniform distribution over $D_X(V_0) \cup \overline{V_1}$. Note that $|D_X(V_0) \cup \overline{V_1}| = 2^n$ and $\Pr[Y' \in \overline{V_1}] \leq \Pr[Y' \notin D_X(V_0)] \leq \nu'(\lambda)$. The statistical distance between Y' and Y'' is

$$\begin{aligned} \Delta(Y', Y'') &= \sum_{a \in D_X(V_0) \cup \overline{V_1}} |\Pr[Y' = a] - \Pr[Y'' = a]| \\ &= \underbrace{\sum_{a \in D_X(V_0)} |\Pr[Y' = a] - \Pr[Y'' = a]|}_{\leq 1 - |V_0| \cdot 2^{-n} \leq 1 - (1 - \nu'(\lambda))} \\ &\quad + \underbrace{\sum_{a \in \overline{V_1}} |\Pr[Y' = a] - \Pr[Y'' = a]|}_{\leq \sum_{a \in \overline{V_1}} (\Pr[Y' = a] + \Pr[Y'' = a]) \leq 2 \cdot \nu'(\lambda)} \leq 3 \cdot \nu'(\lambda) \end{aligned}$$

and hence negligible. Furthermore, since Y'' is the uniform distribution over a set of size 2^n , its min-entropy equals n . Since $\Delta(X, Y') \leq \delta(\lambda)$ and $\Delta(Y', Y'') \leq 3\nu'(\lambda)$ for negligible functions δ and ν' , we have

$$H_{\infty}^{\delta+3\nu'}(X) \geq H_{\infty}(Y'') = n.$$

□

Theorem 17.1 directly yields the following corollary.

COROLLARY 17.1. *If $\text{PREH}_{\equiv_s}^{\text{det}}$ is true for some source X , there exists an $n \leq H_{\infty}^{\epsilon}(X)$ (for some negligible function ϵ), such that X is compressible to length exactly n . The decoding error can then be eliminated applying Lemma 17.1.*

Hence, for $\text{PREH}_{\equiv_s}^{\text{det}}$ to be true for a source X , it is necessary that X is efficiently compressible. However, this is not a sufficient criterion since a compression algorithm can have some structure which makes it easily distinguishable from uniform randomness (e. g., the compression algorithm used in the proof of Lemma 17.1).

The distribution induced by pseudorandom generators can provably not be compressed. However, those distributions have low ϵ -smooth min-entropy.

LEMMA 17.2. *Let iPRG be an injective pseudorandom generator with polynomial stretch $\text{poly}(\cdot)$ and let ϵ be a negligible function. Then, $H_{\infty}^{\epsilon}(\text{iPRG}(U_{\lambda})) \leq \lambda + \delta(\lambda)$ for some negligible function δ .*

Proof. Since $\text{iPRG}(U_{\lambda})$ is a uniform distribution over $\text{iPRG}(\{0, 1\}^{\lambda})$, it suffices to consider all ϵ -close flat distributions over $\{0, 1\}^{\text{poly}(\lambda)}$ to obtain an upper bound on $H_{\infty}^{\epsilon}(\text{iPRG}(U_{\lambda}))$. Let D be the uniform distribution over $\text{iPRG}(\{0, 1\}^{\lambda}) \cup A$, for some set A such that $\text{iPRG}(\{0, 1\}^{\lambda}) \cap A = \emptyset$. Let $|A| = k$. Then,

$$\begin{aligned} \Delta(\text{iPRG}(U_{\lambda}), D) &= \sum_{a \in \text{iPRG}(\{0, 1\}^{\lambda})} \left| \frac{1}{2^{\lambda}} - \frac{1}{2^{\lambda} + k} \right| + \sum_{a \in A} \left| \frac{1}{2^{\lambda} + k} \right| \\ &= 1 - \frac{2^{\lambda}}{2^{\lambda} + k} + \frac{k}{2^{\lambda} + k} = \frac{2k}{2^{\lambda} + k}. \end{aligned}$$

Hence, for $\Delta(\text{iPRG}(U_{\lambda}), D) \leq \epsilon(\lambda)$, k is upper bounded by

$$k \leq \underbrace{k(2 - \epsilon(\lambda))}_{\geq 1} \leq 2^{\lambda} \cdot \epsilon(\lambda).$$

Therefore, the min-entropy of D is at most $H_{\infty}(D) = \log(2^{\lambda} + k) \leq \log(2^{\lambda} + 2^{\lambda}\epsilon) = \lambda + \log(1 + \epsilon)$, which is negligibly close to λ . Hence, $H_{\infty}^{\epsilon}(\text{iPRG}(U_{\lambda})) \leq \lambda + \delta(\lambda)$ for some negligible function δ . □

Thus, if injective PRGs exist (and since the distribution induced by iPRG is clearly compatible, cf. Definition 15.3), then $\text{PREH}_{\equiv_s}^{\text{det}}$ is false. We can further strengthen this result by considering general PRGs.

LEMMA 17.3. *Let PRG be a pseudorandom generator with polynomial stretch $\text{poly}(\cdot)$ and let ϵ be a negligible function. Then, $H_{\infty}^{\epsilon}(\text{PRG}(U_{\lambda})) \leq \lambda + \delta(\lambda)$ for some negligible function δ .*

Proof. Due to Lemma 15.1, there exists a uniform distribution D over a set of size 2^λ such that $\text{PRG}(U_\lambda)$ and D are statistically close. By a similar argument as in Lemma 17.2, in order to upper bound the ϵ -smooth min-entropy of $\text{PRG}(U_\lambda)$, it suffices to consider all ϵ -close flat distributions. Then, by the same computations as in the proof of Lemma 17.2, Lemma 17.3 follows. \square

Theorem 17.1 together with Lemmas 15.1 and 17.3 yield the following corollary.

COROLLARY 17.2. *If one-way functions exist, $\text{PREH}_{\equiv_s}^{\text{det}}$ is false.*

INFORMATION-THEORETIC GUARANTEES WITH SETUP. Intuitively, a common reference string cannot add additional entropy to $E_X(\text{crs}, X)$. That is, if $\text{cPREH}_{\equiv_s}^{\text{det}}$ is true for some source X , the algorithms (E_X, D_X) implied by $\text{cPREH}_{\equiv_s}^{\text{det}}$ compress that source to its ϵ -smooth min-entropy. To prove this, we introduce the following technical lemma.

LEMMA 17.4 (Splitting lemma). *Let X and Y_x be distributions giving rise to a joint distribution (Y, X) . Let $A \subset \text{supp}(Y, X)$ such that $\Pr_{(y,x) \leftarrow (Y,X)}[(y,x) \in A] \geq 1 - \mu(\lambda)$ for some negligible function μ . Further, for all $x \in \text{supp}(X)$, let $p_x := \Pr_{y \leftarrow Y_x}[(y,x) \in A]$. Then, there exist negligible functions ν, ν' and a set $G \subseteq \text{supp}(X)$, such that for all $x \in G$, $p_x \geq 1 - \nu(\lambda)$ and $\Pr_{x \leftarrow X}[x \in G] \geq 1 - \nu'(\lambda)$.*

Proof. Let $\mu(\lambda)$ be some negligible function. Define $G_n := \{x \in \text{supp}(X) : p_x \geq 1 - n \cdot \mu(\lambda)\}$. Then,

$$\begin{aligned} 1 - \mu(\lambda) &\leq \Pr_{(y,x) \leftarrow (Y,X)}[(y,x) \in A] \\ &= \sum_{x \in G_n} \Pr[X = x] \cdot \underbrace{p_x}_{\leq 1} + \sum_{x \notin G_n} \Pr[X = x] \cdot \underbrace{p_x}_{< 1 - n \cdot \mu(\lambda)} \\ &< \Pr_{x \leftarrow X}[x \in G_n] + (1 - n \cdot \mu(\lambda)) \cdot \underbrace{\sum_{x \notin G_n} \Pr[X = x]}_{= 1 - \Pr_{x \leftarrow X}[x \in G_n]} \\ &= \Pr_{x \leftarrow X}[x \in G_n] + 1 - n \cdot \mu(\lambda) - \Pr_{x \leftarrow X}[x \in G_n] \\ &\quad + n \cdot \mu(\lambda) \cdot \Pr_{x \leftarrow X}[x \in G_n] \\ &= 1 - \mu(\lambda) \cdot \left(n - n \cdot \Pr_{x \leftarrow X}[x \in G_n] \right). \end{aligned}$$

Hence,

$$\Pr_{x \leftarrow X}[x \in G_n] \geq 1 - \frac{1}{n}.$$

Without loss of generality, we assume $\mu(\lambda) > 0$ for all $\lambda \in \mathbb{N}$. (If $\mu(\lambda) = 0$ for some λ , then for $G := \text{supp}(X)$, $p_x = 1$ for all $x \in G$, and $\Pr_{x \leftarrow X}[x \in G] = 1$.) Then, $\sqrt{\mu(\lambda)}$ is well-defined and negligible. Let $n := \sqrt{\mu(\lambda)}^{-1}$. Then, by definition of G_n , for all $x \in G_n$, $p_x \geq 1 - n \cdot \mu(\lambda) = 1 - \sqrt{\mu(\lambda)}$. Furthermore, $\Pr_{x \leftarrow X}[x \in G_n] \geq 1 - \frac{1}{n} = 1 - \sqrt{\mu(\lambda)}$. \square

THEOREM 17.2. *If (weak) $\text{cPREH}_{\equiv_s}^{\text{det}}$ is true for X , i. e., there exist a setup algorithm Setup_X and deterministic algorithms (E_X, D_X) with E_X having output length n satisfying correctness and pseudorandomness. Then, the ϵ -smooth min-entropy $H_\infty^\epsilon(X) \geq n$ for some negligible function ϵ .*

Proof. Note that $H_\infty^\epsilon(X) = \tilde{H}_\infty^\epsilon(X | \text{Setup}_X(1^\lambda))$ as the distribution X is independent of the CRS. Hence, it suffices to upper bound the average ϵ -smooth min-entropy $\tilde{H}_\infty^\epsilon(X | \text{Setup}_X(1^\lambda))$.

The proof is similar to the proof of Theorem 17.1. We consider the distribution $Y'_{crs} := D_X(crs, U_n)$. Due to Theorem 16.1, the distributions

$$\left\{ crs \leftarrow \text{Setup}_X(1^\lambda), u \leftarrow U_n : (crs, D_X(crs, u)) \right\} \text{ and } \\ \left\{ crs \leftarrow \text{Setup}_X(1^\lambda), y \leftarrow X : (crs, y) \right\}$$

are statistically indistinguishable. Hence,

$$\Delta((X | \text{Setup}_X(1^\lambda), (Y' | \text{Setup}_X(1^\lambda))) \leq \delta(\lambda)$$

for some negligible function δ .

Further, by a similar argument as in Theorem 17.1, due to correctness and pseudorandomness we have

$$\Pr [crs \leftarrow \text{Setup}_X(1^\lambda), u \leftarrow U_n : E_X(crs, D_X(crs, u)) = u] \geq 1 - \nu'(\lambda)$$

for some negligible function ν' , where the probability is over the choice of crs and u .

The Splitting lemma (Lemma 17.4) implies that there exists a set $G \subseteq \text{supp}(\text{Setup}_X(1^\lambda))$ such that $\Pr[crs \leftarrow \text{Setup}_X(1^\lambda) : crs \in G] \geq 1 - \nu''(\lambda)$ and for all $crs \in G$, $\Pr[u \leftarrow U_n : E_X(crs, D_X(crs, u)) = u] \geq 1 - \nu'''(\lambda)$ for some negligible functions ν'', ν''' . Hence, conditioned on $crs \in G$, $D_X(crs, \cdot)$ operates almost injectively on the set $\{0, 1\}^n$. More formally, let $V_0^{crs} \subset \{0, 1\}^n$ denote the set of all u such that $E_X(crs, D_X(crs, u)) = u$ and let $V_1^{crs} := \{0, 1\}^n \setminus V_0^{crs}$. For $crs \in G$, we have $|V_1^{crs}|/|\{0, 1\}^n| \leq \nu'''(\lambda)$. Let \overline{V}_1^{crs} be some arbitrary subset of $\{0, 1\}^{k(\lambda)} \setminus D_X(crs, V_0^{crs})$ such that $|\overline{V}_1^{crs}| = |V_1^{crs}|$ (for some polynomial k). Let Y'' be the uniform distribution over $D_X(V_0^{crs}) \cup \overline{V}_1^{crs}$. Note that $|D_X(crs, V_0^{crs}) \cup \overline{V}_1^{crs}| = 2^n$, and for $crs \in G$ we have $\Pr[Y'_{crs} \in \overline{V}_1^{crs}] \leq \Pr[Y'_{crs} \notin D_X(crs, V_0^{crs})] \leq \nu'(\lambda)$. We have that

$$\begin{aligned} & \Delta((Y', \text{Setup}_X), (Y'', \text{Setup}_X)) \\ &= \sum_{\substack{crs \in \text{supp}(\text{Setup}_X) \\ a \in \{0, 1\}^{k(\lambda)}}} \left| \Pr[(Y', \text{Setup}_X) = (a, crs)] \right. \\ & \quad \left. - \Pr[(Y'', \text{Setup}_X) = (a, crs)] \right| \\ &= \sum_{crs \in G} \Pr[\text{Setup}_X = crs] \cdot \sum_{a \in \{0, 1\}^{k(\lambda)}} \left| \Pr[Y'_{crs} = a] - \Pr[Y''_{crs} = a] \right| \\ &+ \sum_{crs \notin G} \Pr[\text{Setup}_X = crs] \cdot \sum_{a \in \{0, 1\}^{k(\lambda)}} \left| \Pr[Y'_{crs} = a] - \Pr[Y''_{crs} = a] \right| \end{aligned} \tag{17.1}$$

The separate terms can be upper bounded as follows.

$$\begin{aligned} & \sum_{crs \in G} \Pr[\text{Setup}_X = crs] \cdot \sum_{a \in \{0, 1\}^{k(\lambda)}} \left| \Pr[Y'_{crs} = a] - \Pr[Y''_{crs} = a] \right| \\ &= \sum_{crs \in G} \Pr[\text{Setup}_X = crs] \cdot \\ & \quad \left(\begin{aligned} & \sum_{a \in D_X(crs, V_0^{crs})} \left| \Pr[Y'_{crs} = a] - \Pr[Y''_{crs} = a] \right| \\ & + \sum_{a \in \overline{V}_1^{crs}} \left| \Pr[Y'_{crs} = a] - \Pr[Y''_{crs} = a] \right| \end{aligned} \right) \end{aligned}$$

$$\begin{aligned}
&\leq \sum_{crs \in G} \Pr[\text{Setup}_X = crs] \cdot \left(1 - \frac{|V_0^{crs}|}{2^n} + \Pr[Y'_{crs} \in \overline{V_1^{crs}}] \right. \\
&\quad \left. + \Pr[Y''_{crs} \in \overline{V_1^{crs}}] \right) \\
&\leq \sum_{crs \in G} \Pr[\text{Setup}_X = crs] \cdot \left(\frac{|V_1^{crs}|}{2^n} + \Pr[Y'_{crs} \notin D_X(crs, V_0^{crs})] + \frac{|V_1^{crs}|}{2^n} \right) \\
&\leq 3 \cdot \nu'''(\lambda) \tag{17.2}
\end{aligned}$$

$$\begin{aligned}
&\sum_{crs \notin G} \Pr[\text{Setup}_X = crs] \cdot \sum_{a \in \{0,1\}^{k(n)}} \left| \Pr[Y'_{crs} = a] - \Pr[Y''_{crs} = a] \right| \\
&\leq \sum_{crs \notin G} \Pr[\text{Setup}_X = crs] \cdot 2 \\
&\leq 2 \cdot \nu''(\lambda) \tag{17.3}
\end{aligned}$$

Combining Equations (17.1), (17.2) and (17.3) we have

$$\Delta((Y', \text{Setup}_X), (Y'', \text{Setup}_X)) \leq 2 \cdot \nu''(\lambda) + 3 \cdot \nu'''(\lambda).$$

Furthermore, since for all $crs \in \text{supp}(\text{Setup}_X(1^\lambda))$, Y''_{crs} is the uniform distribution over a set of 2^n elements,

$$\begin{aligned}
\tilde{H}_\infty(Y'' \mid \text{Setup}_X(1^\lambda)) &= -\log \left(\mathbb{E}_{crs \leftarrow \text{Setup}_X(1^\lambda)} \max_a \Pr[Y''_{crs} = a] \right) \\
&= n.
\end{aligned}$$

Therefore, we have

$$H_\infty^{\delta+2\nu''+3\nu'''}(X) = \tilde{H}_\infty^{\delta+2\nu''+3\nu'''}(X \mid \text{Setup}_X(1^\lambda)) \geq n,$$

where δ, ν'', ν''' are negligible functions. \square

Theorem 17.2 together with Lemmas 15.1 and 17.3 yield the following corollary.

COROLLARY 17.3. *If one-way functions exist, $\text{cPREH}_{\equiv_s}^{\text{det}}$ is false.*

17.1.2 Computational Guarantees and Pseudentropy

We study the relation of the pseudorandom encoding hypothesis with deterministic encoding algorithm to HILL and Yao entropy. Interestingly, $\text{PREH}_{\approx_c}^{\text{det}}$ poses a lower bound on the HILL entropy of a source and an upper bound on its Yao entropy.

HILL ENTROPY. The algorithms (E_X, D_X) implied by $\text{PREH}_{\approx_c}^{\text{det}}$ compress a source X to its HILL entropy.

THEOREM 17.3. *If (weak) $\text{PREH}_{\approx_c}^{\text{det}}$ is true for X , i. e., there exist deterministic algorithms (E_X, D_X) with E_X having output length n satisfying correctness and pseudorandomness. Then, $H^{\text{HILL}}(X) \geq n$.*

Proof. We employ a similar strategy as in the proof of Theorem 17.1. Let $Y' := D_X(U_n)$. As already seen in the proof of Theorem 16.1, Y' and X are

computationally indistinguishable due to correctness and pseudorandomness.

Due to correctness, $\Pr[x \leftarrow X: D_X(E_X(x)) = x] \geq 1 - \nu(\lambda)$ for some negligible function ν . Since E_X is required to be deterministic, we get $\Pr[x \leftarrow X: E_X(D_X(E_X(x))) = E_X(x)] \geq 1 - \nu(\lambda)$. Applying pseudorandomness, we have

$$\Pr[u \leftarrow U_n: E_X(D_X(u)) = u] \geq 1 - \nu'(\lambda)$$

for some negligible function ν' , where the probability is only over the choice of u . Let $V_0 \subset \{0, 1\}^n$ denote the set of all u such that $E_X(D_X(u)) = u$ and let $V_1 := \{0, 1\}^n \setminus V_0$. We have that, $|V_1|/|\{0, 1\}^n| \leq \nu'(\lambda)$.

Let $\overline{V_1}$ be some arbitrary subset of $\{0, 1\}^{k(\lambda)} \setminus D_X(V_0)$ such that $|\overline{V_1}| = |V_1|$ (for some polynomial k). Let Y'' be the uniform distribution over $D_X(V_0) \cup \overline{V_1}$. The min-entropy of Y'' equals n . By the same argument as in the proof of Theorem 17.1, the statistical distance between Y' and Y'' is negligible.

Therefore, the distributions X and Y'' are computationally indistinguishable and, hence, $H^{\text{HILL}}(X) \geq n$. \square

This result can be generalized to conditional HILL entropy using general (as opposed to weak) $\text{PREH}_{\approx_c}^{\text{det}}$.

THEOREM 17.4. *Let (X, Z) be a joint distribution. More precisely, let Z be a distribution over words of length λ . For $z \in \text{supp}(Z)$, let X_z denote the conditional distribution when $Z = z$. If (general) $\text{PREH}_{\approx_c}^{\text{det}}$ is true for X , i. e., there exist two deterministic polynomial time algorithms (E_X, D_X) with E_X having output length n satisfying correctness and pseudorandomness. Then, $H^{\text{HILL}}(X | Z) \geq n$.*

Proof. For each $z \in \text{supp}(Z)$, consider the distribution $Y'_z := D_X(z, U_n)$. Due to Theorem 16.1, for all adversarially chosen z , the distributions $\{(Y'_z, z)\}$ and $\{(X_z, z)\}$ are computationally indistinguishable due to correctness and pseudorandomness.

Due to correctness, for all adversarially chosen z , the probability $\Pr[x \leftarrow X_z: D_X(z, E_X(z, x)) = x] \geq 1 - \nu(\lambda)$ for some negligible function ν . Since E_X is required to be deterministic, we have that for all adversarially chosen z , $\Pr[x \leftarrow X_z: E_X(z, D_X(z, E_X(z, x))) = E_X(z, x)] \geq 1 - \nu(\lambda)$. Applying pseudorandomness, we get that there exists a negligible function ν' such that for all adversarially chosen z ,

$$\Pr[u \leftarrow U_n: E_X(z, D_X(z, u)) = u] \geq 1 - \nu'(\lambda),$$

where the probability is only over the choice of u . Let $V_{0,z} \subseteq \{0, 1\}^n$ denote the set of all $u \in \{0, 1\}^n$ such that $E_X(z, D_X(z, u)) = u$ holds and let $V_{1,z} := \{0, 1\}^n \setminus V_{0,z}$. We have that $|V_{1,z}|/|\{0, 1\}^n| \leq \nu'(\lambda)$.

Let $\overline{V_{1,z}}$ be some arbitrary subset of $\{0, 1\}^{k(\lambda)} \setminus D_X(z, V_{0,z})$ such that $|\overline{V_{1,z}}| = |V_{1,z}|$ (for some polynomial k). Let Y''_z be the uniform distribution over $D_X(z, V_{0,z}) \cup \overline{V_{1,z}}$. Clearly,

$$\begin{aligned} \tilde{H}_\infty(Y'' | Z) &= -\log \left(\mathbb{E}_{z \leftarrow Z} \left[\max_{y \in \text{supp}(Y''_z)} \Pr[Y''_z = y] \right] \right) \\ &= n. \end{aligned}$$

By a similar argument as in the proof of Theorem 17.1, for all adversarially chosen z , the statistical distance between Y'_z and Y''_z is negligible:

$$\begin{aligned} \Delta(Y'_z, Y''_z) &= \underbrace{\sum_{a \in D_X(z, V_{z,0})} |\Pr[Y'_z = a] - \Pr[Y''_z = a]|}_{\leq 1 - |V_{0,z}| \cdot 2^{-n} \leq 1 - (1 - \nu'(\lambda))} \\ &+ \underbrace{\sum_{a \in \overline{V}_{1,z}} |\Pr[Y' = a] - \Pr[Y'' = a]|}_{\leq \sum_{a \in \overline{V}_{1,z}} (\Pr[Y' = a] + \Pr[Y'' = a]) \leq 2 \cdot \nu'(\lambda)} \\ &\leq 3 \cdot \nu'(\lambda) \end{aligned}$$

Hence, the (joint) distributions (X, Z) and (Y'', Z) are computationally indistinguishable. Therefore, $H^{\text{HILL}}(X | Z) \geq n$. \square

We recall that in the proof of Theorem 17.2, we used the observation that independent random variables A and B satisfy $H_\infty(A) = H_\infty(A | B)$. However, this does not necessarily extend to the computational case, meaning that in our case $H^{\text{HILL}}(X | \text{Setup}_X(1^\lambda), Z)$ and $H^{\text{HILL}}(X | Z)$ are not equal even though the random variables X and Setup_X are independent. Hence, in order to extend Theorem 17.4 to $\text{cPREH}_{\approx_c}^{\text{det}}$, we need that there exists a “good” CRS crs such that $Y'_{crs,z}$ is almost injective on $\{0, 1\}^n$ and such that $\{(crs, y)\}$ and $\{(crs, D_X(crs, z, u))\}$ are computationally indistinguishable for this fixed crs . However, since the adversary is quantified after the crs , a non-uniform adversary could know the randomness which was used to generate crs compromising all security guarantees. Therefore, we can only hope to obtain an upper bound on n depending on the conditional HILL entropy $H^{\text{HILL}}(X | \text{Setup}_X(1^\lambda), Z)$.

THEOREM 17.5. *Let (X, Z) be a joint distribution. More precisely, let Z be a distribution over words of length λ . For $z \in \text{supp}(Z)$, let X_z denote the conditional distribution when $Z = z$. If (general) $\text{cPREH}_{\approx_c}^{\text{det}}$ is true for X , i. e., there exists a PPT algorithm Setup_X and two deterministic polynomial time algorithms (E_X, D_X) with E_X having output length n satisfying correctness and pseudorandomness. Then, $H^{\text{HILL}}(X | \text{Setup}_X(1^\lambda), Z) \geq n$.*

Proof. The proof is similar to the proof of Theorem 17.4. We consider the distribution $Y'_{crs,z} := D_X(crs, z, U_n)$.

Due to Theorem 16.1, we have that for all adversarially chosen z , the distributions $\{crs \leftarrow \text{Setup}_X(1^\lambda), y \leftarrow X_z : (crs, z, y)\}$ and $\{crs \leftarrow \text{Setup}_X(1^\lambda), u \leftarrow U_n : (crs, z, D_X(crs, z, u))\}$ are computationally indistinguishable.

Due to correctness and pseudorandomness we have that there exists a negligible function ν' , such that for all adversarially chosen z ,

$$\begin{aligned} &\Pr [crs \leftarrow \text{Setup}_X(1^\lambda), u \leftarrow U_n : E_X(crs, z, D_X(crs, z, u)) = u] \\ &\geq 1 - \nu'(\lambda), \end{aligned}$$

where the probability is over the choice of crs and u . The Splitting lemma (Lemma 17.4) implies that for all z ,⁶² there exists a set $G_z \subseteq \text{supp}(\text{Setup}_X(1^\lambda))$ such that $\Pr[crs \leftarrow \text{Setup}_X(1^\lambda) : crs \in G_z] \geq 1 - \nu''_z(\lambda)$ and for all $crs \in G_z$, $\Pr[u \leftarrow U_n : E_X(crs, z, D_X(crs, z, u)) = u] \geq 1 - \nu'''_z(\lambda)$ for some negligible functions ν''_z, ν'''_z . Hence, conditioned on $crs \in G_z$, $D_X(crs, z, \cdot)$ operates almost injectively on the set $\{0, 1\}^n$.

⁶² Note that here we use that the adversary is non-uniform.

Let $V_{0,z}^{crs} \subseteq \{0,1\}^n$ denote the set of all $u \in \{0,1\}^n$ such that $E_X(crs, z, D_X(crs, z, u)) = u$ holds and let $V_{1,z}^{crs} := \{0,1\}^n \setminus V_{0,z}^{crs}$. For $crs \in G_z$, we have $|V_{1,z}^{crs}|/|\{0,1\}^n| \leq \nu_z'''(\lambda)$.

Let $\overline{V}_{1,z}^{crs}$ be some arbitrary subset of $\{0,1\}^{k(\lambda)} \setminus D_X(crs, z, V_{0,z}^{crs})$ such that $|\overline{V}_{1,z}^{crs}| = |V_{1,z}^{crs}|$ (for some polynomial k). Let $Y''_{crs,z}$ be the uniform distribution over $D_X(crs, z, V_{0,z}^{crs}) \cup \overline{V}_{1,z}^{crs}$. Then,

$$\begin{aligned} & \tilde{H}_\infty(Y'' \mid \text{Setup}_X(1^\lambda), Z) \\ &= -\log \left(\mathbb{E}_{\substack{crs \leftarrow \text{Setup}_X(1^\lambda) \\ z \leftarrow Z}} \left[\max_{y \in \text{supp}(Y''_{crs,z})} \Pr[Y''_{crs,z} = y] \right] \right) \\ &= n. \end{aligned}$$

Furthermore, a similar computation as in the proof of Theorem 17.2, for all adversarially chosen z , the statistical distance between $(Y'_z, \text{Setup}_X(1^\lambda))$ and $(Y''_z, \text{Setup}_X(1^\lambda))$ is negligible. Hence, the (joint) distributions $(X, \text{Setup}_X(1^\lambda), Z)$ and $(Y'', \text{Setup}_X(1^\lambda), Z)$ are computationally indistinguishable. Therefore, $H^{\text{HILL}}(X \mid \text{Setup}_X(1^\lambda), Z) \geq n$. \square

YAO ENTROPY. On the other hand, the existence of algorithms (E_X, D_X) as implied by $\text{PREH}_{\approx_c}^{\text{det}}$ give an upper bound for the Yao entropy of a source.

LEMMA 17.5. *If (weak) $\text{PREH}_{\approx_c}^{\text{det}}$ is true for X , i. e., there exist deterministic algorithms (E_X, D_X) with E_X having output length n satisfying correctness and pseudorandomness. Then, $H^{\text{Yao}}(X) < n + \delta(\lambda)$ for some negligible δ .*

Proof. Due to correctness, there exists a pair of efficient algorithms (E_X, D_X) (with E_X having output length n) such that

$$\Pr_{x \leftarrow X} [D_X(E_X(x)) = x] \geq 1 - \nu(\lambda)$$

for some negligible function ν . Hence, by Definition 17.4, $H^{\text{Yao}}(X) < k$, for all k satisfying

$$\begin{aligned} & 1 - \nu(\lambda) \geq 2^{n-k} + \epsilon(\lambda) \\ \iff & 2^k \cdot (1 - \nu(\lambda) - \epsilon(\lambda)) \geq 2^n \\ \iff & k \geq n - \log(1 - \nu(\lambda) - \epsilon(\lambda)). \end{aligned}$$

\square

This result can be generalized to conditional Yao entropy as follows.

LEMMA 17.6. *Let (X, Z) be a joint distribution. More precisely, let Z be a distribution over words of length λ . For $z \in \text{supp}(Z)$, let X_z denote the conditional distribution when $Z = z$. If (general) $\text{PREH}_{\approx_c}^{\text{det}}$ is true for X , i. e., there exist two deterministic polynomial time algorithms (E_X, D_X) with E_X having output length n satisfying correctness and pseudorandomness. Then, $H^{\text{Yao}}(X \mid Z) < n + \delta(\lambda)$ for some negligible δ .*

Proof. Due to correctness, for all adversarially chosen z ,

$$\Pr_{x \leftarrow X_z} [D_X(z, E_X(z, x)) = x] \geq 1 - \nu(\lambda)$$

for some negligible function ν . Hence, by Definition 17.6, $H^{\text{Yao}}(X | Z) < k$, for all k satisfying

$$k \geq n - \log(1 - \nu(\lambda) - \epsilon(\lambda)).$$

□

This result can further be extended to the case of a common setup.

LEMMA 17.7. *Let (X, Z) be a joint distribution. More precisely, let Z be a distribution over words of length λ . For $z \in \text{supp}(Z)$, let X_z denote the conditional distribution when $Z = z$. If (general) $\text{cPREH}_{\approx_c}^{\text{det}}$ is true for X , i. e., there exists a setup algorithm Setup and two deterministic polynomial time algorithms (E_X, D_X) with E_X having output length n satisfying correctness and pseudorandomness. Then, $H^{\text{Yao}}(X | Z) < n + \delta(\lambda)$ for some negligible δ .*

Proof. The Splitting lemma (Lemma 17.4) implies that for all z ,⁶³ there is a set $G_z \subseteq \text{supp}(\text{Setup}_X(1^\lambda))$ such that $\Pr[\text{crs} \leftarrow \text{Setup}_X(1^\lambda) : \text{crs} \in G_z] \geq 1 - \nu_z''(\lambda)$ and for all $\text{crs} \in G_z$,

$$\Pr[\mathbf{u} \leftarrow \mathbf{U}_n : E_X(\text{crs}, z, D_X(\text{crs}, z, \mathbf{u})) = \mathbf{u}] \geq 1 - \nu_z'''(\lambda)$$

for some negligible functions ν_z'', ν_z''' .

We exploit the non-uniformity of the definition (Definitions 17.4 and 17.6). In particular, we define the compression and decompression circuits as $E'_X(y, z; \text{aux}_\lambda) := E_X(\text{aux}_\lambda, z, y)$ and $D'_X(\mathbf{u}, z; \text{aux}_\lambda) := D_X(\text{aux}_\lambda, z, \mathbf{u})$, where aux_λ denotes the non-uniform auxiliary input providing a “good” CRS. □

Due to [HLR07], assuming suitable non-interactive zero-knowledge proof systems and pseudorandom generators, there exists a joint distribution (X, Z) with high conditional Yao entropy but low conditional HILL entropy. The sampler S from [HLR07] takes as input a NIZK common random string σ , samples a seed from \mathbf{U}_λ and outputs $y_1 := \text{PRG}(s)$ together with a proof y_2 that y_1 is in the image of PRG. Due to Lemma 15.2, $S \in \mathcal{S}^{\approx_c}$. This yields the following corollary.

COROLLARY 17.4. *If there exist a pseudorandom generator and a (single-theorem) NIZK proof system such that (i) for an overwhelming fraction of common random strings, the number of accepting proofs for each statement is limited, and (ii) the simulated random string is independent of the statement as in [HLR07]. Then $\text{PREH}_{\approx_c}^{\text{det}}$ is false.*

A Blum integer is a natural number $N = p \cdot q$ such that p, q are primes with $p \equiv q \equiv 3 \pmod{4}$. The quadratic residuosity assumption (QRA) states that for a randomly chosen Blum integer $N = p \cdot q$, the distributions $\{y \leftarrow \mathbb{Z}_N^\times \text{ s.t. } \left(\frac{y}{p}\right) = \left(\frac{y}{q}\right) = 1 : (N, y)\}$ and $\{y \leftarrow \mathbb{Z}_N^\times \text{ s.t. } \left(\frac{y}{N}\right) = 1 : (N, y)\}$ are computationally indistinguishable.

Due to [LMso5], a (single-theorem) NIZK proof system which is suitable for Corollary 17.4 can be instantiated from the quadratic residuosity assumption (QRA).

COROLLARY 17.5. *If the quadratic residuosity assumption is true, then $\text{PREH}_{\approx_c}^{\text{det}}$ is false.*

⁶³ Note that here we use that the adversary is non-uniform.

ON REFUTING $\text{cPREH}_{\approx_c}^{\text{det}}$. Theorem 17.5 only yields an upper bound on the encoding length n depending on $H^{\text{HILL}}(X \mid \text{Setup}_X(1^\lambda), Z)$. Since the Yao entropy of any source $H^{\text{Yao}}(X \mid \text{Setup}_X(1^\lambda), Z)$ is upper bounded by $n + \text{negl}(\lambda)$, we cannot apply the result from [HLR07] to refute $\text{cPREH}_{\approx_c}^{\text{det}}$.

ON REFUTING WEAK $\text{PREH}_{\approx_c}^{\text{det}}$. It is currently not known if plain HILL and plain Yao entropy can be separated as in [HLR07]. Such a separation would refute weak $\text{PREH}_{\approx_c}^{\text{det}}$.

17.2 RANDOMIZED ENCODING ALGORITHM

In the following we prove positive and negative results on the validity of the pseudorandom encoding hypothesis with a randomized encoding algorithm. In particular, in Section 17.2.2 we refute $\text{PREH}_{\approx_s}^{\text{rand}}$ based on subexponential [LWE](#), in Section 17.2.3 we refute $\text{cPREH}_{\approx_c}^{\text{rand}}$ based on extractable one-way functions with *unbounded* auxiliary input. On the positive side, in Section 17.3, we give a construction of perfectly correct $\text{cPREH}_{\approx_c}^{\text{rand}}$ with and without universal setup based on indistinguishability obfuscation and one-way functions following [SW14; DKR15] together with Theorem 16.1. In Section 17.4, we bootstrap $\text{cPREH}_{\approx_c}^{\text{rand}}$ with a common random string from the construction in Chapter 6 additionally assuming *weak* $\text{cPREH}_{\approx_c}^{\text{rand}}$ with a common random string. Since $\text{cPREH}_{\approx_c}^{\text{rand}}$ with a common random string in conjunction with [NIZK](#) proof systems contradicts [EOWFs](#) with common but benign auxiliary information, this refutes even *weak* $\text{cPREH}_{\approx_c}^{\text{rand}}$ with common random string.

17.2.1 (Generalized) Extractable One-way Functions

In this section, we define generalized extractable one-way functions (with respect to common auxiliary input) as in [BCPR16].

DEFINITION 17.9 (Generalized extractable one-way function family ensembles with common auxiliary information, [BCPR16]). A *generalized one-way extractable function (GEOWF) family ensemble with common auxiliary information*, with respect to a relation $\mathcal{R}_{\mathcal{F}}^{\mathcal{F}}$ over triplets $(f, y, x) \in \text{Gen}(1^\lambda) \times \{0, 1\}^{n(\lambda)} \times \{0, 1\}^{m(\lambda)}$, is a function family ensemble $\mathcal{F} = (\text{Gen}, \text{Eval})$ (cf. Definition 2.8) if the following properties are satisfied.

$\mathcal{R}_{\mathcal{F}}^{\mathcal{F}}$ -HARDNESS. For every [PPT](#) adversary \mathcal{A} , for every polynomial b and every $z \in \{0, 1\}^{b(\lambda)}$,

$$\Pr \left[\text{Exp}_{\mathcal{F}, \mathcal{A}, z}^{\text{g-hard-aux}}(\lambda) = 1 \right]$$

is negligible, where $\text{Exp}_{\mathcal{F}, \mathcal{A}, z}^{\text{g-hard-aux}}(\lambda)$ is defined in Figure 17.1a.

$\mathcal{R}_{\mathcal{F}}^{\mathcal{F}}$ -EXTRACTABILITY. For every [PPT](#) adversary \mathcal{X} , there exists a [PPT](#) algorithm $\mathcal{E}_{\mathcal{X}}$ such that for every polynomial b and every $z \in \{0, 1\}^{b(\lambda)}$,

$$\Pr \left[\text{Exp}_{\mathcal{F}, \mathcal{X}, \mathcal{E}_{\mathcal{X}}, z}^{\text{g-ext-aux}}(\lambda) = 1 \right]$$

is overwhelming, where $\text{Exp}_{\mathcal{F}, \mathcal{X}, \mathcal{E}_{\mathcal{X}}, z}^{\text{g-ext-aux}}(\lambda)$ is defined in Figure 17.1b.

We call \mathcal{F}

- *publicly verifiable* if there exists a **PPT** algorithm T such that

$$T(f, f(x), x') = 1 \Leftrightarrow \mathcal{R}_f^{\mathcal{F}}(f(x), x') = 1.$$

- *privately verifiable* if there exists a **PPT** algorithm T such that

$$T(f, x, x') = 1 \Leftrightarrow \mathcal{R}_f^{\mathcal{F}}(f(x), x') = 1.$$

$\text{Exp}_{\mathcal{F}, \mathcal{A}, z}^{\text{g-hard-aux}}(\lambda)$	$\text{Exp}_{\mathcal{F}, \mathcal{X}, \mathcal{E}_{\mathcal{X}}, z}^{\text{g-ext-aux}}(\lambda)$
$f \leftarrow \text{Gen}(1^\lambda)$ $x \leftarrow \{0, 1\}^{n(\lambda)}$ $y := f(x)$ $x' \leftarrow \mathcal{A}(f, y, z)$ return $\mathcal{R}_f^{\mathcal{F}}(y, x')$	$f \leftarrow \text{Gen}(1^\lambda)$ $r_{\mathcal{X}} \leftarrow \{0, 1\}^{p(\lambda)}$ $y \leftarrow \mathcal{X}(f, z; r_{\mathcal{X}})$ $x \leftarrow \mathcal{E}_{\mathcal{X}}(f, z, r_{\mathcal{X}})$ return $(\mathcal{R}_f^{\mathcal{F}}(y, x)) \vee (\forall x': f(x') \neq y)$

(a) The $\mathcal{R}^{\mathcal{F}}$ -hardness game.

(b) The $\mathcal{R}^{\mathcal{F}}$ -extraction game.

Figure 17.1: $\mathcal{R}^{\mathcal{F}}$ -hardness and $\mathcal{R}^{\mathcal{F}}$ -extraction game for **GEOWFs** (with common auxiliary input).

DEFINITION 17.10 (Generalized extractable one-way function family ensembles with b -bounded common auxiliary information, [BCPR16]). Like Definition 17.9 but with a fixed polynomial b determining the length of the common auxiliary information.

Due to [BCPR16], privately verifiable generalized extractable one-way functions with bounded auxiliary input can be instantiated based on falsifiable assumptions.

THEOREM 17.6 ([BCPR16]). *Assuming the learning with errors problem is subexponentially hard, then there exists a $(b(\lambda) - \omega(1))$ -bounded privately verifiable **GEOWF** family ensemble.*

17.2.2 Information-theoretic Guarantees and Privately Verifiable GEOWFs

Assuming a pseudorandom encoding scheme providing information-theoretic guarantees, we can consider unbounded adversaries. In this setting, private verifiability is not a restriction.

THEOREM 17.7. *If privately verifiable generalized extractable one-way functions without auxiliary information exist, then $\text{PREH}_{\equiv_s}^{\text{rand}}$ is false.*

Proof. The proof strategy follows the ideas of [IKOS10]. Let \mathcal{F} be a privately verifiable **GEOWF** (without auxiliary input) with respect to relation $\mathcal{R}^{\mathcal{F}}$.

$\text{PREH}_{\equiv_s}^{\text{rand}}$ implies that for the algorithm S (given in Figure 17.2a) there exist an alternative sampler \bar{S} and a corresponding inverse sampler \bar{S}^{-1} satisfying closeness and invertibility of Definition 16.1 against unbounded adversaries. Since \mathcal{F} is a **GEOWF**, for the algorithm \bar{S} , there exists an extractor $\mathcal{E}_{\bar{S}}$ satisfying extractability from Definition 17.9 (without auxiliary information).



(a) Sampler mapping a random pre-image with f . (b) An adversary breaking one-wayness.

Figure 17.2: Description of the sampler S and of the adversary \mathcal{A} breaking one-wayness of the privately verifiable **GEOWF**.

We prove that for \mathcal{A} given in Figure 17.2b,

$$\Pr \left[\text{Exp}_{\mathcal{F}, \mathcal{A}}^{\text{g-hard}}(\lambda) = 1 \right] \text{ is overwhelming.}$$

Let \tilde{T} be an *unbounded* algorithm that given (f, y, x) , computes the relation $\mathcal{R}_f^{\mathcal{F}}(y, x)$.⁶⁴ We proceed over a series of hybrids, see Figure 17.3.

⁶⁴ Since \mathcal{F} is privately verifiable, an efficient algorithm computing the relation $\mathcal{R}_f^{\mathcal{F}}$ additionally requires the pre-image of y as an input. For our purpose, it suffices to consider an inefficient testing algorithm \tilde{T} .

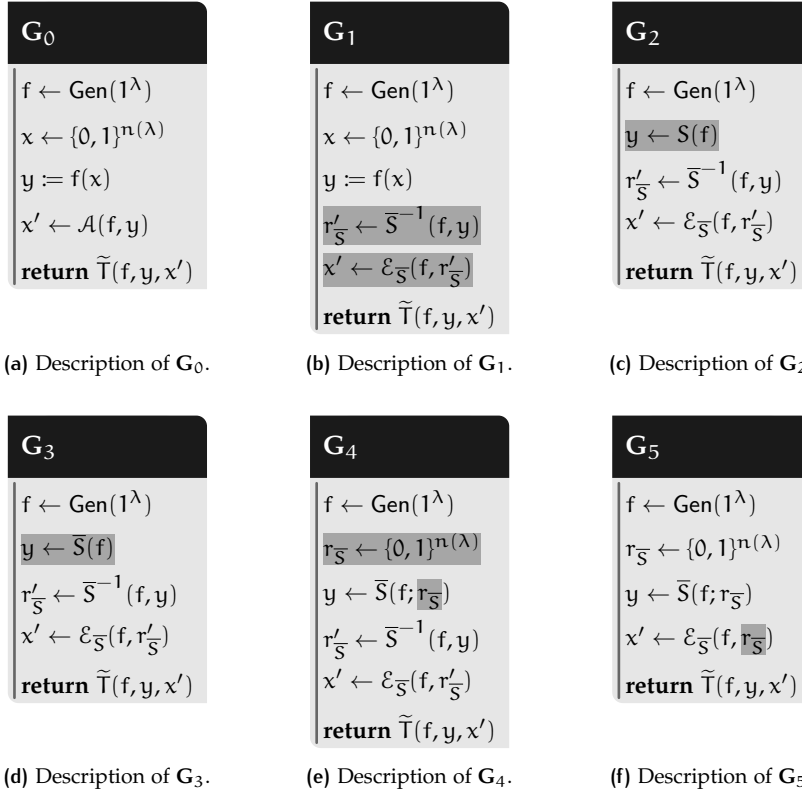


Figure 17.3: Hybrids used in the proof of Theorem 17.7.

GAME G₀. G_0 corresponds to $\text{Exp}_{\mathcal{F}, \mathcal{A}}^{\text{g-hard}}(\lambda)$ using the inefficient algorithm \tilde{T} to test whether $(f(x), x') \in \mathcal{R}_f^{\mathcal{F}}$.

GAME G₁. G_1 is identical to G_0 except that the implementation of \mathcal{A} is explicit in G_1 . Hence, $\Pr[\text{out}_0 = 1] = \Pr[\text{out}_1 = 1]$.

GAME G₂. G_2 is identical to G_1 except that $y := f(x)$ for $x \leftarrow \{0, 1\}^{n(\lambda)}$ is replaced with $y \leftarrow S(f)$ which proceeds identically, see Figure 17.2a. Hence, $\Pr[\text{out}_1 = 1] = \Pr[\text{out}_2 = 1]$.

GAME G_3 . G_3 is identical to G_2 except that G_2 samples y using $S(f)$ whereas G_3 samples y using $\bar{S}(f)$.

CLAIM 17.1. There exists an unbounded adversary \mathcal{B}_3 such that $|\Pr[out_3 = 1] - \Pr[out_2 = 1]| \leq \text{Adv}_{(\bar{S}, \bar{S}^{-1}), \mathcal{B}_3, \bar{f}}^{\text{close}}(\lambda)$ (for some $\bar{f} \in \text{supp}(\text{Gen}(1^\lambda))$).

Proof. Construct an adversary \mathcal{B}_3 breaking statistical closeness. \mathcal{B}_3 receives as input a key \bar{f} and some y that has either been sampled via $S(\bar{f})$ or via $\bar{S}(\bar{f})$. \mathcal{B}_3 computes r'_S and x' as in game G_2 , calls the (inefficient) testing algorithm \tilde{T} on input of (\bar{f}, y, x') and outputs the resulting output. Hence, we have

$$\begin{aligned} \Pr \left[\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{B}_3, \bar{f}}^{(0)\text{-close}}(\lambda) = 1 \right] &= \Pr[out_2 = 1 \mid f = \bar{f}] \text{ and} \\ \Pr \left[\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{B}_3, \bar{f}}^{(1)\text{-close}}(\lambda) = 1 \right] &= \Pr[out_3 = 1 \mid f = \bar{f}], \end{aligned}$$

and, since $\text{Adv}_{(\bar{S}, \bar{S}^{-1}), \mathcal{B}_3, \bar{f}}^{\text{close}}(\lambda)$ is negligible for all $\bar{f} \in \text{supp}(\text{Gen}(1^\lambda))$,

$$\begin{aligned} &|\Pr[out_3 = 1] - \Pr[out_2 = 1]| \\ &= \sum_{\bar{f} \in \text{supp}(\text{Gen}(1^\lambda))} \Pr_{f \leftarrow \text{Gen}(1^\lambda)}[f = \bar{f}] \cdot \text{Adv}_{(\bar{S}, \bar{S}^{-1}), \mathcal{B}_3, \bar{f}}^{\text{close}}(\lambda), \end{aligned}$$

is negligible. \square

GAME G_4 . G_4 is identical to G_3 except that the random coins $r_{\bar{S}}$ used for \bar{S} are made explicit. Hence, $\Pr[out_3 = 1] = \Pr[out_4 = 1]$.

GAME G_5 . G_5 is identical to G_4 except that $\mathcal{E}_{\bar{S}}$ is called using $r_{\bar{S}}$ instead of the inverse sampled random coins $r'_{\bar{S}}$.

CLAIM 17.2. There exists an unbounded adversary \mathcal{B}_5 such that $|\Pr[out_5 = 1] - \Pr[out_4 = 1]| \leq \text{Adv}_{(\bar{S}, \bar{S}^{-1}), \mathcal{B}_5, \bar{f}}^{\text{inv}}(\lambda)$ (for some $\bar{f} \in \text{supp}(\text{Gen}(1^\lambda))$).

Proof. Construct an unbounded adversary \mathcal{B}_5 breaking statistical invertibility. On input of (\bar{f}, r, y) , \mathcal{B}_5 calls $\mathcal{E}_{\bar{S}}$ on input of (\bar{f}, r) and obtains x' . Finally, \mathcal{B}_5 calls the inefficient testing algorithm \tilde{T} on input of (\bar{f}, y, x') and outputs the resulting output. If \mathcal{B}_5 plays the experiment $\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{B}_5, \bar{f}}^{(0)\text{-inv}}(\lambda)$, he perfectly simulates game G_5 and

$$\Pr \left[\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{B}_5, \bar{f}}^{(0)\text{-inv}}(\lambda) = 1 \right] = \Pr[out_5 = 1 \mid f = \bar{f}].$$

If, on the other hand, \mathcal{B}_5 plays the experiment $\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{B}_5, \bar{f}}^{(1)\text{-inv}}(\lambda)$, he perfectly simulates game G_4 and

$$\Pr \left[\text{Exp}_{(\bar{S}, \bar{S}^{-1}), \mathcal{B}_5, \bar{f}}^{(1)\text{-inv}}(\lambda) = 1 \right] = \Pr[out_4 = 1 \mid f = \bar{f}].$$

Thus, the claim follows. \square

Summing up, we have that $|\Pr[out_0 = 1] - \Pr[out_5 = 1]|$ is negligible.

LEMMA 17.8. $\Pr[out_5 = 1]$ is overwhelming.

Proof of Lemma 17.8. Since \mathcal{F} is a **GEOWF**, we have that

$$\Pr_{\mathbf{G}_5} [\mathcal{R}_f^{\mathcal{F}}(\mathbf{y}, \mathbf{x}') = 1 \vee \mathbf{y} \notin \text{image}(f)]$$

is overwhelming. Using a union bound, we get

$$\begin{aligned} \Pr_{\mathbf{G}_5} [\mathcal{R}_f^{\mathcal{F}}(\mathbf{y}, \mathbf{x}') = 1 \vee \mathbf{y} \notin \text{image}(f)] \\ \leq \Pr_{\mathbf{G}_5} [\mathcal{R}_f^{\mathcal{F}}(\mathbf{y}, \mathbf{x}') = 1] + \Pr_{\mathbf{G}_5} [\mathbf{y} \notin \text{image}(f)]. \end{aligned}$$

In the following, we prove that $\Pr_{\mathbf{G}_5} [\mathbf{y} \notin \text{image}(f)]$ is negligible. Consequently, $\Pr_{\mathbf{G}_5} [\mathcal{R}_f^{\mathcal{F}}(\mathbf{y}, \mathbf{x}') = 1] = \Pr[\text{out}_5 = 1]$ is overwhelming.

By construction, S on input of f always produces values $\mathbf{y} \in \text{image}(f)$. If $\Pr_{\mathbf{G}_5} [\mathbf{y} \notin \text{image}(f)] = \Pr[f \leftarrow \text{Gen}(1^\lambda), \mathbf{y} \leftarrow \bar{S}(f): \mathbf{y} \notin \text{image}(f)]$ is non-negligible, then an unbounded adversary can distinguish between outputs of $S(f)$ and outputs of $\bar{S}(f)$ simply by testing all possible pre-images. Let $\bar{\mathcal{B}}$ be the adversary who outputs 1 if and only if $\mathbf{y} \notin \text{image}(f)$. Hence,

$$\Pr_{\mathbf{G}_5} [\mathbf{y} \notin \text{image}(f)] = \sum_{\bar{f} \in \text{supp}(\text{Gen}(1^\lambda))} \Pr[f = \bar{f}] \cdot \text{Adv}_{(\bar{S}, S^{-1}), \bar{\mathcal{B}}, \bar{f}}^{\text{close}}(\lambda).$$

Since for all $f \in \text{supp}(\text{Gen}(1^\lambda))$, $\text{Adv}_{(\bar{S}, S^{-1}), \bar{\mathcal{B}}, \bar{f}}^{\text{close}}(\lambda)$ is negligible, $\Pr_{\mathbf{G}_5} [\mathbf{y} \notin \text{image}(f)]$ is negligible. \square

This concludes the proof of Theorem 17.7. \square

From Theorem 17.7 together with Theorem 17.6, we obtain the following corollary.

COROLLARY 17.6. *Assuming the learning with errors problem is subexponentially hard, then $\text{PREH}_{\equiv_s}^{\text{rand}}$ is false.*

The above proof extends to $\text{cPREH}_{\equiv_s}^{\text{rand}}$ at the cost of assuming privately verifiable **GEOWFs** with unbounded auxiliary information. We omit the details here since no instantiation of this variant of **GEOWFs** is known and refer the reader to the proof of Theorem 17.10 for an example how to extend this strategy to the **CRS** case.

THEOREM 17.8. *If privately verifiable generalized extractable one-way functions with unbounded auxiliary information exist, then $\text{cPREH}_{\equiv_s}^{\text{rand}}$ is false.*

17.2.3 Computational Guarantees and EOWFs with Common Auxiliary Information

Assuming a pseudorandom encoding scheme providing computational guarantees, the above proof strategy must be adapted since the adversary actually uses the alternative sampler of a sampler which outputs an image in the range of the **EOWF**. Computational indistinguishability does not suffice to force the alternative sampler to output an image in the range of the **EOWF**. Hence, for the purpose to force the alternative sampler to output an image of a given extractable one-way function, we follow the lines of [IKOS10] and require the original sampler to additionally provide a non-interactive zero-knowledge proof certifying that the provided image is in the range of the given **EOWF**.

By this strategy [IKOS10] prove that $\text{PREH}_{\approx_c}^{\text{rand}}$ in conjunction with **NIZK** proof systems conflicts with **EOWFs** (without auxiliary input).

THEOREM 17.9 ([IKOS10]). *If extractable one-way functions without auxiliary information and NIZK proof systems for \mathcal{NP} exist, then $\text{PREH}_{\approx_c}^{\text{rand}}$ is false.*

This result can be extended to $\text{cPREH}_{\approx_c}^{\text{rand}}$ at the cost of assuming EOWFs with unbounded common auxiliary inputs.

THEOREM 17.10. *If there exist extractable one-way functions with unbounded common auxiliary information and NIZK proof systems for \mathcal{NP} , then $\text{cPREH}_{\approx_c}^{\text{rand}}$ is false.*

Proof. We adapt the proof from [IKOS10]. Let \mathcal{F} be an extractable one-way function with common auxiliary information. Let

$$L_\lambda := \left\{ (f, y) \in \text{supp}(\text{Gen}(1^\lambda)) \times \{0, 1\}^{m(\lambda)} \mid \begin{array}{l} \exists x \in \{0, 1\}^{n(\lambda)} \text{ s.t.} \\ f(x) = y \end{array} \right\}.$$

Let (Setup, Prove, Verify) be a NIZK proof system for L_λ such that Setup produces uniform random strings from $\{0, 1\}^{n_\pi(\lambda)}$.

$\text{cPREH}_{\approx_c}^{\text{rand}}$ implies that for the PPT algorithm S (see Figure 17.4), there exist a PPT algorithm Setup_S , an alternative sampler \bar{S} and a corresponding inverse sampler \bar{S}^{-1} satisfying closeness and invertibility as in Definition 16.3.

```

S(f, σ)
x ← {0, 1}n(λ)
y := f(x)
π ← Prove(σ, (f, y), x)
return (y, π)
    
```

(a) Sampler mapping a random pre-image with f and proving membership in the range of f .

```

X(f, z := crs; σ || rσ̄)
(y, π) ← S̄(crs, (f, σ); rσ̄)
return y
    
```

(b) An adversary producing an image of f .

```

A(1λ, f, y)
crs ← SetupS(1λ)
(σ, τπ) ← Sim0(1λ)
π ← Sim1(τπ, (f, y))
r'_{σ̄} ← S̄-1(crs, (f, σ), (y, π))
r'_{x̄} := σ || r'_{σ̄}
x' ← KX(f, crs, r'_{x̄})
return x'
    
```

(c) An adversary breaking one-wayness.

Figure 17.4: Description of the sampler S , of the adversary \mathcal{X} against extractability of the EOWF and of the adversary \mathcal{A} against one-wayness of the EOWF. Note that depending on the definition of one-wayness, the CRS crs could also be auxiliary input to \mathcal{A} .

The common auxiliary input is necessary to give the adversary \mathcal{X} and the extractor \mathcal{E}_X access to the same common reference string. Further, it is crucial

to assume *unbounded* auxiliary input since the adversary \mathcal{X} can be considered to be a universal adversary who simply executes the code which is contained in the auxiliary input. Since the size of the adversary cannot be bounded in advance, neither can the size of the auxiliary input.

Since \mathcal{F} is an extractable one-way function with common auxiliary information, we have that for the algorithm \mathcal{X} (see Figure 17.4b), there exists an extractor $\mathcal{E}_{\mathcal{X}}$ such that for every polynomial b and every $crs \in \{0, 1\}^{b(\lambda)}$ (hence, in particular, for every $z := crs$ produced by $\text{Setup}_S(1^\lambda)$),

$$\Pr \left[\text{Exp}_{\mathcal{F}, \mathcal{X}, \mathcal{E}_{\mathcal{X}}, z}^{\text{ext-aux}}(\lambda) = 1 \right]$$

is overwhelming.

We prove that adversary \mathcal{A} given in Figure 17.4c has an overwhelming probability to break the one-wayness of \mathcal{F} . We proceed over a sequence of hybrids, see Figures 17.5 and 17.6.

G₀

```

f ← Gen(1λ)
x ← {0, 1}n(λ)
y := f(x)
x' ← A(f, y)
return f(x') = y

```

(a) Description of G₀.

G₁

```

f ← Gen(1λ)
x ← {0, 1}n(λ)
y := f(x)
crs ← SetupS(1λ)
(σ, τπ) ← Sim0(1λ)
π ← Sim1(τπ, (f, y))
r'_{S̄} ← S̄-1(crs, (f, σ), (y, π))
r'_{X} := σ || r'_{S̄}
x' ← EX(f, crs, r'_{X})
return f(x') = y

```

(b) Description of G₁.

G₂

```

f ← Gen(1λ)
x ← {0, 1}n(λ)
y := f(x)
crs ← SetupS(1λ)
σ ← {0, 1}nπ(λ)
π ← Prove(crs, (f, y), x)
r'_{S̄} ← S̄-1(crs, (f, σ), (y, π))
r'_{X} := σ || r'_{S̄}
x' ← EX(f, crs, r'_{X})
return f(x') = y

```

(c) Description of G₂.

G₃

```

f ← Gen(1λ)
σ ← {0, 1}nπ(λ)
crs ← SetupS(1λ)
(y, π) ← S(f, σ)
r'_{S̄} ← S̄-1(crs, (f, σ), (y, π))
r'_{X} := σ || r'_{S̄}
x' ← EX(f, crs, r'_{X})
return f(x') = y

```

(d) Description of G₃.

Figure 17.5: Hybrids used in the proof of Theorem 17.10.

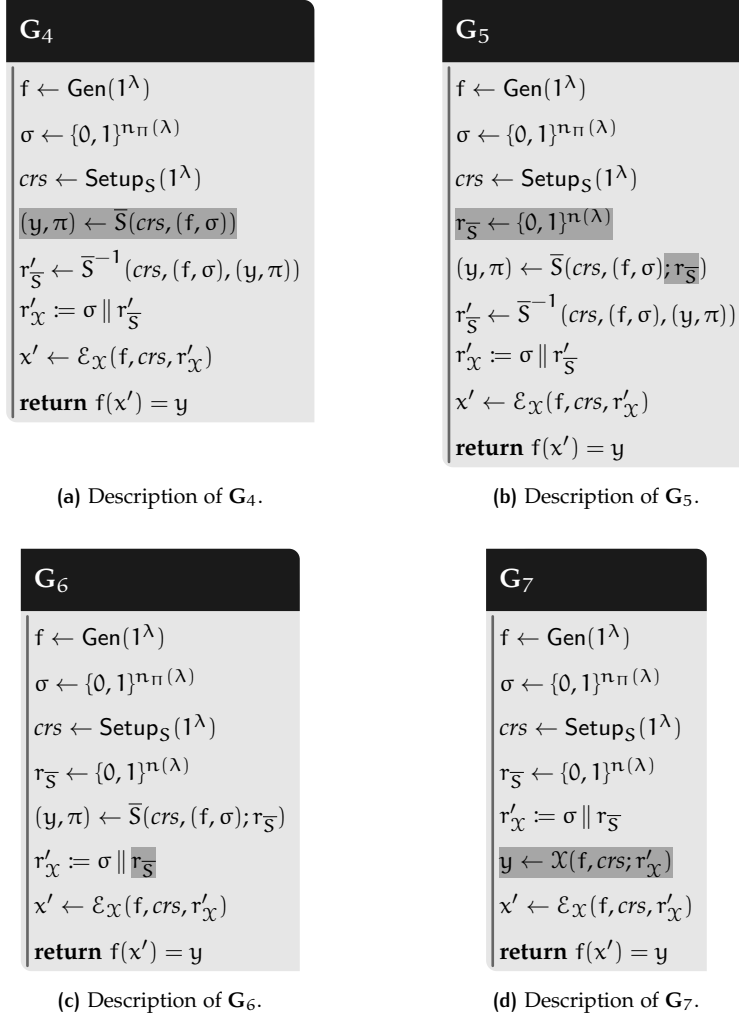


Figure 17.6: Hybrids used in the proof of Theorem 17.10.

GAME G₀. G₀ corresponds to $\text{Exp}_{\mathcal{F}, \mathcal{A}}^{\text{ow}}(\lambda)$.

GAME G₁. G₁ is identical to G₀ except that the implementation of \mathcal{A} is explicit in G₁. Hence, $\Pr[\text{out}_0 = 1] = \Pr[\text{out}_1 = 1]$.

GAME G₂. G₂ is identical to G₁ except that in G₁ σ and π are simulated using $(\text{Sim}_0, \text{Sim}_1)$, whereas in G₂, σ is sampled uniformly at random from $\{0, 1\}^{n_{\pi}(\lambda)}$ and π is produced via $\text{Prove}(\sigma, (f, y), x)$.

CLAIM 17.3. There exists a PPT adversary \mathcal{B}_2 , such that $|\Pr[\text{out}_2 = 1] - \Pr[\text{out}_1 = 1]| \leq \text{Adv}_{\text{NIZK}, \text{Sim}, \mathcal{B}_2}^{\text{zk}}(\lambda)$.

Proof. Construct adversary \mathcal{B}_2 against the zero-knowledge property of NIZK. Initially, \mathcal{B}_2 receives σ as input, produces f , x , y and crs as in G₁ and calls the prove oracle on input of the statement (f, y) together with a witness x . In return, \mathcal{B}_2 receives a proof π and proceeds as in G₁ to produce $r'_{S̄}$, r'_{X} and x' . Finally, \mathcal{B}_2 outputs 1 if $f(x') = y$ and 0 otherwise. If \mathcal{B}_2 receives a simulated CRS produced via $\text{Sim}_0(1^\lambda)$ and a simulated proof produced via $\text{Sim}_1(\sigma, \tau_{\pi}, (f, y))$, \mathcal{B}_2 perfectly simulates G₁. If, on the other hand, σ is uniformly random from $\{0, 1\}^{n_{\pi}(\lambda)}$ and the proof π

is produced by $\text{Prove}(\sigma, (f, y), x)$, then \mathcal{B}_2 perfectly simulates \mathbf{G}_2 . Hence, $|\Pr[out_2 = 1] - \Pr[out_1 = 1]| \leq \text{Adv}_{\text{NIZK, Sim}, \mathcal{B}_2}^{\text{zk}}(\lambda) = 1$. \square

GAME \mathbf{G}_3 . \mathbf{G}_3 and \mathbf{G}_2 are almost identical except for the difference that \mathbf{G}_3 does not produce y and π explicitly but uses $S(f, \sigma)$ to produce them. This difference is only conceptual and $\Pr[out_2 = 1] = \Pr[out_3 = 1]$.

GAME \mathbf{G}_4 . \mathbf{G}_4 is identical to \mathbf{G}_3 except that \mathbf{G}_3 samples (y, π) using $S(f)$ whereas \mathbf{G}_4 samples (y, π) using $\bar{S}(f)$.

CLAIM 17.4. There exists a **PPT** adversary \mathcal{B}_4 , such that $|\Pr[out_4 = 1] - \Pr[out_3 = 1]| \leq \text{Adv}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{B}_4}^{\text{crs-close}}(\lambda)$.

Proof. Construct a **PPT** adversary \mathcal{B}_4 against (static) closeness. Initially, \mathcal{B}_4 produces (f, σ) as in \mathbf{G}_3 and outputs them. In the second phase, \mathcal{B}_4 receives crs and (y, π) , where (y, π) has either been sampled using $S(f, \sigma)$ or using $\bar{S}(\text{crs}, (f, \sigma))$. Further, \mathcal{B}_4 proceeds as in \mathbf{G}_3 producing $r'_{\bar{S}}$, r'_x and x' . Finally, \mathcal{B}_4 outputs $f(x') = y$. Hence,

$$\begin{aligned} \Pr[out_3 = 1] &= \Pr \left[\text{Exp}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{B}_4}^{(0)\text{-crs-close}}(\lambda) = 1 \right] \text{ and} \\ \Pr[out_4 = 1] &= \Pr \left[\text{Exp}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{B}_4}^{(1)\text{-crs-close}}(\lambda) = 1 \right]. \end{aligned}$$

\square

GAME \mathbf{G}_5 . \mathbf{G}_5 is identical to \mathbf{G}_4 except that the random coins $r_{\bar{S}}$ used for \bar{S} are made explicit. This difference is only conceptual and $\Pr[out_4 = 1] = \Pr[out_5 = 1]$.

GAME \mathbf{G}_6 . \mathbf{G}_6 uses $r_{\bar{S}}$ for r'_x instead of the inverse sampled $r'_{\bar{S}}$. Apart from this difference, \mathbf{G}_6 is identical to \mathbf{G}_5 .

CLAIM 17.5. There exists a **PPT** adversary \mathcal{B}_6 , such that $|\Pr[out_6 = 1] - \Pr[out_5 = 1]| \leq \text{Adv}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{B}_6}^{\text{crs-inv}}(\lambda)$.

Proof. Construct a **PPT** adversary \mathcal{B}_6 on (static) invertibility. Initially, \mathcal{B}_6 (statically) produces (f, σ) as in \mathbf{G}_5 and outputs them to the experiment. In the second phase, \mathcal{B}_6 receives $(\text{crs}, r^*, (y, \pi))$, where (y, π) has been sampled using $\bar{S}(\text{crs}, (f, \sigma))$ and r^* either is the actual randomness used for that sampling process, or the inverse sampled randomness produced via $\bar{S}^{-1}(\text{crs}, (f, \sigma), (y, \pi))$. Afterwards, \mathcal{B}_6 proceeds as in \mathbf{G}_5 producing $r'_x := \sigma \parallel r^*$ and x' . Finally, \mathcal{B}_6 outputs $f(x') = y$. Hence,

$$\begin{aligned} \Pr[out_5 = 1] &= \Pr \left[\text{Exp}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{B}_6}^{(1)\text{-crs-inv}}(\lambda) = 1 \right] \text{ and} \\ \Pr[out_6 = 1] &= \Pr \left[\text{Exp}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{B}_6}^{(0)\text{-crs-inv}}(\lambda) = 1 \right]. \end{aligned}$$

\square

GAME \mathbf{G}_7 . \mathbf{G}_7 is identical to \mathbf{G}_6 except for the conceptual difference that \mathbf{G}_7 computes y using $\mathcal{X}(f, \text{crs}; r'_x)$ and does not produce the proof π . Since π is not used anymore, this difference is only conceptual and $\Pr[out_6 = 1] = \Pr[out_7 = 1]$.

Thus, we have that $|\Pr[out_0 = 1] - \Pr[out_7 = 1]|$ is negligible.

LEMMA 17.9. $\Pr[out_7 = 1]$ is overwhelming.

Proof of Lemma 17.9. Since \mathcal{F} is an extractable one-way function, we have that

$$\begin{aligned} \Pr_{\mathbf{G}_7}[f(x') = y \vee (f, y) \notin L_\lambda] \\ \leq \Pr_{\mathbf{G}_7}[f(x') = y] + \Pr_{\mathbf{G}_7}[(f, y) \notin L_\lambda] \end{aligned}$$

is overwhelming. In the following, we prove that $\Pr_{\mathbf{G}_7}[(f, y) \notin L_\lambda]$ is negligible. We proceed over a series of hybrids, see Figure 17.7.

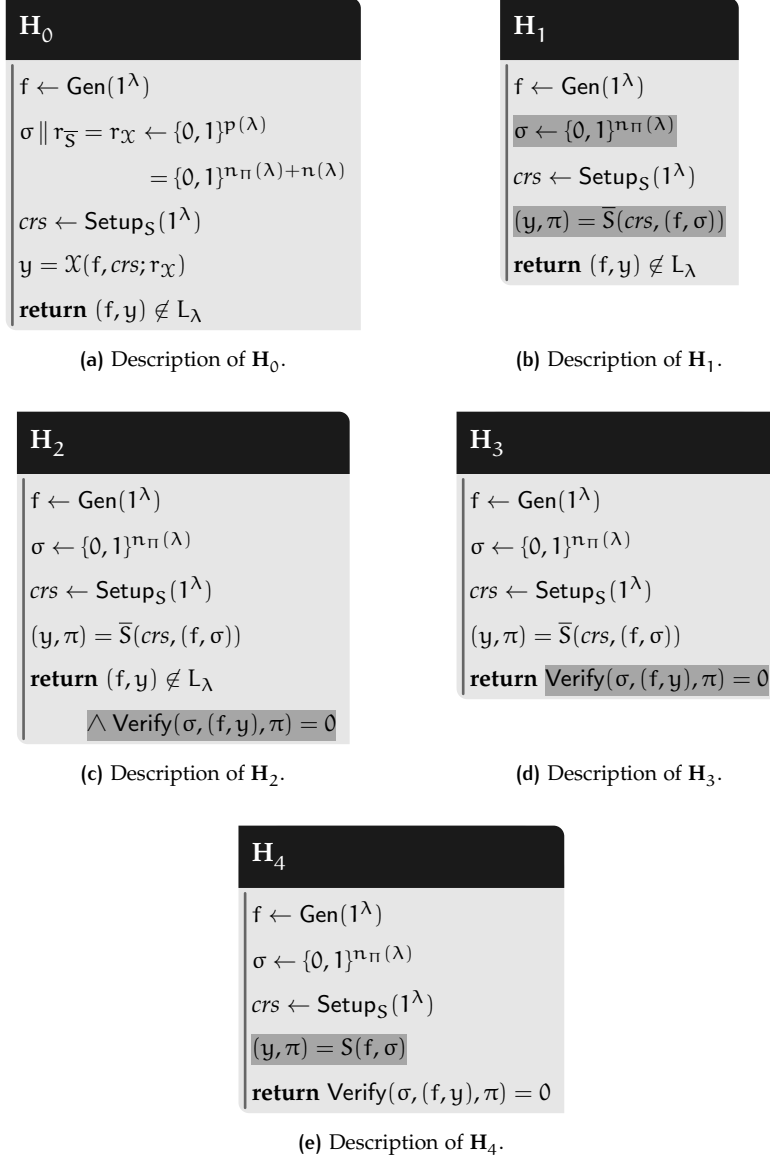


Figure 17.7: Hybrids used in the proof of Lemma 17.9.

GAME \mathbf{H}_0 . \mathbf{H}_0 is identical to \mathbf{G}_7 except that \mathbf{H}_0 returns $(f, y) \notin L_\lambda$. That is, $\Pr[out_{\mathbf{H}_0} = 1] = \Pr_{\mathbf{G}_7}[(f, y) \notin L_\lambda]$.

GAME \mathbf{H}_1 . \mathbf{H}_1 is identical to \mathbf{H}_0 except for the conceptual difference, that the algorithm \mathcal{X} is replaced with its implementation. Hence, $\Pr[out_{\mathbf{H}_0} = 1] = \Pr[out_{\mathbf{H}_1} = 1]$.

GAME H_2 . H_2 is identical to H_1 except that it returns 1 only if $(f, y) \notin L_\lambda$ and $\text{Verify}(\sigma, (f, y), \pi) = 0$.

CLAIM 17.6. There exists a **PPT** adversary \mathcal{B}_{H_2} such that $|\Pr[out_{H_2} = 1] - \Pr[out_{H_1} = 1]| \leq \text{Adv}_{\text{NIZK}, \mathcal{B}_{H_2}}^{\text{sound}}(\lambda)$.

Proof. The only possibility that the output of H_1 and H_2 differ is the case that $(f, y) \notin L_\lambda$ but $\text{Verify}(\sigma, (f, y), \pi) = 1$ (assuming, that the output of Verify is binary). Hence,

$$\begin{aligned} \left| \Pr[out_{H_2} = 1] - \Pr[out_{H_1} = 1] \right| &\leq \Pr_{H_1} \left[\begin{array}{c} (f, y) \notin L_\lambda \wedge \\ \text{Verify}(\sigma, (f, y), \pi) = 1 \end{array} \right] \\ &\leq \text{Adv}_{\text{NIZK}, \mathcal{B}_{H_2}}^{\text{sound}}(\lambda). \end{aligned}$$

The second inequality follows from the statistical soundness of NIZK. More precisely, construct an (unbounded) adversary \mathcal{B}_{H_2} on statistical soundness of the NIZK proof system. On input of σ , \mathcal{B}_{H_2} produces $f, crs, (y, \pi)$ as in H_1 and outputs the statement (f, y) and the proof π . \square

GAME H_3 . The distribution underlying H_2 and H_3 is identical. The only difference between these games is that the condition that H_3 outputs 1 is less restrictive than the one of H_2 , hence, $\Pr[out_{H_2} = 1] \leq \Pr[out_{H_3} = 1]$.

Note that in contrast to the preceding games, the simulation of game H_3 is efficient.

GAME H_4 . H_4 is identical to H_3 except that in H_3 , (y, π) is produced using the alternative sampler \bar{S} and in H_4 , (y, π) is produced using the original sampler S .

CLAIM 17.7. There exists a **PPT** adversary \mathcal{B}_{H_4} such that $|\Pr[out_{H_4} = 1] - \Pr[out_{H_3} = 1]| \leq \text{Adv}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{B}_{H_4}}^{\text{crs-close}}(\lambda)$.

Proof. Construct a **PPT** adversary \mathcal{B}_{H_4} against static closeness. Initially, \mathcal{B}_{H_4} produces (f, σ) as in H_3 and outputs it to the experiment. On input of $(crs, \bar{y} := (y, \pi))$, where \bar{y} is either sampled via $S(f, \sigma)$ or via $\bar{S}(crs, (f, \sigma))$, \mathcal{B}_{H_4} outputs $\text{Verify}(\sigma, (f, y), \pi) = 0$. We have that

$$\begin{aligned} \Pr[out_{H_3} = 1] &= \Pr \left[\text{Exp}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{B}_{H_4}}^{(1)\text{-crs-close}}(\lambda) = 1 \right] \text{ and} \\ \Pr[out_{H_4} = 1] &= \Pr \left[\text{Exp}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{B}_{H_4}}^{(0)\text{-crs-close}}(\lambda) = 1 \right]. \end{aligned}$$

\square

Finally, due to the definition of S and the completeness of the NIZK proof system NIZK,

$$\Pr[out_{H_4} = 1] = 0.$$

This proves that

$$\Pr[out_{H_0} = 1]$$

is negligible concluding the proof of Lemma 17.9. \square

This concludes the proof of Theorem 17.10. \square

By an easy modification of the above proof, we can refute $\text{cPREH}_{\approx_c}^{\text{rand}}$ with universal setup.

THEOREM 17.11. *If there exist extractable one-way functions with unbounded common auxiliary information and NIZK proof systems for \mathcal{NP} , then $\text{cPREH}_{\approx_c}^{\text{rand}}$ with universal setup is false.*

ON THE NEED FOR UNBOUNDED COMMON AUXILIARY INPUT. We stress that assuming *unbounded* common auxiliary input is necessary since the adversary \mathcal{X} can be considered the universal adversary. On a technical level, the common auxiliary input which leads to a contradiction is an output of the universal setup algorithm Setup on input of a sufficiently large bound B on the supported circuit size, or the “special” setup algorithm Setup_S . This bound B as well as the output size of Setup_S depend on the sampler size S and, hence, on the size of the adversary \mathcal{X} .

DISCUSSION. We stress that we are able to prove that EOWFs with unbounded common auxiliary input (in conjunction with NIZK proof systems) implies that $\text{cPREH}_{\approx_c}^{\text{rand}}$ is false. Furthermore, due to [DKR15] and Theorem 16.1 (or alternatively Theorems 17.13 and 17.14 in Section 17.3), $\text{cPREH}_{\approx_c}^{\text{rand}}$ can be instantiated from indistinguishability obfuscation and one-way functions. We restate a theorem from [BCPR16].

THEOREM 17.12 ([BCPR16]). *Assuming indistinguishability obfuscation for all circuits, neither EOWFs nor GEOWFs exist with respect to unbounded common auxiliary information.*

Thus, the above is not a contradiction because due to Theorem 17.12, indistinguishability obfuscation for all polynomial sized circuits does not exist assuming EOWFs with unbounded common auxiliary input.

COMMON BUT BENIGN AUXILIARY INFORMATION. The definition due to [BCPR16] of (G)EOWFs with common auxiliary input requires that extractability holds for all common auxiliary inputs and hence also for a *worst-case* choice of common auxiliary input. This requirement can be weakened such that the common auxiliary input is drawn from some specific distribution. This distribution is called *benign* if it is unlikely that a common auxiliary input sampled according to this distribution encodes a malicious obfuscation⁶⁵. In particular, the uniform distribution over $\{0, 1\}^{b(\lambda)}$ is conjectured to be benign. This notion of (G)EOWFs with common but benign auxiliary information does not contradict IO.

However, by an easy modification of the proof of Theorem 17.10, the existence of (G)EOWFs with common but benign auxiliary input (drawn uniformly at random from $\{0, 1\}^{b(\lambda)}$) contradicts $\text{cPREH}_{\approx_c}^{\text{rand}}$ with common random string.

COROLLARY 17.7. *If there exist extractable one-way functions with common but benign auxiliary information, particularly for auxiliary inputs drawn uniformly at random from $\{0, 1\}^{b(\lambda)}$, and NIZK proof systems for \mathcal{NP} , then $\text{cPREH}_{\approx_c}^{\text{rand}}$, where the setup algorithm produces uniform random strings, is false.*

⁶⁵ By *malicious obfuscation* we mean an obfuscated circuit which renders extraction from an adversary who simply executes the obfuscated circuit on the given key infeasible, see [BCPR16].

KEY-LESS EOWFS. Definitions 2.20, 2.22 and 17.10 can be defined for *key-less* function ensembles, where the key space $\text{supp}(\text{Gen}(1^\lambda))$ contains exactly one element (for every λ). This poses much stronger requirements on one-wayness and extractability since these properties are bound to be met for one fixed key as opposed to the keyed variants.

As observed in [IKOS10], Theorems 17.7 and 17.9 can be adapted to the *weak* variants of $\text{PREH}_{\approx_c}^{\text{rand}}$ and $\text{PREH}_{\equiv_s}^{\text{rand}}$, respectively.⁶⁶ However, key-less (G)EOWFs with unbounded auxiliary input are impossible since the adversary might get a random image of the (fixed) (G)EOWF as auxiliary input. An extractor given the output of the adversary together with his random tape and the same auxiliary input must break one-wayness in order to produce a pre-image. Hence, Theorem 17.10 cannot be adapted to refute weak $\text{cPREH}_{\approx_c}^{\text{rand}}$.

⁶⁶ For weak $\text{PREH}_{\approx_c}^{\text{rand}}$, the NIZK proof system must be replaced by a NIWI proof system (without CRS) and therefore the sampler S samples two images of the EOWF and a NIWI proof using one of the pre-images as witness.

17.3 STATIC PSEUDORANDOM ENCODINGS FROM INDISTINGUISHABILITY OBFUSCATION

(Static) $\text{cISH}_{\approx_c}^{\text{rand}}$ (and hence $\text{cPREH}_{\approx_c}^{\text{rand}}$) is implied by the existence of an explainability compiler [DKR15] which can be built from indistinguishability obfuscation and one-way functions. Using the ideas from [SW14; DKR15], we obtain perfectly correct $\text{cPREH}_{\approx_c}^{\text{rand}}$ with and without universal setup. Note that perfect correctness does not follow from Theorem 16.1. Let S be a PPT algorithm and let $U_B(C, x; r)$ be the universal circuit that accepts any circuit C which can be represented with B bits and evaluates C on input x and randomness r .

Let B be an upper bound on the bitlength which is necessary to describe a sampler. Let $\ell_{\text{in}} = |m|$ be an upper bound on the bitlength of the inputs, ℓ_{out} be an upper bound on bitlength of the outputs and ℓ_r be an upper bound on the bitlength of random tape of such samplers. Let PRG be a PRG that maps $\{0, 1\}^\lambda$ to $\{0, 1\}^{2\lambda}$. Let $|u_1| = \ell_1 = 2B + 2\ell_{\text{in}} + 2\ell_{\text{out}} + 5\lambda$ and $|u_2| = \ell_2 = B + \ell_{\text{in}} + \ell_{\text{out}} + 2\lambda$.

THEOREM 17.13. *Let iO be a perfectly correct indistinguishability obfuscator, F_1, F_2, F_3 be puncturable PRFs satisfying the following additional properties*

- F_1 is extracting when the input min-entropy is greater than $\ell_r + 2(\lambda + 1) + 2$ with error less than $2^{-(\lambda+1)}$ and has input length $\ell_1 + \ell_2 + \ell_{\text{in}} + B$ and output length ℓ_r (such a PRF exists from one-way functions since $\ell_1 + \ell_2 + \ell_{\text{in}} + B \geq \ell_r + 2(\lambda + 1) + 2$),
- F_2 is statistically injective and has input length $B + \ell_{\text{in}} + \ell_{\text{out}} + 2\lambda$ and output length ℓ_1 (such a PRF exists from one-way functions since $\ell_1 \geq 2(B + \ell_{\text{in}} + \ell_{\text{out}} + 2\lambda) + \lambda$),
- F_3 has input length ℓ_1 and output length ℓ_2 .

Then, perfectly correct $\text{cPREH}_{\approx_c}^{\text{rand}}$ with universal setup is true.

Proof. We prove that (Setup, E_S, D_S) defined in Figure 17.8 satisfies perfect correctness and pseudorandomness.

PERFECT CORRECTNESS. Let $\text{crs} = (\Lambda_E, \Lambda_D)$ be parameters in the support of $\text{Setup}(1^\lambda, B)$. Let S be a PPT algorithm represented as a polynomial sized

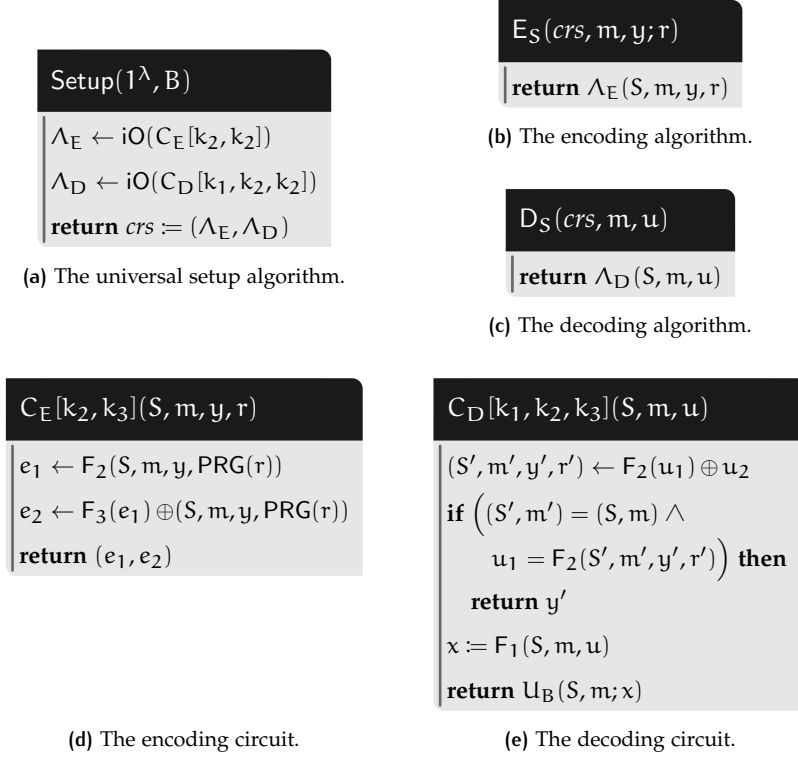


Figure 17.8: Instantiation of perfectly correct $\text{cPREH}_{\approx_c}^{\text{rand}}$ due to [SW14; DKR15] with universal setup.

circuit, $m \in L$ be an input for S and let $y \in \text{supp}(S(m))$. Due to perfect correctness of iO we have

$$\begin{aligned}
 & D_S(crs, m, E_S(crs, m, y; r)) \\
 &= C_D(S, m, C_E(S, m, y, r)) \\
 &= C_D\left(S, m, \underbrace{(F_2(S, m, y, \text{PRG}(r)))}_{=e_1}, \underbrace{(F_3(e_1) \oplus (S, m, y, \text{PRG}(r)))}_{=e_2}\right) = y.
 \end{aligned}$$

PSEUDORANDOMNESS. The proof can be divided into two main steps. In Figure 17.9, we define the corresponding hybrids.

GAME G_0 . G_0 corresponds to $\text{Exp}_{(\text{Setup}, E_S, D_S), \mathcal{A}}^{(b)\text{-crs-pre}}(\lambda)$.

GAME G_1 . G_1 is identical to G_0 except that x^* (i. e., the randomness used by S) is produced differently. In G_0 , x^* is sampled uniformly at random, whereas in G_1 , x^* produced is via $F_1(S^*, m^*, u^*)$ (for uniformly random u^*).

CLAIM 17.8. For all (potentially unbounded) adversaries \mathcal{A} , $|\Pr[out_1 = 1] - \Pr[out_0 = 1]|$ is negligible.

The proof works as the proof of **IND-CPA** security of the deniable encryption scheme from [SW14] or the proof of statistical functional equivalence from [DKR15].

GAME G_2 . G_2 is identical to G_1 except that \mathcal{A} does not receive the “explained” randomness e^* but the uniformly sampled randomness u^* as input.

CLAIM 17.9. For all PPT adversaries \mathcal{A} , $|\Pr[out_2 = 1] - \Pr[out_1 = 1]|$ is negligible.

The proof works as the proof of explainability of the deniable encryption scheme from [SW14] or the proof of explainability from [DKR15]. \square

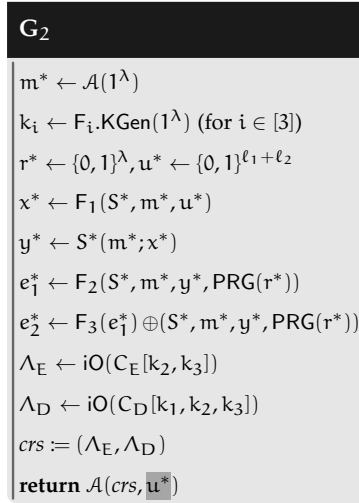
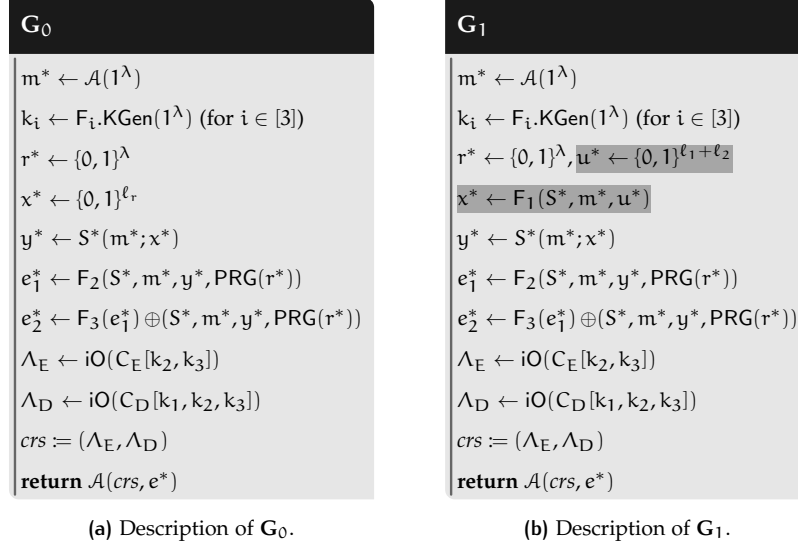


Figure 17.9: Hybrids used in the proof of pseudorandomness for Theorem 17.13.

By allowing the setup algorithm Setup_S to depend on S and replacing the universal circuit in C_D with S , we obtain $\text{cPREH}_{\approx_c}^{\text{rand}}$ without universal setup.

Let S be a PPT sampler. Let $\ell_{\text{in}} = |m|$ be an upper bound on the bitlength of the inputs, ℓ_{out} be an upper bound on bitlength of the outputs and ℓ_r be an upper bound on the bitlength of random tape of S . Let PRG be a PRG that maps $\{0, 1\}^\lambda$ to $\{0, 1\}^{2\lambda}$. Let $|u_1| = \ell_1 = 2\ell_{\text{in}} + 2\ell_{\text{out}} + 5\lambda + \ell_r$ and $|u_2| = \ell_2 = \ell_{\text{in}} + \ell_{\text{out}} + 2\lambda$.

THEOREM 17.14. Let iO be a perfectly correct indistinguishability obfuscator, F_1, F_2, F_3 be puncturable PRFs satisfying the following additional properties

- F_1 is extracting when the input min-entropy is greater than $\ell_r + 2(\lambda + 1) + 2$ with error less than $2^{-(\lambda+1)}$ and has input length $\ell_1 + \ell_2 + \ell_{in}$ and output length ℓ_r (such a PRF exists from one-way functions since $\ell_1 + \ell_2 + \ell_{in} + \ell_r \geq \ell_r + 2(\lambda + 1) + 2$),
- F_2 is statistically injective and has input length $\ell_{in} + \ell_{out} + 2\lambda$ and output length ℓ_1 (such a PRF exists from one-way functions since $\ell_1 \geq 2(\ell_{in} + \ell_{out} + 2\lambda) + \lambda$),
- F_3 has input length ℓ_1 and output length ℓ_2 .

Then, perfectly correct $\text{cPREH}_{\approx_c}^{\text{rand}}$ is true.

The proof of Theorem 17.14 proceeds identically to Theorem 17.13 with $(\text{Setup}_S, E_S, D_S)$ as defined in Figure 17.10. The only differences are changed input and output lengths for the pseudorandom functions.

Theorems 17.13 and 17.14 together with Theorem 15.1 yield the following corollary.

COROLLARY 17.8. Assuming polynomially secure IO for all circuits and one-way functions, then perfectly correct $\text{acPREH}_{\approx_c}^{\text{rand}}$ (with or without universal setup) is true.

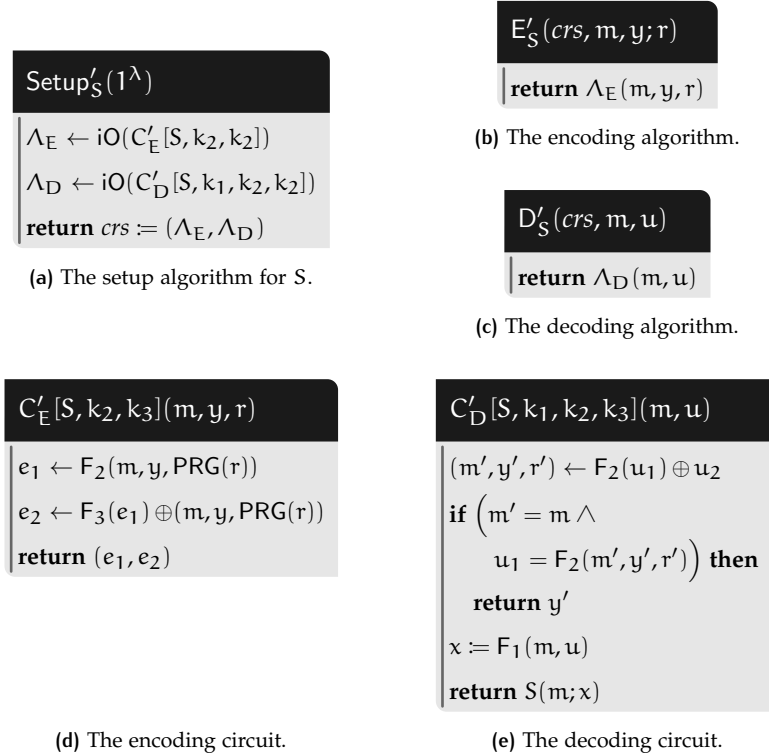


Figure 17.10: Instantiation of perfectly correct $\text{cPREH}_{\approx_c}^{\text{rand}}$ due to [SW14; DKR15] without universal setup.

17.4 BOOTSTRAPPING PSEUDORANDOM ENCODINGS WITH A COMMON RANDOM STRING

Recall that $\text{cPREH}_{\approx_c}^{\text{rand}}$ where the setup algorithm produces uniform random strings (i. e., in the common random string model) conflicts with **EOWFs** with common but benign auxiliary input – an assumption which we believe to be true. However, the very nature of pseudorandom encodings allows to bootstrap pseudorandom encodings in the common random string model from pseudorandom encodings with arbitrary setup if this setup algorithm – interpreted as a sampler – can be pseudorandomly encoded such that the corresponding setup algorithm produces common random strings. Note that this sampler does not expect an input. Via this bootstrapping, we are able to refute even the *weak* pseudorandom encoding hypothesis with common random string (**URC**).

THEOREM 17.15. *Assume (i) $\text{cPREH}_{\approx_c}^{\text{rand}}$ is true for all **PPT** samplers, i. e., for all **PPT** algorithms S , there exists a pseudorandom encoding scheme $(\text{Setup}'_S, E'_S, D'_S)$, and (ii) weak $\text{cPREH}_{\approx_c}^{\text{rand}}$ is true for the class of **PPT** algorithms Setup'_S such that the corresponding setup algorithm $\text{Setup}''_{\text{Setup}'_S}$ produces uniform random strings. Then, $\text{cPREH}_{\approx_c}^{\text{rand}}$ with common random strings is true for all **PPT** samplers.*

Proof. Item **i** implies that for all **PPT** algorithms S , there exists a pseudorandom encoding scheme $(\text{Setup}'_S, E'_S, D'_S)$. Further, Item **ii** guarantees the existence of a pseudorandom encoding scheme $(\text{Setup}''_{\text{Setup}'_S}, E''_{\text{Setup}'_S}, D''_{\text{Setup}'_S})$ such that $\text{Setup}''_{\text{Setup}'_S}$ produces uniform random strings. Let $\{0, 1\}^{n''(\lambda)}$ be the range of $E''_{\text{Setup}'_S}$ (note that $\{0, 1\}^{n''(\lambda)}$ is allowed to depend on the sampler Setup'_S). For notational convenience, we henceforth omit the dependency on the sampler Setup'_S .

Let S be some **PPT** algorithm. We define $(\text{Setup}_S, E_S, D_S)$ corresponding to S as in Figure 17.11.

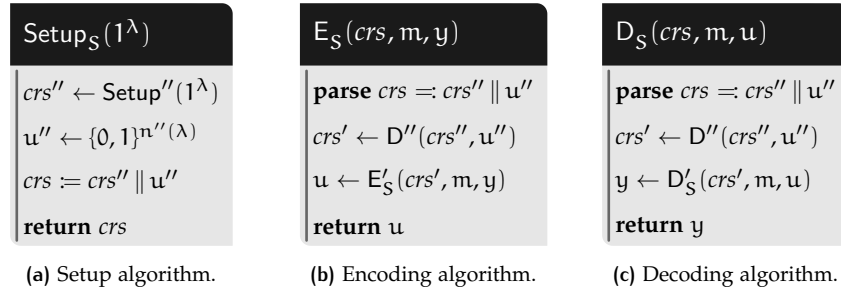


Figure 17.11: Algorithms for $\text{cPREH}_{\approx_c}^{\text{rand}}$ with common random string.

CORRECTNESS. We use a sequence of hybrids, see Figure 17.12.

GAME G_0 . G_0 encodes and decodes a sample y (using a statically chosen input m) using $(\text{Setup}_S, E_S, D_S)$.

GAME G_1 . G_1 is identical to G_0 except that G_1 uses the implementations of Setup_S, E_S and D_S explicitly. This difference is only conceptual and $\Pr[out_0 = 1] = \Pr[out_1 = 1]$.

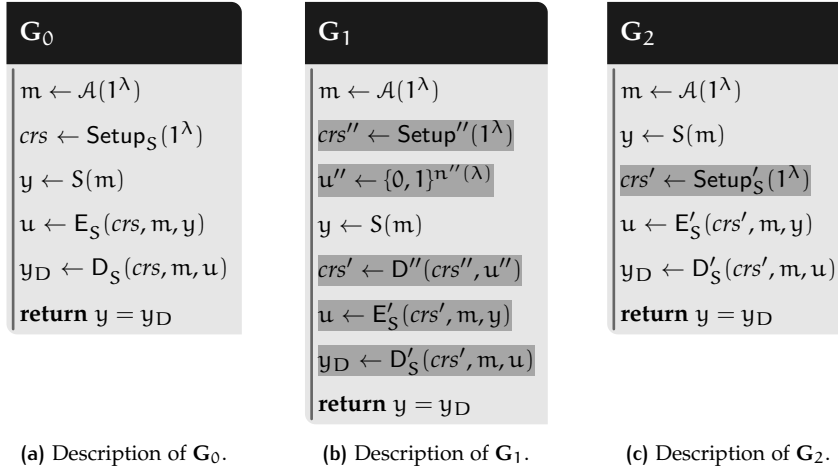


Figure 17.12: Hybrids used in the proof of correctness of Theorem 17.15.

GAME G_2 . G_2 is identical to G_1 except that crs' is sampled using $\text{Setup}'_S(1^\lambda)$ instead of decoding u'' using $D''(crs'', u'')$.

Due to Lemma 16.2, by correctness and pseudorandomness of $(\text{Setup}'', E'', D'')$, the two distributions

$$\left\{ crs'' \leftarrow \text{Setup}''(1^\lambda), crs' \leftarrow \text{Setup}'_S(1^\lambda) : (crs'', crs') \right\} \text{ and}$$

$$\left\{ crs'' \leftarrow \text{Setup}''(1^\lambda), crs' \leftarrow D''(crs'', U_{n''(\lambda)}) : (crs'', crs') \right\}$$

are computationally indistinguishable. Hence, $|\Pr[out_2 = 1] - \Pr[out_1 = 1]|$ is negligible.

Due to correctness of $(\text{Setup}'_S, E'_S, D'_S)$, for all PPT adversaries \mathcal{A} ,

$$\Pr \left[\begin{array}{l} m \leftarrow \mathcal{A}(1^\lambda) \\ crs' \leftarrow \text{Setup}'_S(1^\lambda) : D'_S(crs', m, E'_S(crs', m, y)) = y \\ y \leftarrow S(m) \end{array} \right]$$

and, hence, $\Pr[out_2 = 1]$ are overwhelming.

Before we prove pseudorandomness, we introduce a technical lemma.

LEMMA 17.10. *Let S be a PPT sampler such that $\text{cPREH}_{\approx c}^{\text{rand}}$ is true for S and let $(\text{Setup}_S, E_S, D_S)$ be the corresponding algorithms (such that E_S has output length $n(\lambda)$). Then, for all PPT adversaries \mathcal{A} ,*

$$\left| \Pr \left[\begin{array}{l} m \leftarrow \mathcal{A}(1^\lambda) \\ crs \leftarrow \text{Setup}_S(1^\lambda) \\ y \leftarrow S(m) \\ u \leftarrow E_S(crs, m, y) \end{array} : \mathcal{A}(crs, m, y, u) = 1 \right] \right.$$

$$\left. - \Pr \left[\begin{array}{l} m \leftarrow \mathcal{A}(1^\lambda) \\ crs \leftarrow \text{Setup}_S(1^\lambda) \\ u \leftarrow \{0, 1\}^{n(\lambda)} \\ y \leftarrow D_S(crs, m, u) \end{array} : \mathcal{A}(crs, m, y, u) = 1 \right] \right|$$

is negligible.

Proof of Lemma 17.10. We proceed over a series of hybrids, see Figure 17.13.

GAME H_0 . H_0 corresponds to the left-hand-side of the difference, where \mathcal{A} receives a sample y from $S(m)$ and the corresponding encoded sample $u \leftarrow E_S(crs, m, y)$.

GAME H_1 . H_1 is identical to H_0 except that \mathcal{A} receives the decoded sample y_D instead of the actual sample y_S . If $y_S = y_D$ both games behave identically, hence, by correctness, $|\Pr[out_{H_1} = 1] - \Pr[out_{H_0} = 1]| \leq \epsilon_{(\text{Setup}_S, E_S, D_S), \mathcal{A}}^{\text{c-dec-error}}(\lambda)$.

GAME H_2 . H_2 is identical to H_1 except that u is not produced as an encoding of y , but sampled uniformly at random from $\{0, 1\}^{n(\lambda)}$. This game hop is justified by pseudorandomness $|\Pr[out_{H_2} = 1] - \Pr[out_{H_1} = 1]| \leq \text{Adv}_{(\text{Setup}_S, E_S, D_S), \mathcal{A}}^{\text{crs-pre}}(\lambda)$.

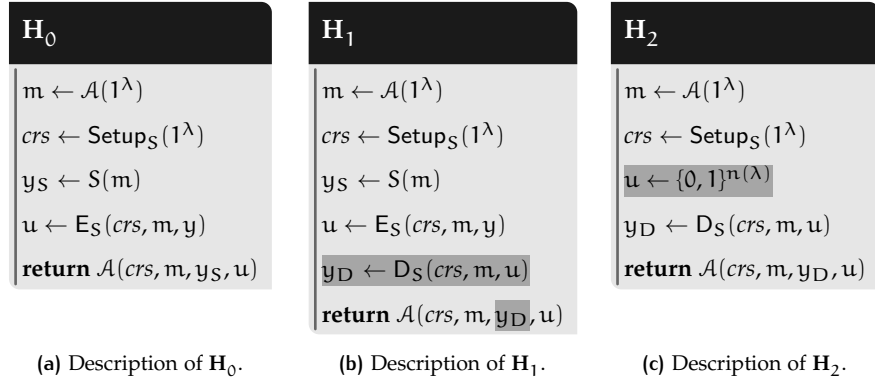


Figure 17.13: Hybrids used in the proof of Lemma 17.10.

This concludes the proof of Lemma 17.10. \square

PSEUDORANDOMNESS. To prove pseudorandomness, we proceed over a sequence of hybrids, see Figure 17.14.

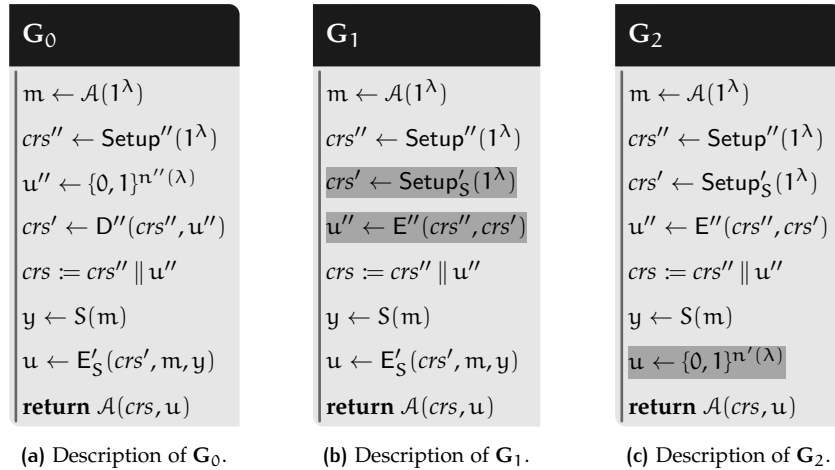


Figure 17.14: Hybrids used in the proof of pseudorandomness of Theorem 17.15.

GAME G_0 . G_0 corresponds to $\text{Exp}_{(\text{Setup}_S, E_S, D_S), \mathcal{A}}^{(0)\text{-crs-pre}}(\lambda)$.

GAME G_1 . G_1 is identical to G_0 except for the following difference. G_1 produces crs' directly via $\text{Setup}'_S(1^\lambda)$ and produces u'' via encoding crs'

with $E''(crs'', crs')$. \mathbf{G}_0 , on the other hand, samples u'' uniformly at random from $\{0, 1\}^{n''(\lambda)}$ and produces crs'' via decoding u'' as $D''(crs'', u'')$. Due to Lemma 17.10, $|\Pr[out_1 = 1] - \Pr[out_0 = 1]|$ is negligible.

GAME \mathbf{G}_2 . \mathbf{G}_2 is identical to \mathbf{G}_1 except that \mathbf{G}_2 samples u uniformly at random from $\{0, 1\}^{n'(\lambda)}$, whereas \mathbf{G}_1 produces u as $E'_S(crs', m, y)$. For all PPT adversaries \mathcal{A} , $|\Pr[out_2 = 1] - \Pr[out_1 = 1]|$ is negligible due to pseudorandomness of $(\text{Setup}'_S, E'_S, D'_S)$. □

Combining Theorem 17.15 and Corollary 17.7, we obtain the following corollary refuting weak $\text{cPREH}_{\approx_c}^{\text{rand}}$ with common random string.

COROLLARY 17.9. *If there exist extractable one-way function family ensembles with common but benign auxiliary information and indistinguishability obfuscation for all circuits, then weak $\text{cPREH}_{\approx_c}^{\text{rand}}$ with common random string (URC) is false.*

REMARK 17.1. If weak $\text{cPREH}_{\approx_c}^{\text{rand}}$ was true such that the setup algorithms always produce uniform random strings, then the global common reference string model could be replaced with the global common random string model (or the non-programmable random oracle model) in several settings in cryptography.

18

APPLICATIONS OF PSEUDORANDOM ENCODINGS

In this chapter we describe further applications of pseudorandom encodings. This unifies several areas in cryptography. In Section 18.1, we show that adaptively secure pseudorandom encodings with setup are equivalent to fully adaptive multi-party computation in the global CRS model for all PPT functionalities (that includes non-adaptively-well-formed functionalities). This extends the results from [IKOS10] and clarifies a claim made in [DKR15]. In Section 18.2, we show that pseudorandom encodings yield honey encryption due to [JR14] for arbitrary message distributions, even such which admit inputs. In Section 18.3, we define a keyless version of steganography. That is, in contrast to symmetric-key or public-key steganography, parties do not need any secret information in order to covertly communicate with each other. In Sections 18.4 and 18.5, we analyze the relation of pseudorandom encodings and certain PKE variants which are known to imply variants of adaptive MPC.

18.1 ADAPTIVELY SECURE MULTI-PARTY COMPUTATION

Due to the equivalence of pseudorandom encodings and invertible sampling, we obtain an equivalence between pseudorandom encodings and fully adaptively secure multi-party computation. Due to [IKOS10], $\text{PREH}_{\approx_c}^{\text{rand}}$ in conjunction with some adaptively secure oblivious transfer protocol implies fully adaptively secure multi-party computation for all randomized functionalities in the plain model.

In Section 18.1.1, we introduce some preliminaries on adaptive multi-party computation, henceforth denoted AMPC. In Sections 18.1.2 and 18.1.3, we extend the results on the relation between the invertible sampling hypothesis and AMPC from [IKOS10] to the global CRS model. This is directly possible since, due to Theorem 15.1 and Corollary 16.1, we have an instantiation of $\text{aclSH}_{\approx_c}^{\text{rand}}$.

18.1.1 Adaptive MPC

The definition of security of multi-party protocols follows the real/ideal model paradigm, [GMW87; Canoo]. A protocol Π is said to be secure if the output of the real execution of the protocol is indistinguishable from the output of an ideal computation, where a trusted third party exists. The behavior of this trusted third party is defined via a functionality \mathcal{F} .

Let \mathcal{F} be a PPT functionality. Let x_i be the input to party P_i . We only consider semi-honest adversaries. Semi-honest adversaries are bound to follow the protocol specification while trying to obtain as much information

as possible. There are general techniques to transform protocols which are secure with respect to semi-honest adversaries to protocols which are secure with respect to malicious adversaries at the cost of a local CRS which is inherent for malicious security, [CLOS02]. In the following, we describe the stand-alone model for adaptive MPC on a very high level. We refer the reader to [Canoo; Lin09] for more details.

THE IDEAL EXECUTION. The ideal execution involves an ideal PPT adversary Sim (also called the simulator), a PPT environment \mathcal{Z} and a trusted third party T . The ideal execution proceeds through several stages. In the *first corruption stage*, Sim adaptively decides to corrupt parties. When Sim corrupts a party, Sim learns that party's input and \mathcal{Z} learns the identity of the corrupted party. In the *computation stage*, the uncorrupted and the corrupted parties send their inputs (which are determined by the environment) to T . The trusted party T evaluates the (possibly randomized) function $(z_1, \dots, z_n) \leftarrow \mathcal{F}(x_1, \dots, x_n)$ and sends the output to the respective parties. In the *second corruption stage*, Sim learns the outputs of the corrupted parties and again adaptively decides to corrupt a party. Upon corruption of a party, Sim learns that party's input *and output* and \mathcal{Z} learns the corrupted party's identity. In the *output stage*, the uncorrupted parties output what they received from T to \mathcal{Z} , the corrupted parties output \perp to \mathcal{Z} and Sim outputs an arbitrary string to \mathcal{Z} . Finally, in the *post-execution corruption stage*, as long as the environment \mathcal{Z} did not halt, \mathcal{Z} sends corruption requests to Sim . Sim generates an arbitrary answer based on its view so far. Furthermore, Sim may additionally corrupt parties as in the second corruption stage.

Let $\mathcal{Z}^{\text{Sim}, \mathcal{F}}(1^\lambda)$ denote the output distribution of the environment \mathcal{Z} when interacting as described above with the simulator Sim and parties P_1, \dots, P_n on inputs x_1, \dots, x_n chosen by \mathcal{Z} .

THE REAL EXECUTION. Initially, the environment \mathcal{Z} (adaptively) chooses inputs x_1, \dots, x_n and sends each party P_i its input. Before the communication rounds start, the adversary \mathcal{A} receives an initial message from \mathcal{Z} . While there exist uncorrupted parties which did not halt, \mathcal{A} may adaptively decide to corrupt new parties. Upon corruption of P_i , \mathcal{Z} learns the identity of P_i , \mathcal{A} learns the input x_i , P_i 's internal state (i. e., its random tape) and all messages received so far. From that time on, \mathcal{A} is in control of the messages this party sends. Furthermore, if a corrupted party receives a message, \mathcal{A} also learns that message. Additionally, \mathcal{A} determines the order in which the uncorrupted parties are activated. If all parties are corrupted or all uncorrupted parties halted, each uncorrupted party and \mathcal{A} produce outputs. The environment \mathcal{Z} learns all of these outputs. Similarly to the post-execution corruption stage in the ideal execution, while \mathcal{Z} did not halt, \mathcal{Z} may instruct \mathcal{A} to corrupt more parties or \mathcal{A} may decide himself to corrupt more parties. Upon corruption, \mathcal{A} learns P_i 's input, randomness and all messages received so far and \mathcal{Z} learns the corrupted party's identity. Additionally, \mathcal{A} sends P_i 's internal state to \mathcal{Z} .

Let $\mathcal{Z}^{\mathcal{A}, \Pi}(1^\lambda)$ denote the output of the environment \mathcal{Z} when interacting with the adversary \mathcal{A} and parties P_1, \dots, P_n running the protocol Π on inputs x_1, \dots, x_n chosen by \mathcal{Z} .

DEFINITION 18.1 (Adaptive security in the stand-alone model, [Canoo]). We say a protocol Π computes functionality \mathcal{F} in the adaptive semi-honest model if for

every PPT adversary \mathcal{A} , there exists a PPT simulator Sim , such that for all PPT environments \mathcal{Z} ,

$$\left| \Pr \left[\mathcal{Z}^{\mathcal{A}, \Pi}(1^\lambda) = 1 \right] - \Pr \left[\mathcal{Z}^{\text{Sim}, \mathcal{F}}(1^\lambda) = 1 \right] \right|$$

is negligible.

COMMON REFERENCE STRING MODEL. In the common reference string (CRS) model, all parties have access to a string which is honestly generated by a trusted third non-participating party. There are two widely used variants of the CRS model. In the *programmable* or *local* CRS model, the simulator may generate the CRS. This enables the simulator to sample the CRS along with corresponding trapdoors which results in an asymmetry between the simulator and the adversary facilitating simulation. However, as soon as two different protocols use the same programmable CRS, all security guarantees related to that CRS break down, see [CDPW07]. In the *non-programmable* or *global* CRS model, the simulator receives the CRS as input and, hence, has no additional power compared to the adversary. A global CRS can be made public and used by any number of protocols without compromising security.

Adaptive security in the global CRS model (with respect to some efficiently samplable CRS distribution Setup) is defined as Definition 18.1 with the difference that all parties including the environment, the adversary and the simulator receive the CRS as input. We stress that the environment may decide on the inputs for the parties adaptively after seeing the CRS.

THE UNIVERSAL COMPOSABILITY FRAMEWORK. The following paragraph describes the universal composability (UC) framework on a very high level. We refer the reader to [Can01] for more details on the model.

In contrast to the stand-alone model for secure computation, the UC model allows the environment to be interactive. In particular, every message that is transmitted between parties is sent to the environment which can arbitrarily decide how to deal with it. The strongest possible adversary in this model is the so-called dummy adversary who simply forwards all instructions he receives from the environment. The UC framework offers a set of security-preserving composition theorems which allow for a modular analysis. If a protocol Π UC-realizes a functionality \mathcal{F} , a protocol Π' which uses protocol Π as a subroutine can be proven to securely realize a functionality \mathcal{F}' in the \mathcal{F} -hybrid model, that is having access to the ideal functionality \mathcal{F} .

DEFINITION 18.2 (Adaptive UC-security, [Can01]). Let \mathcal{A} be the dummy adversary and let θ denote the dummy protocol. We say a protocol Π UC-realizes a functionality \mathcal{F} in the \mathcal{G} -hybrid model in the presence of semi-honest adaptive adversaries if for all PPT environments \mathcal{Z} , there exists a simulator Sim , such that

$$\left| \Pr \left[\text{real}[\mathcal{Z}, \mathcal{A}, \pi^{\mathcal{G}}](\lambda) = 1 \right] - \Pr \left[\text{ideal}[\mathcal{Z}, \text{Sim}, \theta^{\mathcal{F}}](\lambda) = 1 \right] \right|$$

is negligible, where $\text{real}[\mathcal{Z}, \mathcal{A}, \pi^{\mathcal{G}}](\lambda)$ denotes the output distribution of \mathcal{Z} when interacting in the real world with \mathcal{A} and $\pi^{\mathcal{G}}$ and $\text{ideal}[\mathcal{Z}, \text{Sim}, \theta^{\mathcal{F}}](\lambda)$ denotes the output distribution of \mathcal{Z} when interacting in the ideal world with Sim and $\theta^{\mathcal{F}}$.

UC-security is a stronger notion than stand-alone security. More precisely, if there exists a protocol which UC-realizes a functionality according to

Definition 18.2, then there exists a protocol which securely computes that functionality according to Definition 18.1.

In the UC framework, the global CRS model corresponds to the $\mathcal{F}_{\text{Setup}}^{\text{CRS}}$ -hybrid model. The ideal functionality $\mathcal{F}_{\text{Setup}}^{\text{CRS}}$ can be queried by any party and the adversary. Upon receiving such a query, $\mathcal{F}_{\text{Setup}}^{\text{CRS}}$ gives the CRS crs to the querying party. If crs is not initialized, $\mathcal{F}_{\text{Setup}}^{\text{CRS}}$ samples crs according to $\text{Setup}(1^\lambda)$.

18.1.2 Pseudorandom Encodings Imply Adaptive MPC

AMPC for deterministic functionalities is possible assuming adaptively secure oblivious transfer due to [Kil88].

THEOREM 18.1 (informal, [Kil88; IPS08]). *Any deterministic functionality can be UC-realized in the OT-hybrid model in the presence of semi-honest adversaries adaptively corrupting any number of parties.*

Ishai et al. [IKOS10] prove that adaptively secure oblivious transfer in conjunction with $\text{ISH}_{\approx_c}^{\text{rand}}$ yields AMPC for all randomized functionalities in the plain model.

Assuming $\text{acPREH}_{\approx_c}^{\text{rand}}$, the strategy from [IKOS10] can be applied to show that every randomized functionality can be UC-realized in the (global CRS, OT)-hybrid model in the presence of semi-honest adversaries adaptively corrupting any number of parties.

THEOREM 18.2. *Assume $\text{acPREH}_{\approx_c}^{\text{rand}}$ holds, then any functionality \mathcal{F} can be UC-realized in the (global CRS, \mathcal{F}_{OT})-hybrid-model, in the presence of semi-honest adaptive adversaries corrupting any number of parties.*

Proof. The proof follows the strategy from [IKOS10]. We consider functionalities \mathcal{F} giving the same output to all parties. (This is without loss of generality since evaluating the function \mathcal{F}' which distributes blinded individual outputs to the parties, i. e., $\mathcal{F}'((x_1, k_1), (x_2, k_2)) := \mathcal{F}_1(x_1, x_2) \oplus k_1 \parallel \mathcal{F}_2(x_1, x_2) \oplus k_2$, yields the general case, see [GL91].) That is, the functionality \mathcal{F} takes as input x_1, x_2 (the inputs of the parties P_1 and P_2 , respectively) and internal randomness ρ , and outputs z to both parties. We consider the case of two parties (the proof easily extends to the case of many parties).

We view \mathcal{F} as a PPT algorithm S . Due to $\text{acPREH}_{\approx_c}^{\text{rand}}$, there exist Setup_S, \bar{S} and \bar{S}^{-1} satisfying adaptive closeness and adaptive invertibility. Define a deterministic functionality \mathcal{G} (in the $\mathcal{F}_{\text{g-crs}}^{\text{Setup}_S}$ -hybrid model) which takes inputs (x_1, ρ_1) from party P_1 and (x_2, ρ_2) from party P_2 , and evaluates $\bar{S}(\text{crs}, (x_1, x_2); \rho_1 \oplus \rho_2)$, where crs comes from $\mathcal{F}_{\text{g-crs}}^{\text{Setup}_S}$. Due to Theorem 18.1, \mathcal{G} can be UC-realized in the OT-hybrid model in the presence of semi-honest adversaries adaptively corrupting any number of parties.

Now we turn to realize \mathcal{F} in the $(\mathcal{G}, \mathcal{F}_{\text{g-crs}}^{\text{Setup}_S})$ -hybrid model. Party P_i chooses randomness ρ_i , feeds (x_i, ρ_i) into \mathcal{G} and waits for the output. Without loss of generality, we assume that \mathcal{Z} first observes the entire protocol before corrupting parties.⁶⁷ The simulator Sim works as follows. Sim initially only receives the output z of the ideal functionality \mathcal{F} as input and calls $\mathcal{F}_{\text{g-crs}}^{\text{Setup}_S}$ to obtain crs . If both parties are corrupted (first P_1 , then P_2), Sim outputs a uniformly random string ρ_1 to explain the internal randomness of party P_1 , computes $\rho \leftarrow \bar{S}^{-1}(\text{crs}, (x_1, x_2), z)$ and outputs $\rho_2 := \rho \oplus \rho_1$ to explain the

⁶⁷ This is because we consider semi-honest adversaries. Such adversaries collect as much information as possible without deviating from the protocol description. An adversary who corrupts parties earlier (or not at all) obtains strictly less information and hence simulation is easier.

internal randomness of party P_2 (note that Sim learns the inputs x_1 and x_2 at the time of corrupting P_1 and P_2 , respectively).

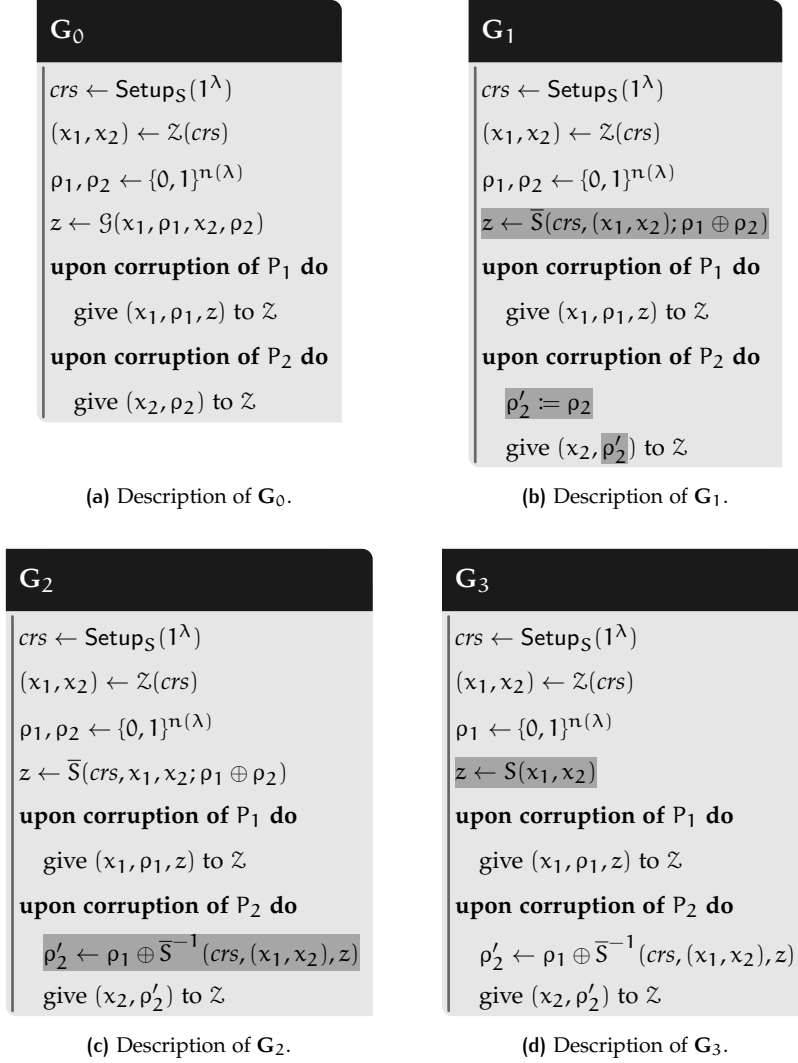


Figure 18.1: Hybrids used in proof of Theorem 18.2.

We proceed over a series of hybrids, see Figure 18.1.

GAME G_0 . G_0 is identical to $\text{real}[\mathcal{Z}, \mathcal{A}, \pi^{\mathcal{G}}](\lambda)$.

GAME G_1 . G_1 is identical to G_0 except for the conceptual change that \mathcal{G} is replaced with its implementation and ρ_2 is renamed to ρ'_2 before being given to \mathcal{Z} . Hence, $|\Pr[out_0 = 1] - \Pr[out_1 = 1]|$.

GAME G_2 . G_2 is identical to G_1 except that in G_1 , \mathcal{Z} receives the real random tape ρ_2 used by party P_2 , whereas in G_2 , \mathcal{Z} receives ρ'_2 which is computed as the logical XOR of ρ_1 and inverse sampled random coins $\bar{S}^{-1}(crs, (x_1, x_2), z)$.

CLAIM 18.1. For all PPT environments \mathcal{Z} , there exists a PPT adversary \mathcal{B}_2 , such that $|\Pr[out_2 = 1] - \Pr[out_1 = 1]| \leq \text{Adv}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{B}_2}^{\text{a-crs-inv}}(\lambda)$.

Proof. Construct an adversary \mathcal{B}_2 breaking adaptive closeness as follows. \mathcal{B}_2 receives crs sampled from $\text{Setup}_S(1^\lambda)$ and calls \mathcal{Z} on input of crs . In

return, \mathcal{B}_2 receives (x_1, x_2) and samples $\rho_1 \leftarrow \{0, 1\}^{n(\lambda)}$. Then, \mathcal{B}_2 outputs $m := (x_1, x_2)$ and receives (r, y) from the invertibility experiment. \mathcal{B}_2 uses $z := y$ and $\rho'_2 := \rho_1 \oplus r$ and continues as in \mathbf{G}_1 .

If (r, y) were produced such that r was sampled uniformly from $\{0, 1\}^{n(\lambda)}$ and $y \leftarrow \bar{S}(crs, (x_1, x_2); r)$, then \mathcal{B}_2 perfectly simulates \mathbf{G}_1 . If (r, y) were produced such that y was sampled from $\bar{S}(crs, (x_1, x_2))$ (using some uniform randomness) and r was inversely sampled as $r \leftarrow \bar{S}^{-1}(crs, (m_1, m_2), y)$, then \mathcal{B}_2 perfectly simulates \mathbf{G}_2 . \square

GAME \mathbf{G}_3 . \mathbf{G}_3 is identical to \mathbf{G}_2 except that in \mathbf{G}_2 , z is sampled using the alternative sampler $\bar{S}(crs, (x_1, x_2))$, whereas in \mathbf{G}_3 , z is sampled using the original sampler $S(x_1, x_2)$. \mathbf{G}_3 is identical to $\text{ideal}[\mathcal{Z}, \mathcal{A}, \pi^{\mathcal{G}}](\lambda)$.

CLAIM 18.2. For all PPT environments \mathcal{Z} , there exists a PPT adversary \mathcal{B}_3 , such that $|\Pr[out_3 = 1] - \Pr[out_2 = 1]| \leq \text{Adv}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{B}_3}^{\text{a-crs-close}}(\lambda)$.

Proof. Construct an adversary \mathcal{B}_3 breaking adaptive closeness as follows. \mathcal{B}_3 receives crs sampled from $\text{Setup}_S(1^\lambda)$ and calls \mathcal{Z} on input of crs . In return, \mathcal{B}_3 receives (x_1, x_2) , samples $\rho_1 \leftarrow \{0, 1\}^{n(\lambda)}$ and outputs $m := (x_1, x_2)$ to the adaptive closeness experiment to receive y . \mathcal{B}_3 uses $z := y$ and continues as in \mathbf{G}_2 . Note that \mathbf{G}_2 can be simulated without knowing the random coins used for \bar{S} .

If y was sampled using $\bar{S}(crs, (x_1, x_2))$, then \mathcal{B}_3 perfectly simulates \mathbf{G}_2 . If y was sampled using $S(x_1, x_2)$, then \mathcal{B}_3 perfectly simulates \mathbf{G}_3 . \square

Thus, due to adaptive invertibility and adaptive closeness, for all environments \mathcal{Z} ,

$$\left| \Pr \left[\text{real}[\mathcal{Z}, \mathcal{A}, \pi^{\mathcal{G}}](\lambda) = 1 \right] - \Pr \left[\text{ideal}[\mathcal{Z}, \text{Sim}, \theta^{\mathcal{F}}](\lambda) = 1 \right] \right|$$

is negligible. \square

Note that due to [DKR15] it is possible to realize UC-AMPC for every functionality only assuming the static version $\text{cPREH}_{\approx_c}^{\text{rand}}$ at the cost of additionally assuming the existence of an adaptively secure two-round oblivious transfer protocol. This can be avoided due to Corollary 16.1.

18.1.3 Adaptive MPC Implies Pseudorandom Encodings

In the following, we consider the standalone model [Canoo]. Note that in the standalone model, the environment is strictly weaker than in the UC-framework.

THEOREM 18.3. *If for all (two party) PPT functionalities \mathcal{F} , there exists a protocol Π that securely implements \mathcal{F} in the global CRS model in the presence of adaptive semi-honest adversaries with post-execution corruption corrupting any number of parties, then $\text{acPREH}_{\approx_c}^{\text{rand}}$ (without universal setup) is true.*

Proof. The proof follows the ideas of [IKOS10]. Let S be an arbitrary PPT algorithm. Let x_1 denote the input of party P_1 and let x_2 denote the input of party P_2 .

Consider the randomized functionality \mathcal{F} which computes $z := S(x_1, x_2; \rho)$ and outputs z to P_1 and \perp to P_2 , where ρ is the internal random tape of the

functionality. We denote the message space of S by L . Let Π be a protocol which securely realizes \mathcal{F} in the global CRS model (where the common reference string is drawn according some distribution $\text{Setup}_{\mathcal{F}}$ possibly depending on the functionality \mathcal{F}) in the presence of adaptive semi-honest adversaries (with post-execution corruption). Further, let p_1 and p_2 be stateful PPT algorithms modeling the protocol messages exchanged between the parties P_1 and P_2 . More formally, $(m_{2i+1}, z_i) \leftarrow p_1(\text{crs}, x_1, m_{2i}; r_1)$ and $m_{2j} \leftarrow p_2(\text{crs}, x_2, m_{2j-1}; r_2)$ for $i \geq 0, j \geq 1$ and $r_1, r_2 \leftarrow \{0, 1\}^{p(\lambda)}$ (for some sufficiently large polynomial p). (Without loss of generality, we let p_1 additionally output z_i which equals \perp during the protocol execution and contains the output of party P_1 after the execution.)

We define the invertible sampling scheme $(\text{Setup}_S := \text{Setup}_{\mathcal{F}}, \bar{S}, \bar{S}^{-1})$ as in Figure 18.2.

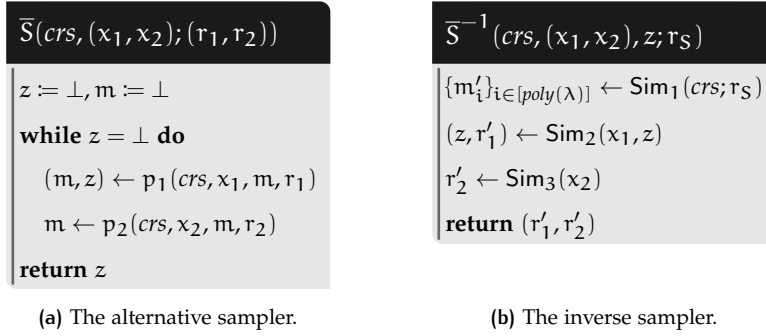


Figure 18.2: Definition of the alternative sampler \bar{S} and the corresponding inverse sampler \bar{S}^{-1} .

LEMMA 18.1. $(\text{Setup}_S := \text{Setup}_{\mathcal{F}}, \bar{S}, \bar{S}^{-1})$ defined in Figure 18.2 satisfies adaptive closeness.

Proof. Since we only consider protocols Π with at most polynomially many rounds, \bar{S} is a PPT algorithm.

Intuitively, the output distributions of S and \bar{S} are computationally close because of “correctness” of the protocol Π .

Let \mathcal{B} be the adversary who only observes the protocol Π and does not interfere at all. By Definition 18.1, there exists a simulator Sim such that for all environments \mathcal{Z} ,

$$\left| \Pr \left[\text{crs} \leftarrow \text{Setup}_{\mathcal{F}}(1^\lambda): \mathcal{Z}^{\mathcal{A}, \Pi}(\text{crs}) = 1 \right] - \Pr \left[\text{crs} \leftarrow \text{Setup}_{\mathcal{F}}(1^\lambda): \mathcal{Z}^{\text{Sim}, \mathcal{F}}(\text{crs}) = 1 \right] \right|$$

is negligible. Note that since \mathcal{B} corrupts no parties, Sim also corrupts no parties.

Let \mathcal{A} be an adversary breaking closeness. We construct an environment \mathcal{Z} which distinguishes between the real and the ideal world as follows. Initially, \mathcal{Z} receives crs as input, calls \mathcal{A} on input of crs and obtains $(x_1, x_2) \in L$. \mathcal{Z} uses x_1 as input for party P_1 and x_2 as input for P_2 . Then, \mathcal{Z} executes the protocol (without any interference) and finally receives the output (z, \perp) . \mathcal{Z} calls \mathcal{A} on input z and outputs \mathcal{A} 's output.

In the ideal world, \mathcal{Z} receives the outputs z and \perp for P_1 and P_2 from the trusted third party T (since no party is corrupted) which are computed by \mathcal{F} .

Hence, if \mathcal{Z} is in the ideal world, it simulates $\text{Exp}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{A}}^{(0)\text{-a-crs-close}}(\lambda)$. In the real world, \mathcal{Z} interacts with the actual parties running protocol Π and, hence, simulates $\text{Exp}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{A}}^{(1)\text{-a-crs-close}}(\lambda)$. Therefore,

$$\text{Adv}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{A}}^{\text{a-crs-close}}(\lambda) \leq \left| \Pr \left[\text{crs} \leftarrow \text{Setup}_{\mathcal{F}}(1^\lambda) : \mathcal{Z}^{\mathcal{A}, \Pi}(\text{crs}) = 1 \right] - \Pr \left[\text{crs} \leftarrow \text{Setup}_{\mathcal{F}}(1^\lambda) : \mathcal{Z}^{\text{Sim}, \mathcal{F}}(\text{crs}) = 1 \right] \right|$$

for the environment \mathcal{Z} and the adversary \mathcal{B} which do not interfere with the protocol execution. \square

Let \mathcal{B} be an adversary that operates as follows. During the protocol, \mathcal{B} records the transcripts $\{m_i\}_{i \in [\text{poly}'(\lambda)]}$ and outputs them. At the end of the protocol, \mathcal{B} corrupts P_1 obtaining the input x_1 , output z and random tape r_1 . \mathcal{B} directly outputs (z, r_1) . Afterwards, \mathcal{B} (post-execution) corrupts P_2 and obtains the input x_2 and the random tape r_2 and outputs r_2 .

Since we assume Π securely realizes \mathcal{F} (in particular in the presence of \mathcal{B}), there exists a (stateful) simulator $\text{Sim} := (\text{Sim}_1, \text{Sim}_2, \text{Sim}_3)$ producing outputs that are computationally indistinguishable from the outputs \mathcal{B} produces. Sim_1 simulates the execution before a corruption occurs. On input of crs and a random tape r_S , Sim_1 produces $\{m'_i\}_{i \in [\text{poly}'(\lambda)]}$. On corruption of P_1 (in the post-execution corruption stage), Sim_2 obtains x_1, z and outputs a string r'_1 corresponding to the randomness that P_1 used to produce its messages in the protocol (using input x_1) and outputs (z, r'_1) . On corruption of P_2 (in the post-execution corruption stage), Sim_3 obtains x_2 and outputs a string r'_2 corresponding to the randomness that P_2 used to produce its messages (using input x_2).

LEMMA 18.2. ($\text{Setup}_S := \text{Setup}_{\mathcal{F}, \bar{S}, \bar{S}^{-1}}$) defined in Figure 18.2 satisfies adaptive invertibility.

Proof. Since $\text{Sim}_1, \text{Sim}_2$ and Sim_3 are PPT algorithms, \bar{S}^{-1} is a PPT algorithm.

Recall that $\mathcal{F}(x_1, x_2; \rho) := S(x_1, x_2; \rho)$. Define a modified functionality $\mathcal{F}'(x_1, x_2; \rho') := \bar{S}(x_1, x_2; \rho')$. We consider an intermediate game, where \mathcal{Z} interacts with the simulator Sim and the modified ideal functionality \mathcal{F}' .

CLAIM 18.3. For all PPT environments \mathcal{Z} ,

$$\left| \Pr \left[\text{crs} \leftarrow \text{Setup}_{\mathcal{F}}(1^\lambda) : \mathcal{Z}^{\text{Sim}, \mathcal{F}}(\text{crs}) = 1 \right] - \Pr \left[\text{crs} \leftarrow \text{Setup}_{\mathcal{F}}(1^\lambda) : \mathcal{Z}^{\text{Sim}, \mathcal{F}'}(\text{crs}) = 1 \right] \right|$$

is negligible.

Proof of Claim 18.3. Let \mathcal{Z} be a PPT environment. We construct an adversary \mathcal{B} breaking adaptive closeness. \mathcal{B} simulates the entire interaction between \mathcal{Z} , the simulator Sim and the trusted party T , which evaluates the functionality \mathcal{F} (or \mathcal{F}'). In particular, \mathcal{B} takes control of the party T . On input of crs , \mathcal{B} simulates the interaction between \mathcal{Z} and Sim (given crs) until the computation stage. In the computation stage, the parties P_1 and P_2 send their inputs (which may have been adaptively chosen based on crs) to \mathcal{B} . \mathcal{B} outputs (x_1, x_2) to $\text{Exp}_{(\text{Setup}_S, \bar{S}, \bar{S}^{-1}), \mathcal{A}}^{(b)\text{-a-crs-close}}(\lambda)$ and receives an output $z := y_b$, where $y_0 \leftarrow S(x_1, x_2)$ and $y_1 \leftarrow \bar{S}(x_1, x_2)$. \mathcal{B} sends z to P_1 and P_2 and continues to simulate the

interaction between \mathcal{Z} and Sim . This is possible, because the randomness which is actually used by \mathcal{T} is necessary for simulation. Finally, \mathcal{Z} halts and outputs a bit b' . \mathcal{B} outputs b' . \square

Let \mathcal{A} be an adversary breaking adaptive invertibility. We construct an environment \mathcal{Z} which distinguishes between the real execution of Π with adversary \mathcal{B} and the ideal execution with Sim and the ideal functionality \mathcal{F}' . Initially, \mathcal{Z} receives crs and calls \mathcal{A} on input of crs to obtain (x_1, x_2) . \mathcal{Z} uses x_1 as input for \mathcal{P}_1 and x_2 as input for \mathcal{P}_2 . After the execution of the protocol and the post-execution corruptions, \mathcal{Z} receives $(r_1, r_2), z$ (either by the adversary \mathcal{B} or the simulator Sim). Finally, \mathcal{Z} calls \mathcal{A} on input of $((r_1, r_2), z)$ and outputs \mathcal{A} 's output.

In the real world, z is the output of the real protocol and r_1, r_2 is the actual randomness used by the parties. Hence, by definition of $\bar{\mathcal{S}}$, (r_1, r_2) is the randomness actually used by $\bar{\mathcal{S}}$ to produce the output z . Therefore,

$$\Pr \left[\text{Exp}_{(\text{Setup}_{\mathcal{S}}, \bar{\mathcal{S}}, \bar{\mathcal{S}}^{-1}), \mathcal{A}}^{(0)\text{-a-crs-inv}}(\lambda) = 1 \right] = \Pr_{\text{crs}} \left[\mathcal{Z}^{\mathcal{B}, \Pi}(\text{crs}) = 1 \right].$$

In the ideal world, z is produced by the functionality \mathcal{F}' (hence, by $\bar{\mathcal{S}}$) and $(r_1, r_2) := (r'_1, r'_2)$ is produced by $\bar{\mathcal{S}}^{-1}((x_1, x_2), z)$. Hence,

$$\Pr \left[\text{Exp}_{(\text{Setup}_{\mathcal{S}}, \bar{\mathcal{S}}, \bar{\mathcal{S}}^{-1}), \mathcal{A}}^{(1)\text{-a-crs-inv}}(\lambda) = 1 \right] = \Pr_{\text{crs}} \left[\mathcal{Z}^{\text{Sim}, \mathcal{F}'}(\text{crs}) = 1 \right].$$

Hence,

$$\begin{aligned} & \text{Adv}_{(\text{Setup}_{\mathcal{S}}, \bar{\mathcal{S}}, \bar{\mathcal{S}}^{-1}), \mathcal{A}}^{\text{a-crs-inv}}(\lambda) \\ &= \left| \Pr_{\text{crs}} \left[\mathcal{Z}^{\mathcal{B}, \Pi}(\text{crs}) = 1 \right] - \Pr_{\text{crs}} \left[\mathcal{Z}^{\text{Sim}, \mathcal{F}'}(\text{crs}) = 1 \right] \right| \\ &\leq \left| \Pr_{\text{crs}} \left[\mathcal{Z}^{\mathcal{B}, \Pi}(\text{crs}) = 1 \right] - \Pr_{\text{crs}} \left[\mathcal{Z}^{\text{Sim}, \mathcal{F}}(\text{crs}) = 1 \right] \right| \\ &\quad + \left| \Pr_{\text{crs}} \left[\mathcal{Z}^{\text{Sim}, \mathcal{F}}(\text{crs}) = 1 \right] - \Pr_{\text{crs}} \left[\mathcal{Z}^{\text{Sim}, \mathcal{F}'}(\text{crs}) = 1 \right] \right| \end{aligned}$$

which is negligible by assumption and Claim 18.3. \square

Therefore, for every PPT algorithm \mathcal{S} , there exists an algorithm $\text{Setup}_{\mathcal{S}}$, an alternative sampler $\bar{\mathcal{S}}$ and an inverse sampler $\bar{\mathcal{S}}^{-1}$ such that adaptive closeness and adaptive invertibility hold. This concludes the proof. \square

Combining Theorems 18.2 and 18.3 and since semi-honest adaptive MPC in the non-programmable common random string model implies semi-honest adaptive MPC in the plain model, we obtain the following corollary improving Corollary 17.9.

COROLLARY 18.1. *If there exist extractable one-way function family ensembles (without auxiliary information) and indistinguishability obfuscation for all circuits, then weak $\text{cPREH}_{\approx_c}^{\text{rand}}$ with common random string (URC) is false.*

18.2 HONEY ENCRYPTION

Honey encryption schemes [JR14] are secret-key encryption schemes which offer security even against unbounded adversaries, i. e., even against adversaries who are able to perform an exhaustive search over all secret-keys.

This is a particularly useful notion when the key comes from a distribution with low min-entropy, like passwords. On decryption with a wrong key, a honey encryption scheme behaves indistinguishably from decryption with the actually used key.

Let D_1 be an efficiently samplable key distribution over the key space \mathcal{K}_λ and let D_2 be an efficiently samplable message distribution over the message space \mathcal{M}_λ .

DEFINITION 18.3 (Honey encryption, [JR14]). A secret-key encryption (SKE) scheme (Enc, Dec) is called *honey encryption (HE) scheme for key distribution D_1 and plaintext distribution D_2* if it satisfies the following property.

SECURITY AGAINST MESSAGE RECOVERY. The encryption scheme (Enc, Dec) is secure against message recovery (with respect to D_1, D_2) if for all unbounded adversaries \mathcal{A} , the advantage

$$\text{Adv}_{(\text{Enc}, \text{Dec}), \mathcal{A}, D_1, D_2}^{\text{mr}}(\lambda) = \Pr \left[\text{Exp}_{(\text{Enc}, \text{Dec}), \mathcal{A}, D_1, D_2}^{\text{mr}}(\lambda) = 1 \right]$$

is negligibly close to $2^{-\mu_1}$, where μ_1 is the min-entropy of D_1 (i. e., $\mu_1 = H_\infty(D_1)$), and $\text{Exp}_{(\text{Enc}, \text{Dec}), \mathcal{A}, D_1, D_2}^{\text{mr}}(\lambda)$ is defined as in Figure 18.3.

```

Expmr(Enc, Dec), A, D1, D2(λ)
k ← D1(1λ)
m* ← D2(1λ)
c* ← Enc(k, m*)
m ← A(c*)
return m = m*

```

Figure 18.3: Message recovery game.

DEFINITION 18.4 (Distribution-transforming encoder, [JR14; JRT16]). A *distribution-transforming encoder (DTE)* for a distribution sampled by a sampler S over $\mathcal{Y}_{S, \lambda}$ is a tuple of efficient algorithms (E_S, D_S) , where D_S is deterministic, such that the following properties are satisfied.

PERFECT CORRECTNESS. For all $y \in \mathcal{Y}_{S, \lambda}$, $\Pr [D_S(E_S(m)) = m] = 1$, where the probability is over the randomness of E_S .

DTE SECURITY. For all unbounded adversaries,

$$\text{Adv}_{(E_S, D_S), \mathcal{A}}^{\text{dte}}(\lambda) := \left| \Pr \left[\text{Exp}_{(E_S, D_S), \mathcal{A}}^{(0)\text{-dte}}(\lambda) = 1 \right] - \Pr \left[\text{Exp}_{(E_S, D_S), \mathcal{A}}^{(1)\text{-dte}}(\lambda) = 1 \right] \right|$$

is negligible, where $\text{Exp}_{(E_S, D_S), \mathcal{A}}^{(0)\text{-dte}}(\lambda)$ and $\text{Exp}_{(E_S, D_S), \mathcal{A}}^{(1)\text{-dte}}(\lambda)$ are defined in Figure 18.4.

The above notions are defined with respect to unbounded adversaries. This is a simplification meant to capture the fact that an adversary is able to perform work proportional to $2^{-\mu_1}$. For low-entropic key distributions, a computational variant of the above definition suffices. We refer to this as *computational honey encryption*.



(a) DTE game using an encoding of a sample. (b) DTE game using a uniform random string.

Figure 18.4: DTE games.

Relaxing the perfect correctness requirement from Definition 18.4 recovers the definition of (weak) $\text{PREH}_{\equiv_s}^{\text{rand}}$ (or weak $\text{PREH}_{\approx_c}^{\text{rand}}$ in the computational case) by Lemma 17.10 and still suffices to imply honey encryption via the DTE-then-Encrypt framework due to [JR14] given a secret-key encryption scheme (with message space matching the range of E_S). Applying the DTE-then-Encrypt framework on $\text{acPREH}_{\equiv_s}^{\text{rand}}$ and $\text{acPREH}_{\approx_c}^{\text{rand}}$ yields honey encryption in a CRS model for high-entropic and low-entropic key distributions, respectively. Note that the adaptive versions are only necessary when considering message distributions D_2 with input.

THEOREM 18.4. *If $\text{acPREH}_{\equiv_s}^{\text{rand}}$ (or $\text{acPREH}_{\approx_c}^{\text{rand}}$) is true and a suitable secret-key encryption scheme (modeled as random cipher) exists, then (computational) honey encryption exists for all efficiently samplable message distributions in the CRS model.*

The proof directly follows from [JR14]. Together with Theorems 17.13 and 17.14, we obtain the following corollary.

COROLLARY 18.2. *If polynomially secure IO exists and a (suitable) secret-key encryption scheme (modeled as random cipher) exists, then computational honey encryption exists for all efficiently samplable message distributions in the CRS model.*

18.3 KEYLESS STEGANOGRAPHY

Pseudorandom encodings yield a notion of keyless steganography. In the case of $\text{acPREH}_{\approx_c}^{\text{rand}}$ all parties need access to some public parameters, but none of the parties needs access to a secret key. We adopt the notation from [BL18b]. Since decoding does not involve any secret information, any attack corresponds to an equivalent of SS-CCA-security for public-key stegosystems (PKStS), where the decoding may even be applied on the challenge. Hence, the definition of SS-CCA-security must be adapted such that the message to be hidden is not chosen by the adversary but sampled according to a predefined message distribution. Let $dl(\lambda)$ be the document length and $ol(\lambda)$ be the output length. Let $\mathcal{C}_{\lambda, dl(\lambda)}$ be a channel which defines a probability distribution over $\Sigma^{dl(\lambda)}$ depending on the history $(\Sigma^{dl(\lambda)})^*$.

DEFINITION 18.5 (Keyless stegosystem for distribution S). *A keyless stegosystem (KlStS) for message distribution S is a triple of PPT algorithms (KlStS.Setup, KlStS. E_S , KlStS. D_S), where*

- KlStS.Setup(1^λ) produces public parameters pp (without corresponding secret information),

```

ExpKIStS, $\mathcal{A}$ , $\mathcal{C}$ (b)-klsts-sec( $\lambda$ )
pp  $\leftarrow$  KIStS.Setup( $1^\lambda$ )
m  $\leftarrow$   $\mathcal{A}$ (pp)
m*  $\leftarrow$  S(m)
hist*  $\leftarrow$   $\mathcal{A}$ ( $1^\lambda$ , pp, m*)
d*0  $\leftarrow$  KIStS.ES $\mathcal{C}$ (pp, m*, hist*)
d*1  $\leftarrow$   $\mathcal{C}_{\lambda, dl(\lambda), hist^*}^{ol(\lambda)}$ 
return  $\mathcal{A}$ (d*b)

```

Figure 18.5: Definition of the security game for keyless stegosystems.

- KIStS.E_S on input of pp , a message y sampled from S , a history $hist \in (\Sigma^{dl(\lambda)})^*$ and some state information $s \in \{0, 1\}^*$, produces a document $d \in \Sigma^{dl}$ (by being able to sample from $C_{\lambda, dl(\lambda)}$). KIStS.E_S ^{\mathcal{C}} ($pp, m, hist$) denotes sampling $ol(\lambda)$ documents using KIStS.E_S one-by-one.
- KIStS.D_S on input of pp , and a sequence of documents $d_1, \dots, d_{ol(\lambda)}$, outputs a message m' .

We require KIStS to meet the following properties.

UNIVERSALITY. KIStS works on any channel without prior knowledge of the distribution of the channel.

RELIABILITY. The probability

$$\Pr [\text{KIStS.D}_S(pp, \text{KIStS.E}_S(pp, m, hist), hist) \neq m]$$

is negligible, where the probability is over the choice of pp , m and the random coins of KIStS.E_S.

SECURITY. KIStS is secure (on channel \mathcal{C}), if for all PPT adversaries \mathcal{A} ,

$$\text{Adv}_{\text{KIStS}, \mathcal{A}, \mathcal{C}}^{\text{klsts-sec}}(\lambda) := \left| \Pr \left[\text{Exp}_{\text{KIStS}, \mathcal{A}, \mathcal{C}}^{(0)\text{-klsts-sec}}(\lambda) = 1 \right] - \Pr \left[\text{Exp}_{\text{KIStS}, \mathcal{A}, \mathcal{C}}^{(1)\text{-klsts-sec}}(\lambda) = 1 \right] \right|$$

is negligible, where $\text{Exp}_{\text{KIStS}, \mathcal{A}, \mathcal{C}}^{(0)\text{-klsts-sec}}(\lambda)$ and $\text{Exp}_{\text{KIStS}, \mathcal{A}, \mathcal{C}}^{(1)\text{-klsts-sec}}(\lambda)$ are defined in Figure 18.5.

Applying a similar strategy as in [vHo4; Hop05], we obtain the following theorem.

THEOREM 18.5. *Let S be an efficiently samplable message distribution and let \mathcal{H} be a family of pairwise independent hash functions. If $\text{acPREH}_{\approx_c}^{\text{rand}}$ is true for S , then $(\text{KIStS.Setup}_S, \text{KIStS.E}_S, \text{KIStS.D}_S)$ defined in Figure 18.6 is a keyless stegosystem for the message distribution S .*

The result basically follows from the Leftover Hash Lemma [HILL99] and the ability to embed the message distribution into the uniform distribution due to $\text{acPREH}_{\approx_c}^{\text{rand}}$. Note that since in $\text{Exp}_{\text{KIStS}, \mathcal{A}, \mathcal{C}}^{(b)\text{-klsts-sec}}(\lambda)$, the adversary does not know the challenge message m^* , it is not necessary that each encoding of a message corresponds to exactly one stegotext, see [Hop05]. The system is universal since it does not assume any knowledge on the channel.

```

KIStS.SetupS(1λ)
crs ← SetupS(1λ)
H ← ℋ
return pp := (crs, H)

```

(a) Setup algorithm.

```

KIStS.ES((crs, H), m)
u ← ES(crs, m)
u1 || ... || uℓ =: u
for i ∈ [ℓ] do
  do
    di ← ℂ(λ, hist)
  until prefixm(H(di)) = ui
return (d1, ..., dℓ)

```

(b) Encode algorithm.

```

KIStS.DS((crs, H), (d1, ..., dℓ))
for i ∈ [ℓ] do
  ui ← prefixm(H(di))
u ← u1 || ... || uℓ
return DS(crs, u)

```

(c) Decode algorithm.

Figure 18.6: Description of a keyless stegosystem from pseudorandom encodings inspired by [vHo4; Hop05].

18.4 DENIABLE ENCRYPTION

As noted in [CDNO97], sender deniable encryption is related to adaptively secure MPC (for adaptively well-formed functionalities). In this section, we analyze if $\text{cPREH}_{\approx_c}^{\text{rand}}$ together with the existence of a PKE scheme suffices to obtain deniable encryption.

Recall, that the explainability compiler of [DKR15] which is based on the deniable encryption scheme of [SW14] corresponds to $\text{cISH}_{\approx_c}^{\text{rand}}$, where closeness actually holds information-theoretically. Let cISH' denote this variant of $\text{cISH}_{\approx_c}^{\text{rand}}$, where closeness holds information-theoretically.

DEFINITION 18.6 (Publicly deniable encryption, [SW14]). A publicly deniable encryption (DE) scheme for message space \mathcal{M} is a tuple $\text{DE} = (\text{KGen}, \text{Enc}, \text{Dec}, \text{Explain})$ such that $(\text{KGen}, \text{Enc}, \text{Dec})$ is an IND-CPA secure PKE scheme and the following property is satisfied.

INDISTINGUISHABILITY OF EXPLANATION. For all PPT adversaries \mathcal{A} ,

$$\text{Adv}_{\text{DE}, \mathcal{A}}^{\text{ind-expl}}(\lambda) := \left| \Pr \left[\text{Exp}_{\text{DE}, \mathcal{A}}^{(0)\text{-ind-expl}}(\lambda) = 1 \right] - \Pr \left[\text{Exp}_{\text{DE}, \mathcal{A}}^{(1)\text{-ind-expl}}(\lambda) = 1 \right] \right|$$

is negligible, where $\text{Exp}_{\text{DE}, \mathcal{A}}^{(0)\text{-ind-expl}}(\lambda)$ and $\text{Exp}_{\text{DE}, \mathcal{A}}^{(1)\text{-ind-expl}}(\lambda)$ are defined in Figure 18.7 and the randomness space of Enc is $\{0, 1\}^{n(\lambda)}$.

We only consider publicly deniable encryption schemes with message space $\{0, 1\}$, since a deniable encryption scheme with message space $\{0, 1\}$ implies a publicly deniable encryption scheme for message space $\{0, 1\}^n$ (for a polynomial n in λ).

```

ExpDE,A(0)-ind-expl(λ)
(pk, sk) ← KGen(1λ)
m* ← A(pk)
u* ← {0, 1}n(λ)
c* ← Enc(pk, m*; u*)
return A(c*, u*)

```

(a) Adaptive indistinguishability of explanation game using true randomness.

```

ExpDE,A(1)-ind-expl(λ)
(pk, sk) ← KGen(1λ)
m* ← A(pk)
u* ← {0, 1}n(λ)
c* ← Enc(pk, m*; u*)
e* ← Explain(pk, c*)
return A(c*, e*)

```

(b) Adaptive indistinguishability of explanation game using explained randomness.

```

ExpDE,A(0)-ind-expl'(λ)
m* ← A(1λ)
(pk, sk) ← KGen(1λ)
u* ← {0, 1}n(λ)
c* ← Enc(pk, m*; u*)
return A(pk, c*, u*)

```

(c) Static indistinguishability of explanation game using true randomness.

```

ExpDE,A(1)-ind-expl'(λ)
m* ← A(1λ)
(pk, sk) ← KGen(1λ)
u* ← {0, 1}n(λ)
c* ← Enc(pk, m*; u*)
e* ← Explain(pk, c*)
return A(pk, c*, e*)

```

(d) Static indistinguishability of explanation game using explained randomness.

Figure 18.7: Indistinguishability of explanation games in its adaptive (Figures 18.7a and 18.7b) and static (Figures 18.7c and 18.7d) variant.

For message space $\{0, 1\}$, indistinguishability of explanation is equivalent to static indistinguishability of explanation, i. e., indistinguishability between $\text{Exp}_{\text{DE},A}^{(0)\text{-ind-expl}'}$ (λ) and $\text{Exp}_{\text{DE},A}^{(1)\text{-ind-expl}'}$ (λ) as defined in Figure 18.7.

clSH' in conjunction with a **PKE** scheme yields (publicly) sender deniable encryption.

THEOREM 18.6. *Let $(\text{KGen}, \text{Enc}, \text{Dec})$ be an **IND-CPA**-secure public-key encryption scheme for message space $\{0, 1\}$. If clSH' holds, then there exists a publicly deniable encryption scheme for message space $\{0, 1\}$.*

Proof. clSH' implies that for the **PPT** algorithm Enc there exist a setup algorithm $\text{Setup}_{\text{Enc}'}$, an alternative sampler $\overline{\text{Enc}}$ and an inverse sampler $\overline{\text{Enc}}^{-1}$ satisfying statistical closeness and computational invertibility. We define a publicly deniable encryption scheme $(\text{KGen}', \text{Enc}', \text{Dec}', \text{Explain}')$ in Figure 18.8.

CORRECTNESS. Let $m \in \{0, 1\}$ be a plaintext.

$$\epsilon_1 := \Pr \left[\begin{array}{l} ((\text{crs}, pk), sk) = (pk', sk') \leftarrow \text{KGen}'(1^\lambda) \\ c \leftarrow \text{Enc}(pk, m) \end{array} : \text{Dec}(sk, c) \neq m \right]$$

$$\epsilon_2 := \Pr \left[\begin{array}{l} ((\text{crs}, pk), sk) = (pk', sk') \leftarrow \text{KGen}'(1^\lambda) \\ c \leftarrow \overline{\text{Enc}}(pk, m) \end{array} : \text{Dec}(sk, c) \neq m \right]$$

```

KGen'( $1^\lambda$ )
   $crs \leftarrow \text{Setup}_{\text{Enc}}(1^\lambda)$ 
   $(pk, sk) \leftarrow \text{KGen}(1^\lambda)$ 
   $pk' := (crs, pk), sk' := sk$ 
  return  $(pk', sk')$ 

```

(a) The key generation algorithm.

```

Enc'( $pk', m$ )
  parse  $(crs, pk) = pk'$ 
   $c \leftarrow \overline{\text{Enc}}(crs, (pk, m))$ 
  return  $c$ 

```

(b) The encryption algorithm.

```

Dec'( $sk, c$ )
   $m \leftarrow \text{Dec}(sk, c)$ 
  return  $m$ 

```

(c) The decryption algorithm.

```

Explain'( $pk', m$ )
  parse  $(crs, pk) = pk'$ 
   $u \leftarrow \overline{\text{Enc}}^{-1}(crs, (pk, m), c)$ 
  return  $u$ 

```

(d) The explain algorithm.

Figure 18.8: Publicly deniable encryption scheme from cSH'.

Consider an *unbounded* adversary \mathcal{A} on closeness which on input of (crs, pk, c) , computes sk (using exhaustive search) and outputs 1 if $\text{Dec}(sk, c) \neq m$ and 0 otherwise. The advantage of this adversary is

$$\text{Adv}_{(\text{Setup}_{\text{Enc}}, \overline{\text{Enc}}, \overline{\text{Enc}}^{-1}), \mathcal{A}}^{\text{crs-close}}(\lambda) = |\epsilon_1 - \epsilon_2|.$$

Hence,

$$\epsilon_2 \leq \text{Adv}_{(\text{Setup}_{\text{Enc}}, \overline{\text{Enc}}, \overline{\text{Enc}}^{-1}), \mathcal{A}}^{\text{crs-close}}(\lambda) + \epsilon_1$$

and therefore negligible due to statistical closeness and correctness of $(\text{KGen}, \text{Enc}, \text{Dec})$.

IND-CPA SECURITY. For message space $\{0, 1\}$, **IND-CPA** security of $(\text{KGen}', \text{Enc}', \text{Dec}')$ is equivalent to the indistinguishability between the games \mathbf{G}_0 and \mathbf{G}_3 of Figure 18.9.

GAME \mathbf{G}_0 . \mathbf{G}_0 generates a key pair as in $\text{DE.KGen}(1^\lambda)$, encrypts 0 using DE.Enc and gives the public key and the ciphertext to the adversary \mathcal{A} .

GAME \mathbf{G}_1 . \mathbf{G}_1 is identical to \mathbf{G}_0 except that c^* is produced using the original sampler Enc instead of the alternative sampler $\overline{\text{Enc}}$.

CLAIM 18.4. For all (unbounded) adversaries \mathcal{A} there exists an (unbounded) adversary \mathcal{B}_1 such that

$$|\Pr[out_1 = 1] - \Pr[out_0 = 1]| \leq \text{Adv}_{(\text{Setup}_{\text{Enc}}, \overline{\text{Enc}}, \overline{\text{Enc}}^{-1}), \mathcal{B}_1}^{\text{crs-close}}(\lambda).$$

Proof. Let \mathcal{A} be an unbounded adversary distinguishing \mathbf{G}_0 and \mathbf{G}_1 . Construct an adversary \mathcal{B}_1 breaking closeness. Initially, \mathcal{B}_1 samples $(pk, sk) \leftarrow \text{KGen}(1^\lambda)$ and outputs $m := (pk, 0)$ to the experiment. In return, \mathcal{B}_1 receives

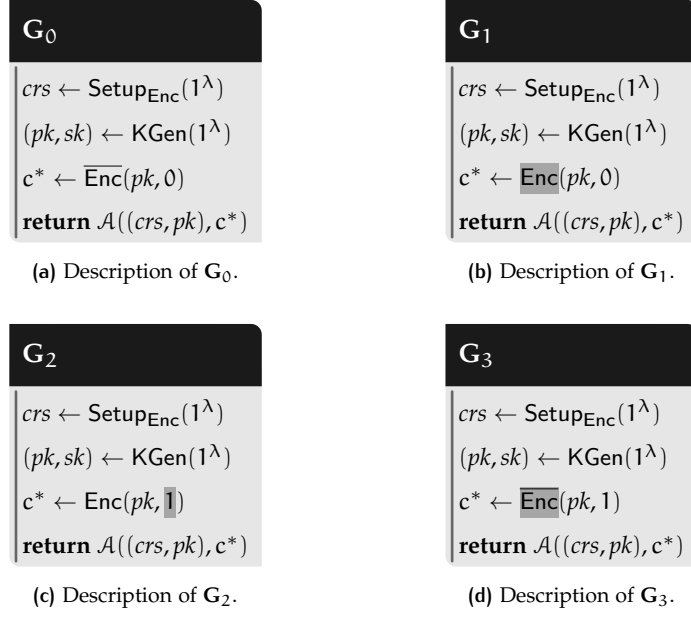


Figure 18.9: Hybrids used in the proof of IND-CPA security for Theorem 18.6.

(crs, y) , where y is either sampled using $\text{Enc}(m)$ or $\overline{\text{Enc}}(m)$. \mathcal{B}_1 calls \mathcal{A} on input of $((crs, pk), y)$. Hence,

$$\Pr[out_0 = 1] = \Pr \left[\text{Exp}_{(\text{Setup}_{\text{Enc}}, \overline{\text{Enc}}, \overline{\text{Enc}}^{-1})}^{(1)\text{-crs-close}}(\lambda) = 1 \right],$$

$$\Pr[out_1 = 1] = \Pr \left[\text{Exp}_{(\text{Setup}_{\text{Enc}}, \overline{\text{Enc}}, \overline{\text{Enc}}^{-1})}^{(0)\text{-crs-close}}(\lambda) = 1 \right].$$

□

GAME G_2 . G_2 is identical to G_1 except that Enc is used to encrypt 1 instead of 0. This game hop is justified by the IND-CPA security of $(\text{KGen}, \text{Enc}, \text{Dec})$.

GAME G_3 . G_3 is identical to G_2 except that c^* is produced using the alternative sampler $\overline{\text{Enc}}$ instead of the original sampler Enc . This game hop is justified by statistical closeness like the game hop from G_0 to G_1 .

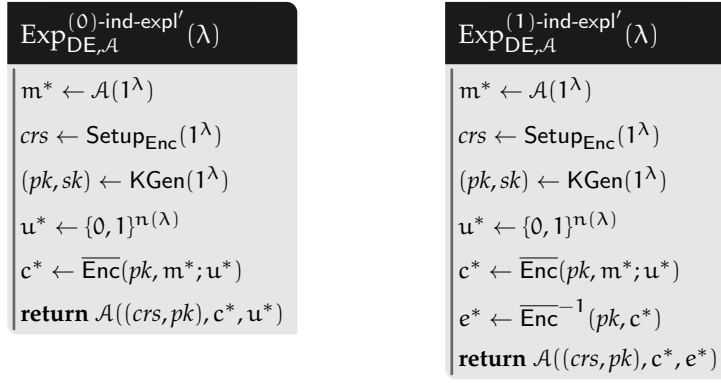
INDISTINGUISHABILITY OF EXPLANATION. We prove indistinguishability between the games described in Figure 18.10.

CLAIM 18.5. For all PPT adversaries \mathcal{A} , $\text{Adv}_{\text{DE}, \mathcal{A}}^{\text{ind-expl}'}(\lambda)$ is negligible.

Proof. Let \mathcal{A} be an adversary distinguishing between the two experiments $\text{Exp}_{\text{DE}, \mathcal{A}}^{(0)\text{-ind-expl}'}$ and $\text{Exp}_{\text{DE}, \mathcal{A}}^{(1)\text{-ind-expl}'}$ as above. Construct an adversary \mathcal{B} breaking invertibility. Initially, \mathcal{B} calls \mathcal{A} , obtains m^* , samples $(pk, sk) \leftarrow \text{KGen}(1^\lambda)$ and outputs $m := (pk, m^*)$ to the experiment. In return, \mathcal{B} receives (crs, r, y) from the experiment, where y is sampled using $\overline{\text{Enc}}(pk, m^*)$ and r is either the randomness used by $\overline{\text{Enc}}$ or is sampled via $\overline{\text{Enc}}^{-1}(pk, m^*, y)$. \mathcal{B} calls \mathcal{A} on input of $((crs, pk), y, r)$. Hence,

$$\Pr \left[\text{Exp}_{\text{DE}, \mathcal{A}}^{(b)\text{-ind-expl}'}(\lambda) = 1 \right] = \Pr \left[\text{Exp}_{(\text{Setup}_{\text{Enc}}, \overline{\text{Enc}}, \overline{\text{Enc}}^{-1})}^{(b)\text{-crs-inv}}(\lambda) = 1 \right]$$

for $b \in \{0, 1\}$. □



(a) Indistinguishability of explanation game using true randomness. (b) Indistinguishability of explanation game using explained randomness.

Figure 18.10: Unwrapped static indistinguishability of explanation games used in the proof of indistinguishability of explanation of Theorem 18.6.

□

Whether $\text{cISH}_{\approx_c}^{\text{rand}}$ (without statistical closeness) implies deniable encryption remains open.

18.5 NON-COMMITTING ENCRYPTION

Non-committing encryption is a powerful notion which is known to imply adaptive MPC for well-formed functionalities [CFGN96; CLOS02].

On a high level, if $\text{acPREH}_{\approx_c}^{\text{rand}}$ is true, then any IND-CPA secure PKE scheme is a simulatable PKE scheme [DN00] (in the CRS model). The strategy of [CDMW09] translates to the CRS model (where inputs may be chosen depending on the CRS) directly.

DEFINITION 18.7 (Non-committing bit encryption in the global CRS model). A non-committing bit encryption (NCE) scheme in the global CRS model is a tuple of PPT algorithms (Setup, KGen, Enc, Dec, Sim), such that (KGen, Enc, Dec) is a PKE scheme and the following distributions are computationally indistinguishable:

$$\left\{ \begin{array}{l} crs \leftarrow \text{Setup}(1^\lambda) \\ (pk, sk) \leftarrow \text{KGen}(1^\lambda, crs; r_{\text{KGen}}) : (crs, pk, c, r_{\text{KGen}}, r_{\text{Enc}}) \\ c \leftarrow \text{Enc}(crs, pk, b; r_{\text{Enc}}) \end{array} \right\},$$

$$\left\{ \begin{array}{l} crs \leftarrow \text{Setup}(1^\lambda) \\ (pk, c, r_{\text{KGen}}^0, r_{\text{KGen}}^1, r_{\text{Enc}}^0, r_{\text{Enc}}^1) \leftarrow \text{Sim}(1^\lambda, crs) : (crs, pk, c, r_{\text{KGen}}^b, r_{\text{Enc}}^b) \end{array} \right\}.$$

Following the lines of [CDMW09], if $\text{acPREH}_{\approx_c}^{\text{rand}}$ is true and there exists an IND-CPA secure PKE scheme, then there exists a non-committing encryption scheme.

THEOREM 18.7. Let $(\text{KGen}', \text{Enc}', \text{Dec}')$ be an IND-CPA secure PKE scheme for message space $\{0, 1\}^\lambda$. If $\text{acPREH}_{\approx_c}^{\text{rand}}$ holds, then there exists a non-committing encryption scheme.

```

KGen( $1^\lambda, crs$ )
 $M_0, M_1 \leftarrow \{0, 1\}^\lambda$ 
 $T \leftarrow [4\lambda]$  s.t.  $|T| = \lambda$ 
 $(pk_i, sk_i) \leftarrow \begin{cases} \text{KGen}'(1^\lambda) & \text{if } i \in T \\ \overline{\text{OGen}}(crs) & \text{otherwise} \end{cases}$ 
 $pk := (M_0, M_1, pk_1, \dots, pk_{4\lambda})$ 
 $sk := (T, (sk_i)_{i \in T})$ 
return  $(pk, sk)$ 

```

(a) The key generation algorithm.

```

Enc( $crs, pk, b$ )
 $S \leftarrow [4\lambda]$  s.t.  $|S| = \lambda$ 
 $c_i \leftarrow \begin{cases} \text{Enc}'(pk_i, M_b) & \text{if } i \in S \\ \overline{\text{OEnc}}(crs, pk_i) & \text{otherwise} \end{cases}$ 
return  $c := (c_i)_{i \in [4\lambda]}$ 

```

(b) The encryption algorithm.

```

Dec( $crs, sk, c$ )
 $J := \{\text{Dec}'(sk_i, c_i) \mid i \in T\}$ 
if  $M_0 \in J$  then
  return 0
else
  return 1

```

(c) The decryption algorithm.

Figure 18.11: Non-committing encryption scheme in the global CRS model based on [CDMW09].

Proof sketch. Let OGen be the algorithm which on input of 1^λ calls $\text{KGen}'(1^\lambda)$ and outputs only (pk, \perp) . Further let OEnc be the algorithm which on input of $(1^\lambda, pk)$ samples $m \leftarrow \{0, 1\}^\lambda$ and outputs $c \leftarrow \text{Enc}'(1^\lambda, pk, m)$. If $\text{acPREH}_{\approx_c}^{\text{rand}}$ is true, then there exist alternative and inverse samplers for OGen and OEnc , denoted by $(\overline{\text{OGen}}, \overline{\text{OGen}}^{-1})$ and $(\overline{\text{OEnc}}, \overline{\text{OEnc}}^{-1})$, respectively. Note that the alternative sampler and the inverse sampler need access to the CRS. If the setup algorithm Setup is trivial (i.e., outputs the empty string), this corresponds to the notion of simulatable encryption due to [DN00] and yields non-committing encryption directly due to [CDMW09]. Figures 18.11 and 18.12 show the construction of non-committing encryption in the global CRS model. Note that assuming *adaptive* $\text{acPREH}_{\approx_c}^{\text{rand}}$ is necessary since the inputs to the sampler OEnc are sampled via $\text{OGen}(crs)$ during the simulation and, hence, depend on the CRS. The indistinguishability between the real and the simulated distribution follows the same ideas as in [CDMW09]. \square

18.6 SUPER-POLYNOMIAL ENCODING

Extremely lossy functions due to [Zha16] are functions which can be set up in two computationally indistinguishable modes – an injective mode and a extremely lossy mode, where the range of the function is merely polynomial, see Definition 4.2 in Section 4.2. A slight relaxation of this notion is what we call very lossy functions (VLFs). The difference to ELFs is that we require


```

Sim(crs)
 $M_0, M_1 \leftarrow \{0, 1\}^\lambda$ 
 $S_0, T_0 \leftarrow [4\lambda]$  s.t.  $|S_0| = |T_0| = \lambda$ 
 $S_1, T_1 \leftarrow [4\lambda] \setminus (S_0 \cup T_0)$  s.t.  $|S_0 \cap T_0| = |S_1 \cap T_1|$ 
 $(pk_i, sk_i) \leftarrow \begin{cases} \text{KGen}'(1^\lambda; r_{\text{KGen}'})^{(i)} & \text{if } i \in T_0 \cup S_0 \cup T_1 \cup S_1 \\ \overline{\text{OGen}}(crs; r_{\text{OGen}})^{(i)} & \text{otherwise} \end{cases}$ 
 $c_i \leftarrow \begin{cases} \text{Enc}'(pk_i, M_0; r_{\text{Enc}'})^{(i)} & \text{if } i \in S_0 \\ \text{Enc}'(pk_i, M_1; r_{\text{Enc}'})^{(i)} & \text{if } i \in S_1 \\ \overline{\text{OEnc}}(crs, pk_i; r_{\text{OEnc}})^{(i)} & \text{otherwise} \end{cases}$ 
define  $r_{\text{KGen}}^b := (T_b, (u_{\text{KGen}'}^{(b,i)})_{i \in [4\lambda]})$  and  $r_{\text{Enc}}^b := (S_b, (u_{\text{Enc}'}^{(b,i)})_{i \in [4\lambda]})$ 
 $u_{\text{KGen}'}^{(b,i)} \leftarrow \begin{cases} r_{\text{KGen}'}^{(i)} & \text{if } i \in T_b \\ \overline{\text{OGen}}^{-1}(crs, pk_i) & \text{if } i \in (T_0 \cup T_1 \cup S_0 \cup S_1) \setminus T_b \\ r_{\text{OGen}}^{(i)} & \text{otherwise} \end{cases}$ 
 $u_{\text{Enc}'}^{(b,i)} \leftarrow \begin{cases} r_{\text{Enc}'}^{(i)} & \text{if } i \in S_b \\ \overline{\text{OEnc}}^{-1}(crs, (pk_i, M_{1-b}), c_i) & \text{if } i \in S_{1-b} \\ r_{\text{OEnc}}^{(i)} & \text{otherwise} \end{cases}$ 
return  $(pk := (M_0, M_1, (pk_i)_{i \in [4\lambda]}), c := (c_i)_{i \in [4\lambda]}, r_{\text{KGen}}^0, r_{\text{KGen}}^1, r_{\text{Enc}}^0, r_{\text{Enc}}^1)$ 

```

Figure 18.12: The simulator Sim for the non-committing encryption scheme in the global CRS model described in Figure 18.11. Sim is based on the simulator given in [CDMW09].

indistinguishability between functions with exponential range and *super-polynomial* range. The existence of these functions implies a relaxation of $\text{cPREH}_{\approx_c}^{\text{rand}}$. The definition of VLFs is inspired by the definition of bounded adversary ELFs from [Zha16].

DEFINITION 18.8 (Very lossy functions). A *very lossy function* (VLF) is a tuple of PPT algorithms $\text{VLF} = (\text{Gen}, \text{Eval})$ and a computable function $N(M)$ (such that $\log N$ is polynomial in $\log M$) such that the following properties are satisfied.

- Gen on input of M , $r \in [M]$ and a flag $b \in \{\text{inj}, \text{lossy}\}$, outputs a function description $G: [M] \rightarrow [N]$.
- If $b = \text{inj}$, G is injective with overwhelming probability (in $\log M$).
- For all $r \in [M]$, if $b = \text{lossy}$, $|G([M])| \leq r$ with overwhelming probability (in $\log M$).
- $\text{Eval}(G, \cdot)$ evaluates the function G in polynomial time. A very lossy function is required to satisfy the following property.

INDISTINGUISHABILITY. There exists a *super-polynomial* function q such that for all PPT adversaries \mathcal{A} and any $r \in [q(\log M), M]$,

$$\left| \Pr[\mathcal{A}(\text{VLF.Gen}(M, r, \text{inj})) = 1] - \Pr[\mathcal{A}(\text{VLF.Gen}(M, r, \text{lossy})) = 1] \right|$$

is negligible.

As a shorthand notation, we often write $G(x)$ to denote $\text{VLF.Eval}(G, x)$.

DEFINITION 18.9 (Strong regularity, [Zha16]). A very lossy function VLF is *strongly regular* if for all $r \in [M]$, with overwhelming probability over the choice of $G \leftarrow \text{VLF.Gen}(M, r)$, the distribution $\{x \leftarrow [M] : G(x)\}$ is statistically close to the uniform distribution over $G([M])$.

A VLF is called *strongly efficiently enumerable*, if there exists a (potentially randomized) algorithm running in polynomial time in $\log M$ and r , which given $G \leftarrow \text{VLF.Gen}(M, r, \text{lossy})$ and r outputs a set $S \subseteq [N]$ such that with overwhelming probability (over the choice of G and the randomness of the algorithm), we have $S = G([M])$. If VLF is strongly regular, then it also is strongly efficiently enumerable, [Zha16]. Zhandry [Zha16] shows that strongly regular ELFs are implied by the exponential decisional Diffie-Hellman (DDH) assumption.

We note that in contrast to ELFs, for VLFs an adversary who learns r does not harm security.

Assuming the subexponential hardness of the decisional Diffie-Hellman (DDH) problem, the bounded adversary extremely lossy function instantiation from [Zha16] is a strongly regular very lossy function according to Definitions 18.8 and 18.9.

THEOREM 18.8. *If strongly regular very lossy functions exist, then $\text{acPREH}_{\approx_c}^{\text{rand}}$ with super-polynomial encoding is true.*

Proof. Let S be a PPT sampler with input space L . For $m \in L$, let $\ell_r(|m|)$ denote the polynomial which upper bounds the number of random bits $S(m)$ takes, i. e., the random tape of S is uniform over $\{0, 1\}^{\ell_r} = [2^{\ell_r}]$. Further, let ξ be a super-polynomial function. The setup, the encoding and the decoding algorithms are defined in Figure 18.13. Let $\bar{S}(crs, m; r) := D_S(crs, m, r)$ and $\bar{S}^{-1}(crs, m, y) := E_S(crs, m, y)$ be the corresponding alternative and inverse sampler as in the proof of Theorem 16.1. We prove the equivalent properties closeness and invertibility.

Setup($1^\lambda, 2^{\ell_r}$)

$G \leftarrow \text{VLF.Gen}(2^{\ell_r}, \xi, \text{lossy})$
return $crs := G$

(a) The setup algorithm.

$D_S(crs, m, r)$

$y \leftarrow S(m; G(r))$
return y

(b) The decoding algorithm.

$E_S(crs, m, y)$

$\mathcal{R} := \emptyset$
for $r' \in \text{image}(G)$ **do**
 if $S(m; G(r')) = y$ **then**
 $\mathcal{R} := \mathcal{R} \cup r'$
 $r \leftarrow \mathcal{R}$
return r

(c) The encoding algorithm.

Figure 18.13: Description of the pseudorandom encoding scheme with super-polynomial time encoding algorithm.

INVERTIBILITY. The encoding algorithm produces perfectly distributed inverse sampled random tapes. Hence, adaptive invertibility follows.

CLOSENESS. We start from the game $\text{Exp}_{(\text{Setup}, \overline{S}, \overline{S}^{-1}), \mathcal{A}}^{(b)\text{-a-crs-close}}(\lambda)$ and switch the VLF to injective mode $G \leftarrow \text{VLF.Gen}(2^{\ell_r}, \xi, \text{inj})$. This is computationally indistinguishable for the adversary. The inverse sampler will not work anymore, but since the adversary is polynomially bounded, he cannot call the inverse sampler anyway. Strong regularity implies that for $G \leftarrow \text{VLF.Gen}(2^{\ell_r}, \xi, \text{inj})$ the distribution $\{x \leftarrow [2^{\ell_r}]: G(x)\}$ is statistically close to uniform distribution over $G([2^{\ell_r}])$. Hence, $S(m; G(r))$ and $S(m; r)$ for uniform r from $[2^{\ell_r}]$ are statistically close. Then, the view of the adversary is identical regardless if we started from $\text{Exp}_{(\text{Setup}, \overline{S}, \overline{S}^{-1}), \mathcal{A}}^{(0)\text{-a-crs-close}}(\lambda)$ or $\text{Exp}_{(\text{Setup}, \overline{S}, \overline{S}^{-1}), \mathcal{A}}^{(1)\text{-a-crs-close}}(\lambda)$. \square

Unfortunately, $\text{acPREH}_{\approx_c}^{\text{rand}}$ with super-polynomial encoding algorithm (or, equivalently $\text{acISH}_{\approx_c}^{\text{rand}}$ with super-polynomial inverse sampler) does not suffice to imply adaptive MPC even if the simulator is allowed to run in super-polynomial time, [Paso3]. This is because plugging $\text{cPREH}_{\approx_c}^{\text{rand}}$ with super-polynomial encoding algorithm into the proof of Theorem 18.2, the game hop from G_2 to G_3 cannot be made, since the simulation of these games requires super-polynomial time and, hence, a reduction to closeness against PPT adversaries is not possible. We do, however, obtain a non-standard notion of adaptive MPC with super-polynomial simulation, namely for any functionality \mathcal{F} we are able to adaptively realize a functionality $\overline{\mathcal{F}}$ which produces a computationally (against PPT adversaries) indistinguishable output distribution to the original functionality. We view it as an interesting problem to further study this notion of adaptive MPC.

OUTLOOK AND FURTHER DIRECTIONS

Finally, we provide an outlook on further interesting research directions.

USE OUR DP10 FRAMEWORK IN MORE SETTINGS. In Part I and [ACH20], we identify fully homomorphic encryption and spooky encryption to fit our dp10 framework and thereby improve the corresponding security proofs removing the need for subexponentially secure IO. Later, Döttling and Nishimaki [DN18] utilize our framework to improve the security reduction of their novel notion of universal proxy re-encryption. It seems plausible that our framework can be used to improve the security reduction of further primitives which require probabilistic IO. For instance, the IND-CCA₁-secure fully homomorphic encryption scheme of [CRRV17] could be a good starting point.

FIND MORE APPLICATIONS FOR THE ALGEBRAIC WRAPPER. Since its introduction in [FKL18], the algebraic group model (AGM) enjoys an increasing popularity [FPS20; BFL20; RS20; KLRX20; GT20]. In Part II and [AHK20], we demonstrate how to implement AGM-based proofs from [FKL18] and [FPS20] in the standard (group) model. Further, we identify properties used in some proofs which exceed the capabilities of the algebraic wrapper. It would be interesting to identify further instances, where the extraction properties provided by the algebraic wrapper suffice. This would transport the corresponding proofs into the standard (group) model while avoiding knowledge-type assumptions.

NOVEL CONSTRUCTIONS OF IO. In Part III which is based on [ACI+20], we obtain partial evidence that indistinguishability obfuscation (or at least an assumption which stands in conflict with extractable one-way functions with unbounded auxiliary input) is necessary for pseudorandom encoding schemes for all efficiently samplable distributions. Hence, we raised the question whether IO is indeed necessary for pseudorandom encodings.

Inspired by [ACI+20], Wee and Wichs [WW20] answer this question in the affirmative assuming the subexponential hardness of the learning-with-errors (LWE) problem. This opens up new opportunities towards novel constructions of indistinguishability obfuscation. On the one hand, one can seek to improve the construction of [WW20]. On the other hand, the question whether IO is necessary for pseudorandom encoding schemes remains open relative to other hardness assumptions such as the LPN assumption or assumptions in cryptographic groups. We believe this will give rise to some new roads towards plausible IO candidates.

EXTEND COVERT COMPUTATION – HOW FAR CAN WE GO? Our notion of pseudorandom encodings from Part III has connections to variants of steganography and covert computation [ACI+20]. Covert computation, introduced in [vHL05; CGOS07], allows to hide the very fact that a multi-party

protocol is being executed and only reveals this fact (together with the protocol output) if an admissibility function is satisfied and all present parties actually participated. This admissibility function models that all participating parties are satisfied with the outcome of the protocol.

If a covert multi-party protocol terminates successfully, all participating parties learn not only their output but also the fact that all other parties actively participated. A natural question is whether we can transcend this notion of multi-party computation to hide participation even after a successful run. Whether MPC with deniable participation is feasible is an intriguing question currently under investigation. As a first step we identify “anonymous (message) transfer” (AT) as an interesting special case. An AT protocol considers three players: one receiver, one sender and one non-participant who only outputs random noise. The goal of AT is to deliver a message from the sender to the receiver in a way that leaves the receiver unable to determine which of the two other parties was the sender. Whether such an anonymous transfer protocol can exist is a very interesting question.

This is ongoing work together with Geoffroy Couteau, Sven Maier and Rafael Pass.

BIBLIOGRAPHY

- [ABR01] Michel Abdalla, Mihir Bellare, and Phillip Rogaway. “The Oracle Diffie-Hellman Assumptions and an Analysis of DHIES.” In: *Topics in Cryptology – CT-RSA 2001*. Ed. by David Naccache. Vol. 2020. Lecture Notes in Computer Science. San Francisco, CA, USA: Springer, Heidelberg, Germany, Apr. 2001, pp. 143–158. DOI: [10.1007/3-540-45353-9_12](https://doi.org/10.1007/3-540-45353-9_12) (cit. on p. 129).
- [Agr19] Shweta Agrawal. “Indistinguishability Obfuscation Without Multilinear Maps: New Methods for Bootstrapping and Instantiation.” In: *Advances in Cryptology – EUROCRYPT 2019, Part I*. Ed. by Yuval Ishai and Vincent Rijmen. Vol. 11476. Lecture Notes in Computer Science. Darmstadt, Germany: Springer, Heidelberg, Germany, May 2019, pp. 191–225. DOI: [10.1007/978-3-030-17653-2_7](https://doi.org/10.1007/978-3-030-17653-2_7) (cit. on pp. 2, 41).
- [AP20] Shweta Agrawal and Alice Pellet-Mary. “Indistinguishability Obfuscation Without Maps: Attacks and Fixes for Noisy Linear FE.” In: *Advances in Cryptology – EUROCRYPT 2020, Part I*. Ed. by Anne Canteaut and Yuval Ishai. Vol. 12105. Lecture Notes in Computer Science. Zagreb, Croatia: Springer, Heidelberg, Germany, May 2020, pp. 110–140. DOI: [10.1007/978-3-030-45721-1_5](https://doi.org/10.1007/978-3-030-45721-1_5) (cit. on pp. 2, 41).
- [Ajt96] Miklós Ajtai. “Generating Hard Instances of Lattice Problems (Extended Abstract).” In: *28th Annual ACM Symposium on Theory of Computing*. Philadelphia, PA, USA: ACM Press, May 1996, pp. 99–108. DOI: [10.1145/237814.237838](https://doi.org/10.1145/237814.237838) (cit. on p. 2).
- [AMP20] Navid Alamati, Hart Montgomery, and Sikhar Patranabis. *Stronger Multilinear Maps from Indistinguishability Obfuscation*. Cryptology ePrint Archive, Report 2020/610. <https://eprint.iacr.org/2020/610>. 2020 (cit. on pp. 91, 92, 110).
- [AFH+16] Martin R. Albrecht, Pooya Farshim, Dennis Hofheinz, Enrique Larraia, and Kenneth G. Paterson. “Multilinear Maps from Obfuscation.” In: *TCC 2016-A: 13th Theory of Cryptography Conference, Part I*. Ed. by Eyal Kushilevitz and Tal Malkin. Vol. 9562. Lecture Notes in Computer Science. Tel Aviv, Israel: Springer, Heidelberg, Germany, Jan. 2016, pp. 446–473. DOI: [10.1007/978-3-662-49096-9_19](https://doi.org/10.1007/978-3-662-49096-9_19) (cit. on pp. 11, 90–92, 95, 110, 111, 115, 116, 120).
- [ABG+13] Prabhanjan Ananth, Dan Boneh, Sanjam Garg, Amit Sahai, and Mark Zhandry. *Differing-Inputs Obfuscation and Applications*. Cryptology ePrint Archive, Report 2013/689. <http://eprint.iacr.org/2013/689>. 2013 (cit. on pp. 27, 41, 45).
- [AJL+19] Prabhanjan Ananth, Aayush Jain, Huijia Lin, Christian Matt, and Amit Sahai. “Indistinguishability Obfuscation Without Multilinear Maps: New Paradigms via Low Degree Weak Pseudorandomness and Security Amplification.” In: *Advances*

- in Cryptology – CRYPTO 2019, Part III*. Ed. by Alexandra Boldyreva and Daniele Micciancio. Vol. 11694. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2019, pp. 284–332. DOI: [10.1007/978-3-030-26954-8_10](https://doi.org/10.1007/978-3-030-26954-8_10) (cit. on pp. 2, 41).
- [A]15] Prabhanjan Ananth and Abhishek Jain. “Indistinguishability Obfuscation from Compact Functional Encryption.” In: *Advances in Cryptology – CRYPTO 2015, Part I*. Ed. by Rosario Gennaro and Matthew J. B. Robshaw. Vol. 9215. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2015, pp. 308–326. DOI: [10.1007/978-3-662-47989-6_15](https://doi.org/10.1007/978-3-662-47989-6_15) (cit. on p. 41).
- [AS17] Prabhanjan Ananth and Amit Sahai. “Projective Arithmetic Functional Encryption and Indistinguishability Obfuscation from Degree-5 Multilinear Maps.” In: *Advances in Cryptology – EUROCRYPT 2017, Part I*. Ed. by Jean-Sébastien Coron and Jesper Buus Nielsen. Vol. 10210. Lecture Notes in Computer Science. Paris, France: Springer, Heidelberg, Germany, Apr. 2017, pp. 152–181. DOI: [10.1007/978-3-319-56620-7_6](https://doi.org/10.1007/978-3-319-56620-7_6) (cit. on p. 41).
- [ADGM17] Daniel Apon, Nico Döttling, Sanjam Garg, and Pratyay Mukherjee. “Cryptanalysis of Indistinguishability Obfuscations of Circuits over GGH_{13} .” In: *ICALP 2017: 44th International Colloquium on Automata, Languages and Programming*. Ed. by Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl. Vol. 80. LIPIcs. Warsaw, Poland: Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, July 2017, 38:1–38:16. DOI: [10.4230/LIPIcs.ICALP.2017.38](https://doi.org/10.4230/LIPIcs.ICALP.2017.38) (cit. on p. 2).
- [BMSZ16] Saikrishna Badrinarayanan, Eric Miles, Amit Sahai, and Mark Zhandry. “Post-zeroizing Obfuscation: New Mathematical Tools, and the Case of Evasive Circuits.” In: *Advances in Cryptology – EUROCRYPT 2016, Part II*. Ed. by Marc Fischlin and Jean-Sébastien Coron. Vol. 9666. Lecture Notes in Computer Science. Vienna, Austria: Springer, Heidelberg, Germany, May 2016, pp. 764–791. DOI: [10.1007/978-3-662-49896-5_27](https://doi.org/10.1007/978-3-662-49896-5_27) (cit. on pp. 2, 5).
- [Bar01] Boaz Barak. “How to Go Beyond the Black-Box Simulation Barrier.” In: *42nd Annual Symposium on Foundations of Computer Science*. Las Vegas, NV, USA: IEEE Computer Society Press, Oct. 2001, pp. 106–115. DOI: [10.1109/SFCS.2001.959885](https://doi.org/10.1109/SFCS.2001.959885) (cit. on p. 32).
- [BGI+01] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. “On the (Im)possibility of Obfuscating Programs.” In: *Advances in Cryptology – CRYPTO 2001*. Ed. by Joe Kilian. Vol. 2139. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2001, pp. 1–18. DOI: [10.1007/3-540-44647-8_1](https://doi.org/10.1007/3-540-44647-8_1) (cit. on pp. 25, 41, 45).

- [BGI+12] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. “On the (im)possibility of obfuscating programs.” In: *J. ACM* 59.2 (2012), 6:1–6:48. DOI: [10.1145/2160158.2160159](https://doi.org/10.1145/2160158.2160159). URL: <https://doi.org/10.1145/2160158.2160159> (cit. on pp. 2, 25, 27).
- [BHHI10] Boaz Barak, Iftach Haitner, Dennis Hofheinz, and Yuval Ishai. “Bounded Key-Dependent Message Security.” In: *Advances in Cryptology – EUROCRYPT 2010*. Ed. by Henri Gilbert. Vol. 6110. Lecture Notes in Computer Science. French Riviera: Springer, Heidelberg, Germany, May 2010, pp. 423–444. DOI: [10.1007/978-3-642-13190-5_22](https://doi.org/10.1007/978-3-642-13190-5_22) (cit. on p. 46).
- [BSW03] Boaz Barak, Ronen Shaltiel, and Avi Wigderson. “Computational Analogues of Entropy.” In: *Approximation, Randomization, and Combinatorial Optimization: Algorithms and Techniques, 6th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems, APPROX 2003 and 7th International Workshop on Randomization and Approximation Techniques in Computer Science, RANDOM 2003, Princeton, NJ, USA, August 24–26, 2003, Proceedings*. Ed. by Sanjeev Arora, Klaus Jansen, José D. P. Rolim, and Amit Sahai. Vol. 2764. Lecture Notes in Computer Science. Springer, 2003, pp. 200–215. DOI: [10.1007/978-3-540-45198-3_18](https://doi.org/10.1007/978-3-540-45198-3_18). URL: https://doi.org/10.1007/978-3-540-45198-3_18 (cit. on p. 190).
- [BJ+20] James Bartusek, Yuval Ishai, Aayush Jain, Fermi Ma, Amit Sahai, and Mark Zhandry. “Affine Determinant Programs: A Framework for Obfuscation and Witness Encryption.” In: *ITCS 2020: 11th Innovations in Theoretical Computer Science Conference*. Ed. by Thomas Vidick. Vol. 151. Seattle, WA, USA: LIPIcs, Jan. 2020, 82:1–82:39. DOI: [10.4230/LIPIcs.ITCS.2020.82](https://doi.org/10.4230/LIPIcs.ITCS.2020.82) (cit. on pp. 3, 5, 41).
- [BFL20] Balthazar Bauer, Georg Fuchsbauer, and Julian Loss. “A Classification of Computational Assumptions in the Algebraic Group Model.” In: *Advances in Cryptology – CRYPTO 2020, Part II*. Ed. by Daniele Micciancio and Thomas Ristenpart. Vol. 12171. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2020, pp. 121–151. DOI: [10.1007/978-3-030-56880-1_5](https://doi.org/10.1007/978-3-030-56880-1_5) (cit. on p. 245).
- [Bea97] Donald Beaver. “Plug and Play Encryption.” In: *Advances in Cryptology – CRYPTO’97*. Ed. by Burton S. Kaliski Jr. Vol. 1294. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 1997, pp. 75–89. DOI: [10.1007/BFb0052228](https://doi.org/10.1007/BFb0052228) (cit. on pp. 8, 16).
- [BH93] Donald Beaver and Stuart Haber. “Cryptographic Protocols Provably Secure Against Dynamic Adversaries.” In: *Advances in Cryptology – EUROCRYPT’92*. Ed. by Rainer A. Rueppel. Vol. 658. Lecture Notes in Computer Science. Balatonfüred, Hungary: Springer, Heidelberg, Germany, May 1993, pp. 307–323. DOI: [10.1007/3-540-47555-9_26](https://doi.org/10.1007/3-540-47555-9_26) (cit. on pp. 16, 154).

- [BP04] Mihir Bellare and Adriana Palacio. “The Knowledge-of-Exponent Assumptions and 3-Round Zero-Knowledge Protocols.” In: *Advances in Cryptology – CRYPTO 2004*. Ed. by Matthew Franklin. Vol. 3152. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2004, pp. 273–289. DOI: [10.1007/978-3-540-28628-8_17](https://doi.org/10.1007/978-3-540-28628-8_17) (cit. on pp. [32](#), [87](#)).
- [BR93] Mihir Bellare and Phillip Rogaway. “Random Oracles are Practical: A Paradigm for Designing Efficient Protocols.” In: *ACM CCS 93: 1st Conference on Computer and Communications Security*. Ed. by Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, Ravi S. Sandhu, and Victoria Ashby. Fairfax, Virginia, USA: ACM Press, Nov. 1993, pp. 62–73. DOI: [10.1145/168588.168596](https://doi.org/10.1145/168588.168596) (cit. on p. [10](#)).
- [BR95] Mihir Bellare and Phillip Rogaway. “Optimal Asymmetric Encryption.” In: *Advances in Cryptology – EUROCRYPT’94*. Ed. by Alfredo De Santis. Vol. 950. Lecture Notes in Computer Science. Perugia, Italy: Springer, Heidelberg, Germany, May 1995, pp. 92–111. DOI: [10.1007/BFb0053428](https://doi.org/10.1007/BFb0053428) (cit. on pp. [91](#), [129](#)).
- [BSW16] Mihir Bellare, Igors Stepanovs, and Brent Waters. “New Negative Results on Differing-Inputs Obfuscation.” In: *Advances in Cryptology – EUROCRYPT 2016, Part II*. Ed. by Marc Fischlin and Jean-Sébastien Coron. Vol. 9666. Lecture Notes in Computer Science. Vienna, Austria: Springer, Heidelberg, Germany, May 2016, pp. 792–821. DOI: [10.1007/978-3-662-49896-5_28](https://doi.org/10.1007/978-3-662-49896-5_28) (cit. on p. [45](#)).
- [BM92] Steven M. Bellovin and Michael Merritt. “Encrypted Key Exchange: Password-Based Protocols Secure against Dictionary Attacks.” In: *1992 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, May 1992, pp. 72–84. DOI: [10.1109/RISP.1992.213269](https://doi.org/10.1109/RISP.1992.213269) (cit. on pp. [13](#), [147](#), [150](#)).
- [BL18a] Fabrice Benhamouda and Huijia Lin. “k-Round Multiparty Computation from k-Round Oblivious Transfer via Garbled Interactive Circuits.” In: *Advances in Cryptology – EUROCRYPT 2018, Part II*. Ed. by Jesper Buus Nielsen and Vincent Rijmen. Vol. 10821. Lecture Notes in Computer Science. Tel Aviv, Israel: Springer, Heidelberg, Germany, Apr. 2018, pp. 500–532. DOI: [10.1007/978-3-319-78375-8_17](https://doi.org/10.1007/978-3-319-78375-8_17) (cit. on pp. [3](#), [41](#)).
- [BL18b] Sebastian Berndt and Maciej Liskiewicz. “On the Gold Standard for Security of Universal Steganography.” In: *Advances in Cryptology – EUROCRYPT 2018, Part I*. Ed. by Jesper Buus Nielsen and Vincent Rijmen. Vol. 10820. Lecture Notes in Computer Science. Tel Aviv, Israel: Springer, Heidelberg, Germany, Apr. 2018, pp. 29–60. DOI: [10.1007/978-3-319-78381-9_2](https://doi.org/10.1007/978-3-319-78381-9_2) (cit. on p. [231](#)).
- [Bit17] Nir Bitansky. “Verifiable Random Functions from Non-interactive Witness-Indistinguishable Proofs.” In: *TCC 2017: 15th Theory of Cryptography Conference, Part II*. Ed. by Yael Kalai

- and Leonid Reyzin. Vol. 10678. Lecture Notes in Computer Science. Baltimore, MD, USA: Springer, Heidelberg, Germany, Nov. 2017, pp. 567–594. DOI: [10.1007/978-3-319-70503-3_19](https://doi.org/10.1007/978-3-319-70503-3_19) (cit. on pp. 3, 41).
- [BCC+17] Nir Bitansky, Ran Canetti, Alessandro Chiesa, Shafi Goldwasser, Huijia Lin, Aviad Rubinfeld, and Eran Tromer. “The Hunting of the SNARK.” In: *Journal of Cryptology* 30.4 (Oct. 2017), pp. 989–1066. DOI: [10.1007/s00145-016-9241-9](https://doi.org/10.1007/s00145-016-9241-9) (cit. on p. 32).
- [BCKP14] Nir Bitansky, Ran Canetti, Yael Tauman Kalai, and Omer Paneth. “On Virtual Grey Box Obfuscation for General Circuits.” In: *Advances in Cryptology – CRYPTO 2014, Part II*. Ed. by Juan A. Garay and Rosario Gennaro. Vol. 8617. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2014, pp. 108–125. DOI: [10.1007/978-3-662-44381-1_7](https://doi.org/10.1007/978-3-662-44381-1_7) (cit. on p. 41).
- [BCPR16] Nir Bitansky, Ran Canetti, Omer Paneth, and Alon Rosen. “On the Existence of Extractable One-Way Functions.” In: *SIAM J. Comput.* 45.5 (2016), pp. 1910–1952. DOI: [10.1137/140975048](https://doi.org/10.1137/140975048). URL: <https://doi.org/10.1137/140975048> (cit. on pp. 10, 31–33, 88, 89, 94, 109, 152, 153, 158, 200, 201, 211).
- [BGL+15] Nir Bitansky, Sanjam Garg, Huijia Lin, Rafael Pass, and Sidharth Telang. “Succinct Randomized Encodings and their Applications.” In: *47th Annual ACM Symposium on Theory of Computing*. Ed. by Rocco A. Servedio and Ronitt Rubinfeld. Portland, OR, USA: ACM Press, June 2015, pp. 439–448. DOI: [10.1145/2746539.2746574](https://doi.org/10.1145/2746539.2746574) (cit. on pp. 3, 80).
- [BP13] Nir Bitansky and Omer Paneth. “On the impossibility of approximate obfuscation and applications to resettable cryptography.” In: *45th Annual ACM Symposium on Theory of Computing*. Ed. by Dan Boneh, Tim Roughgarden, and Joan Feigenbaum. Palo Alto, CA, USA: ACM Press, June 2013, pp. 241–250. DOI: [10.1145/2488608.2488639](https://doi.org/10.1145/2488608.2488639) (cit. on p. 41).
- [BP15a] Nir Bitansky and Omer Paneth. “ZAPs and Non-Interactive Witness Indistinguishability from Indistinguishability Obfuscation.” In: *TCC 2015: 12th Theory of Cryptography Conference, Part II*. Ed. by Yevgeniy Dodis and Jesper Buus Nielsen. Vol. 9015. Lecture Notes in Computer Science. Warsaw, Poland: Springer, Heidelberg, Germany, Mar. 2015, pp. 401–427. DOI: [10.1007/978-3-662-46497-7_16](https://doi.org/10.1007/978-3-662-46497-7_16) (cit. on pp. 3, 36, 61, 95).
- [BPR15] Nir Bitansky, Omer Paneth, and Alon Rosen. “On the Cryptographic Hardness of Finding a Nash Equilibrium.” In: *56th Annual Symposium on Foundations of Computer Science*. Ed. by Venkatesan Guruswami. Berkeley, CA, USA: IEEE Computer Society Press, Oct. 2015, pp. 1480–1498. DOI: [10.1109/FOCS.2015.94](https://doi.org/10.1109/FOCS.2015.94) (cit. on p. 3).

- [BPW16] Nir Bitansky, Omer Paneth, and Daniel Wichs. “Perfect Structure on the Edge of Chaos - Trapdoor Permutations from Indistinguishability Obfuscation.” In: *TCC 2016-A: 13th Theory of Cryptography Conference, Part I*. Ed. by Eyal Kushilevitz and Tal Malkin. Vol. 9562. Lecture Notes in Computer Science. Tel Aviv, Israel: Springer, Heidelberg, Germany, Jan. 2016, pp. 474–502. DOI: [10.1007/978-3-662-49096-9_20](https://doi.org/10.1007/978-3-662-49096-9_20) (cit. on p. 3).
- [BV15] Nir Bitansky and Vinod Vaikuntanathan. “Indistinguishability Obfuscation from Functional Encryption.” In: *56th Annual Symposium on Foundations of Computer Science*. Ed. by Venkatesan Guruswami. Berkeley, CA, USA: IEEE Computer Society Press, Oct. 2015, pp. 171–190. DOI: [10.1109/FOCS.2015.20](https://doi.org/10.1109/FOCS.2015.20) (cit. on p. 41).
- [BRS03] John Black, Phillip Rogaway, and Thomas Shrimpton. “Encryption-Scheme Security in the Presence of Key-Dependent Messages.” In: *SAC 2002: 9th Annual International Workshop on Selected Areas in Cryptography*. Ed. by Kaisa Nyberg and Howard M. Heys. Vol. 2595. Lecture Notes in Computer Science. St. John’s, Newfoundland, Canada: Springer, Heidelberg, Germany, Aug. 2003, pp. 62–75. DOI: [10.1007/3-540-36492-7_6](https://doi.org/10.1007/3-540-36492-7_6) (cit. on p. 45).
- [BFLK94] Avrim Blum, Merrick L. Furst, Michael J. Kearns, and Richard J. Lipton. “Cryptographic Primitives Based on Hard Learning Problems.” In: *Advances in Cryptology – CRYPTO’93*. Ed. by Douglas R. Stinson. Vol. 773. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 1994, pp. 278–291. DOI: [10.1007/3-540-48329-2_24](https://doi.org/10.1007/3-540-48329-2_24) (cit. on p. 2).
- [BFM88] Manuel Blum, Paul Feldman, and Silvio Micali. “Non-Interactive Zero-Knowledge and Its Applications (Extended Abstract).” In: *20th Annual ACM Symposium on Theory of Computing*. Chicago, IL, USA: ACM Press, May 1988, pp. 103–112. DOI: [10.1145/62212.62222](https://doi.org/10.1145/62212.62222) (cit. on p. 35).
- [BM82] Manuel Blum and Silvio Micali. “How to Generate Cryptographically Strong Sequences of Pseudo Random Bits.” In: *23rd Annual Symposium on Foundations of Computer Science*. Chicago, Illinois: IEEE Computer Society Press, Nov. 1982, pp. 112–117. DOI: [10.1109/SFCS.1982.72](https://doi.org/10.1109/SFCS.1982.72) (cit. on p. 21).
- [BB04a] Dan Boneh and Xavier Boyen. “Efficient Selective-ID Secure Identity Based Encryption Without Random Oracles.” In: *Advances in Cryptology – EUROCRYPT 2004*. Ed. by Christian Cachin and Jan Camenisch. Vol. 3027. Lecture Notes in Computer Science. Interlaken, Switzerland: Springer, Heidelberg, Germany, May 2004, pp. 223–238. DOI: [10.1007/978-3-540-24676-3_14](https://doi.org/10.1007/978-3-540-24676-3_14) (cit. on p. 43).
- [BB04b] Dan Boneh and Xavier Boyen. “Secure Identity Based Encryption Without Random Oracles.” In: *Advances in Cryptology – CRYPTO 2004*. Ed. by Matthew Franklin. Vol. 3152. Lecture Notes in Computer Science. Santa Barbara, CA, USA:

- Springer, Heidelberg, Germany, Aug. 2004, pp. 443–459. DOI: [10.1007/978-3-540-28628-8_27](https://doi.org/10.1007/978-3-540-28628-8_27) (cit. on p. 44).
- [BF01] Dan Boneh and Matthew K. Franklin. “Identity-Based Encryption from the Weil Pairing.” In: *Advances in Cryptology – CRYPTO 2001*. Ed. by Joe Kilian. Vol. 2139. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2001, pp. 213–229. DOI: [10.1007/3-540-44647-8_13](https://doi.org/10.1007/3-540-44647-8_13) (cit. on p. 8).
- [BHHO08] Dan Boneh, Shai Halevi, Michael Hamburg, and Rafail Ostrovsky. “Circular-Secure Encryption from Decision Diffie-Hellman.” In: *Advances in Cryptology – CRYPTO 2008*. Ed. by David Wagner. Vol. 5157. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2008, pp. 108–125. DOI: [10.1007/978-3-540-85174-5_7](https://doi.org/10.1007/978-3-540-85174-5_7) (cit. on p. 45).
- [BL96] Dan Boneh and Richard J. Lipton. “Algorithms for Black-Box Fields and their Application to Cryptography (Extended Abstract).” In: *Advances in Cryptology – CRYPTO’96*. Ed. by Neal Koblitz. Vol. 1109. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 1996, pp. 283–297. DOI: [10.1007/3-540-68697-5_22](https://doi.org/10.1007/3-540-68697-5_22) (cit. on p. 87).
- [BLS04] Dan Boneh, Ben Lynn, and Hovav Shacham. “Short Signatures from the Weil Pairing.” In: *Journal of Cryptology* 17.4 (Sept. 2004), pp. 297–319. DOI: [10.1007/s00145-004-0314-9](https://doi.org/10.1007/s00145-004-0314-9) (cit. on pp. 12, 89).
- [BV98] Dan Boneh and Ramarathnam Venkatesan. “Breaking RSA May Not Be Equivalent to Factoring.” In: *Advances in Cryptology – EUROCRYPT’98*. Ed. by Kaisa Nyberg. Vol. 1403. Lecture Notes in Computer Science. Espoo, Finland: Springer, Heidelberg, Germany, May 1998, pp. 59–71. DOI: [10.1007/BFb0054117](https://doi.org/10.1007/BFb0054117) (cit. on p. 87).
- [BW13] Dan Boneh and Brent Waters. “Constrained Pseudorandom Functions and Their Applications.” In: *Advances in Cryptology – ASIACRYPT 2013, Part II*. Ed. by Kazue Sako and Palash Sarkar. Vol. 8270. Lecture Notes in Computer Science. Bangalore, India: Springer, Heidelberg, Germany, Dec. 2013, pp. 280–300. DOI: [10.1007/978-3-642-42045-0_15](https://doi.org/10.1007/978-3-642-42045-0_15) (cit. on pp. 23, 24).
- [BZ14] Dan Boneh and Mark Zhandry. “Multiparty Key Exchange, Efficient Traitor Tracing, and More from Indistinguishability Obfuscation.” In: *Advances in Cryptology – CRYPTO 2014, Part I*. Ed. by Juan A. Garay and Rosario Gennaro. Vol. 8616. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2014, pp. 480–499. DOI: [10.1007/978-3-662-44371-2_27](https://doi.org/10.1007/978-3-662-44371-2_27) (cit. on p. 3).
- [BMNo1] Colin Boyd, Paul Montague, and Khanh Quoc Nguyen. “Elliptic Curve Based Password Authenticated Key Exchange Protocols.” In: *ACISP 01: 6th Australasian Conference on Information Security and Privacy*. Ed. by Vijay Varadharajan and Yi Mu. Vol. 2119. Lecture Notes in Computer Science. Sydney, NSW,

- Australia: Springer, Heidelberg, Germany, July 2001, pp. 487–501. DOI: [10.1007/3-540-47719-5_38](https://doi.org/10.1007/3-540-47719-5_38) (cit. on pp. 13, 150).
- [Boyo8] Xavier Boyen. “The Uber-Assumption Family.” In: *Pairing-Based Cryptography - Pairing 2008, Second International Conference, Egham, UK, September 1-3, 2008. Proceedings*. 2008, pp. 39–56. DOI: [10.1007/978-3-540-85538-5_3](https://doi.org/10.1007/978-3-540-85538-5_3). URL: https://doi.org/10.1007/978-3-540-85538-5_3 (cit. on p. 92).
- [BCP14] Elette Boyle, Kai-Min Chung, and Rafael Pass. “On Extractability Obfuscation.” In: *TCC 2014: 11th Theory of Cryptography Conference*. Ed. by Yehuda Lindell. Vol. 8349. Lecture Notes in Computer Science. San Diego, CA, USA: Springer, Heidelberg, Germany, Feb. 2014, pp. 52–73. DOI: [10.1007/978-3-642-54242-8_3](https://doi.org/10.1007/978-3-642-54242-8_3) (cit. on pp. 27, 41).
- [BCG+19a] Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, Peter Rindal, and Peter Scholl. “Efficient Two-Round OT Extension and Silent Non-Interactive Secure Computation.” In: *ACM CCS 2019: 26th Conference on Computer and Communications Security*. Ed. by Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz. ACM Press, Nov. 2019, pp. 291–308. DOI: [10.1145/3319535.3354255](https://doi.org/10.1145/3319535.3354255) (cit. on p. 3).
- [BCG+19b] Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, and Peter Scholl. “Efficient Pseudorandom Correlation Generators: Silent OT Extension and More.” In: *Advances in Cryptology – CRYPTO 2019, Part III*. Ed. by Alexandra Boldyreva and Daniele Micciancio. Vol. 11694. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2019, pp. 489–518. DOI: [10.1007/978-3-030-26954-8_16](https://doi.org/10.1007/978-3-030-26954-8_16) (cit. on p. 3).
- [BGI17] Elette Boyle, Niv Gilboa, and Yuval Ishai. “Group-Based Secure Computation: Optimizing Rounds, Communication, and Computation.” In: *Advances in Cryptology – EUROCRYPT 2017, Part II*. Ed. by Jean-Sébastien Coron and Jesper Buus Nielsen. Vol. 10211. Lecture Notes in Computer Science. Paris, France: Springer, Heidelberg, Germany, Apr. 2017, pp. 163–193. DOI: [10.1007/978-3-319-56614-6_6](https://doi.org/10.1007/978-3-319-56614-6_6) (cit. on p. 3).
- [BGI14] Elette Boyle, Shafi Goldwasser, and Ioana Ivan. “Functional Signatures and Pseudorandom Functions.” In: *PKC 2014: 17th International Conference on Theory and Practice of Public Key Cryptography*. Ed. by Hugo Krawczyk. Vol. 8383. Lecture Notes in Computer Science. Buenos Aires, Argentina: Springer, Heidelberg, Germany, Mar. 2014, pp. 501–519. DOI: [10.1007/978-3-642-54631-0_29](https://doi.org/10.1007/978-3-642-54631-0_29) (cit. on pp. 23, 24).
- [BP15b] Elette Boyle and Rafael Pass. “Limits of Extractability Assumptions with Distributional Auxiliary Input.” In: *Advances in Cryptology – ASIACRYPT 2015, Part II*. Ed. by Tetsu Iwata and Jung Hee Cheon. Vol. 9453. Lecture Notes in Computer Science. Auckland, New Zealand: Springer, Heidelberg, Germany, Nov. 2015, pp. 236–261. DOI: [10.1007/978-3-662-48800-3_10](https://doi.org/10.1007/978-3-662-48800-3_10) (cit. on p. 27).

- [BDGM20a] Zvika Brakerski, Nico Döttling, Sanjam Garg, and Giulio Malavolta. “Candidate iO from Homomorphic Encryption Schemes.” In: *Advances in Cryptology – EUROCRYPT 2020, Part I*. Ed. by Anne Canteaut and Yuval Ishai. Vol. 12105. Lecture Notes in Computer Science. Zagreb, Croatia: Springer, Heidelberg, Germany, May 2020, pp. 79–109. DOI: [10.1007/978-3-030-45721-1_4](https://doi.org/10.1007/978-3-030-45721-1_4) (cit. on pp. 2, 5, 41, 158).
- [BDGM20b] Zvika Brakerski, Nico Döttling, Sanjam Garg, and Giulio Malavolta. *Factoring and Pairings are not Necessary for iO: Circular-Secure LWE Suffices*. Cryptology ePrint Archive, Report 2020/1024. <https://eprint.iacr.org/2020/1024>. 2020 (cit. on pp. 2, 5, 41).
- [BV11] Zvika Brakerski and Vinod Vaikuntanathan. “Fully Homomorphic Encryption from Ring-LWE and Security for Key Dependent Messages.” In: *Advances in Cryptology – CRYPTO 2011*. Ed. by Phillip Rogaway. Vol. 6841. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2011, pp. 505–524. DOI: [10.1007/978-3-642-22792-9_29](https://doi.org/10.1007/978-3-642-22792-9_29) (cit. on p. 34).
- [BV14a] Zvika Brakerski and Vinod Vaikuntanathan. “Efficient Fully Homomorphic Encryption from (Standard) LWE.” In: *SIAM J. Comput.* 43.2 (2014), pp. 831–871. DOI: [10.1137/120868669](https://doi.org/10.1137/120868669). URL: <https://doi.org/10.1137/120868669> (cit. on p. 34).
- [BV14b] Zvika Brakerski and Vinod Vaikuntanathan. “Lattice-based FHE as secure as PKE.” In: *ITCS 2014: 5th Conference on Innovations in Theoretical Computer Science*. Ed. by Moni Naor. Princeton, NJ, USA: Association for Computing Machinery, Jan. 2014, pp. 1–12. DOI: [10.1145/2554797.2554799](https://doi.org/10.1145/2554797.2554799) (cit. on p. 34).
- [Can97] Ran Canetti. “Towards Realizing Random Oracles: Hash Functions That Hide All Partial Information.” In: *Advances in Cryptology – CRYPTO’97*. Ed. by Burton S. Kaliski Jr. Vol. 1294. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 1997, pp. 455–469. DOI: [10.1007/BFb0052255](https://doi.org/10.1007/BFb0052255) (cit. on p. 41).
- [Can00] Ran Canetti. “Security and Composition of Multiparty Cryptographic Protocols.” In: *Journal of Cryptology* 13.1 (Jan. 2000), pp. 143–202. DOI: [10.1007/s001459910006](https://doi.org/10.1007/s001459910006) (cit. on pp. 221, 222, 226).
- [Can01] Ran Canetti. “Universally Composable Security: A New Paradigm for Cryptographic Protocols.” In: *42nd Annual Symposium on Foundations of Computer Science*. Las Vegas, NV, USA: IEEE Computer Society Press, Oct. 2001, pp. 136–145. DOI: [10.1109/SFCS.2001.959888](https://doi.org/10.1109/SFCS.2001.959888) (cit. on p. 223).
- [CC17] Ran Canetti and Yilei Chen. “Constraint-Hiding Constrained PRFs for NC^1 from LWE.” In: *Advances in Cryptology – EUROCRYPT 2017, Part I*. Ed. by Jean-Sébastien Coron and Jesper Buus Nielsen. Vol. 10210. Lecture Notes in Computer Science. Paris, France: Springer, Heidelberg, Germany, Apr. 2017,

- pp. 446–476. DOI: [10.1007/978-3-319-56620-7_16](https://doi.org/10.1007/978-3-319-56620-7_16) (cit. on p. 3).
- [CCH+19] Ran Canetti, Yilei Chen, Justin Holmgren, Alex Lombardi, Guy N. Rothblum, Ron D. Rothblum, and Daniel Wichs. “Fiat-Shamir: from practice to theory.” In: *51st Annual ACM Symposium on Theory of Computing*. Ed. by Moses Charikar and Edith Cohen. Phoenix, AZ, USA: ACM Press, June 2019, pp. 1082–1090. DOI: [10.1145/3313276.3316380](https://doi.org/10.1145/3313276.3316380) (cit. on p. 3).
- [CCRR18] Ran Canetti, Yilei Chen, Leonid Reyzin, and Ron D. Rothblum. “Fiat-Shamir and Correlation Intractability from Strong KDM-Secure Encryption.” In: *Advances in Cryptology – EUROCRYPT 2018, Part I*. Ed. by Jesper Buus Nielsen and Vincent Rijmen. Vol. 10820. Lecture Notes in Computer Science. Tel Aviv, Israel: Springer, Heidelberg, Germany, Apr. 2018, pp. 91–122. DOI: [10.1007/978-3-319-78381-9_4](https://doi.org/10.1007/978-3-319-78381-9_4) (cit. on pp. 3, 41).
- [CD08] Ran Canetti and Ronny Ramzi Dakdouk. “Extractable Perfectly One-Way Functions.” In: *ICALP 2008: 35th International Colloquium on Automata, Languages and Programming, Part II*. Ed. by Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfssdóttir, and Igor Walukiewicz. Vol. 5126. Lecture Notes in Computer Science. Reykjavik, Iceland: Springer, Heidelberg, Germany, July 2008, pp. 449–460. DOI: [10.1007/978-3-540-70583-3_37](https://doi.org/10.1007/978-3-540-70583-3_37) (cit. on p. 31).
- [CD09] Ran Canetti and Ronny Ramzi Dakdouk. “Towards a Theory of Extractable Functions.” In: *TCC 2009: 6th Theory of Cryptography Conference*. Ed. by Omer Reingold. Vol. 5444. Lecture Notes in Computer Science. Springer, Heidelberg, Germany, Mar. 2009, pp. 595–613. DOI: [10.1007/978-3-642-00457-5_35](https://doi.org/10.1007/978-3-642-00457-5_35) (cit. on pp. 31, 32).
- [CDPW07] Ran Canetti, Yevgeniy Dodis, Rafael Pass, and Shabsi Walfish. “Universally Composable Security with Global Setup.” In: *TCC 2007: 4th Theory of Cryptography Conference*. Ed. by Salil P. Vadhan. Vol. 4392. Lecture Notes in Computer Science. Amsterdam, The Netherlands: Springer, Heidelberg, Germany, Feb. 2007, pp. 61–85. DOI: [10.1007/978-3-540-70936-7_4](https://doi.org/10.1007/978-3-540-70936-7_4) (cit. on p. 223).
- [CDNO97] Ran Canetti, Cynthia Dwork, Moni Naor, and Rafail Ostrovsky. “Deniable Encryption.” In: *Advances in Cryptology – CRYPTO’97*. Ed. by Burton S. Kaliski Jr. Vol. 1294. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 1997, pp. 90–104. DOI: [10.1007/BFb0052229](https://doi.org/10.1007/BFb0052229) (cit. on pp. 155, 233).
- [CFGN96] Ran Canetti, Uriel Feige, Oded Goldreich, and Moni Naor. “Adaptively Secure Multi-Party Computation.” In: *28th Annual ACM Symposium on Theory of Computing*. Philadelphia, PA, USA: ACM Press, May 1996, pp. 639–648. DOI: [10.1145/237814.238015](https://doi.org/10.1145/237814.238015) (cit. on pp. 16, 17, 151, 154, 155, 159, 237).

- [CL18] Ran Canetti and Amit Lichtenberg. “Certifying Trapdoor Permutations, Revisited.” In: *TCC 2018: 16th Theory of Cryptography Conference, Part I*. Ed. by Amos Beimel and Stefan Dziembowski. Vol. 11239. Lecture Notes in Computer Science. Panaji, India: Springer, Heidelberg, Germany, Nov. 2018, pp. 476–506. DOI: [10.1007/978-3-030-03807-6_18](https://doi.org/10.1007/978-3-030-03807-6_18) (cit. on p. 3).
- [CLTV15] Ran Canetti, Huijia Lin, Stefano Tessaro, and Vinod Vaikuntanathan. “Obfuscation of Probabilistic Circuits and Applications.” In: *TCC 2015: 12th Theory of Cryptography Conference, Part II*. Ed. by Yevgeniy Dodis and Jesper Buus Nielsen. Vol. 9015. Lecture Notes in Computer Science. Warsaw, Poland: Springer, Heidelberg, Germany, Mar. 2015, pp. 468–497. DOI: [10.1007/978-3-662-46497-7_19](https://doi.org/10.1007/978-3-662-46497-7_19) (cit. on pp. [iii](#), [3](#), [5](#), [8](#), [27–30](#), [34](#), [41–46](#), [56](#), [60](#), [67](#), [68](#), [73](#), [74](#), [80](#), [81](#), [89](#), [97](#), [99–105](#), [107](#)).
- [CLOS02] Ran Canetti, Yehuda Lindell, Rafail Ostrovsky, and Amit Sahai. “Universally composable two-party and multi-party secure computation.” In: *34th Annual ACM Symposium on Theory of Computing*. Montréal, Québec, Canada: ACM Press, May 2002, pp. 494–503. DOI: [10.1145/509907.509980](https://doi.org/10.1145/509907.509980) (cit. on pp. [16](#), [154](#), [222](#), [237](#)).
- [CPR17] Ran Canetti, Oxana Poburinnaya, and Mariana Raykova. “Optimal-Rate Non-Committing Encryption.” In: *Advances in Cryptology – ASIACRYPT 2017, Part III*. Ed. by Tsuyoshi Takagi and Thomas Peyrin. Vol. 10626. Lecture Notes in Computer Science. Hong Kong, China: Springer, Heidelberg, Germany, Dec. 2017, pp. 212–241. DOI: [10.1007/978-3-319-70700-6_8](https://doi.org/10.1007/978-3-319-70700-6_8) (cit. on pp. [150](#), [161](#), [171](#)).
- [CPV17] Ran Canetti, Oxana Poburinnaya, and Muthuramakrishnan Venkitasubramaniam. “Better Two-Round Adaptive Multi-party Computation.” In: *PKC 2017: 20th International Conference on Theory and Practice of Public Key Cryptography, Part II*. Ed. by Serge Fehr. Vol. 10175. Lecture Notes in Computer Science. Amsterdam, The Netherlands: Springer, Heidelberg, Germany, Mar. 2017, pp. 396–427. DOI: [10.1007/978-3-662-54388-7_14](https://doi.org/10.1007/978-3-662-54388-7_14) (cit. on pp. [16](#), [150](#), [161](#), [171](#)).
- [CRRV17] Ran Canetti, Srinivasan Raghuraman, Silas Richelson, and Vinod Vaikuntanathan. “Chosen-Ciphertext Secure Fully Homomorphic Encryption.” In: *PKC 2017: 20th International Conference on Theory and Practice of Public Key Cryptography, Part II*. Ed. by Serge Fehr. Vol. 10175. Lecture Notes in Computer Science. Amsterdam, The Netherlands: Springer, Heidelberg, Germany, Mar. 2017, pp. 213–240. DOI: [10.1007/978-3-662-54388-7_8](https://doi.org/10.1007/978-3-662-54388-7_8) (cit. on pp. [46](#), [245](#)).
- [Cha69] Gregory J. Chaitin. “On the Simplicity and Speed of Programs for Computing Infinite Sets of Natural Numbers.” In: *J. ACM* 16.3 (1969), pp. 407–422. DOI: [10.1145/321526.321530](https://doi.org/10.1145/321526.321530). URL: <https://doi.org/10.1145/321526.321530> (cit. on pp. [13](#), [147](#)).

- [CGOS07] Nishanth Chandran, Vipul Goyal, Rafail Ostrovsky, and Amit Sahai. “Covert Multi-Party Computation.” In: *48th Annual Symposium on Foundations of Computer Science*. Providence, RI, USA: IEEE Computer Society Press, Oct. 2007, pp. 238–248. DOI: [10.1109/FOCS.2007.21](https://doi.org/10.1109/FOCS.2007.21) (cit. on pp. 13, 147, 155, 245).
- [CKLM12] Melissa Chase, Markulf Kohlweiss, Anna Lysyanskaya, and Sarah Meiklejohn. “Malleable Proof Systems and Applications.” In: *Advances in Cryptology – EUROCRYPT 2012*. Ed. by David Pointcheval and Thomas Johansson. Vol. 7237. Lecture Notes in Computer Science. Cambridge, UK: Springer, Heidelberg, Germany, Apr. 2012, pp. 281–300. DOI: [10.1007/978-3-642-29011-4_18](https://doi.org/10.1007/978-3-642-29011-4_18) (cit. on p. 117).
- [CGH17] Yilei Chen, Craig Gentry, and Shai Halevi. “Cryptanalyses of Candidate Branching Program Obfuscators.” In: *Advances in Cryptology – EUROCRYPT 2017, Part III*. Ed. by Jean-Sébastien Coron and Jesper Buus Nielsen. Vol. 10212. Lecture Notes in Computer Science. Paris, France: Springer, Heidelberg, Germany, Apr. 2017, pp. 278–307. DOI: [10.1007/978-3-319-56617-7_10](https://doi.org/10.1007/978-3-319-56617-7_10) (cit. on p. 2).
- [CHVW19] Yilei Chen, Minki Hhan, Vinod Vaikuntanathan, and Hoeteck Wee. “Matrix PRFs: Constructions, Attacks, and Applications to Obfuscation.” In: *TCC 2019: 17th Theory of Cryptography Conference, Part I*. Ed. by Dennis Hofheinz and Alon Rosen. Vol. 11891. Lecture Notes in Computer Science. Nuremberg, Germany: Springer, Heidelberg, Germany, Dec. 2019, pp. 55–80. DOI: [10.1007/978-3-030-36030-6_3](https://doi.org/10.1007/978-3-030-36030-6_3) (cit. on pp. 2, 5, 41).
- [CVW18] Yilei Chen, Vinod Vaikuntanathan, and Hoeteck Wee. “GGH15 Beyond Permutation Branching Programs: Proofs, Attacks, and Candidates.” In: *Advances in Cryptology – CRYPTO 2018, Part II*. Ed. by Hovav Shacham and Alexandra Boldyreva. Vol. 10992. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2018, pp. 577–607. DOI: [10.1007/978-3-319-96881-0_20](https://doi.org/10.1007/978-3-319-96881-0_20) (cit. on pp. 2, 5, 41).
- [CHL+15] Jung Hee Cheon, Kyoohyung Han, Changmin Lee, Hansol Ryu, and Damien Stehlé. “Cryptanalysis of the Multilinear Map over the Integers.” In: *Advances in Cryptology – EUROCRYPT 2015, Part I*. Ed. by Elisabeth Oswald and Marc Fischlin. Vol. 9056. Lecture Notes in Computer Science. Sofia, Bulgaria: Springer, Heidelberg, Germany, Apr. 2015, pp. 3–12. DOI: [10.1007/978-3-662-46800-5_1](https://doi.org/10.1007/978-3-662-46800-5_1) (cit. on p. 2).
- [CDG+17] Chongwon Cho, Nico Döttling, Sanjam Garg, Divya Gupta, Peihan Miao, and Antigoni Polychroniadou. “Laconic Oblivious Transfer and Its Applications.” In: *Advances in Cryptology – CRYPTO 2017, Part II*. Ed. by Jonathan Katz and Hovav Shacham. Vol. 10402. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2017, pp. 33–65. DOI: [10.1007/978-3-319-63715-0_2](https://doi.org/10.1007/978-3-319-63715-0_2) (cit. on p. 3).

- [CDMW09] Seung Geol Choi, Dana Dachman-Soled, Tal Malkin, and Hoeteck Wee. “Improved Non-committing Encryption with Applications to Adaptively Secure Protocols.” In: *Advances in Cryptology – ASIACRYPT 2009*. Ed. by Mitsuru Matsui. Vol. 5912. Lecture Notes in Computer Science. Tokyo, Japan: Springer, Heidelberg, Germany, Dec. 2009, pp. 287–302. DOI: [10.1007/978-3-642-10366-7_17](https://doi.org/10.1007/978-3-642-10366-7_17) (cit. on pp. 153, 237–239).
- [CLP15] Kai-Min Chung, Huijia Lin, and Rafael Pass. “Constant-Round Concurrent Zero-Knowledge from Indistinguishability Obfuscation.” In: *Advances in Cryptology – CRYPTO 2015, Part I*. Ed. by Rosario Gennaro and Matthew J. B. Robshaw. Vol. 9215. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2015, pp. 287–307. DOI: [10.1007/978-3-662-47989-6_14](https://doi.org/10.1007/978-3-662-47989-6_14) (cit. on p. 3).
- [CsW19] Ran Cohen, abhi shelat, and Daniel Wichs. “Adaptively Secure MPC with Sublinear Communication Complexity.” In: *Advances in Cryptology – CRYPTO 2019, Part II*. Ed. by Alexandra Boldyreva and Daniele Micciancio. Vol. 11693. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2019, pp. 30–60. DOI: [10.1007/978-3-030-26951-7_2](https://doi.org/10.1007/978-3-030-26951-7_2) (cit. on pp. 4, 149, 153, 154, 187).
- [CTL98a] Christian S. Collberg, Clark D. Thomborson, and Douglas Low. “Breaking Abstractions and Unstructuring Data Structures.” In: *Proceedings of the 1998 International Conference on Computer Languages, ICCL 1998, Chicago, IL, USA, May 14-16, 1998*. IEEE Computer Society, 1998, pp. 28–38. DOI: [10.1109/ICCL.1998.674154](https://doi.org/10.1109/ICCL.1998.674154). URL: <https://doi.org/10.1109/ICCL.1998.674154> (cit. on p. 2).
- [CTL98b] Christian S. Collberg, Clark D. Thomborson, and Douglas Low. “Manufacturing Cheap, Resilient, and Stealthy Opaque Constructs.” In: *POPL ’98, Proceedings of the 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, San Diego, CA, USA, January 19-21, 1998*. Ed. by David B. MacQueen and Luca Cardelli. ACM, 1998, pp. 184–196. DOI: [10.1145/268946.268962](https://doi.org/10.1145/268946.268962). URL: <https://doi.org/10.1145/268946.268962> (cit. on p. 2).
- [Coroo] Jean-Sébastien Coron. “On the Exact Security of Full Domain Hash.” In: *Advances in Cryptology – CRYPTO 2000*. Ed. by Mihir Bellare. Vol. 1880. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2000, pp. 229–235. DOI: [10.1007/3-540-44598-6_14](https://doi.org/10.1007/3-540-44598-6_14) (cit. on p. 89).
- [CLLT16] Jean-Sébastien Coron, Moon Sung Lee, Tancrede Lepoint, and Mehdi Tibouchi. “Cryptanalysis of GGH15 Multilinear Maps.” In: *Advances in Cryptology – CRYPTO 2016, Part II*. Ed. by Matthew Robshaw and Jonathan Katz. Vol. 9815. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2016, pp. 607–628. DOI: [10.1007/978-3-662-53008-5_21](https://doi.org/10.1007/978-3-662-53008-5_21) (cit. on p. 2).

- [CLLT17] Jean-Sébastien Coron, Moon Sung Lee, Tancrede Lepoint, and Mehdi Tibouchi. “Zeroizing Attacks on Indistinguishability Obfuscation over CLT_{13} .” In: *PKC 2017: 20th International Conference on Theory and Practice of Public Key Cryptography, Part I*. Ed. by Serge Fehr. Vol. 10174. Lecture Notes in Computer Science. Amsterdam, The Netherlands: Springer, Heidelberg, Germany, Mar. 2017, pp. 41–58. DOI: [10.1007/978-3-662-54365-8_3](https://doi.org/10.1007/978-3-662-54365-8_3) (cit. on p. 2).
- [CH19] Geoffroy Couteau and Dennis Hofheinz. “Designated-Verifier Pseudorandom Generators, and Their Applications.” In: *Advances in Cryptology – EUROCRYPT 2019, Part II*. Ed. by Yuval Ishai and Vincent Rijmen. Vol. 11477. Lecture Notes in Computer Science. Darmstadt, Germany: Springer, Heidelberg, Germany, May 2019, pp. 562–592. DOI: [10.1007/978-3-030-17656-3_20](https://doi.org/10.1007/978-3-030-17656-3_20) (cit. on p. 36).
- [CS02] Ronald Cramer and Victor Shoup. “Universal Hash Proofs and a Paradigm for Adaptive Chosen Ciphertext Secure Public-Key Encryption.” In: *Advances in Cryptology – EUROCRYPT 2002*. Ed. by Lars R. Knudsen. Vol. 2332. Lecture Notes in Computer Science. Amsterdam, The Netherlands: Springer, Heidelberg, Germany, Apr. 2002, pp. 45–64. DOI: [10.1007/3-540-46035-7_4](https://doi.org/10.1007/3-540-46035-7_4) (cit. on p. 94).
- [CS03] Ronald Cramer and Victor Shoup. “Design and Analysis of Practical Public-Key Encryption Schemes Secure against Adaptive Chosen Ciphertext Attack.” In: *SIAM J. Comput.* 33.1 (2003), pp. 167–226. DOI: [10.1137/S0097539702403773](https://doi.org/10.1137/S0097539702403773). URL: <https://doi.org/10.1137/S0097539702403773> (cit. on p. 129).
- [Dac16] Dana Dachman-Soled. “Towards Non-Black-Box Separations of Public Key Encryption and One Way Function.” In: *TCC 2016-B: 14th Theory of Cryptography Conference, Part II*. Ed. by Martin Hirt and Adam D. Smith. Vol. 9986. Lecture Notes in Computer Science. Beijing, China: Springer, Heidelberg, Germany, Oct. 2016, pp. 169–191. DOI: [10.1007/978-3-662-53644-5_7](https://doi.org/10.1007/978-3-662-53644-5_7) (cit. on p. 33).
- [DKR15] Dana Dachman-Soled, Jonathan Katz, and Vanishree Rao. “Adaptively Secure, Universally Composable, Multiparty Computation in Constant Rounds.” In: *TCC 2015: 12th Theory of Cryptography Conference, Part II*. Ed. by Yevgeniy Dodis and Jesper Buus Nielsen. Vol. 9015. Lecture Notes in Computer Science. Warsaw, Poland: Springer, Heidelberg, Germany, Mar. 2015, pp. 586–613. DOI: [10.1007/978-3-662-46497-7_23](https://doi.org/10.1007/978-3-662-46497-7_23) (cit. on pp. 16, 152–154, 187, 189, 200, 211–215, 221, 226, 233).
- [Dam92] Ivan Damgård. “Towards Practical Public Key Systems Secure Against Chosen Ciphertext Attacks.” In: *Advances in Cryptology – CRYPTO’91*. Ed. by Joan Feigenbaum. Vol. 576. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 1992, pp. 445–456. DOI: [10.1007/3-540-46766-1_36](https://doi.org/10.1007/3-540-46766-1_36) (cit. on pp. 9, 30, 87, 88, 92).

- [DJ01] Ivan Damgård and Mats Jurik. “A Generalisation, a Simplification and Some Applications of Paillier’s Probabilistic Public-Key System.” In: *PKC 2001: 4th International Workshop on Theory and Practice in Public Key Cryptography*. Ed. by Kwangjo Kim. Vol. 1992. Lecture Notes in Computer Science. Cheju Island, South Korea: Springer, Heidelberg, Germany, Feb. 2001, pp. 119–136. DOI: [10.1007/3-540-44586-2_9](https://doi.org/10.1007/3-540-44586-2_9) (cit. on p. 97).
- [DN00] Ivan Damgård and Jesper Buus Nielsen. “Improved Non-committing Encryption Schemes Based on a General Complexity Assumption.” In: *Advances in Cryptology – CRYPTO 2000*. Ed. by Mihir Bellare. Vol. 1880. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2000, pp. 432–450. DOI: [10.1007/3-540-44598-6_27](https://doi.org/10.1007/3-540-44598-6_27) (cit. on pp. 17, 151, 159, 177, 237, 238).
- [DPR16] Ivan Damgård, Antigoni Polychroniadou, and Vanishree Rao. “Adaptively Secure Multi-Party Computation from LWE (via Equivocal FHE).” In: *PKC 2016: 19th International Conference on Theory and Practice of Public Key Cryptography, Part II*. Ed. by Chen-Mou Cheng, Kai-Min Chung, Giuseppe Persiano, and Bo-Yin Yang. Vol. 9615. Lecture Notes in Computer Science. Taipei, Taiwan: Springer, Heidelberg, Germany, Mar. 2016, pp. 208–233. DOI: [10.1007/978-3-662-49387-8_9](https://doi.org/10.1007/978-3-662-49387-8_9) (cit. on p. 16).
- [Deno2] Alexander W. Dent. “Adapting the Weaknesses of the Random Oracle Model to the Generic Group Model.” In: *Advances in Cryptology – ASIACRYPT 2002*. Ed. by Yuliang Zheng. Vol. 2501. Lecture Notes in Computer Science. Queenstown, New Zealand: Springer, Heidelberg, Germany, Dec. 2002, pp. 100–109. DOI: [10.1007/3-540-36178-2_6](https://doi.org/10.1007/3-540-36178-2_6) (cit. on pp. 9, 87).
- [Deno6] Alexander W. Dent. “The Cramer-Shoup Encryption Scheme Is Plaintext Aware in the Standard Model.” In: *Advances in Cryptology – EUROCRYPT 2006*. Ed. by Serge Vaudenay. Vol. 4004. Lecture Notes in Computer Science. St. Petersburg, Russia: Springer, Heidelberg, Germany, May 2006, pp. 289–307. DOI: [10.1007/11761679_18](https://doi.org/10.1007/11761679_18) (cit. on p. 87).
- [DH76] Whitfield Diffie and Martin E. Hellman. “New directions in cryptography.” In: *IEEE Trans. Inf. Theory* 22.6 (1976), pp. 644–654. DOI: [10.1109/TIT.1976.1055638](https://doi.org/10.1109/TIT.1976.1055638). URL: <https://doi.org/10.1109/TIT.1976.1055638> (cit. on pp. 2, 8, 23).
- [DHRW16] Yevgeniy Dodis, Shai Halevi, Ron D. Rothblum, and Daniel Wichs. “Spooky Encryption and Its Applications.” In: *Advances in Cryptology – CRYPTO 2016, Part III*. Ed. by Matthew Robshaw and Jonathan Katz. Vol. 9816. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2016, pp. 93–122. DOI: [10.1007/978-3-662-53015-3_4](https://doi.org/10.1007/978-3-662-53015-3_4) (cit. on pp. 30, 41, 42, 46, 47).
- [DORS08] Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam D. Smith. “Fuzzy Extractors: How to Generate Strong Keys from Biometrics and Other Noisy Data.” In: *SIAM J. Comput.*

- 38.1 (2008), pp. 97–139. DOI: [10.1137/060651380](https://doi.org/10.1137/060651380). URL: <https://doi.org/10.1137/060651380> (cit. on pp. 20, 189).
- [DRV12] Yevgeniy Dodis, Thomas Ristenpart, and Salil P. Vadhan. “Randomness Condensers for Efficiently Samplable, Seed-Dependent Sources.” In: *TCC 2012: 9th Theory of Cryptography Conference*. Ed. by Ronald Cramer. Vol. 7194. Lecture Notes in Computer Science. Taormina, Sicily, Italy: Springer, Heidelberg, Germany, Mar. 2012, pp. 618–635. DOI: [10.1007/978-3-642-28914-9_35](https://doi.org/10.1007/978-3-642-28914-9_35) (cit. on pp. 13, 147).
- [DG17] Nico Döttling and Sanjam Garg. “Identity-Based Encryption from the Diffie-Hellman Assumption.” In: *Advances in Cryptology – CRYPTO 2017, Part I*. Ed. by Jonathan Katz and Hovav Shacham. Vol. 10401. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2017, pp. 537–569. DOI: [10.1007/978-3-319-63688-7_18](https://doi.org/10.1007/978-3-319-63688-7_18) (cit. on p. 3).
- [DGI+19] Nico Döttling, Sanjam Garg, Yuval Ishai, Giulio Malavolta, Tamer Mour, and Rafail Ostrovsky. “Trapdoor Hash Functions and Their Applications.” In: *Advances in Cryptology – CRYPTO 2019, Part III*. Ed. by Alexandra Boldyreva and Daniele Micciancio. Vol. 11694. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2019, pp. 3–32. DOI: [10.1007/978-3-030-26954-8_1](https://doi.org/10.1007/978-3-030-26954-8_1) (cit. on p. 3).
- [DN18] Nico Döttling and Ryo Nishimaki. *Universal Proxy Re-Encryption*. Cryptology ePrint Archive, Report 2018/840. <https://eprint.iacr.org/2018/840>. 2018 (cit. on pp. 8, 46, 245).
- [ElG85] Taher ElGamal. “A public key cryptosystem and a signature scheme based on discrete logarithms.” In: *IEEE Trans. Information Theory* 31.4 (1985), pp. 469–472. DOI: [10.1109/TIT.1985.1057074](https://doi.org/10.1109/TIT.1985.1057074). URL: <https://doi.org/10.1109/TIT.1985.1057074> (cit. on pp. 8, 33, 97, 129).
- [EGL85] Shimon Even, Oded Goldreich, and Abraham Lempel. “A Randomized Protocol for Signing Contracts.” In: *Commun. ACM* 28.6 (1985), pp. 637–647. DOI: [10.1145/3812.3818](https://doi.org/10.1145/3812.3818). URL: <http://doi.acm.org/10.1145/3812.3818> (cit. on p. 159).
- [FHHL18] Pooya Farshim, Julia Hesse, Dennis Hofheinz, and Enrique Larraia. “Graded Encoding Schemes from Obfuscation.” In: *PKC 2018: 21st International Conference on Theory and Practice of Public Key Cryptography, Part II*. Ed. by Michel Abdalla and Ricardo Dahab. Vol. 10770. Lecture Notes in Computer Science. Rio de Janeiro, Brazil: Springer, Heidelberg, Germany, Mar. 2018, pp. 371–400. DOI: [10.1007/978-3-319-76581-5_13](https://doi.org/10.1007/978-3-319-76581-5_13) (cit. on pp. 41, 91, 92, 110, 111, 115, 116).
- [FLS90] Uriel Feige, Dror Lapidot, and Adi Shamir. “Multiple Non-Interactive Zero Knowledge Proofs Based on a Single Random String (Extended Abstract).” In: *31st Annual Symposium on Foundations of Computer Science*. St. Louis, MO, USA: IEEE

- Computer Society Press, Oct. 1990, pp. 308–317. DOI: [10.1109/FSCS.1990.89549](https://doi.org/10.1109/FSCS.1990.89549) (cit. on p. 36).
- [FS90] Uriel Feige and Adi Shamir. “Witness Indistinguishable and Witness Hiding Protocols.” In: *22nd Annual ACM Symposium on Theory of Computing*. Baltimore, MD, USA: ACM Press, May 1990, pp. 416–426. DOI: [10.1145/100216.100272](https://doi.org/10.1145/100216.100272) (cit. on p. 95).
- [FS87] Amos Fiat and Adi Shamir. “How to Prove Yourself: Practical Solutions to Identification and Signature Problems.” In: *Advances in Cryptology – CRYPTO’86*. Ed. by Andrew M. Odlyzko. Vol. 263. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 1987, pp. 186–194. DOI: [10.1007/3-540-47721-7_12](https://doi.org/10.1007/3-540-47721-7_12) (cit. on p. 10).
- [FGJ18] Nils Fleischhacker, Vipul Goyal, and Abhishek Jain. “On the Existence of Three Round Zero-Knowledge Proofs.” In: *Advances in Cryptology – EUROCRYPT 2018, Part III*. Ed. by Jesper Buus Nielsen and Vincent Rijmen. Vol. 10822. Lecture Notes in Computer Science. Tel Aviv, Israel: Springer, Heidelberg, Germany, Apr. 2018, pp. 3–33. DOI: [10.1007/978-3-319-78372-7_1](https://doi.org/10.1007/978-3-319-78372-7_1) (cit. on p. 32).
- [FJS19] Nils Fleischhacker, Tibor Jager, and Dominique Schröder. “On Tight Security Proofs for Schnorr Signatures.” In: *Journal of Cryptology* 32.2 (Apr. 2019), pp. 566–599. DOI: [10.1007/s00145-019-09311-5](https://doi.org/10.1007/s00145-019-09311-5) (cit. on p. 89).
- [Flo64] William B. Floyd. “Review of ‘Information Theory and Coding’ (Abramson, N.; 1963).” In: *IEEE Trans. Inf. Theory* 10.4 (1964), p. 392. DOI: [10.1109/TIT.1964.1053709](https://doi.org/10.1109/TIT.1964.1053709). URL: <https://doi.org/10.1109/TIT.1964.1053709> (cit. on p. 12).
- [FJNT15] Tore Kasper Frederiksen, Thomas P. Jakobsen, Jesper Buus Nielsen, and Roberto Trifiletti. *TinyLEGO: An Interactive Garbling Scheme for Maliciously Secure Two-Party Computation*. Cryptology ePrint Archive, Report 2015/309. <http://eprint.iacr.org/2015/309>. 2015 (cit. on p. 3).
- [FHPS13] Eduarda S. V. Freire, Dennis Hofheinz, Kenneth G. Paterson, and Christoph Striecks. “Programmable Hash Functions in the Multilinear Setting.” In: *Advances in Cryptology – CRYPTO 2013, Part I*. Ed. by Ran Canetti and Juan A. Garay. Vol. 8042. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2013, pp. 513–530. DOI: [10.1007/978-3-642-40041-4_28](https://doi.org/10.1007/978-3-642-40041-4_28) (cit. on pp. 10, 90).
- [FKL18] Georg Fuchsbauer, Eike Kiltz, and Julian Loss. “The Algebraic Group Model and its Applications.” In: *Advances in Cryptology – CRYPTO 2018, Part II*. Ed. by Hovav Shacham and Alexandra Boldyreva. Vol. 10992. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2018, pp. 33–62. DOI: [10.1007/978-3-319-96881-0_2](https://doi.org/10.1007/978-3-319-96881-0_2) (cit. on pp. iii, 9, 10, 12, 87–89, 91, 93, 94, 111, 245).

- [FPS20] Georg Fuchsbauer, Antoine Plouviez, and Yannick Seurin. “Blind Schnorr Signatures and Signed ElGamal Encryption in the Algebraic Group Model.” In: *Advances in Cryptology – EUROCRYPT 2020, Part II*. Ed. by Anne Canteaut and Yuval Ishai. Vol. 12106. Lecture Notes in Computer Science. Zagreb, Croatia: Springer, Heidelberg, Germany, May 2020, pp. 63–95. DOI: [10.1007/978-3-030-45724-2_3](https://doi.org/10.1007/978-3-030-45724-2_3) (cit. on pp. [12](#), [88](#), [89](#), [91](#), [92](#), [129–133](#), [139–141](#), [245](#)).
- [GGH13] Sanjam Garg, Craig Gentry, and Shai Halevi. “Candidate Multilinear Maps from Ideal Lattices.” In: *Advances in Cryptology – EUROCRYPT 2013*. Ed. by Thomas Johansson and Phong Q. Nguyen. Vol. 7881. Lecture Notes in Computer Science. Athens, Greece: Springer, Heidelberg, Germany, May 2013, pp. 1–17. DOI: [10.1007/978-3-642-38348-9_1](https://doi.org/10.1007/978-3-642-38348-9_1) (cit. on pp. [90](#), [109](#), [110](#)).
- [GGHR14] Sanjam Garg, Craig Gentry, Shai Halevi, and Mariana Raykova. “Two-Round Secure MPC from Indistinguishability Obfuscation.” In: *TCC 2014: 11th Theory of Cryptography Conference*. Ed. by Yehuda Lindell. Vol. 8349. Lecture Notes in Computer Science. San Diego, CA, USA: Springer, Heidelberg, Germany, Feb. 2014, pp. 74–94. DOI: [10.1007/978-3-642-54242-8_4](https://doi.org/10.1007/978-3-642-54242-8_4) (cit. on pp. [3](#), [41](#)).
- [GGH+13] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. “Candidate Indistinguishability Obfuscation and Functional Encryption for all Circuits.” In: *54th Annual Symposium on Foundations of Computer Science*. Berkeley, CA, USA: IEEE Computer Society Press, Oct. 2013, pp. 40–49. DOI: [10.1109/FOCS.2013.13](https://doi.org/10.1109/FOCS.2013.13) (cit. on pp. [2](#), [3](#), [25](#), [26](#), [41](#)).
- [GGHW17] Sanjam Garg, Craig Gentry, Shai Halevi, and Daniel Wichs. “On the Implausibility of Differing-Inputs Obfuscation and Extractable Witness Encryption with Auxiliary Input.” In: *Algorithmica* 79.4 (2017), pp. 1353–1373. DOI: [10.1007/s00453-017-0276-6](https://doi.org/10.1007/s00453-017-0276-6). URL: <https://doi.org/10.1007/s00453-017-0276-6> (cit. on pp. [27](#), [28](#), [45](#), [60](#)).
- [GMM+16] Sanjam Garg, Eric Miles, Pratyay Mukherjee, Amit Sahai, Akshayaram Srinivasan, and Mark Zhandry. “Secure Obfuscation in a Weak Multilinear Map Model.” In: *TCC 2016-B: 14th Theory of Cryptography Conference, Part II*. Ed. by Martin Hirt and Adam D. Smith. Vol. 9986. Lecture Notes in Computer Science. Beijing, China: Springer, Heidelberg, Germany, Oct. 2016, pp. 241–268. DOI: [10.1007/978-3-662-53644-5_10](https://doi.org/10.1007/978-3-662-53644-5_10) (cit. on pp. [2](#), [5](#)).
- [GPS16] Sanjam Garg, Omkant Pandey, and Akshayaram Srinivasan. “Revisiting the Cryptographic Hardness of Finding a Nash Equilibrium.” In: *Advances in Cryptology – CRYPTO 2016, Part II*. Ed. by Matthew Robshaw and Jonathan Katz. Vol. 9815. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2016, pp. 579–604. DOI: [10.1007/978-3-662-53008-5_20](https://doi.org/10.1007/978-3-662-53008-5_20) (cit. on pp. [3](#), [5](#), [42](#)).

- [GPSZ17] Sanjam Garg, Omkant Pandey, Akshayaram Srinivasan, and Mark Zhandry. “Breaking the Sub-Exponential Barrier in Obfuscation.” In: *Advances in Cryptology – EUROCRYPT 2017, Part III*. Ed. by Jean-Sébastien Coron and Jesper Buus Nielsen. Vol. 10212. Lecture Notes in Computer Science. Paris, France: Springer, Heidelberg, Germany, Apr. 2017, pp. 156–181. DOI: [10.1007/978-3-319-56617-7_6](https://doi.org/10.1007/978-3-319-56617-7_6) (cit. on pp. 5, 42, 44).
- [GS12] Sanjam Garg and Amit Sahai. “Adaptively Secure Multi-Party Computation with Dishonest Majority.” In: *Advances in Cryptology – CRYPTO 2012*. Ed. by Reihaneh Safavi-Naini and Ran Canetti. Vol. 7417. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2012, pp. 105–123. DOI: [10.1007/978-3-642-32009-5_8](https://doi.org/10.1007/978-3-642-32009-5_8) (cit. on pp. 16, 154).
- [GS16] Sanjam Garg and Akshayaram Srinivasan. “Single-Key to Multi-Key Functional Encryption with Polynomial Loss.” In: *TCC 2016-B: 14th Theory of Cryptography Conference, Part II*. Ed. by Martin Hirt and Adam D. Smith. Vol. 9986. Lecture Notes in Computer Science. Beijing, China: Springer, Heidelberg, Germany, Oct. 2016, pp. 419–442. DOI: [10.1007/978-3-662-53644-5_16](https://doi.org/10.1007/978-3-662-53644-5_16) (cit. on p. 42).
- [GS18] Sanjam Garg and Akshayaram Srinivasan. “Two-Round Multiparty Secure Computation from Minimal Assumptions.” In: *Advances in Cryptology – EUROCRYPT 2018, Part II*. Ed. by Jesper Buus Nielsen and Vincent Rijmen. Vol. 10821. Lecture Notes in Computer Science. Tel Aviv, Israel: Springer, Heidelberg, Germany, Apr. 2018, pp. 468–499. DOI: [10.1007/978-3-319-78375-8_16](https://doi.org/10.1007/978-3-319-78375-8_16) (cit. on pp. 3, 41).
- [GJLS20] Romain Gay, Aayush Jain, Huijia Lin, and Amit Sahai. *Indistinguishability Obfuscation from Simple-to-State Hard Problems: New Assumptions, New Techniques, and Simplification*. Cryptology ePrint Archive, Report 2020/764. <https://eprint.iacr.org/2020/764>. 2020 (cit. on pp. 2, 5, 41).
- [GP20] Romain Gay and Rafael Pass. *Indistinguishability Obfuscation from Circular Security*. Cryptology ePrint Archive, Report 2020/1010. <https://eprint.iacr.org/2020/1010>. 2020 (cit. on pp. 2, 5, 41).
- [Gen09] Craig Gentry. “Fully homomorphic encryption using ideal lattices.” In: *41st Annual ACM Symposium on Theory of Computing*. Ed. by Michael Mitzenmacher. Bethesda, MD, USA: ACM Press, May 2009, pp. 169–178. DOI: [10.1145/1536414.1536440](https://doi.org/10.1145/1536414.1536440) (cit. on pp. 34, 73).
- [GGH15] Craig Gentry, Sergey Gorbunov, and Shai Halevi. “Graph-Induced Multilinear Maps from Lattices.” In: *TCC 2015: 12th Theory of Cryptography Conference, Part II*. Ed. by Yevgeniy Dodis and Jesper Buus Nielsen. Vol. 9015. Lecture Notes in Computer Science. Warsaw, Poland: Springer, Heidelberg, Germany, Mar. 2015, pp. 498–527. DOI: [10.1007/978-3-662-46497-7_20](https://doi.org/10.1007/978-3-662-46497-7_20) (cit. on pp. 2, 3, 5).

- [GJK18] Craig Gentry, Charanjit S. Jutla, and Daniel Kane. *Obfuscation Using Tensor Products*. Cryptology ePrint Archive, Report 2018/756. <https://eprint.iacr.org/2018/756>. 2018 (cit. on pp. 3, 5, 41).
- [GKM+00] Yael Gertner, Sampath Kannan, Tal Malkin, Omer Reingold, and Mahesh Viswanathan. “The Relationship between Public Key Encryption and Oblivious Transfer.” In: *41st Annual Symposium on Foundations of Computer Science*. Redondo Beach, CA, USA: IEEE Computer Society Press, Nov. 2000, pp. 325–335. DOI: [10.1109/SFCS.2000.892121](https://doi.org/10.1109/SFCS.2000.892121) (cit. on pp. 17, 151, 159).
- [GT20] Ashrujit Ghoshal and Stefano Tessaro. *Tight State-Restoration Soundness in the Algebraic Group Model*. Cryptology ePrint Archive, Report 2020/1351. <https://eprint.iacr.org/2020/1351>. 2020 (cit. on p. 245).
- [GM59] E. N. Gilbert and E. F. Moore. “Variable-length binary encodings.” In: *The Bell System Technical Journal* 38.4 (1959), pp. 933–967. DOI: [10.1002/j.1538-7305.1959.tb01583.x](https://doi.org/10.1002/j.1538-7305.1959.tb01583.x) (cit. on p. 12).
- [GS85] Andrew V. Goldberg and Michael Sipser. “Compression and Ranking.” In: *17th Annual ACM Symposium on Theory of Computing*. Providence, RI, USA: ACM Press, May 1985, pp. 440–448. DOI: [10.1145/22145.22194](https://doi.org/10.1145/22145.22194) (cit. on pp. 12, 13, 147, 150, 156).
- [GGM84a] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. “How to Construct Random Functions (Extended Abstract).” In: *25th Annual Symposium on Foundations of Computer Science*. Singer Island, Florida: IEEE Computer Society Press, Oct. 1984, pp. 464–479. DOI: [10.1109/SFCS.1984.715949](https://doi.org/10.1109/SFCS.1984.715949) (cit. on p. 61).
- [GGM84b] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. “On the Cryptographic Applications of Random Functions.” In: *Advances in Cryptology – CRYPTO’84*. Ed. by G. R. Blakley and David Chaum. Vol. 196. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 1984, pp. 276–288 (cit. on p. 10).
- [GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. “How to construct random functions.” In: *J. ACM* 33.4 (1986), pp. 792–807. DOI: [10.1145/6490.6503](https://doi.org/10.1145/6490.6503). URL: <https://doi.org/10.1145/6490.6503> (cit. on pp. 10, 23, 24).
- [GK96] Oded Goldreich and Hugo Krawczyk. “On the Composition of Zero-Knowledge Proof Systems.” In: *SIAM J. Comput.* 25.1 (1996), pp. 169–192. DOI: [10.1137/S0097539791220688](https://doi.org/10.1137/S0097539791220688). URL: <https://doi.org/10.1137/S0097539791220688> (cit. on p. 32).
- [GMW86] Oded Goldreich, Silvio Micali, and Avi Wigderson. “Proofs that Yield Nothing But their Validity and a Methodology of Cryptographic Protocol Design (Extended Abstract).” In: *27th Annual Symposium on Foundations of Computer Science*. Toronto, Ontario, Canada: IEEE Computer Society Press, Oct. 1986, pp. 174–187. DOI: [10.1109/SFCS.1986.47](https://doi.org/10.1109/SFCS.1986.47) (cit. on p. 36).

- [GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. “How to Play any Mental Game or A Completeness Theorem for Protocols with Honest Majority.” In: *19th Annual ACM Symposium on Theory of Computing*. Ed. by Alfred Aho. New York City, NY, USA: ACM Press, May 1987, pp. 218–229. DOI: [10.1145/28395.28420](https://doi.org/10.1145/28395.28420) (cit. on pp. 16, 221).
- [GO94] Oded Goldreich and Yair Oren. “Definitions and Properties of Zero-Knowledge Proof Systems.” In: *Journal of Cryptology* 7.1 (Dec. 1994), pp. 1–32. DOI: [10.1007/BF00195207](https://doi.org/10.1007/BF00195207) (cit. on p. 35).
- [GGG+14] Shafi Goldwasser, S. Dov Gordon, Vipul Goyal, Abhishek Jain, Jonathan Katz, Feng-Hao Liu, Amit Sahai, Elaine Shi, and Hong-Sheng Zhou. “Multi-input Functional Encryption.” In: *Advances in Cryptology – EUROCRYPT 2014*. Ed. by Phong Q. Nguyen and Elisabeth Oswald. Vol. 8441. Lecture Notes in Computer Science. Copenhagen, Denmark: Springer, Heidelberg, Germany, May 2014, pp. 578–602. DOI: [10.1007/978-3-642-55220-5_32](https://doi.org/10.1007/978-3-642-55220-5_32) (cit. on p. 3).
- [GK05] Shafi Goldwasser and Yael Tauman Kalai. “On the Impossibility of Obfuscation with Auxiliary Input.” In: *46th Annual Symposium on Foundations of Computer Science*. Pittsburgh, PA, USA: IEEE Computer Society Press, Oct. 2005, pp. 553–562. DOI: [10.1109/SFCS.2005.60](https://doi.org/10.1109/SFCS.2005.60) (cit. on p. 41).
- [GL91] Shafi Goldwasser and Leonid A. Levin. “Fair Computation of General Functions in Presence of Immoral Majority.” In: *Advances in Cryptology – CRYPTO’90*. Ed. by Alfred J. Menezes and Scott A. Vanstone. Vol. 537. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 1991, pp. 77–93. DOI: [10.1007/3-540-38424-3_6](https://doi.org/10.1007/3-540-38424-3_6) (cit. on p. 224).
- [GM82] Shafi Goldwasser and Silvio Micali. “Probabilistic Encryption and How to Play Mental Poker Keeping Secret All Partial Information.” In: *14th Annual ACM Symposium on Theory of Computing*. San Francisco, CA, USA: ACM Press, May 1982, pp. 365–377. DOI: [10.1145/800070.802212](https://doi.org/10.1145/800070.802212) (cit. on p. 21).
- [GM84] Shafi Goldwasser and Silvio Micali. “Probabilistic Encryption.” In: *J. Comput. Syst. Sci.* 28.2 (1984), pp. 270–299. DOI: [10.1016/0022-0000\(84\)90070-9](https://doi.org/10.1016/0022-0000(84)90070-9). URL: [https://doi.org/10.1016/0022-0000\(84\)90070-9](https://doi.org/10.1016/0022-0000(84)90070-9) (cit. on pp. 33, 97).
- [GMR85] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. “The Knowledge Complexity of Interactive Proof-Systems (Extended Abstract).” In: *17th Annual ACM Symposium on Theory of Computing*. Providence, RI, USA: ACM Press, May 1985, pp. 291–304. DOI: [10.1145/22145.22178](https://doi.org/10.1145/22145.22178) (cit. on p. 35).
- [GR07] Shafi Goldwasser and Guy N. Rothblum. “On Best-Possible Obfuscation.” In: *TCC 2007: 4th Theory of Cryptography Conference*. Ed. by Salil P. Vadhan. Vol. 4392. Lecture Notes in Computer Science. Amsterdam, The Netherlands: Springer, Heidelberg, Germany, Feb. 2007, pp. 194–213. DOI: [10.1007/978-3-540-70936-7_11](https://doi.org/10.1007/978-3-540-70936-7_11) (cit. on pp. 2, 25, 41).

- [GHKW17] Rishab Goyal, Susan Hohenberger, Venkata Koppula, and Brent Waters. “A Generic Approach to Constructing and Proving Verifiable Random Functions.” In: *TCC 2017: 15th Theory of Cryptography Conference, Part II*. Ed. by Yael Kalai and Leonid Reyzin. Vol. 10678. Lecture Notes in Computer Science. Baltimore, MD, USA: Springer, Heidelberg, Germany, Nov. 2017, pp. 537–566. DOI: [10.1007/978-3-319-70503-3_18](https://doi.org/10.1007/978-3-319-70503-3_18) (cit. on pp. 3, 41).
- [GKW17] Rishab Goyal, Venkata Koppula, and Brent Waters. “Separating Semantic and Circular Security for Symmetric-Key Bit Encryption from the Learning with Errors Assumption.” In: *Advances in Cryptology – EUROCRYPT 2017, Part II*. Ed. by Jean-Sébastien Coron and Jesper Buus Nielsen. Vol. 10211. Lecture Notes in Computer Science. Paris, France: Springer, Heidelberg, Germany, Apr. 2017, pp. 528–557. DOI: [10.1007/978-3-319-56614-6_18](https://doi.org/10.1007/978-3-319-56614-6_18) (cit. on p. 3).
- [GOS06] Jens Groth, Rafail Ostrovsky, and Amit Sahai. “Perfect Non-interactive Zero Knowledge for NP.” In: *Advances in Cryptology – EUROCRYPT 2006*. Ed. by Serge Vaudenay. Vol. 4004. Lecture Notes in Computer Science. St. Petersburg, Russia: Springer, Heidelberg, Germany, May 2006, pp. 339–358. DOI: [10.1007/11761679_21](https://doi.org/10.1007/11761679_21) (cit. on pp. 35, 36).
- [GOS12] Jens Groth, Rafail Ostrovsky, and Amit Sahai. “New Techniques for Noninteractive Zero-Knowledge.” In: *J. ACM* 59.3 (2012), 11:1–11:35. DOI: [10.1145/2220357.2220358](https://doi.org/10.1145/2220357.2220358). URL: <https://doi.org/10.1145/2220357.2220358> (cit. on p. 8).
- [GS08] Jens Groth and Amit Sahai. “Efficient Non-interactive Proof Systems for Bilinear Groups.” In: *Advances in Cryptology – EUROCRYPT 2008*. Ed. by Nigel P. Smart. Vol. 4965. Lecture Notes in Computer Science. Istanbul, Turkey: Springer, Heidelberg, Germany, Apr. 2008, pp. 415–432. DOI: [10.1007/978-3-540-78967-3_24](https://doi.org/10.1007/978-3-540-78967-3_24) (cit. on pp. 95, 96).
- [Hadoo] Satoshi Hada. “Zero-Knowledge and Code Obfuscation.” In: *Advances in Cryptology – ASIACRYPT 2000*. Ed. by Tatsuaki Okamoto. Vol. 1976. Lecture Notes in Computer Science. Kyoto, Japan: Springer, Heidelberg, Germany, Dec. 2000, pp. 443–457. DOI: [10.1007/3-540-44448-3_34](https://doi.org/10.1007/3-540-44448-3_34) (cit. on pp. 2, 25, 27, 41).
- [HT98] Satoshi Hada and Toshiaki Tanaka. “On the Existence of 3-Round Zero-Knowledge Protocols.” In: *Advances in Cryptology – CRYPTO’98*. Ed. by Hugo Krawczyk. Vol. 1462. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 1998, pp. 408–423. DOI: [10.1007/BFb0055744](https://doi.org/10.1007/BFb0055744) (cit. on pp. 32, 87).
- [HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. “A Pseudorandom Generator from any One-way Function.” In: *SIAM J. Comput.* 28.4 (1999), pp. 1364–1396. DOI: [10.1137/S0097539793244708](https://doi.org/10.1137/S0097539793244708). URL: <https://doi.org/10.1137/S0097539793244708> (cit. on pp. 61, 103, 156, 157, 162, 189, 190, 232).

- [HP14] Carmit Hazay and Arpita Patra. “One-Sided Adaptively Secure Two-Party Computation.” In: *TCC 2014: 11th Theory of Cryptography Conference*. Ed. by Yehuda Lindell. Vol. 8349. Lecture Notes in Computer Science. San Diego, CA, USA: Springer, Heidelberg, Germany, Feb. 2014, pp. 368–393. DOI: [10.1007/978-3-642-54242-8_16](https://doi.org/10.1007/978-3-642-54242-8_16) (cit. on p. 16).
- [HLOV11] Brett Hemenway, Benoît Libert, Rafail Ostrovsky, and Damien Vergnaud. “Lossy Encryption: Constructions from General Assumptions and Efficient Selective Opening Chosen Ciphertext Security.” In: *Advances in Cryptology – ASIACRYPT 2011*. Ed. by Dong Hoon Lee and Xiaoyun Wang. Vol. 7073. Lecture Notes in Computer Science. Seoul, South Korea: Springer, Heidelberg, Germany, Dec. 2011, pp. 70–88. DOI: [10.1007/978-3-642-25385-0_4](https://doi.org/10.1007/978-3-642-25385-0_4) (cit. on p. 97).
- [HJK+16] Dennis Hofheinz, Tibor Jager, Dakshita Khurana, Amit Sahai, Brent Waters, and Mark Zhandry. “How to Generate and Use Universal Samplers.” In: *Advances in Cryptology – ASIACRYPT 2016, Part II*. Ed. by Jung Hee Cheon and Tsuyoshi Takagi. Vol. 10032. Lecture Notes in Computer Science. Hanoi, Vietnam: Springer, Heidelberg, Germany, Dec. 2016, pp. 715–744. DOI: [10.1007/978-3-662-53890-6_24](https://doi.org/10.1007/978-3-662-53890-6_24) (cit. on pp. 150, 161, 171).
- [HKo8] Dennis Hofheinz and Eike Kiltz. “Programmable Hash Functions and Their Applications.” In: *Advances in Cryptology – CRYPTO 2008*. Ed. by David Wagner. Vol. 5157. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2008, pp. 21–38. DOI: [10.1007/978-3-540-85174-5_2](https://doi.org/10.1007/978-3-540-85174-5_2) (cit. on p. 44).
- [HK12] Dennis Hofheinz and Eike Kiltz. “Programmable Hash Functions and Their Applications.” In: *Journal of Cryptology* 25.3 (July 2012), pp. 484–527. DOI: [10.1007/s00145-011-9102-5](https://doi.org/10.1007/s00145-011-9102-5) (cit. on pp. 10, 90).
- [HMS07] Dennis Hofheinz, John Malone-Lee, and Martijn Stam. “Obfuscation for Cryptographic Purposes.” In: *TCC 2007: 4th Theory of Cryptography Conference*. Ed. by Salil P. Vadhan. Vol. 4392. Lecture Notes in Computer Science. Amsterdam, The Netherlands: Springer, Heidelberg, Germany, Feb. 2007, pp. 214–232. DOI: [10.1007/978-3-540-70936-7_12](https://doi.org/10.1007/978-3-540-70936-7_12) (cit. on p. 41).
- [HRW16] Dennis Hofheinz, Vanishree Rao, and Daniel Wichs. “Standard Security Does Not Imply Indistinguishability Under Selective Opening.” In: *TCC 2016-B: 14th Theory of Cryptography Conference, Part II*. Ed. by Martin Hirt and Adam D. Smith. Vol. 9986. Lecture Notes in Computer Science. Beijing, China: Springer, Heidelberg, Germany, Oct. 2016, pp. 121–145. DOI: [10.1007/978-3-662-53644-5_5](https://doi.org/10.1007/978-3-662-53644-5_5) (cit. on p. 41).
- [HU19] Dennis Hofheinz and Bogdan Ursu. “Dual-Mode NIZKs from Obfuscation.” In: *Advances in Cryptology – ASIACRYPT 2019, Part I*. Ed. by Steven D. Galbraith and Shiho Moriai. Vol. 11921. Lecture Notes in Computer Science. Kobe, Japan: Springer,

- Heidelberg, Germany, Dec. 2019, pp. 311–341. DOI: [10.1007/978-3-030-34578-5_12](https://doi.org/10.1007/978-3-030-34578-5_12) (cit. on pp. [89](#), [96](#)).
- [HRsV07] Susan Hohenberger, Guy N. Rothblum, abhi shelat, and Vinod Vaikuntanathan. “Securely Obfuscating Re-encryption.” In: *TCC 2007: 4th Theory of Cryptography Conference*. Ed. by Salil P. Vadhan. Vol. 4392. Lecture Notes in Computer Science. Amsterdam, The Netherlands: Springer, Heidelberg, Germany, Feb. 2007, pp. 233–252. DOI: [10.1007/978-3-540-70936-7_13](https://doi.org/10.1007/978-3-540-70936-7_13) (cit. on p. [41](#)).
- [HSW13] Susan Hohenberger, Amit Sahai, and Brent Waters. “Full Domain Hash from (Leveled) Multilinear Maps and Identity-Based Aggregate Signatures.” In: *Advances in Cryptology – CRYPTO 2013, Part I*. Ed. by Ran Canetti and Juan A. Garay. Vol. 8042. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2013, pp. 494–512. DOI: [10.1007/978-3-642-40041-4_27](https://doi.org/10.1007/978-3-642-40041-4_27) (cit. on pp. [10](#), [90](#)).
- [HSW14] Susan Hohenberger, Amit Sahai, and Brent Waters. “Replacing a Random Oracle: Full Domain Hash from Indistinguishability Obfuscation.” In: *Advances in Cryptology – EUROCRYPT 2014*. Ed. by Phong Q. Nguyen and Elisabeth Oswald. Vol. 8441. Lecture Notes in Computer Science. Copenhagen, Denmark: Springer, Heidelberg, Germany, May 2014, pp. 201–220. DOI: [10.1007/978-3-642-55220-5_12](https://doi.org/10.1007/978-3-642-55220-5_12) (cit. on pp. [10](#), [44](#), [90](#)).
- [HW09] Susan Hohenberger and Brent Waters. “Short and Stateless Signatures from the RSA Assumption.” In: *Advances in Cryptology – CRYPTO 2009*. Ed. by Shai Halevi. Vol. 5677. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2009, pp. 654–670. DOI: [10.1007/978-3-642-03356-8_38](https://doi.org/10.1007/978-3-642-03356-8_38) (cit. on p. [44](#)).
- [HL18] Justin Holmgren and Alex Lombardi. “Cryptographic Hashing from Strong One-Way Functions (Or: One-Way Product Functions and Their Applications).” In: *59th Annual Symposium on Foundations of Computer Science*. Ed. by Mikkel Thorup. Paris, France: IEEE Computer Society Press, Oct. 2018, pp. 850–858. DOI: [10.1109/FOCS.2018.00085](https://doi.org/10.1109/FOCS.2018.00085) (cit. on p. [3](#)).
- [Hop05] Nicholas Hopper. “On Steganographic Chosen Coverttext Security.” In: *ICALP 2005: 32nd International Colloquium on Automata, Languages and Programming*. Ed. by Luís Caires, Giuseppe F. Italiano, Luís Monteiro, Catuscia Palamidessi, and Moti Yung. Vol. 3580. Lecture Notes in Computer Science. Lisbon, Portugal: Springer, Heidelberg, Germany, July 2005, pp. 311–323. DOI: [10.1007/11523468_26](https://doi.org/10.1007/11523468_26) (cit. on pp. [232](#), [233](#)).
- [HPRV19] Thibaut Horel, Sunoo Park, Silas Richelson, and Vinod Vaikuntanathan. “How to Subvert Backdoored Encryption: Security Against Adversaries that Decrypt All Ciphertexts.” In: *ITCS 2019: 10th Innovations in Theoretical Computer Science Conference*. Ed. by Avrim Blum. Vol. 124. San Diego, CA, USA: LIPIcs, Jan. 2019, 42:1–42:20. DOI: [10.4230/LIPIcs.ITCS.2019.42](https://doi.org/10.4230/LIPIcs.ITCS.2019.42) (cit. on pp. [13](#), [147](#)).

- [HLR07] Chun-Yuan Hsiao, Chi-Jen Lu, and Leonid Reyzin. “Conditional Computational Entropy, or Toward Separating Pseudentropy from Compressibility.” In: *Advances in Cryptology – EUROCRYPT 2007*. Ed. by Moni Naor. Vol. 4515. Lecture Notes in Computer Science. Barcelona, Spain: Springer, Heidelberg, Germany, May 2007, pp. 169–186. DOI: [10.1007/978-3-540-72540-4_10](https://doi.org/10.1007/978-3-540-72540-4_10) (cit. on pp. [12](#), [20](#), [147](#), [153](#), [157](#), [163](#), [189](#), [190](#), [199](#), [200](#)).
- [HW15] Pavel Hubacek and Daniel Wichs. “On the Communication Complexity of Secure Function Evaluation with Long Output.” In: *ITCS 2015: 6th Conference on Innovations in Theoretical Computer Science*. Ed. by Tim Roughgarden. Rehovot, Israel: Association for Computing Machinery, Jan. 2015, pp. 163–172. DOI: [10.1145/2688073.2688105](https://doi.org/10.1145/2688073.2688105) (cit. on p. [3](#)).
- [Huf52] David A Huffman. “A method for the construction of minimum-redundancy codes.” In: *Proceedings of the IRE* 40.9 (1952), pp. 1098–1101 (cit. on p. [12](#)).
- [Imp95] Russell Impagliazzo. “A Personal View of Average-Case Complexity.” In: *Proceedings of the Tenth Annual Structure in Complexity Theory Conference, Minneapolis, Minnesota, USA, June 19-22, 1995*. IEEE Computer Society, 1995, pp. 134–147. DOI: [10.1109/SCT.1995.514853](https://doi.org/10.1109/SCT.1995.514853). URL: <https://doi.org/10.1109/SCT.1995.514853> (cit. on p. [156](#)).
- [IL90] Russell Impagliazzo and Leonid A. Levin. “No Better Ways to Generate Hard NP Instances than Picking Uniformly at Random.” In: *31st Annual Symposium on Foundations of Computer Science*. St. Louis, MO, USA: IEEE Computer Society Press, Oct. 1990, pp. 812–821. DOI: [10.1109/FSCS.1990.89604](https://doi.org/10.1109/FSCS.1990.89604) (cit. on p. [21](#)).
- [ILL89] Russell Impagliazzo, Leonid A. Levin, and Michael Luby. “Pseudo-random Generation from one-way functions (Extended Abstracts).” In: *21st Annual ACM Symposium on Theory of Computing*. Seattle, WA, USA: ACM Press, May 1989, pp. 12–24. DOI: [10.1145/73007.73009](https://doi.org/10.1145/73007.73009) (cit. on p. [20](#)).
- [IR89] Russell Impagliazzo and Steven Rudich. “Limits on the Provable Consequences of One-Way Permutations.” In: *21st Annual ACM Symposium on Theory of Computing*. Seattle, WA, USA: ACM Press, May 1989, pp. 44–61. DOI: [10.1145/73007.73012](https://doi.org/10.1145/73007.73012) (cit. on p. [33](#)).
- [IKOS10] Yuval Ishai, Abishek Kumarasubramanian, Claudio Orlandi, and Amit Sahai. “On Invertible Sampling and Adaptive Security.” In: *Advances in Cryptology – ASIACRYPT 2010*. Ed. by Masayuki Abe. Vol. 6477. Lecture Notes in Computer Science. Singapore: Springer, Heidelberg, Germany, Dec. 2010, pp. 466–482. DOI: [10.1007/978-3-642-17373-8_27](https://doi.org/10.1007/978-3-642-17373-8_27) (cit. on pp. [16](#), [17](#), [151–154](#), [158](#), [159](#), [177](#), [201](#), [204](#), [205](#), [212](#), [221](#), [224](#), [226](#)).

- [IPS15] Yuval Ishai, Omkant Pandey, and Amit Sahai. “Public-Coin Differing-Inputs Obfuscation and Its Applications.” In: *TCC 2015: 12th Theory of Cryptography Conference, Part II*. Ed. by Yevgeniy Dodis and Jesper Buus Nielsen. Vol. 9015. Lecture Notes in Computer Science. Warsaw, Poland: Springer, Heidelberg, Germany, Mar. 2015, pp. 668–697. DOI: [10.1007/978-3-662-46497-7_26](https://doi.org/10.1007/978-3-662-46497-7_26) (cit. on p. 41).
- [IPSo8] Yuval Ishai, Manoj Prabhakaran, and Amit Sahai. “Founding Cryptography on Oblivious Transfer - Efficiently.” In: *Advances in Cryptology – CRYPTO 2008*. Ed. by David Wagner. Vol. 5157. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2008, pp. 572–591. DOI: [10.1007/978-3-540-85174-5_32](https://doi.org/10.1007/978-3-540-85174-5_32) (cit. on pp. 16, 224).
- [JRT16] Joseph Jaeger, Thomas Ristenpart, and Qiang Tang. “Honey Encryption Beyond Message Recovery Security.” In: *Advances in Cryptology – EUROCRYPT 2016, Part I*. Ed. by Marc Fischlin and Jean-Sébastien Coron. Vol. 9665. Lecture Notes in Computer Science. Vienna, Austria: Springer, Heidelberg, Germany, May 2016, pp. 758–788. DOI: [10.1007/978-3-662-49890-3_29](https://doi.org/10.1007/978-3-662-49890-3_29) (cit. on p. 230).
- [JLMS19] Aayush Jain, Huijia Lin, Christian Matt, and Amit Sahai. “How to Leverage Hardness of Constant-Degree Expanding Polynomials over \mathbb{R} to build iO .” In: *Advances in Cryptology – EUROCRYPT 2019, Part I*. Ed. by Yuval Ishai and Vincent Rijmen. Vol. 11476. Lecture Notes in Computer Science. Darmstadt, Germany: Springer, Heidelberg, Germany, May 2019, pp. 251–281. DOI: [10.1007/978-3-030-17653-2_9](https://doi.org/10.1007/978-3-030-17653-2_9) (cit. on pp. 2, 41).
- [JLS20] Aayush Jain, Huijia Lin, and Amit Sahai. *Indistinguishability Obfuscation from Well-Founded Assumptions*. Cryptology ePrint Archive, Report 2020/1003. <https://eprint.iacr.org/2020/1003>. 2020 (cit. on pp. 2, 5, 27, 41).
- [Jak98] Markus Jakobsson. “A Practical Mix.” In: *Advances in Cryptology – EUROCRYPT’98*. Ed. by Kaisa Nyberg. Vol. 1403. Lecture Notes in Computer Science. Espoo, Finland: Springer, Heidelberg, Germany, May 1998, pp. 448–461. DOI: [10.1007/BFb0054145](https://doi.org/10.1007/BFb0054145) (cit. on pp. 89, 129).
- [JR14] Ari Juels and Thomas Ristenpart. “Honey Encryption: Security Beyond the Brute-Force Bound.” In: *Advances in Cryptology – EUROCRYPT 2014*. Ed. by Phong Q. Nguyen and Elisabeth Oswald. Vol. 8441. Lecture Notes in Computer Science. Copenhagen, Denmark: Springer, Heidelberg, Germany, May 2014, pp. 293–310. DOI: [10.1007/978-3-642-55220-5_17](https://doi.org/10.1007/978-3-642-55220-5_17) (cit. on pp. 13, 147, 153, 156, 158, 221, 229–231).
- [KLMR18] Yuan Kang, Chengyu Lin, Tal Malkin, and Mariana Raykova. “Obfuscation from Polynomial Hardness: Beyond Decomposable Obfuscation.” In: *SCN 18: 11th International Conference on Security in Communication Networks*. Ed. by Dario Catalano and Roberto De Prisco. Vol. 11035. Lecture Notes in Computer Science. Amalfi, Italy: Springer, Heidelberg, Germany, Sept. 2018,

- pp. 407–424. DOI: [10.1007/978-3-319-98113-0_22](https://doi.org/10.1007/978-3-319-98113-0_22) (cit. on p. 5).
- [KLRX20] Julia Kastner, Julian Loss, Michael Rosenberg, and Jiayu Xu. *On Pairing-Free Blind Signature Schemes in the Algebraic Group Model*. Cryptology ePrint Archive, Report 2020/1071. <https://eprint.iacr.org/2020/1071>. 2020 (cit. on p. 245).
- [KP19] Julia Kastner and Jiaxin Pan. *Towards Instantiating the Algebraic Group Model*. Cryptology ePrint Archive, Report 2019/1018. <https://eprint.iacr.org/2019/1018>. 2019 (cit. on pp. 9, 88, 92, 94, 109, 110).
- [Kato8] Jonathan Katz. “Which Languages Have 4-Round Zero-Knowledge Proofs?” In: *TCC 2008: 5th Theory of Cryptography Conference*. Ed. by Ran Canetti. Vol. 4948. Lecture Notes in Computer Science. San Francisco, CA, USA: Springer, Heidelberg, Germany, Mar. 2008, pp. 73–88. DOI: [10.1007/978-3-540-78524-8_5](https://doi.org/10.1007/978-3-540-78524-8_5) (cit. on p. 32).
- [KO04] Jonathan Katz and Rafail Ostrovsky. “Round-Optimal Secure Two-Party Computation.” In: *Advances in Cryptology – CRYPTO 2004*. Ed. by Matthew Franklin. Vol. 3152. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2004, pp. 335–354. DOI: [10.1007/978-3-540-28628-8_21](https://doi.org/10.1007/978-3-540-28628-8_21) (cit. on p. 16).
- [KPTZ13] Aggelos Kiayias, Stavros Papadopoulos, Nikos Triandopoulos, and Thomas Zacharias. “Delegatable pseudorandom functions and applications.” In: *ACM CCS 2013: 20th Conference on Computer and Communications Security*. Ed. by Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung. Berlin, Germany: ACM Press, Nov. 2013, pp. 669–684. DOI: [10.1145/2508859.2516668](https://doi.org/10.1145/2508859.2516668) (cit. on pp. 23, 24).
- [Kil88] Joe Kilian. “Founding Cryptography on Oblivious Transfer.” In: *20th Annual ACM Symposium on Theory of Computing*. Chicago, IL, USA: ACM Press, May 1988, pp. 20–31. DOI: [10.1145/62212.62215](https://doi.org/10.1145/62212.62215) (cit. on p. 224).
- [Kol68] Andrei Nikolaevich Kolmogorov. “Three approaches to the quantitative definition of information.” In: *International journal of computer mathematics* 2.1-4 (1968), pp. 157–168 (cit. on pp. 13, 147).
- [KNY14] Ilan Komargodski, Moni Naor, and Eylon Yogev. “Secret-Sharing for NP.” In: *Advances in Cryptology – ASIACRYPT 2014, Part II*. Ed. by Palash Sarkar and Tetsu Iwata. Vol. 8874. Lecture Notes in Computer Science. Kaoshiung, Taiwan, R.O.C.: Springer, Heidelberg, Germany, Dec. 2014, pp. 254–273. DOI: [10.1007/978-3-662-45608-8_14](https://doi.org/10.1007/978-3-662-45608-8_14) (cit. on p. 3).
- [KLW15] Venkata Koppula, Allison Bishop Lewko, and Brent Waters. “Indistinguishability Obfuscation for Turing Machines with Unbounded Memory.” In: *47th Annual ACM Symposium on Theory of Computing*. Ed. by Rocco A. Servedio and Ronitt Rubinfeld. Portland, OR, USA: ACM Press, June 2015, pp. 419–428. DOI: [10.1145/2746539.2746614](https://doi.org/10.1145/2746539.2746614) (cit. on p. 3).

- [LMs05] Matt Lepinski, Silvio Micali, and abhi shelat. “Fair-Zero Knowledge.” In: *TCC 2005: 2nd Theory of Cryptography Conference*. Ed. by Joe Kilian. Vol. 3378. Lecture Notes in Computer Science. Cambridge, MA, USA: Springer, Heidelberg, Germany, Feb. 2005, pp. 245–263. DOI: [10.1007/978-3-540-30576-7_14](https://doi.org/10.1007/978-3-540-30576-7_14) (cit. on pp. [153](#), [163](#), [199](#)).
- [Lev86] Leonid A. Levin. “Average Case Complete Problems.” In: *SIAM J. Comput.* 15.1 (1986), pp. 285–286. DOI: [10.1137/0215020](https://doi.org/10.1137/0215020). URL: <https://doi.org/10.1137/0215020> (cit. on pp. [1](#), [21](#)).
- [LM16] Baiyu Li and Daniele Micciancio. “Compactness vs Collusion Resistance in Functional Encryption.” In: *TCC 2016-B: 14th Theory of Cryptography Conference, Part II*. Ed. by Martin Hirt and Adam D. Smith. Vol. 9986. Lecture Notes in Computer Science. Beijing, China: Springer, Heidelberg, Germany, Oct. 2016, pp. 443–468. DOI: [10.1007/978-3-662-53644-5_17](https://doi.org/10.1007/978-3-662-53644-5_17) (cit. on p. [42](#)).
- [LV90] Ming Li and Paul M. B. Vitányi. “Handbook of Theoretical Computer Science (Vol. A).” In: ed. by Jan van Leeuwen. Cambridge, MA, USA: MIT Press, 1990. Chap. Kolmogorov Complexity and Its Applications, pp. 187–254. ISBN: 0-444-88071-2. URL: <http://dl.acm.org/citation.cfm?id=114872.114876> (cit. on pp. [13](#), [147](#)).
- [LV19] Ming Li and Paul M. B. Vitányi. *An Introduction to Kolmogorov Complexity and Its Applications, 4th Edition*. Texts in Computer Science. Springer, 2019. ISBN: 978-3-030-11297-4. DOI: [10.1007/978-3-030-11298-1](https://doi.org/10.1007/978-3-030-11298-1) (cit. on pp. [13](#), [147](#)).
- [Lin16] Huijia Lin. “Indistinguishability Obfuscation from Constant-Degree Graded Encoding Schemes.” In: *Advances in Cryptology – EUROCRYPT 2016, Part I*. Ed. by Marc Fischlin and Jean-Sébastien Coron. Vol. 9665. Lecture Notes in Computer Science. Vienna, Austria: Springer, Heidelberg, Germany, May 2016, pp. 28–57. DOI: [10.1007/978-3-662-49890-3_2](https://doi.org/10.1007/978-3-662-49890-3_2) (cit. on p. [2](#)).
- [Lin17] Huijia Lin. “Indistinguishability Obfuscation from SXDH on 5-Linear Maps and Locality-5 PRGs.” In: *Advances in Cryptology – CRYPTO 2017, Part I*. Ed. by Jonathan Katz and Hovav Shacham. Vol. 10401. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2017, pp. 599–629. DOI: [10.1007/978-3-319-63688-7_20](https://doi.org/10.1007/978-3-319-63688-7_20) (cit. on pp. [2](#), [41](#)).
- [LT17] Huijia Lin and Stefano Tessaro. “Indistinguishability Obfuscation from Trilinear Maps and Block-Wise Local PRGs.” In: *Advances in Cryptology – CRYPTO 2017, Part I*. Ed. by Jonathan Katz and Hovav Shacham. Vol. 10401. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2017, pp. 630–660. DOI: [10.1007/978-3-319-63688-7_21](https://doi.org/10.1007/978-3-319-63688-7_21) (cit. on pp. [2](#), [41](#)).

- [LV16] Huijia Lin and Vinod Vaikuntanathan. “Indistinguishability Obfuscation from DDH-Like Assumptions on Constant-Degree Graded Encodings.” In: *57th Annual Symposium on Foundations of Computer Science*. Ed. by Irit Dinur. New Brunswick, NJ, USA: IEEE Computer Society Press, Oct. 2016, pp. 11–20. DOI: [10.1109/FOCS.2016.11](https://doi.org/10.1109/FOCS.2016.11) (cit. on p. 2).
- [Lin09] Andrew Y. Lindell. “Adaptively Secure Two-Party Computation with Erasures.” In: *Topics in Cryptology – CT-RSA 2009*. Ed. by Marc Fischlin. Vol. 5473. Lecture Notes in Computer Science. San Francisco, CA, USA: Springer, Heidelberg, Germany, Apr. 2009, pp. 117–132. DOI: [10.1007/978-3-642-00862-7_8](https://doi.org/10.1007/978-3-642-00862-7_8) (cit. on p. 222).
- [LZ17] Qipeng Liu and Mark Zhandry. “Decomposable Obfuscation: A Framework for Building Applications of Obfuscation from Polynomial Hardness.” In: *TCC 2017: 15th Theory of Cryptography Conference, Part I*. Ed. by Yael Kalai and Leonid Reyzin. Vol. 10677. Lecture Notes in Computer Science. Baltimore, MD, USA: Springer, Heidelberg, Germany, Nov. 2017, pp. 138–169. DOI: [10.1007/978-3-319-70500-2_6](https://doi.org/10.1007/978-3-319-70500-2_6) (cit. on pp. 5, 42, 44).
- [LPS04] Ben Lynn, Manoj Prabhakaran, and Amit Sahai. “Positive Results and Techniques for Obfuscation.” In: *Advances in Cryptology – EUROCRYPT 2004*. Ed. by Christian Cachin and Jan Camenisch. Vol. 3027. Lecture Notes in Computer Science. Interlaken, Switzerland: Springer, Heidelberg, Germany, May 2004, pp. 20–39. DOI: [10.1007/978-3-540-24676-3_2](https://doi.org/10.1007/978-3-540-24676-3_2) (cit. on p. 41).
- [Mau05] Ueli M. Maurer. “Abstract Models of Computation in Cryptography (Invited Paper).” In: *10th IMA International Conference on Cryptography and Coding*. Ed. by Nigel P. Smart. Vol. 3796. Lecture Notes in Computer Science. Cirencester, UK: Springer, Heidelberg, Germany, Dec. 2005, pp. 1–12 (cit. on pp. 8, 87, 93).
- [MW98] Ueli M. Maurer and Stefan Wolf. “Lower Bounds on Generic Algorithms in Groups.” In: *Advances in Cryptology – EUROCRYPT’98*. Ed. by Kaisa Nyberg. Vol. 1403. Lecture Notes in Computer Science. Espoo, Finland: Springer, Heidelberg, Germany, May 1998, pp. 72–84. DOI: [10.1007/BFb0054118](https://doi.org/10.1007/BFb0054118) (cit. on p. 87).
- [MG02] Daniele Micciancio and Shafi Goldwasser. *Complexity of lattice problems - a cryptographic perspective*. Vol. 671. The Kluwer international series in engineering and computer science. Springer, 2002. ISBN: 978-0-7923-7688-0 (cit. on p. 106).
- [MSZ16] Eric Miles, Amit Sahai, and Mark Zhandry. “Annihilation Attacks for Multilinear Maps: Cryptanalysis of Indistinguishability Obfuscation over GGH_{13} .” In: *Advances in Cryptology – CRYPTO 2016, Part II*. Ed. by Matthew Robshaw and Jonathan Katz. Vol. 9815. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2016,

- pp. 629–658. DOI: [10.1007/978-3-662-53008-5_22](https://doi.org/10.1007/978-3-662-53008-5_22) (cit. on p. 2).
- [Na003] Moni Naor. “On Cryptographic Assumptions and Challenges (Invited Talk).” In: *Advances in Cryptology – CRYPTO 2003*. Ed. by Dan Boneh. Vol. 2729. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2003, pp. 96–109. DOI: [10.1007/978-3-540-45146-4_6](https://doi.org/10.1007/978-3-540-45146-4_6) (cit. on pp. 10, 22, 27, 32, 88, 94).
- [Nec94] V. I. Nechaev. “Complexity of a Determinate Algorithm for the Discrete Logarithm.” In: *Mathematical Notes* 55.2 (1994), pp. 165–172 (cit. on p. 87).
- [NZ96] Noam Nisan and David Zuckerman. “Randomness is Linear in Space.” In: *J. Comput. Syst. Sci.* 52.1 (1996), pp. 43–52. DOI: [10.1006/jcss.1996.0004](https://doi.org/10.1006/jcss.1996.0004). URL: <https://doi.org/10.1006/jcss.1996.0004> (cit. on p. 20).
- [OPWW15] Tatsuaki Okamoto, Krzysztof Pietrzak, Brent Waters, and Daniel Wichs. “New Realizations of Somewhere Statistically Binding Hashing and Positional Accumulators.” In: *Advances in Cryptology – ASIACRYPT 2015, Part I*. Ed. by Tetsu Iwata and Jung Hee Cheon. Vol. 9452. Lecture Notes in Computer Science. Auckland, New Zealand: Springer, Heidelberg, Germany, Nov. 2015, pp. 121–145. DOI: [10.1007/978-3-662-48797-6_6](https://doi.org/10.1007/978-3-662-48797-6_6) (cit. on p. 3).
- [OW93] Rafail Ostrovsky and Avi Wigderson. “One-Way Functions are Essential for Non-Trivial Zero-Knowledge.” In: *Second Israel Symposium on Theory of Computing Systems, ISTCS 1993, Natanya, Israel, June 7-9, 1993, Proceedings*. IEEE Computer Society, 1993, pp. 3–17. DOI: [10.1109/ISTCS.1993.253489](https://doi.org/10.1109/ISTCS.1993.253489). URL: <https://doi.org/10.1109/ISTCS.1993.253489> (cit. on p. 35).
- [Pai99] Pascal Paillier. “Public-Key Cryptosystems Based on Composite Degree Residuosity Classes.” In: *Advances in Cryptology – EUROCRYPT’99*. Ed. by Jacques Stern. Vol. 1592. Lecture Notes in Computer Science. Prague, Czech Republic: Springer, Heidelberg, Germany, May 1999, pp. 223–238. DOI: [10.1007/3-540-48910-X_16](https://doi.org/10.1007/3-540-48910-X_16) (cit. on p. 97).
- [PV05] Pascal Paillier and Damien Vergnaud. “Discrete-Log-Based Signatures May Not Be Equivalent to Discrete Log.” In: *Advances in Cryptology – ASIACRYPT 2005*. Ed. by Bimal K. Roy. Vol. 3788. Lecture Notes in Computer Science. Chennai, India: Springer, Heidelberg, Germany, Dec. 2005, pp. 1–20. DOI: [10.1007/11593447_1](https://doi.org/10.1007/11593447_1) (cit. on p. 87).
- [Pas77] Richard C. Pasco. “Source coding algorithms for fast data compression (Ph.D. Thesis abstr.)” In: *IEEE Trans. Inf. Theory* 23.4 (1977), p. 548. DOI: [10.1109/TIT.1977.1055739](https://doi.org/10.1109/TIT.1977.1055739). URL: <https://doi.org/10.1109/TIT.1977.1055739> (cit. on p. 12).

- [Pas03] Rafael Pass. “Simulation in Quasi-Polynomial Time, and Its Application to Protocol Composition.” In: *Advances in Cryptology – EUROCRYPT 2003*. Ed. by Eli Biham. Vol. 2656. Lecture Notes in Computer Science. Warsaw, Poland: Springer, Heidelberg, Germany, May 2003, pp. 160–176. DOI: [10.1007/3-540-39200-9_10](https://doi.org/10.1007/3-540-39200-9_10) (cit. on p. 241).
- [PST14] Rafael Pass, Karn Seth, and Sidharth Telang. “Indistinguishability Obfuscation from Semantically-Secure Multilinear Encodings.” In: *Advances in Cryptology – CRYPTO 2014, Part I*. Ed. by Juan A. Garay and Rosario Gennaro. Vol. 8616. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2014, pp. 500–517. DOI: [10.1007/978-3-662-44371-2_28](https://doi.org/10.1007/978-3-662-44371-2_28) (cit. on pp. 41, 43).
- [Ps16] Rafael Pass and abhi shelat. “Impossibility of VBB Obfuscation with Ideal Constant-Degree Graded Encodings.” In: *TCC 2016-A: 13th Theory of Cryptography Conference, Part I*. Ed. by Eyal Kushilevitz and Tal Malkin. Vol. 9562. Lecture Notes in Computer Science. Tel Aviv, Israel: Springer, Heidelberg, Germany, Jan. 2016, pp. 3–17. DOI: [10.1007/978-3-662-49096-9_1](https://doi.org/10.1007/978-3-662-49096-9_1) (cit. on p. 42).
- [PS19] Chris Peikert and Sina Shiehian. “Noninteractive Zero Knowledge for NP from (Plain) Learning with Errors.” In: *Advances in Cryptology – CRYPTO 2019, Part I*. Ed. by Alexandra Boldyreva and Daniele Micciancio. Vol. 11692. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2019, pp. 89–114. DOI: [10.1007/978-3-030-26948-7_4](https://doi.org/10.1007/978-3-030-26948-7_4) (cit. on pp. 3, 36, 96).
- [PR07] Manoj Prabhakaran and Mike Rosulek. “Rerandomizable RCCA Encryption.” In: *Advances in Cryptology – CRYPTO 2007*. Ed. by Alfred Menezes. Vol. 4622. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2007, pp. 517–534. DOI: [10.1007/978-3-540-74143-5_29](https://doi.org/10.1007/978-3-540-74143-5_29) (cit. on p. 97).
- [RS92] Charles Rackoff and Daniel R. Simon. “Non-Interactive Zero-Knowledge Proof of Knowledge and Chosen Ciphertext Attack.” In: *Advances in Cryptology – CRYPTO’91*. Ed. by Joan Feigenbaum. Vol. 576. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 1992, pp. 433–444. DOI: [10.1007/3-540-46766-1_35](https://doi.org/10.1007/3-540-46766-1_35) (cit. on p. 129).
- [RR99] Ran Raz and Omer Reingold. “On Recycling the Randomness of States in Space Bounded Computation.” In: *31st Annual ACM Symposium on Theory of Computing*. Atlanta, GA, USA: ACM Press, May 1999, pp. 159–168. DOI: [10.1145/301250.301294](https://doi.org/10.1145/301250.301294) (cit. on pp. 13, 147).
- [Rey11] Leonid Reyzin. “Some Notions of Entropy for Cryptography - (Invited Talk).” In: *ICITS 11: 5th International Conference on Information Theoretic Security*. Ed. by Serge Fehr. Vol. 6673. Lecture Notes in Computer Science. Amsterdam, The Netherlands:

- Springer, Heidelberg, Germany, May 2011, pp. 138–142. DOI: [10.1007/978-3-642-20728-0_13](https://doi.org/10.1007/978-3-642-20728-0_13) (cit. on p. 189).
- [Ris76] Jorma Rissanen. “Generalized Kraft Inequality and Arithmetic Coding.” In: *IBM J. Res. Dev.* 20.3 (1976), pp. 198–203. DOI: [10.1147/rd.203.0198](https://doi.org/10.1147/rd.203.0198). URL: <https://doi.org/10.1147/rd.203.0198> (cit. on p. 12).
- [RJ79] Jorma Rissanen and Glen G. Langdon Jr. “Arithmetic Coding.” In: *IBM J. Res. Dev.* 23.2 (1979), pp. 149–162. DOI: [10.1147/rd.232.0149](https://doi.org/10.1147/rd.232.0149). URL: <https://doi.org/10.1147/rd.232.0149> (cit. on p. 12).
- [RAD78] Ronald L. Rivest, Leonard M. Adleman, and Michael L. Derouzos. “On data banks and privacy homomorphisms.” In: *Foundations of secure computation* 4.11 (1978), pp. 169–180 (cit. on p. 34).
- [RSA78] Ronald L. Rivest, Adi Shamir, and Leonard M. Adleman. “A Method for Obtaining Digital Signatures and Public-Key Cryptosystems.” In: *Commun. ACM* 21.2 (1978), pp. 120–126. DOI: [10.1145/359340.359342](https://doi.org/10.1145/359340.359342). URL: <http://doi.acm.org/10.1145/359340.359342> (cit. on p. 2).
- [RS20] Lior Rotem and Gil Segev. “Algebraic Distinguishers: From Discrete Logarithms to Decisional Uber Assumptions.” In: (2020) (cit. on p. 245).
- [SW14] Amit Sahai and Brent Waters. “How to use indistinguishability obfuscation: deniable encryption, and more.” In: *46th Annual ACM Symposium on Theory of Computing*. Ed. by David B. Shmoys. New York, NY, USA: ACM Press, May 2014, pp. 475–484. DOI: [10.1145/2591796.2591825](https://doi.org/10.1145/2591796.2591825) (cit. on pp. 3, 24, 41, 43, 103, 104, 152, 153, 155, 189, 200, 212–215, 233).
- [San87] Miklos Santha. “On Using Deterministic Functions to Reduce Randomness in Probabilistic Algorithms.” In: *Inf. Comput.* 74.3 (1987), pp. 241–249. DOI: [10.1016/0890-5401\(87\)90023-X](https://doi.org/10.1016/0890-5401(87)90023-X). URL: [https://doi.org/10.1016/0890-5401\(87\)90023-X](https://doi.org/10.1016/0890-5401(87)90023-X) (cit. on p. 20).
- [Sch72] J. Pieter M. Schalkwijk. “An algorithm for source coding.” In: *IEEE Trans. Inf. Theory* 18.3 (1972), pp. 395–399. DOI: [10.1109/TIT.1972.1054832](https://doi.org/10.1109/TIT.1972.1054832). URL: <https://doi.org/10.1109/TIT.1972.1054832> (cit. on p. 12).
- [Sch91] Claus-Peter Schnorr. “Efficient Signature Generation by Smart Cards.” In: *Journal of Cryptology* 4.3 (Jan. 1991), pp. 161–174. DOI: [10.1007/BF00196725](https://doi.org/10.1007/BF00196725) (cit. on p. 89).
- [SJ00] Claus-Peter Schnorr and Markus Jakobsson. “Security of Signed ElGamal Encryption.” In: *Advances in Cryptology – ASIACRYPT 2000*. Ed. by Tatsuaki Okamoto. Vol. 1976. Lecture Notes in Computer Science. Kyoto, Japan: Springer, Heidelberg, Germany, Dec. 2000, pp. 73–89. DOI: [10.1007/3-540-44448-3_7](https://doi.org/10.1007/3-540-44448-3_7) (cit. on p. 129).

- [Sha48] Claude E. Shannon. “A mathematical theory of communication.” In: *Bell Syst. Tech. J.* 27.3 (1948), pp. 379–423. DOI: [10.1002/j.1538-7305.1948.tb01338.x](https://doi.org/10.1002/j.1538-7305.1948.tb01338.x). URL: <https://doi.org/10.1002/j.1538-7305.1948.tb01338.x> (cit. on pp. 12, 162, 190).
- [TUZ01] Amnon Ta-Shma, Christopher Umans, and David Zuckerman. “Loss-less condensers, unbalanced expanders, and extractors.” In: *33rd Annual ACM Symposium on Theory of Computing*. Crete, Greece: ACM Press, July 2001, pp. 143–152. DOI: [10.1145/380752.380790](https://doi.org/10.1145/380752.380790) (cit. on pp. 13, 147).
- [Sho97] Victor Shoup. “Lower Bounds for Discrete Logarithms and Related Problems.” In: *Advances in Cryptology – EUROCRYPT’97*. Ed. by Walter Fumy. Vol. 1233. Lecture Notes in Computer Science. Konstanz, Germany: Springer, Heidelberg, Germany, May 1997, pp. 256–266. DOI: [10.1007/3-540-69053-0_18](https://doi.org/10.1007/3-540-69053-0_18) (cit. on pp. 8, 10, 87, 93).
- [Sho04] Victor Shoup. *Sequences of games: a tool for taming complexity in security proofs*. Cryptology ePrint Archive, Report 2004/332. <http://eprint.iacr.org/2004/332>. 2004 (cit. on pp. 172, 185, 186).
- [Sip88] Michael Sipser. “Expanders, Randomness, or Time versus Space.” In: *J. Comput. Syst. Sci.* 36.3 (1988), pp. 379–383. DOI: [10.1016/0022-0000\(88\)90035-9](https://doi.org/10.1016/0022-0000(88)90035-9). URL: [https://doi.org/10.1016/0022-0000\(88\)90035-9](https://doi.org/10.1016/0022-0000(88)90035-9) (cit. on p. 20).
- [Sol64] Ray J. Solomonoff. “A Formal Theory of Inductive Inference. Part I.” In: *Inf. Control.* 7.1 (1964), pp. 1–22. DOI: [10.1016/S0019-9958\(64\)90223-2](https://doi.org/10.1016/S0019-9958(64)90223-2). URL: [https://doi.org/10.1016/S0019-9958\(64\)90223-2](https://doi.org/10.1016/S0019-9958(64)90223-2) (cit. on pp. 13, 147).
- [TV00] Luca Trevisan and Salil P. Vadhan. “Extracting Randomness from Samplable Distributions.” In: *41st Annual Symposium on Foundations of Computer Science*. Redondo Beach, CA, USA: IEEE Computer Society Press, Nov. 2000, pp. 32–42. DOI: [10.1109/SFCS.2000.892063](https://doi.org/10.1109/SFCS.2000.892063) (cit. on pp. 13, 147).
- [TVZ05] Luca Trevisan, Salil P. Vadhan, and David Zuckerman. “Compression of Samplable Sources.” In: *Computational Complexity* 14.3 (2005), pp. 186–227. DOI: [10.1007/s00037-005-0198-6](https://doi.org/10.1007/s00037-005-0198-6). URL: <https://doi.org/10.1007/s00037-005-0198-6> (cit. on pp. 12, 13, 147, 156, 162, 168, 190, 191).
- [TY98] Yiannis Tsiounis and Moti Yung. “On the Security of ElGamal Based Encryption.” In: *PKC’98: 1st International Workshop on Theory and Practice in Public Key Cryptography*. Ed. by Hideki Imai and Yuliang Zheng. Vol. 1431. Lecture Notes in Computer Science. Pacifico Yokohama, Japan: Springer, Heidelberg, Germany, Feb. 1998, pp. 117–134. DOI: [10.1007/BFb0054019](https://doi.org/10.1007/BFb0054019) (cit. on pp. 89, 129).
- [Tur50] AM Turing. “Mind.” In: *Mind* 59.236 (1950), pp. 433–460 (cit. on p. vii).

- [Val76] Leslie G. Valiant. "Relative Complexity of Checking and Evaluating." In: *Inf. Process. Lett.* 5.1 (1976), pp. 20–23. DOI: [10.1016/0020-0190\(76\)90097-1](https://doi.org/10.1016/0020-0190(76)90097-1). URL: [https://doi.org/10.1016/0020-0190\(76\)90097-1](https://doi.org/10.1016/0020-0190(76)90097-1) (cit. on p. 94).
- [vHo4] Luis von Ahn and Nicholas J. Hopper. "Public-Key Steganography." In: *Advances in Cryptology – EUROCRYPT 2004*. Ed. by Christian Cachin and Jan Camenisch. Vol. 3027. Lecture Notes in Computer Science. Interlaken, Switzerland: Springer, Heidelberg, Germany, May 2004, pp. 323–341. DOI: [10.1007/978-3-540-24676-3_20](https://doi.org/10.1007/978-3-540-24676-3_20) (cit. on pp. 232, 233).
- [vHL05] Luis von Ahn, Nicholas J. Hopper, and John Langford. "Covert two-party computation." In: *37th Annual ACM Symposium on Theory of Computing*. Ed. by Harold N. Gabow and Ronald Fagin. Baltimore, MA, USA: ACM Press, May 2005, pp. 513–522. DOI: [10.1145/1060590.1060668](https://doi.org/10.1145/1060590.1060668) (cit. on pp. 13, 147, 155, 245).
- [Wato5] Brent R. Waters. "Efficient Identity-Based Encryption Without Random Oracles." In: *Advances in Cryptology – EUROCRYPT 2005*. Ed. by Ronald Cramer. Vol. 3494. Lecture Notes in Computer Science. Aarhus, Denmark: Springer, Heidelberg, Germany, May 2005, pp. 114–127. DOI: [10.1007/11426639_7](https://doi.org/10.1007/11426639_7) (cit. on p. 44).
- [Wee04] Hoeteck Wee. "On Pseudoentropy versus Compressibility." In: *19th Annual IEEE Conference on Computational Complexity (CCC 2004), 21-24 June 2004, Amherst, MA, USA*. IEEE Computer Society, 2004, pp. 29–41. DOI: [10.1109/CCC.2004.1313782](https://doi.org/10.1109/CCC.2004.1313782). URL: <https://doi.org/10.1109/CCC.2004.1313782> (cit. on pp. 12, 147, 157, 162, 190).
- [Wee05] Hoeteck Wee. "On obfuscating point functions." In: *37th Annual ACM Symposium on Theory of Computing*. Ed. by Harold N. Gabow and Ronald Fagin. Baltimore, MA, USA: ACM Press, May 2005, pp. 523–532. DOI: [10.1145/1060590.1060669](https://doi.org/10.1145/1060590.1060669) (cit. on p. 41).
- [WW20] Hoeteck Wee and Daniel Wichs. *Candidate Obfuscation via Oblivious LWE Sampling*. Cryptology ePrint Archive, Report 2020/1042. <https://eprint.iacr.org/2020/1042>. 2020 (cit. on pp. 2, 17, 41, 158, 159, 245).
- [WZ17] Daniel Wichs and Giorgos Zirdelis. "Obfuscating Compute-and-Compare Programs under LWE." In: *58th Annual Symposium on Foundations of Computer Science*. Ed. by Chris Umans. Berkeley, CA, USA: IEEE Computer Society Press, Oct. 2017, pp. 600–611. DOI: [10.1109/FOCS.2017.61](https://doi.org/10.1109/FOCS.2017.61) (cit. on p. 3).
- [WS07] J. Wu and D.R. Stinson. *An Efficient Identification Protocol and the Knowledge-of-Exponent Assumption*. Cryptology ePrint Archive, Report 2007/479. <http://eprint.iacr.org/2007/479>. 2007 (cit. on pp. 9, 88, 109).

- [Yao82] Andrew Chi-Chih Yao. “Theory and Applications of Trapdoor Functions (Extended Abstract).” In: *23rd Annual Symposium on Foundations of Computer Science*. Chicago, Illinois: IEEE Computer Society Press, Nov. 1982, pp. 80–91. DOI: [10.1109/SFCS.1982.45](https://doi.org/10.1109/SFCS.1982.45) (cit. on pp. [21](#), [33](#), [157](#), [162](#), [190](#)).
- [Zha16] Mark Zhandry. “The Magic of ELFs.” In: *Advances in Cryptology – CRYPTO 2016, Part I*. Ed. by Matthew Robshaw and Jonathan Katz. Vol. 9814. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 2016, pp. 479–508. DOI: [10.1007/978-3-662-53018-4_18](https://doi.org/10.1007/978-3-662-53018-4_18) (cit. on pp. [7](#), [42](#), [44](#), [45](#), [47–49](#), [153](#), [238–240](#)).
- [Zim15] Joe Zimmerman. “How to Obfuscate Programs Directly.” In: *Advances in Cryptology – EUROCRYPT 2015, Part II*. Ed. by Elisabeth Oswald and Marc Fischlin. Vol. 9057. Lecture Notes in Computer Science. Sofia, Bulgaria: Springer, Heidelberg, Germany, Apr. 2015, pp. 439–467. DOI: [10.1007/978-3-662-46803-6_15](https://doi.org/10.1007/978-3-662-46803-6_15) (cit. on p. [41](#)).
- [ZL77] Jacob Ziv and Abraham Lempel. “A universal algorithm for sequential data compression.” In: *IEEE Trans. Inf. Theory* 23.3 (1977), pp. 337–343. DOI: [10.1109/TIT.1977.1055714](https://doi.org/10.1109/TIT.1977.1055714). URL: <https://doi.org/10.1109/TIT.1977.1055714> (cit. on p. [12](#)).
- [ZL78] Jacob Ziv and Abraham Lempel. “Compression of individual sequences via variable-rate coding.” In: *IEEE Trans. Inf. Theory* 24.5 (1978), pp. 530–536. DOI: [10.1109/TIT.1978.1055934](https://doi.org/10.1109/TIT.1978.1055934). URL: <https://doi.org/10.1109/TIT.1978.1055934> (cit. on p. [12](#)).

COLOPHON

This document was typeset with $\text{\LaTeX 2}_{\epsilon}$ using the typographical look-and-feel `classicthesis` developed by André Miede and Ivo Pletikosić using Hermann Zapf's *Palatino* and *Euler* type faces (Type 1 PostScript fonts *URW Palladio L* and *FPL* were used). The style was inspired by Robert Bringhurst's seminal book on typography "*The Elements of Typographic Style*".

Final version as of June 15, 2021.