

Simulation and measurement of the dynamics of ultra-short electron bunch profiles for the generation of coherent THz radiation

Zur Erlangung des akademischen Grades eines
DOKTORS DER NATURWISSENSCHAFTEN (Dr. rer. nat.)

von der Fakultät für Physik
des Karlsruher Institutes für Technologie (KIT)
angenommene


DISSERTATION

von

Dipl.-Phys.
Patrik Schönfeldt
aus Oldenburg

Tag der mündlichen Prüfung: 13.07.2018

Referent: **Prof. Dr. A.-S. Müller**
Korreferent: **Prof. Dr. G. Quast**

 This work is licensed under the Creative Commons Attribution-ShareAlike 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-sa/4.0/> or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.

Contents

1. Introduction	1
2. Fundamentals of Accelerator Physics	3
2.1. Charged, Relativistic Particles in Electro-Magnetic Fields	3
2.2. Magnet Optics	4
2.3. Radiation and Fields from Charged Particles	5
2.3.1. Coulomb Field	6
2.3.2. Synchrotron Radiation	6
2.3.3. Radiation Damping	8
2.3.4. Beam Line Cutoff	9
2.4. Energy Deviation and Longitudinal Motion	9
3. Self-Interaction in the Longitudinal Phase-Space	12
3.1. Vlasov-Fokker-Planck Equation	12
3.2. Wake Fields and Impedances	13
3.2.1. Resistive Wall Impedance	15
3.2.2. Impedance of a Collimator	15
3.2.3. Free Space CSR Impedance	16
3.2.4. Parallel Plates CSR Impedance	16
3.3. Perturbed Vlasov-Fokker-Planck Equation	17
3.3.1. Potential Well Distortion	17
3.3.2. Micro-Bunching Instability	18
4. Beam Diagnostics	20
4.1. Streak Camera	20
4.2. THz Detectors	22
4.3. Electro-Optical Methods	23
4.3.1. The Pockels-Effect	23
4.3.2. The Electric Field Inside the EO Crystal	23
4.3.3. Electro-Optical Spectral Decoding	25
5. Optimization of the EO Setup	27
5.1. Turn-By-Turn Detector Readout	27
5.2. Reconstruction of the Measured Wake Fields	30
5.3. Optimization of the EO Arm Geometry	32
5.4. Test Measurements	34
6. Bunch-Profile Measurements	36
6.1. Streak Camera Measurements	36
6.2. Electro-Optical Measurements	38
6.3. Comparison of Bunch Profile Measurement Methods	39
6.4. Correlation with THz-Bursts	40

7. Parallelized Computing	42
7.1. Parallelization Categories	43
7.1.1. Single Instruction, Single Data	43
7.1.2. Multiple Instruction, Multiple Data	45
7.1.3. Single Instruction, Multiple Data	46
7.2. Example Hardware Implementations	47
7.3. OpenCL	48
7.4. Applicability	50
8. Numerical Solution of the VFPE	51
8.1. Discretization	51
8.1.1. Time Steps	51
8.1.2. The Grid Coordinates	52
8.1.3. Interpolation	53
8.1.4. Integration	54
8.2. Source-Mapping	55
8.3. Manhattan-Rotation	57
8.4. Calculating the Wake Potential	58
8.5. Single-particle dynamics	58
9. Influences of Numerical Effects	60
9.1. Explicit and Implicit Discretization	60
9.1.1. General Considerations	60
9.1.2. Data Types	61
9.1.3. Simple Example	61
9.2. Numerical Artifacts	62
9.2.1. Interpolation	62
9.2.2. Numerical Differentiation	64
9.2.3. Example Studies featuring Physical Cases	65
9.3. Benchmark Scenarios	69
9.3.1. Convergence of Unperturbed VFPE	69
9.3.2. Conservation of Probability	70
9.4. Starting Conditions to Simulate Unstable Beams	75
9.4.1. Well above Instability Threshold	75
9.4.2. Near the Instability Threshold	76
9.5. Error Estimation for Bursting Thresholds	77
10. Simulation of the Longitudinal Phase-Space	80
10.1. Current-Dependent, Static Phenomena	80
10.1.1. Bunch Deformation	80
10.1.2. Shift and Spread of the Synchrotron Tune	81
10.1.3. Growing Bunch Length and Energy Spread	83
10.1.4. Synchrotron Radiation	85
10.2. Dynamics of Typical CSR Bursts	86
10.2.1. Regular Bursting Regime	86
10.2.2. Sawtooth Bursting Regime	89
10.3. Contribution of Additional Impedances	91
10.4. Wavelength Dependency of CSR Power Fluctuation	93
10.5. Further Interpretation of the Results	96
10.5.1. Fast Bursting Frequencies	96
10.5.2. Growth and Decay of Substructures	99

11. Explicit Comparison of Simulation and Measurement	102
11.1. Bursting Thresholds	102
11.2. Influence of a Geometrical Impedance	103
11.3. Bursting Spectrogram	104
11.4. Bunch Lengthening	105
12. Implications and Outlook	107
12.1. Radiation Losses and Phase Shift	107
12.2. Methods to Determine the Micro-Bunching Threshold	108
12.3. Influencing the Bursting CSR	110
12.4. Verification of Non-Material Charge Modulation	110
13. Summary	111
Appendix	113
A. Complete Lists of Used Simulation Settings	113
B. Mathematical Definitions and Discussions	113
B.1. Airy and Bessel Functions	113
B.2. Fourier Transforms	114
B.3. Rounding	115
C. Complementary Data	115
C.1. Prediction of EOS v2 wake	115
C.2. Implicit Discretization: Differences in Convergence	116
C.3. Contributions to the Form Factor	116
C.4. Contribution of Additional Impedances	118
C.5. Fluctuation of the CSR Spectrum	119
C.6. Difficulty to Obtain the Synchrotron Frequency	120
D. Inovesa Quick Reference	121
D.1. Program Information	121
D.2. General Program Parameters	121
D.3. Numerical Parameters for Simulation	121
D.4. Physical Parameters for Simulation	122
E. List of Symbols	124
F. Acronyms	125
Bibliography	127
Acknowledgments	141

List of Figures

2.1.	Accelerator Coordinate System	4
2.2.	Electric Field of a Relativistic Electron Bunch	6
2.3.	Synchrotron radiation spectrum	8
2.4.	Phase Focusing	10
3.1.	Resistive wall impedance	15
3.2.	Unshielded (free space) and shielded (parallel plates) CSR impedance	16
3.3.	Azimuthal modes of synchrotron motion	18
4.1.	Working Principle of a Streak Camera	21
4.2.	EO Crystal Orientation	24
4.3.	Simulated Fields inside the EO Crystal and at the Electron Orbit	25
4.4.	EOSD Working Principle	26
5.1.	Schematic Drawing of KALYPSO	27
5.2.	Histogram of KALYPSO Background Measurement	28
5.3.	Electronic Contributions to the EOSD Background Signal	29
5.4.	Design of the EOS v1 EO Arm	30
5.5.	EO sampling signal using EOS v1	31
5.6.	Comparison of Bunch Profile and EO Signal	32
5.7.	Design of the EOS v2 EO Arm	33
5.8.	Wake Loss of EOS v1 and EOS v2	33
5.9.	Wake impedance of EOS v1 and EOS v2	34
5.10.	EOS wake field measurements using EOS v1 and EOS v2	35
5.11.	Single bunch profile recorded using EOSD	35
6.1.	Raw Image Recorded using a Streak Camera.	37
6.2.	Bunch Length Histogram (Using Streak Camera Measurements)	38
6.3.	EOSD measurements using EOS v1 and EOS v2	39
6.4.	Longitudinal Bunch Profiles measured using SC and EOSD	40
6.5.	Synchronous EOSD and THz Measurements	41
7.1.	Flynn's taxonomy of computer architectures	44
7.2.	Visualization of an instruction pipeline	45
8.1.	An example grid	53
8.2.	Computational Performance Using Source-Mapping	57
9.1.	Gaussian distribution sampled at multiples of 0.1	62
9.2.	Example for Spline Interpolation featuring Overshoots	63
9.3.	Comparison of different interpolation types	64
9.4.	Damping of a test pattern carried out in multiple time steps	65
9.5.	Self-amplifying interpolation artifacts	67
9.6.	Suspected amplification of quantization noise	68

9.7.	Working principle of the Convergence Test	70
9.8.	Simulated damping using different data types	71
9.9.	Reconstructed damping time and energy spread	72
9.10.	Influence of Charge Renormalization	74
9.11.	Evolution of the RMS energy spread over	75
9.12.	Convergence Using Different Starting Conditions	76
9.13.	Response of a Meta-Stable State to Numerical Changes	77
9.14.	Bursting Threshold obtained with Different Numerical Settings	78
10.1.	Example Deformed Bunch Profiles	81
10.2.	Particle Movement in the Longitudinal Phase Space	82
10.3.	Synchronous Phase Compared to Bunch Position	83
10.4.	Current-dependency of Bunch Length and Energy Spread	84
10.5.	Scaling of the Bunch Length with Current	85
10.6.	Averaged CSR Spectrum	86
10.7.	Integrated CSR Power and its Fluctuation	87
10.8.	The “Regular” Bursting Regime	88
10.9.	The “Sawtooth” Bursting Regime	90
10.10.	Influence of Additional Impedances	92
10.11.	CSR Spectra and their Fluctuation	94
10.12.	CSR Fluctuation Averaging Out	95
10.13.	Rotation of Substructures in Phase Space	97
10.14.	Formation of Substructures in the Longitudinal Phase Space	100
11.1.	CSR Power Fluctuation: Measurement and Simulation	103
11.2.	Measured Influence of a Scaper	104
11.3.	Measured and Simulated Bursting Spectrogram	105
11.4.	Distributions of Measured and Simulated Bunch Lengths	106
C.1.	Simulated EOS v2 wake-field compared to measurements	116
C.2.	Convergence based on Data Type	117
C.3.	Contributions to the Form-Factor	117
C.4.	Influence of Additional Impedances	118
C.5.	Example CSR Spectrum and its detailed Fluctuation	119
C.6.	Difficulty to Obtain the Synchrotron Frequency	120

List of Tables

8.1.	Runtime of different implementations of the rotation	56
9.1.	Parameters for an example run to check for numerical artifacts.	66
10.1.	Overview of bursting frequencies	98
A.1.	Simulation settings used for the results obtained in this thesis.	113

1. Introduction

Synchrotron radiation has proven to be a unique tool for the study of various phenomena for samples coming from different research areas. Example experimental methods include – each named with just one possible application: X-ray topography can image crystal imperfections [Moo95]. Time-resolved spectroscopy allows to track excited states of matter [Bon+80]. Fourier transform IR spectroscopy and microscopy show molecular vibrations in cancer cells [BS12]. X-Ray crystallography can map the structure of viruses [Liu+09]. Angle-resolved photoemission spectroscopy unveils superconducting gap quasi particles [Nor+98; Kam+00]. Broadband THz spectroscopy [Bar+13] allows gas-phase studies of unstable molecules.

Coherent synchrotron radiation (CSR) enhances the potential of the methods, as it can increase the photon flux by several orders of magnitude. Emission is coherent for wavelength longer than the emitting ensemble of electrons – typically an electron bunch. When the emitting electrons form substructures, so called micro-bunches, the CSR emission can shift to frequencies higher than for smooth bunches. The principle is used in free-electron lasers (FELs) [Dea+77] but is also applicable for storage-ring based synchrotron light sources. It however, took a long time until these potentials were utilized. In the beginning, synchrotron radiation and in particular CSR, as well as micro-bunching were regarded as undesired at particle accelerators.

A first theory on the radiation of charged particles was formulated by J. Larmor [Lar97] in 1897. One year later, in 1898, A. Liénard [Lié98] extended this work for relativistic particles moving on a curved path. These authors, however, could not correctly account for the relativity: The needed Lorentz transform was not formulated until 1904 [PB16]. Shortly thereafter, a first theory of synchrotron radiation [Sch07] that still holds in its original formulation was created by G. A. Schott – in 1907. The first experimental observation of the radiation [Eld+47] at a synchrotron in 1947 led to the naming. Just a few years before, the spectrum of the radiation had been first calculated by J. Schwinger (in 1945). His theory already included CSR [Sch98], however, CSR did not seem to be relevant, so only parts were published in 1949 [Sch49]. As particle accelerators at that time focused on high energy physics, this radiation was mainly an energy loss – with CSR just bringing additional losses. Early publications considering the CSR suggested to shield it by parallel plates [NS54]. It is interesting that this simple model still holds as a description for the beam vacuum pipe.

Over the time, more and more use was made of the emitted radiation, first parasitically at high-energy particle experiments [TH56; CM65], later also at dedicated synchrotron light sources [RM73]. In 1984, the first beam line designed to couple out light with long wavelengths from the accelerator vacuum chamber came into operation, the infra-red beam line in Daresbury [Yar+84]. As it turned out later, the light was CSR emitted due to the micro-bunching instability. First explicit observations of the increase of coherent emission were made at the NSLS VUV (**N**ational **S**ynchrotron **L**ight **S**ource **V**acuum **U**ltraviolet)

ring [Car+99; Car+01]. Shortly thereafter, stable CSR was produced using short bunches, first at BESSY [Abo+02], later also at other machines (e.g. ANKA [Mül+05]). At that time, the micro-bunching instability was still mostly seen as problem limiting the maximum bunch charge. On the other hand, the term “bursting CSR” was coined, to emphasize on the additional emission due to the micro-bunching. In the following, experimental studies were conducted to map out the parameters governing the bursting behavior [Byr+02; Abo+03; Tak+05; Kar+10; Mül10a; Wüs+10; Eva+12; Shi+12], and theoretic models were formulated to predict thresholds [VW02; SH02; BCS10; Cai11] and to simulate the dynamics. For simulations, there are two approaches: One is to solve the Vlasov-Fokker-Planck equation for the longitudinal phase space density [WE00; Ven+05; Rou+14a]. Alternatively, using supercomputers, particle tracking with one million macro particles or more [Eva+12; Rou+14b] can be performed. A good estimation for the threshold has been found by [BCS10]. In the very same paper, they still proposed methods to counter-act the micro-bunching. It is only since a couple of years, that also the potential due to emission shorter than bunch length has been considered. There are suggestions for controlled (seeded) micro-bunching [RCH11; She+11] in FELs, and beams have been modulated intentionally using a laser [Ung+17]. However, until now there is no method that can control the self-amplifying process that naturally forms micro-bunches at storage rings.

The aim of this thesis is to better understand the dynamics of the micro-bunching – following the long-term goal to find possibilities to use and influence it. The scope spans both, simulation and measurements: Current measurement equipment cannot fulfill the requirements needed to to successfully proof or reject the predictions of the model. So, experimental methods have to be advanced. Simulation, on the other hand, is demanding for both, hardware and user of the simulation tool. And as there is no intuitive model, the resulting data is sometimes hard to interpret and to visualize in a meaningful manner.

First (Chap. 2), the working principles of synchrotron light sources are discussed, including the process of CSR emission. The micro-bunching then is described in a Vlasov-Fokker-Planck picture (Chap. 3), that has already proven to be an efficient model when it comes to simulations. While the focus of this thesis is always on the motion of electrons emitting CSR, for some chapters it branches into a more experimental part and a simulation-centered one: Different kinds of beam diagnostics relevant for this work are shortly presented (Chap. 4), while the focus here is on electro-optical (EO) methods, which are also optimized (Chap. 5). A comparison of two different methods to measure bunch profiles – namely a streak camera and electro-optical spectral decoding (EOSD) – closes this experimental part (Chap. 6). The simulation part is started by introducing concepts of parallelized computing (Chap. 7). Those are used in the implementation of a Vlasov-Fokker-Planck solver (VFPS) (Chap. 8) which is extensively benchmarked to show different numerical effects (Chap. 9) that could otherwise be misinterpreted as physics. The actual numerical studies (Chap. 10) include detailed descriptions of different beam behaviors as well as a study of ways how to influence or utilize the CSR emission. Both parts merge in a comparison of measurement and simulation results (Chap. 11) before a final outlook (Chap. 12) is given, on how the numerical results can be utilized in experimental setups and how new predictions could be verified.

2. Fundamentals of Accelerator Physics

“Time dilates as the speed of light approaches. To the extent that light consists of particles, it is in its own way, timeless.”

Prokhor Zakharov

This chapter will introduce the fundamental quantities of accelerator physics that are important for the generation of coherent synchrotron radiation at storage rings. This focus narrows down the topic quite a bit. More comprehensive deductions as well as broader views that do also include e.g. other types of accelerators can be found in almost any standard work on accelerator physics, e.g. [Wil01; Hin08; Wie15]. Besides one of the standard works [Wil96], this chapter is strongly influenced by [Ng06] which puts the focus on beam instabilities.

Section 2.1 briefly comments on deflection and acceleration and introduces the fundamental tools to theoretically describe an abstract model of a storage ring. Deflection will be elaborated more in Section 2.2. For a detailed view on this particular topic, the reader may consult e.g. [Wie15, Chap. 5], [Hin08, Chap. 6], or [Wil01, Chap. 3]. Here only the aspects are discussed that are crucial to understand how synchrotron radiation (see Section 2.3) can be influenced. To do so, mostly longitudinal beam dynamics has to be discussed. This is done in section 2.4.

2.1. Charged, Relativistic Particles in Electro-Magnetic Fields

In high energy particle accelerators, a magnet optics is used to influence the beam path, while the energy is increased using electric fields. The physical reason for this technical implementation is the nature of the Lorentz force

$$\vec{F}_L = Q (\vec{v} \times \vec{B} + \vec{E}), \quad (2.1)$$

where \vec{E} denotes the electric field, \vec{B} the magnetic field, Q the charge, and \vec{v} the velocity of the particle. The equation directly implies that for large \vec{v} , it is easier to achieve large forces using the magnetic component. However, because $\vec{F}_L \perp \vec{B}$, a static magnetic field will only deflect the particle and not change its energy

$$E = \sqrt{m^2 c^4 + |\vec{p}|^2 c^2}. \quad (2.2)$$

The process of increasing the energy is usually handled by (a chain of) pre-accelerators and will be omitted here. Instead, only acceleration of particles near the accelerator’s design energy is considered. In fact, typically $\beta = v/c \approx 1$, so the relevant changes are also more




Figure 2.1.: Accelerator Coordinate System

The reference particle moves along s with velocity v_0 . An arbitrary particle's position is described in terms of the particle's offsets x , y and z from this reference. Notice that the time axis τ is inverted with respect to z . (Based on [Wil96].)

visible in momentum or energy. For a fixed energy E_0 , the synchronous or reference particle has velocity v_0 . With the Lorentz factor $\gamma = 1/\sqrt{1-\beta^2}$, its momentum can be expressed as $|\vec{p}_0| = \gamma m |\vec{v}_0|$. The momentum deviation is then typically given relative to p_0 , as

$$\delta := \frac{|\Delta\vec{p}|}{|\vec{p}_0|} = \frac{|\vec{p} - \vec{p}_0|}{|\vec{p}_0|} = \frac{\Delta E}{\beta^2 E_0}. \quad (2.3)$$

Electro-magnetic fields can differ in shape, strength, orientation, and location: Motion can follow complicated paths; but – for small beams – it is similar for almost all particles. So, to describe the position of an arbitrary particle, it is convenient to do so from the perspective of a *reference particle*, which marks the origin of a co-propagating coordinate system (see Fig. 2.1). With this coordinate system in mind, the bunch head is pointing in positive direction along the z -axis, motion along this axis is called *longitudinal motion*, while *transverse motion* is in the x and y directions¹. The time axis τ is oriented in opposite direction to z because a stationary observer will first (small time) see things at large z .

For a circular accelerator with a circumference of $L_0 = \oint ds$, also the revolution time and revolution frequency are defined based on the synchronous particle

$$1/f_{\text{rev}} := T_0 := L_0/v_0. \quad (2.4)$$

Sometimes, also an effective radius $\tilde{R} := L_0/(2\pi)$ is defined. This definition accounts for the fact that accelerators are not perfectly round but are typically designed to have alternating deflecting and straight sections (cf. Sec. 2.2). The latter can house the accelerating structures that compensate energy losses (cf. Sec. 2.3.2). Their technical realizations uses a voltage that oscillates with f_{RF} . The number of field oscillations per revolution is called the harmonic number

$$h := f_{\text{RF}}/f_{\text{rev}} \quad (2.5)$$

and is always an integer. As section 2.4 shows, this number also denotes the number of positions where on-momentum particles can be at. Because of this, h is also referred to as the number of (RF) buckets.

2.2. Magnet Optics

To describe the effect of the accelerator magnets, it is useful to expand their fields into multi-pole components

$$B_y(x) = B_{y,0} + \sum_{N=1}^{\infty} \frac{1}{N!} \frac{\partial^N B_y}{\partial x^N} \times x^N, \quad (2.6)$$

¹In this thesis, transverse beam dynamics are neglected. Typically, x and y will stand for grid coordinates.

and respectively for the x -components. Accordingly, the strength of different magnet types is defined

$$1/R = \frac{e}{|\vec{p}|} B_y, \quad K = \frac{e}{|\vec{p}|} \frac{\partial B_y}{\partial x}, \quad S = \frac{e}{2|\vec{p}|} \frac{\partial^2 B_y}{\partial x^2} \quad (2.7)$$

for dipole, quadrupole and sextupole magnets. The strength of the bending magnet here is defined as the inverse radius of the deflection, the so-called ‘‘bending radius’’ R . It should be noted that $2\pi R \leq L_0$. The equality is an edge-case where the particle accelerator has a constant $B_y \equiv B_0$. This case is called an iso-magnetic ring, which is a theoretical model that is not technically realized. In contrast, storage rings are typically designed [San69] to have an iso-magnetic guide field, so $B_y(s) \in \{0, B_0\}$. Strictly speaking, any location where $B_y = 0$ is a straight section. However, typically parts of the accelerator that feature any magnets and even space between the magnets that is required for technical reasons is not called that way. In fact, the usage of the term ‘straight section’ most often implies a purpose why there are no magnets placed at that specific location.

When only considering linear beam dynamics (esp. $S \equiv 0$) the linear equations of motion

$$x''(s) + x(s) \left(\frac{1}{R^2(s)} - K(s) \right) = \frac{1}{R(s)} \delta \quad (2.8a)$$

$$y''(s) + y(s) K(s) = 0 \quad (2.8b)$$

can be derived [Wil96] directly from Eq. (2.7). Note that both, $R(s)$ and $K(s)$, are defined by the magnets and thus periodic with the orbit length L_0 . In the linear model, there is no coupling between horizontal and vertical plain. For $1/R \neq 0$ and $\delta = 0$ the solution is

$$u(s) = A_u(s) \cos(\Psi_u(s) + \phi_u), \quad (2.9)$$

where $A_u(s) = A_u(s + L_0)$ is the oscillation amplitude. It can be expressed as a function of the (constant) emittance ϵ , and the (position-dependent) beta-function β

$$A(s) = \sqrt{\epsilon \times \beta(s)}. \quad (2.10)$$

Inside gradient-free bending magnets ($1/R \neq 0$, $K = 0$), it is convenient to define the path $D(s)$ (called dispersion) of a particle with $\delta = 1$. It has the form

$$D''(s) + \frac{1}{R^2} D(s) = \frac{1}{R} \quad (2.11)$$

and – for linear optics – gives the only influence of longitudinal beam dynamics to the transverse plains. The resulting path can be described as [Wil96]

$$x(s) = \sqrt{\epsilon_x \times \beta_x(s)} \cos(\Psi_x(s) + \phi_x) + D(s)\delta. \quad (2.12)$$

Notice that this expression is purely horizontal. To have coupling between x and y , sextupole magnets have to be considered.

2.3. Radiation and Fields from Charged Particles

At rest, the static electric field of a charged particle is going out radially in all directions. When the particle is moving, a magnetic field forms according to Ampère’s circuital law; but also – when relativity is considered – the shape of the fields changes (*Liénard-Wiechert potential*). When $\beta \rightarrow 1$, the fields get a pancake-like shape with an opening angle of about $1/\gamma$. This boosted Coulomb field is always perpendicular to the direction of motion and cannot not lag behind (see Sec. 2.3.1). When the particle’s path is curved, parts of the field are radiated off, leading to synchrotron radiation (Sec. 2.3.2).




Figure 2.2.: Electric Field of a Relativistic Electron Bunch

Radial electric field measured by an observer at a distance of $r = 10$ mm from a relativistic charge passing by. The field of the point-charge is a significantly broadened but for the charge distribution with $\sigma_z = 0.1$ ps only a slight difference is visible when comparing to the case with $\gamma = \infty$. [Hil13, Fig. 2.4]

When placed inside a beam pipe, the fields cannot reach infinity but are trimmed at a certain distance. Where the field lines touch the wall, image charges travel parallel to the original particle, creating the so-called *image current*. When the wall is not perfectly smooth or not perfectly conducting, these image charges account for so-called *wake fields* [Ng06, p. 1]: Fields are left behind when the co-propagating image charges are slowed down. (For a more comprehensive description of different wake fields, see Sec. 3.2.)

2.3.1. Coulomb Field

The Coulomb field of a point charge Q moving with constant velocity $v = \beta c$ is always co-propagating its source. While at rest, the field is equally distributed in all directions, the shape changes to a more compressed shape. When the movement is along the z -axis, the electric field of a point charge at $r = \sqrt{x^2 + y^2}$ can be expressed as [Jac98, p. 595]

$$E_r(r, t) = \frac{Q}{4\pi\epsilon_0} \times \frac{\gamma r}{(r^2 + \gamma^2 v^2 t^2)^{3/2}}. \quad (2.13)$$

For the field of a line charge distribution $Q_b \times \rho(z)$ this means

$$E_{r,\rho}(r, t) = Q_b \int_{-\infty}^{\infty} E_r(r, t - t') \rho(t) dt'. \quad (2.14)$$

So, for $\gamma \rightarrow \infty$ the electric field has the same shape as the charge distribution, but for $\gamma < \infty$ a point spread is introduced. Figure 2.2 illustrates the effect for a point charge and a Gaussian charge distribution with $\sigma_z = 0.1$ ps, both moving with $\gamma = 2544$ and at a distance of $r = 10$ mm from the observer. While for the point-charge, the spread is a significant, for the charge distribution it is almost not visible. This means that the shape of a relativistically moving charge distribution can be measured by observing the electric field in its proximity. A technical implementation using electro-optical sampling is discussed in Sec. 4.3.

2.3.2. Synchrotron Radiation

When a charged particle is moving on a curved trajectory with radius R , part of its Coulomb field can no longer co-propagate. It is irradiated as so-called synchrotron radiation.

Using classical electromagnetism theory, the total power emitted by one electron can be derived [Sch49] to be [Hin08, p. 347]

$$P_1 = \frac{Q^2 c \gamma^4}{6\pi\epsilon_0 R^2}. \quad (2.15)$$

This power is emitted in a spectrum that depends on the emission angle. When averaging over space and a sufficiently long time [San69, p. 113] of $\Delta t > R/(\gamma c)$ an “instantaneous” radiated power spectrum

$$P(\omega) = \frac{P_1}{\omega_c} S\left(\frac{\omega}{\omega_c}\right) \quad (2.16)$$

can be defined, where

$$\omega_c = \frac{3 c \gamma^3}{2 R} \quad (2.17)$$

is the critical frequency. The photons are emitted in forward direction, but there is a uncertainty of

$$\Theta_{\text{RMS}} = 1/\gamma, \quad (2.18)$$

that leads to an opening angle of the emitted radiation with the same value [Hin08, p. 348]. It is noteworthy that photons emitted in this angle to the inner side cross the electron path in a distance of

$$s = 2R \sin(1/\gamma). \quad (2.19)$$

This will be of relevance when self-interaction of the electron-bunch is considered (see Chap. 3). Further, as

$$P_1 = \int_0^\infty P(\omega) d\omega, \quad (2.20)$$

the function S in Eq. (2.16) has to be normalized. It can be expressed as

$$S(\xi) = \frac{9\sqrt{3}}{8\pi} \xi \int_\xi^\infty K_{5/3}(\hat{\xi}) d\hat{\xi}, \quad (2.21)$$

where $K_{5/3}$ is a modified Bessel function (see Appendix B.1).

With N particles, the total power is not only amplified by a factor of N , additionally there is coherent amplification for wavelength larger than the emitting ensemble of electrons. It can be expressed as

$$\begin{aligned} P(\omega) &= N \left[1 + (N-1) |\tilde{\varrho}(\omega)|^2 \right] \times P_1(\omega) \\ &= N \times P_1(\omega) + N(N-1) |\tilde{\varrho}(\omega)|^2 \times P_1(\omega) \\ &= P_{\text{ISR}}(\omega) + P_{\text{CSR}}(\omega), \end{aligned} \quad (2.22)$$

where

$$\tilde{\varrho}(\omega) := \int_{-\infty}^{\infty} \varrho(t) \times e^{-i\omega t} dt \quad (2.23)$$

denotes the Fourier transform of the normalized longitudinal bunch profile $\varrho(t)$. The first summand (P_{ISR}) gives the power of the incoherent synchrotron radiation (ISR) and is linear to the number of particles. The CSR on the other hand increases quadratically with the number of particles but also depends on the bunch’s form-factor $F(\omega) := |\tilde{\varrho}(\omega)|^2$. Example spectra are shown in Fig. 2.3.

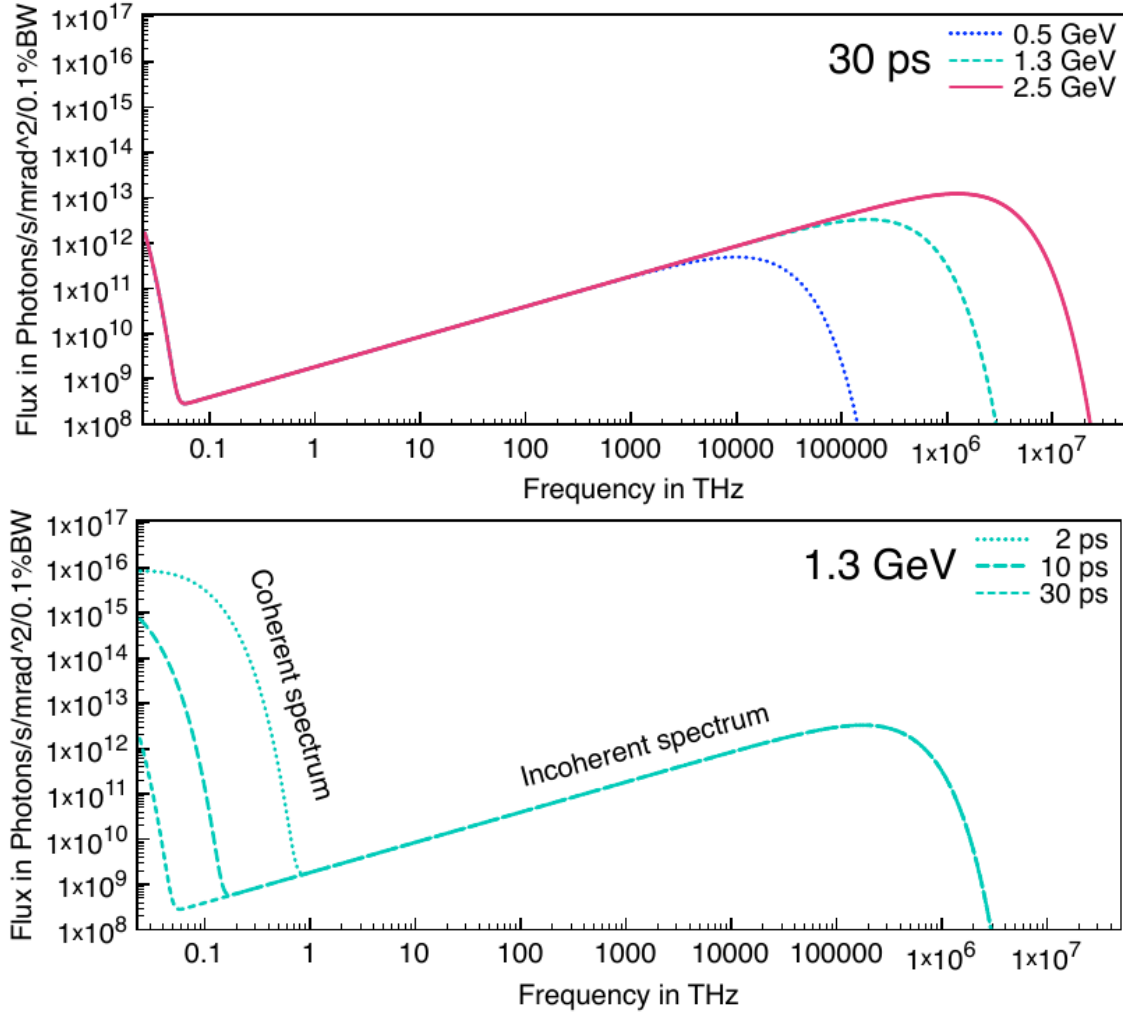


Figure 2.3.: Synchrotron radiation spectrum calculated for 10^8 electrons in a bunch of Gaussian shape at the **K**arlsruhe **R**esearch **A**ccelerator (KARA). [Hil13, p. 19]

2.3.3. Radiation Damping

When only constant deflection inside the bending magnets is considered, the total energy lost by an electron in one revolution can be computed by integrating Eq. (2.15). The result then is

$$W_0 = \oint P_1 dt = \frac{Q^2 \gamma_0^4}{3\epsilon_0 R}. \quad (2.24)$$

From this, the longitudinal damping time can be derived. The calculation can be found in [San69], or in [Hin08, p. 351] which is based on the first source. Here, just the result

$$1/\tau_d := \frac{1}{T_0} \frac{W_0 \mathcal{J}_E}{E_0} \quad (2.25)$$

is given, where $\mathcal{J}_E = 2 + d \approx 2$ is the longitudinal damping partition number. This definition of the damping time is based on the orbit change of the single particles. As the beam spread is a second order effect, its damping is twice as fast. Also, the beam spread will not damp down to zero. Instead, due to the quantum nature of the emission, there is always a non-zero spread of the energy distribution. For equilibrium cases, the distribution is Gaussian with a spread of [San69, p. 123]

$$\sigma_\delta = \gamma \times \sqrt{\frac{C_q}{\mathcal{J}_E R}}, \quad (2.26)$$

where $C_q = 55\hbar/(32\sqrt{3}mc) \approx 3.84 \times 10^{-13}$ m may be called the quantum constant. Non-equilibrium cases will be discussed later, starting from Sec. 3.3.2.

2.3.4. Beam Line Cutoff

Light cannot travel in structures smaller than its wavelength. For synchrotron radiation this fact can be considered already for the emission, this will be done in Sec. 3.2.4. When the emission frequencies of question are not completely suppressed, also the path between the electron beam and the observer – the beam line – has to be considered. This, of course, is the fact when unshielded emission models like the one in Eq. (2.16) are used. A beam line model acts just on the synchrotron light already irradiated, so it has no influence on the beam. However, to compare calculated results (e.g. from simulations) and measurements, it is important not to be dominated by CSR in the very low frequency range. An easy model just implements an exponential suppression. The spectrum $P_{\text{BL}}(f)$ that arrives at a beam line then is [Mül10b]

$$P_{\text{BL}}(f) = P_{\text{SR}}(f) \times \left[1 - \exp(-f/f_c)\right]^n, \quad (2.27)$$

where the cutoff is described by the cutoff frequency f_c and the strength of the suppression can be chosen by the order n .

2.4. Energy Deviation and Longitudinal Motion

As described in Section 2.2, particles with higher energy are deflected less by the accelerator magnets. Thus a momentum offset δ usually introduces different path lengths. The relation between deviations in energy and path length is called the momentum compaction factor.

$$\alpha_c := \sum_{n=0}^{\infty} \alpha_n \delta^n := \frac{\Delta L/L_0}{\delta}. \quad (2.28)$$

Usually one has $\delta \ll 1$, thus often only α_0 is considered to contribute significantly. Just considering linear optics leads to another commonly used approximation, namely

$$\alpha_c \approx \frac{1}{L_0} \oint \frac{D(s)}{R(s)} ds. \quad (2.29)$$

The difference in path length generally leads to a difference in arrival time, which can be seen as a motion relative to the synchronous particle and is called synchrotron motion². It is convenient to also define the slip factor, which gives the change in arrival time due to a momentum offset³

$$\eta_c := \frac{\Delta T/T_0}{\delta} = \alpha_c - \frac{1}{\gamma^2}. \quad (2.30)$$

As the time scale of the synchrotron motion is slow compared to the revolution time, it is possible to assume $d\Psi_z/dt = \Delta\Psi_z/T_0$. Now, expressing arrival times as phases, the definition above directly yields an equation for the average phase shift per turn due to an energy offset:

$$\dot{\Psi}_z = -2\pi f_{\text{rev}} h \eta_c \delta \quad (2.31)$$

In the accelerating structure, the field is typically higher at earlier times and thus gives more energy to the particles with less energy, and vice versa. This case ($\eta_c > 0$) is also depicted in Fig. 2.4. The resulting mechanism is called phase focusing and also leads to the bunching of the beam. It is noteworthy that such a mechanism is needed for stable




Figure 2.4.: Phase Focusing

Particles with less momentum (red) move on a shorter path, so that they reach the RF cavity at an earlier time, and thus at a smaller phase. This means they see a higher voltage and are accelerated more than the synchronous particle (black). For higher momentum particles (green) the same argument is reversed. (Based on [Wil96].)

operation of a storage ring. To allow $\eta_c < 0$, additional measures (like active feedback) have to be implemented.

Due to synchrotron radiation it is guaranteed that the energy will not be constantly fixed. In a storage ring, exactly the radiation losses W_0 of the synchronous particle will be compensated by the accelerating voltage. Typically this voltage is sinusoidal with $V_{\text{RF}}(t) = \hat{V} \sin(\Psi_z(t))$, where $\Psi_z(t)$ is the arrival time dependent phase. For the synchronous particle this leads to the synchronous phase

$$\sin(\Psi_{z,0}) := \frac{W_0}{e\hat{V}}. \quad (2.32)$$

When only incoherent synchrotron radiation in bending magnets is considered, Eq. (2.24) holds for W_0 . Additional losses (see Sec. 3.2) lead to a different phase. If the losses are current dependent, this phenomenon is called *synchronous phase shift*. At this point, it should be noted that the original definition of the synchronous phase is only valid for a single (virtual) particle. Many authors (e.g. [PS98; RBT11; Che+17]) relax it to the average phase of a whole bunch of particles. In this thesis, however, the synchronous phase will always refer to the position of the (virtual) *particle* that does gain exactly the energy that it has lost before⁴. The shift of the average phase will be called *coherent phase shift* or just *phase shift*.

For off-momentum particles, the energy changes by

$$\Delta E = e\hat{V} \sin \Psi_z - W_0 \quad (2.33)$$

when passing the accelerating voltage. This directly leads to the second equation of motion

$$\dot{\delta} = \frac{f_{\text{rev}} \Delta E}{\beta^2 E_0} = \frac{f_{\text{rev}}}{\beta^2 E_0} e\hat{V} [\sin(\Psi_z) - \sin(\Psi_{z,0})]. \quad (2.34)$$

It also implies that the maximum positive/negative energy offset as well as the maximum phase offset that can be compensated are dependent on $\Psi_{z,0}$. The trajectory along this positions in phase space is called the separatrix. It marks the borders of the RF buckets.

²This gives the (implicit) definition that the synchronous particle does not show any synchrotron motion.

³For ultra-relativistic particles we have $\gamma \rightarrow \infty$, hence $\eta_c \approx \alpha_c$.

⁴Section 10.1.2 will show the difference, while Sec. 12.1 will comment on the implications.

For small phase variations

$$\phi := \Psi_z - \Psi_{z,0}, \quad (2.35)$$

the effect of the RF potential can be approximated as

$$e\hat{V} \sin(\Psi_z) = e\hat{V} \sin(\Psi_{z,0} + \phi) \quad (2.36a)$$

$$= e\hat{V} [\sin(\Psi_{z,0}) \cos(\phi) + \sin(\phi) \cos(\Psi_{z,0})] \quad (2.36b)$$

$$\approx e\hat{V} [\sin(\Psi_{z,0}) + \phi \cos(\Psi_{z,0})]. \quad (2.36c)$$

The voltage seen by the particles is called effective voltage⁵ $V := \hat{V} \cos(\Psi_{z,0})$. Using this, Eq. (2.34) simplifies to

$$\dot{\delta} = \frac{eV f_{\text{rev}}}{\beta^2 E_0} \phi. \quad (2.37)$$

From the combination of Eqs. (2.31), (2.35), and (2.37), it follows that the change of energy and phase follows a harmonic oscillation, called synchrotron oscillation, with

$$\ddot{\phi} = -f_{s,0}^2 \phi \quad (2.38a)$$

$$\ddot{\delta} = -f_{s,0}^2 \delta, \quad (2.38b)$$

where

$$f_{s,0} = f_{\text{rev}} \sqrt{\frac{h\eta_c eV}{2\pi\beta^2 E_0}} \quad (2.39)$$

is the nominal synchrotron frequency. It holds true for a single electron at $E \approx E_0$. For example the oscillation frequency is lower for larger oscillations [cf. Eq. (2.36b)], and finally approach 0 at the separatrix. The Hamiltonian of this system reads

$$H = \pi h f_{\text{rev}} \eta_c \delta^2 + \frac{1}{2} \frac{eV f_{\text{rev}}}{\beta^2 E_0} \phi^2. \quad (2.40)$$

For now, any possible additional interaction of the particles is neglected. (It will be discussed as perturbations to the solution presented here, see Section 3.3.) Because of the the random nature of quantum radiation and excitations, the particles are distributed Gaussian in δ , where the RMS energy spread $\sigma_{\delta,0}$ is determined by synchrotron radiation as described in Eq. (2.26). As the solution to this harmonic oscillator is a rotation in phase space, the spacial distribution is Gaussian as well. With the longitudinal coordinate

$$z = -\frac{\beta c T_0}{2\pi h} \phi, \quad (2.41)$$

the RMS bunch length is

$$\sigma_{z,0} = \frac{c|\eta_c|}{2\pi f_{s,0} \beta} \sigma_{\delta,0}. \quad (2.42)$$

To reflect the fact that this equation only holds true without any interactions, $\sigma_{z,0}$ is called the ‘natural’ or ‘zero-current’ bunch length.

⁵The definition $V := \sqrt{\hat{V}^2 - U_0^2}$ with $U_0 := W_0/e$ is a fully equivalent.

3. Self-Interaction in the Longitudinal Phase-Space

“Truthful movements spontaneously attract to themselves all manner of pure and disinterested help.”

Mahatma Gandhi

This chapter first introduces an equation that can be used to describe the temporal evolution of a charge density, the Vlasov-Fokker-Planck Equation (Sec. 3.1). Self-interaction of the distribution is discussed as a perturbation. To do so, the concept of wake fields and impedances (Sec. 3.2) is used and also some example impedances are named. At last (Sec. 3.3), a brief theoretical treatment of the effect of the perturbation is given for both, equilibrium and non-equilibrium cases. These formulations make a rich toolbox that can be used to describe most of the known longitudinal beam dynamics.

3.1. Vlasov-Fokker-Planck Equation

The longitudinal phase space is spanned by position z and particle energy E . For large numbers of particles the normalized charge density $\psi(z, E, t)$ of electrons in a storage ring can be assumed to be smooth. When working with densities of finite spread, it is convenient to define the dimensionless, generalized coordinates [WE00]

$$q := \frac{z}{\sigma_{z,0}} \tag{3.1a}$$

and

$$p := \frac{E - E_0}{\sigma_{E,0}}. \tag{3.1b}$$

By giving it in multiples of synchrotron periods¹, also the time can be made dimensionless

$$\theta := f_{s,0} \times t. \tag{3.1c}$$

Assuming that the transverse dimensions of the charge distribution are $\sigma_{x,y} \ll \sigma_z$ [DRS95], using this set of coordinates, Eq. (2.40) simplifies to

$$H = \frac{1}{2} (q^2 + p^2). \tag{3.2}$$

¹Other authors prefer to define $\theta := \omega_{s,0} \times t$, so that θ is an angle (in radian).

In a conservative system, the area in the phase space is constant. It can then be described by the Vlasov Equation [Vla45] (cited from [Ng06, p. 48])²

$$\frac{\partial \psi}{\partial \theta} + \frac{\partial H}{\partial p} \frac{\partial \psi}{\partial q} - \frac{\partial H}{\partial q} \frac{\partial \psi}{\partial p} = 0. \quad (3.3)$$

Due to synchrotron radiation the system of course is not conservative. However, if the energy is sufficiently low, the time range is short, or in equilibrium cases, the Vlasov equation can be a valid approximation.

To account for radiation losses, in the image of a charge density their quantum nature usually is not considered. This is a good assumption for the large numbers of particles that are investigated in this thesis but should be reconsidered in other cases. Another simplification follows from $\sigma_E \ll E_0$: It means that for the single particle radiation the fact that $P(E) \propto E^4$ can be neglected, hence $P(E)_1 \approx \text{const}$. The evolution of $\psi(q, p, \theta)$ can then be described by the Vlasov-Fokker-Planck equation (VFPE). Following the notation of [BCS10], it reads

$$\frac{d\psi}{d\theta} = \frac{\partial \psi}{\partial \theta} + \frac{\partial H}{\partial p} \frac{\partial \psi}{\partial q} - \frac{\partial H}{\partial q} \frac{\partial \psi}{\partial p} = \beta_d \frac{\partial}{\partial p} \left(p\psi + \frac{\partial \psi}{\partial p} \right), \quad (3.4)$$

where $\beta_d = 1/(f_{s,0} \times \tau_d)$ is the dimensionless parameter for the longitudinal damping time τ_d . In this picture, collective effects can be described as a perturbation to the Hamiltonian

$$H(q, p, t) = \underbrace{H_e(q, p, t)}_{\text{external fields}} + \underbrace{H_c(q, t)}_{\text{collective effects}} \quad (3.5a)$$

$$= \frac{1}{2} (q^2 + p^2) + \frac{e f_{\text{rev}}}{\sigma_E f_{s,0}} \int_q^{\infty} V_c(q', t) dq'. \quad (3.5b)$$

Here, the perturbation is expressed using the potential V_c of one full revolution. This time (T_0) is also reflected in the factor f_{rev} . To calculate the energy from the potential, the factor e is needed, $1/\sigma_E$ and $f_{s,0}$ are to normalize energy and time. How the potential can be derived is discussed in the next section. It is noteworthy that Eq. (3.2) which makes the first term in Eq. (3.5b) assumes the unperturbed Hamiltonian to be linear. In general, also higher order terms of both, the momentum compaction factor α_c [cf. Eq. (2.28)] and the accelerating voltage V_{RF} [cf. Eq. (2.34)] can be included (see [Sch+18a] for details).

3.2. Wake Fields and Impedances

Wake fields are electromagnetic fields that are left behind by moving particles. As the deployed field has an energy, wake fields directly correspond to energy losses. A difference can be drawn between primary and secondary fields. Primary fields can be considered to be in proximity to the bunch $E(x, y, z \approx 0, t)$, secondary fields may lag behind. The latter may include influences like reflections that do not correspond to (further) energy-losses of the particle beam. Forces from an on-axis field can influence trailing particles. This influence does not have to refer to bunch-bunch interaction. In most cases, the trailing particles are still in the same bunch. However, if a wake persists for long enough, it can also cause long-range interaction.

To avoid confusion, in this thesis an explicit difference is made between the three distinct quantities *wake field* E , *wake function* W , and *wake potential* V . The wake function $W(z)$ is caused by a single particle, whereas the actual wake field $E(z)$ is typically caused by an electron bunch. However, although some authors speak of the wake function as the “single

²Especially when talking about the transverse plains, it is also referred to as Liouville theorem.

particle wake” [Kle12], or “wake potential” [WE00; Ven+05], W is a charge-normalized voltage with $[W] = 1 \text{ V}/c$. $E(z)$, in contrast, really is a (wake) *field* with $[E] = 1 \text{ V}/\text{m}$, and V a (wake) *potential* with $[V] = 1 \text{ V}$. All three quantities are connected³ by [WE00, p. 7]

$$V(z) = \int_0^L E(z, s) ds \quad (3.6a)$$

$$= eN_b \int_{-\infty}^{\infty} W(z - z') \varrho(z') dz'. \quad (3.6b)$$

To emphasize causality, it is possible to use different integration limits in Eq. (3.6b). The option preferred here is to express causality using a Heaviside step function in the formulation of the wake, so $W(z) \propto \Theta(z)$. When this is done, the longitudinal wake function $W_z(z)$ can be written as an impedance [Ng06, p. 13]⁴

$$Z_z(\omega) = \int_{-\infty}^{\infty} W_z(z) e^{-i\omega z/v} \frac{dz}{v}. \quad (3.7)$$

Furthermore, $E(s)$ is defined at any position, while W gives the integrated effect of the whole structure that has the impedance⁵ Z . Notice that $\Im[W_z(z)] \equiv 0$, so it directly follows that the impedance is Hermitian

$$Z(-\omega) = \Re[Z(\omega)] - i \times \Im[Z(\omega)]. \quad (3.8)$$

To study beam dynamics caused by the wake field, the energy lost by a beam particle has to be known. Eq. (3.6b) directly gives the representation as

$$\begin{aligned} \Delta E(z) &= e \times V(z) \\ &= e^2 N_b \int_{-\infty}^{\infty} W(z - z') \varrho(z') dz'. \end{aligned} \quad (3.9a)$$

Alternatively, using Eq. (3.7), it can also be worked in the frequency domain⁶

$$\Delta E(z) = \frac{e^2 N_b}{2\pi} \int_{-\infty}^{\infty} \tilde{\varrho}(\omega) Z(\omega) e^{i\omega z/v} d\omega. \quad (3.9b)$$

While impedances are typically derived starting from the wake function, the impedance representation has one great advantage: The effect of an impedance can be directly seen from its value. Inductive impedances [$\Im(Z) > 0$] counter-act changes in the current. As the instantaneous current rises on the arrival of the bunch and drops at its departure this means that the bunch length will be increased. For capacitive impedances [$\Im(Z) < 0$] the effect is the opposite, the bunch shortens. The resistive part of the impedance [$\Re(Z)$] expresses energy losses but also leads to asymmetric distortions. [Ng06, p. 228]

³It has to be noted that this formulation neglects the finite propagation speed of the fields. This assumption is valid as long as $\varrho(z')$ can be seen as constant at those time scales.

⁴Notice that the definitions for the Fourier transform differ between authors. See Appendix B.2 for details.

⁵As in this work mostly longitudinal effects will be considered, in the following, qualifiers for the impedance will be disregarded, $Z \equiv Z_z$.

⁶In the implementation, the factor of $1/(2\pi)$ is barely visible due to the usage of $f = \omega/(2\pi)$.




Figure 3.1.: Real part of the resistive wall impedance calculated for a 35 m long Aluminum pipe with a diameter of 32 mm. (These values approximate the part of KARA where dipole magnets are installed.) The imaginary part has the same value but the opposite sign, see Eq. (3.10).

3.2.1. Resistive Wall Impedance

The first impedance presented here is due to the resistivity of the vacuum chamber. It can be explained by the effect of mirror-charges moving inside the resistive wall of the beam pipe. The impedance can be calculated as [Ng06, p. 18]

$$Z_{RW}(f) = [1 - i \operatorname{sign}(f)] \sqrt{\frac{\mu_r Z_0}{\pi c \sigma_c} \frac{L}{2b}} \sqrt{|f|}, \quad (3.10)$$

where σ_c is the conductivity, μ_r the permeability, L the length and b the radius of a round beam pipe. Z_0 is the vacuum impedance

$$Z_0 := \mu_0 \times c \quad (3.11)$$

$$= 2.99792458 \times 40\pi \Omega \quad (3.12)$$

$$\approx 377 \Omega.$$

As one would expect, the impedance vanishes for superconducting or large beam pipes. A non-trivial example with $\sigma_c = 37.7 \times 10^6$ S/m, $\mu_r = 1 + 2.1 \times 10^{-5}$ (both Aluminum [WebE18]), $L = 2\pi \times 5.559$ m = 35 m, and $b = 16$ mm is displayed in Fig. 3.1.

3.2.2. Impedance of a Collimator

A collimator is an aperture smaller than the beam pipe. It is typically brought to the beam to protect sensitive structures from being hit by particle losses. It is of special interest because it can easily be modified, e.g. in case an iris is used. In high frequency limit the impedance of such a device is [SBZ07, p. 14]

$$Z_{\text{Coll}} = \frac{Z_0}{\pi} \ln \frac{b}{a}, \quad (3.13)$$

where $a < b$ is the inner radius of the collimator, and b the outer radius (typically the one of the beam pipe). For example sizes of $a = 2$ mm and $b = 16$ mm, the impedance is $Z = 249.36 \Omega$. Note that it is $\Im(Z) \equiv 0$ for infinitesimally short collimators only. If two steps in the beam pipe can be considered separately, the imaginary part does not vanish.

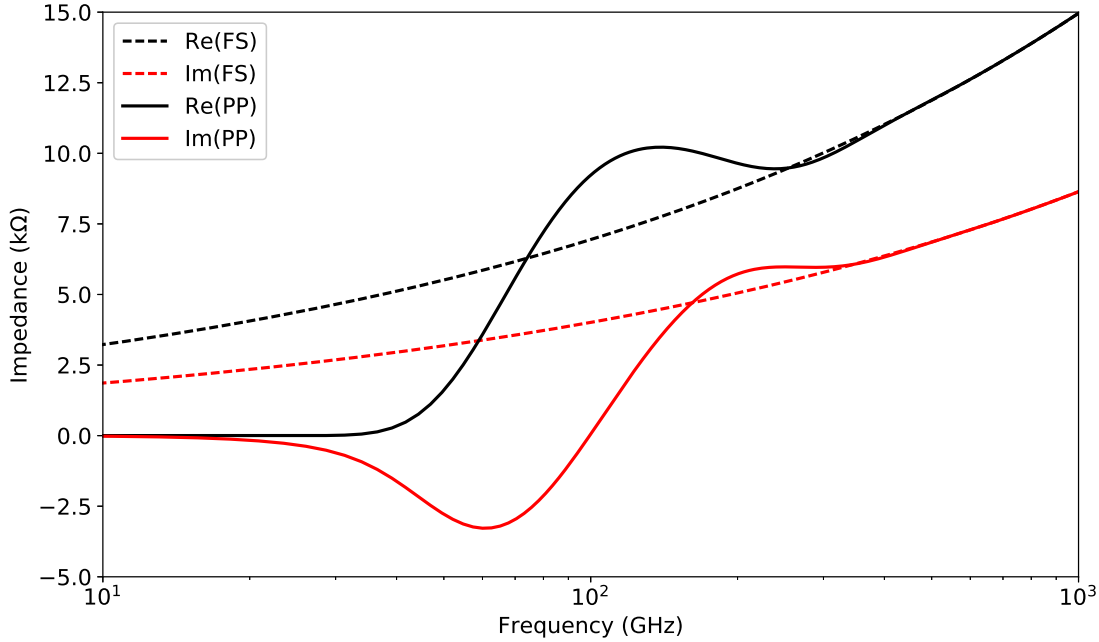


Figure 3.2.: Unshielded (free space) CSR impedance (FS) and CSR impedance shielded by parallel plates (PP), calculated for an example accelerator with $R = 5.559$ m, and (for the shielded case) $g = 32$ mm. For high frequencies both impedances converge, as short wavelength are not effected by the shielding. For $f \rightarrow 0$ Hz both impedances approach $Z = 0 \Omega$. (First published in [Sch+17c].)

3.2.3. Free Space CSR Impedance

The free space CSR impedance describes the effect of coherent synchrotron radiation of particles traveling on a curved path in vacuum. Even though both, the particles and the emitted photons, travel with the speed of light ($\beta \approx 1$), photons can overtake the bunch due to the curvature. Assuming that $\sigma_{x,y} \ll \sigma_z$, Derbenev, Rossbach and Saldin ([DRS95]) derived a “longitudinal overtake potential”. A good approximation [MKG97] is found to be

$$Z_{\text{CSR0}}(n) \approx Z_0 \frac{\Gamma(2/3)}{3^{1/3}} \left(\frac{\sqrt{3}}{2} + \frac{i}{2} \right) n^{1/3}, \quad (3.14)$$

where $n = f/f_0$ is frequency expressed in multiples of the revolution frequency f_0 in an iso-magnetic ring of length $2\pi R$.

3.2.4. Parallel Plates CSR Impedance

A more advanced CSR impedance adds the shielding effect of the beam pipe – being approximated by two parallel plates with distance g . In terms of Bessel functions, it can be approximated [Cai11] with the Airy functions Ai and Bi ⁷:

$$Z_{\text{PP}}(n, R/g) \approx \frac{4\pi^2 2^{1/3}}{\epsilon_0 c} \left(\frac{R}{g} \right) n^{-1/3} \times \sum_p \text{Ai}'(u_p) \text{Ci}'(u_p) + u_p \text{Ai}(u_p) \text{Ci}(u_p), \quad (3.15)$$

where the prime marks the derivative, R is the beam path’s radius, $\text{Ci} := \text{Ai} - i \text{Bi}$, and

$$u_p := \frac{\pi^2 (2p+1)^2}{2^{2/3}} \left(\frac{R}{g} \right)^2 n^{-4/3}. \quad (3.16)$$

⁷See Appendix B.1 for the function definitions.

This impedance has already proven to describe the micro-bunching instability threshold within the uncertainty of the measurements [Bro+16b]. For the limit $g \rightarrow \infty$ it converges to the free space impedance. Figure 3.2 shows these two impedances, for an example accelerator with $R = 5.559$ m, and (for the shielded case) $g = 32$ mm. An interesting feature of this impedance is that it is capacitive for low frequencies (here: below 100 GHz) and becomes inductive for higher ones.

3.3. Perturbed Vlasov-Fokker-Planck Equation

When an impedance is considered in the Hamiltonian, the additional potential modifies the solution of the VFPE and the bunch profile deforms. This deformation might be stable in time (Sec. 3.3.1), but when the perturbation term is large enough, the beam becomes unstable (Sec. 3.3.2).

3.3.1. Potential Well Distortion

As discussed in Sec. 3.1, the effect of collective effects is treated as a perturbation to the Hamiltonian. In the equilibrium case, a stationary solution is reached. In this case, the particle distribution has to be a function of the Hamiltonian [Ng02, Eq. (3.2)]

$$\psi(q, p) = \psi(H). \quad (3.17)$$

As only position dependent energies are altered [see Eq. (3.5b)], only the potential well has changed. Thus, the change in the particle distribution does not affect the energy distribution

$$\psi(q, p) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{p}{2}\right) \times \varrho(q). \quad (3.18)$$

The bunch shape $\varrho(q)$ that provides the solution to Eq. (3.17) is called the Haïssinski distribution. It was first derived in [Hai73], but deductions can also be found e.g. in [Ng02]. It is

$$\varrho(q) = \xi \times \exp\left(-\frac{q^2}{2} + \frac{e^2 N_b f_{\text{rev}}}{2\pi \sigma_E f_{s,0}} \int_q^\infty \int_{-\infty}^\infty \tilde{\varrho}(\omega) Z(\omega) e^{i\omega q' \sigma_{z,0}/v} d\omega dq'\right) \quad (3.19a)$$

$$= \xi \times \exp\left(-\frac{q^2}{2} + Q_c \int_q^\infty \frac{e}{2\pi} \int_{-\infty}^\infty \tilde{\varrho}(\omega) Z(\omega) e^{i\omega q'} d\omega dq'\right) \quad (3.19b)$$

$$= \xi \times \exp\left(-\frac{q^2}{2} + \frac{Q_c}{N_b} \int_q^\infty V_c(q') dq'\right), \quad (3.19c)$$

where the constant Q_c represents the charge. Notice that in Eq. (3.19c), both Q_c as well as V_c include the total charge, so that a corrective division by N_b has to be applied. To obtain the factor ξ , normalization is done, so that $\int_{-\infty}^\infty \varrho(q) dq \stackrel{!}{=} 1$. Examples will be shown in Sec. 10.1.1. Notice that the equation is consistent with the expectations for an unperturbed beam, as in the limit $N_b \rightarrow 0$, the distribution still is Gaussian.

The potential also changes the synchrotron frequency: The particles do not only see the accelerating voltage, as assumed in the derivation of Eq. (2.39) but additionally the wake voltage [Eq. (3.6)]. As $\partial V/\partial z \neq 0$ this also means that f_s will be different on different paths through phase space. This phenomenon is called ‘‘tune-spread’’ and will be illustrated based on simulation data in Sec. 10.1.2. To express this position dependent motion of the particles in phase space, the parameter

$$J(q, p, t) = \sqrt{\left(\frac{dq}{dt}\right)^2 + \left(\frac{dp}{dt}\right)^2} \quad (3.20)$$

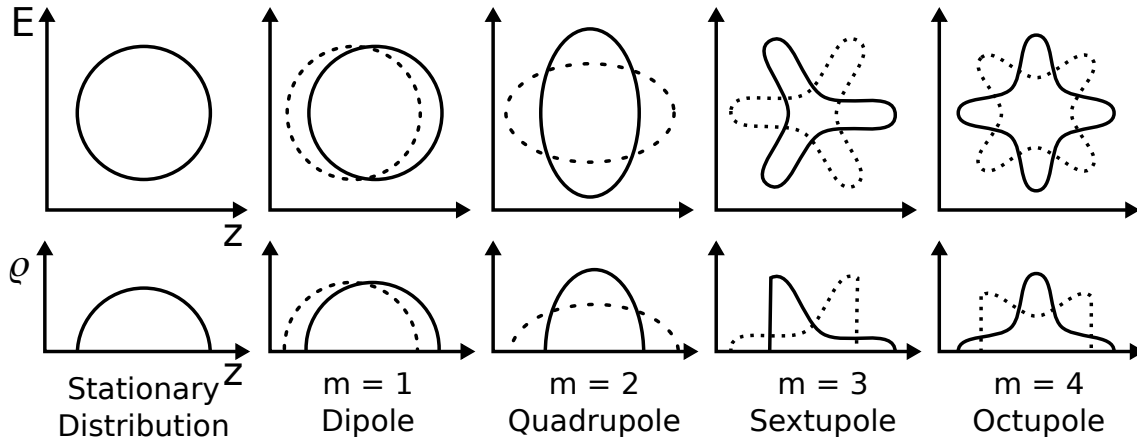


Figure 3.3.: The top row shows a visualization of azimuthal synchrotron modes of a bunch in the longitudinal phase space. In the bottom row, the projection of the density to the space axis and is sketched as a linear density. Also a combination of modes is possible, e.g. a stationary distribution with small sub-structures caused by an octupole motion. (Based on [Ng02, Fig. 7.1].)

can be used. It has the dimension of a frequency but should rather be seen as a “velocity” in phase space, and will be called *instantaneous synchrotron motion*. For example, in the unperturbed case, where all synchrotron motion can be expressed by a constant angular frequency, it is

$$J_0(q, p) = 2\pi f_{s,0} \sqrt{q^2 + p^2}. \quad (3.21)$$

3.3.2. Micro-Bunching Instability

For large number of particles Eq. (3.19) does not converge. This is not a mathematical or numerical problem, there really is no stable charge distribution for such cases. A different ansatz gives more insight. Extending the theory of the microwave instability for coasting beams (see e.g. [Ng02]) to bunched beams a small structure starts to grow in intensity due to positive feedback loop: The spectrum of the structure samples a part of the impedance that supports its growth. The stability criterion for a bunched beam with a Gaussian energy spread⁸ was derived by Krinsky and Wang [KW85] to be

$$\left| \frac{Z}{n} \right| < \frac{2\pi |\eta_c| E_0}{e I_{\text{peak}}} \delta^2, \quad (3.22)$$

where $I_{\text{peak}} = \max[\partial Q_b(z)/\partial z]/(\beta c)$ is the peak current of the electron bunch. This simple treatment assumes [Ng02, p. 7-1] that the bunch is much longer than the wavelength of the perturbation and that the growth time is much shorter than a synchrotron period. For electron beams, however, this is not necessarily given.

If synchrotron oscillation cannot be neglected, different oscillation eigenmodes (see Fig. 3.3) can be considered [Ng02]. For example the $m = 2$ mode corresponds to a length fluctuation with a frequency of $2f_s$. As the bunch spectrum plays a vital role in its emission [see Eq. (2.22)], the synchrotron oscillation modes are also connected to fluctuations in the emitted synchrotron radiation. In this view, focused on the emission of synchrotron light, the micro-bunching instability is referred to as “bursting CSR”.

As will be shown in Sec. 10.2.2, the instability may grow the bunch length beyond the conditions where substructures are formed. In that case, a breakdown with shrinking

⁸The Krinsky-Wang criterion is an elaborated version of the Keil-Schnell criterion [KS69]. The latter was first expanded to bunched beams by Boussard [Bou75].

bunch length and decreased CSR emission follows its growth. Based on observations of this particular periodic changes of the bunch length, the term “sawtooth instability” was coined [Kre+93]. Experiments at the Stanford Linear Collider (SLC) positron damping ring correlated streak camera images with different phases of the (coherent) synchrotron oscillation. Afterwards, changes in the bunch profile were identified with specific modes: Different kinds of substructures were observed [Pod99] and attributed to quadrupole and sextupole motion, respectively. When the beam current is varied, not only rapid growth and shrinking of the electron bunch are observed but also an overall bunch lengthening. Based on calculations using the different modes of the synchrotron motion, Chao and Gareyte [CG90] derived a scaling law [Ng02, Eq. 7.17] for arbitrary microwave instabilities. When the part of the impedance sampled by the bunch behaves like

$$Z \propto \omega^a \quad (3.23)$$

the bunch length scales with

$$\sigma_z \propto I_b^{1/(2+a)}. \quad (3.24)$$

Notice that this expression (implicitly) assumes that there is no fluctuation of the bunch length. [Ng06, p. 225] explicitly states:

“As soon as the threshold is passed, the bunch will be lengthened (...) to such an extent that stability is regained again. Unlike proton bunches no overshoot is observed in electron bunches, probably because of the smoothing effect of synchrotron oscillation. Probably radiation damping also plays a role.”

So in the special case of the micro-bunching instability with a fluctuating length, Eq. (3.24) can only give an estimate.

Until now, there is no closed formulation to describe the beam dynamics during the micro-bunching instability. Reasons are the tune spread due to the potential well distortion – in this case even with $\partial V/\partial t \neq 0$ – and that mode-mixing of different azimuthal and radial modes is possible. To get an insight, the dynamics have to be simulated, e.g. using particle tracking or by iteratively solving the VFPE for a series of points in time (see Chap. 8).

In a bunched beam theory (BBT), when only the CSR impedance shielded by parallel plates (see Sec. 3.2.4) is considered, and radiation damping can be neglected, two generalized quantities can be used to describe the dynamics, namely the “CSR strength” and the “shielding”. The first can be seen as a scaled current and is defined to be

$$\xi = \frac{Q_b f_{\text{rev}}}{\alpha_c \gamma (\sigma_{E,0}/E_0)^2 I_A} \left(\frac{R}{\sigma_{z,0}} \right)^{1/3}, \quad (3.25)$$

where $I_A = 17045$ A is the Alfvén current. The shielding is the natural bunch length scaled with the geometric sizes of the parallel plates approximating the beam pipe

$$\Pi = \frac{\sigma_{z,0} R^{1/2}}{(g/2)^{3/2}}. \quad (3.26)$$

Based on agreeing simulations using both, a VFPS and a linearized Vlasov solver, the relation, [BCS10] found that the beam becomes unstable for

$$\xi_{\text{th}}(\Pi) \gtrsim 0.5 + 0.12 \Pi. \quad (3.27)$$

This model is valid as long as the bunch profile does not change on shorter time scales than the fields need to propagate.

4. Beam Diagnostics

“One doesn’t discover new lands
without losing sight of the shore.”

André Gide

The measurements carried out in the framework of this thesis were performed at the KARA, that is formerly known as Ångströmquelle Karlsruhe (ANKA). This chapter will shortly introduce two kinds of well tested measurement devices that were used for reference measurements to verify the applicability of the model used for simulation (see Chap. 11): Streak cameras (Sec. 4.1) for the longitudinal bunch profile and THz detectors (Sec. 4.2) for CSR emission. To overcome limitations of the streak camera, also EO methods were used. They are introduced in Sec. 4.3. As near-field EO at a storage ring itself is a cutting-edge research topic, the complete Chap. 5 is dedicated to the optimization of the setup. Measurements will be shown in Chap. 6 and later (Chap. 11) be compared to simulation results.

4.1. Streak Camera

Provided that the bunch length is much larger than its height and width or that the opening angle of the emission of the synchrotron light is small, for simple geometrical reasons, the temporal profile of the emitting electron bunch is the same as the temporal probability density function of detecting synchrotron light. The first condition is typically given but for sufficiently high beam energy, the second is always true [cf. Eq. (2.18)]. This fact can be utilized to indirectly measure the longitudinal bunch profile. As the pulse, however, can be in the range of picoseconds or even shorter, detection of the temporal photon pulse shape is a hard task. To do so, a streak camera (see Fig. 4.1) can be used. It first sends the photons onto a photo cathode, the photo-electrons are then deflected by a time-varying voltage. They are detected e.g. using a fluorescent screen. Using this method, the time distribution is transformed to a spacial distribution that can easily be measured. Using two time-varying voltages, a second time axis can be added. Thus, it is possible to monitor the pulse shape over time.

The main drawbacks of streak cameras are the influence of shot noise, the lacking ability to pick pulses of a pulse train, and the limited time range – to measure the bunch profile, the coherent synchrotron oscillation has to be resolved. Pulse picking would be needed when bunches have different shapes, e.g. because of current dependent bunch deformation. To overcome this, often only a single bunch is stored in the accelerator. Shot noise especially plays a role when the bunch charge is low. Under these conditions only a very limited number of photons arrives at the streak camera. To still reconstruct the beam profile, it has to be integrated over a long time. However, when the phase of the (coherent) synchrotron motion is not perfectly known or the bunch profile varies in time, this integration introduces




Figure 4.1.: Working Principle of a Streak Camera

A streak camera firsts turns an incident light pulse into photoelectrons. Those than can be swept over a fluorescent screen using a time-varying voltage. This way, the original temporal distribution is turned into a spacial distribution that can easily recorded using a camera. Also, a second time varying voltage can be used to add a second (typically slower) time axis. ([Hil13], based on [MPG04].)

significant errors. [Pod99] attacked this problem by only integrating over periods of time where the shape should look similar. If the underlying model is correct, this approach works, but it is of course not feasible for exploratory studies.

Using a dual sweep, a second time axis can be added. On this (slower) time axis, time range can be traded for temporal resolution. The fact that lower bunch charges are of higher importance when the synchrotron frequency is also low¹ allows to partly compensate the low intensity: In this case, longer integration times still allow to reconstruct low $f_{s,0}$. Still, using this approach, faster dynamical changes of the bunch profile will be averaged out.

Streak cameras are commercially available [Axi17; BIF17; Ham17; Opt17; Syd17] and can be placed at a beam line of the accelerator. They are comparably robust and user-friendly, and can be considered *the* standard diagnostics to measure longitudinal bunch profiles at electron storage rings. At KARA, a Hamamatsu C5680 is used. It features a dual sweep unit.

For measurements of the bunch length, several factors have to be considered. First of all, on the fluorescent screen every incoming light is detected as spots with a macroscopic diameter. The effect is intensified when the photomultiplier is tuned to detect even single incoming photons. As a counter measure, the resulting profiles may be unfolded with the point spread function. Secondly, windows in the beam line act as dispersive elements. That way, the light pulse is elongated. To minimize this effect, a filter is used to only pick light around a central region of 550 nm, which is the region where the streak camera is most sensitive. At last, noise on the recorded data can be hardly distinguishable from signal. This fact has most impact, when the bunch length has to be determined. As the bunch shape is not known a priori, no fits can be used. So, to minimize the effect of outliers, usually the full

¹When neglecting the changing form factor, e.g. Eq. (3.19a) does not change if the ratio $N_b/f_{s,0}$ is constant. Eq. (3.22) contains the ratio $|\eta_c|/I_{peak} \propto f_{s,0}^2/N_b$, see Sec. 3.3 for details.

width half maximum (FWHM) should be preferred over the RMS as a measure to quantify the bunch lengths. The results presented in this thesis are obtained using a fully automated data evaluation by a custom made software [Sch13] that implements these guidelines.

An edge-case is the determination of the natural bunch lengths $\sigma_{z,0}$: In the needed limit of very low bunch currents, collective effects become unimportant, so multiple bunches – although with different current – have a sufficiently similar shape. So, when it is guaranteed that phase² and amplitude of the coherent synchrotron oscillation is also the same for all bunches, streak camera measurements should be favorably carried out in a multi-bunch environment. In this special case, averaging the (noisy) signal of the different bunches comes as a feature rather than a problem. Furthermore, for very low current the bunch form is known to be Gaussian, so a fit may be applied to reduce the impact of noise.

4.2. THz Detectors

To detect the synchrotron radiation, the spectral sensitivity has to be high enough to detect pulse powers in the μJ range. The most sensitive detector available at KARA is a hot electron bolometer (HEB). In such detectors, a superconductor is cooled exactly to the critical temperature to be sensitive to very small temperature changes. The HEB can even see ISR, however, it has a very limited dynamic range that has to be adjusted using absorbers in front of the detector. The detector's antenna is designed for a broad band spectral sensitivity between 200 GHz and 4 THz with a peak at the lower end of this range.

For the sake of flexibility, mostly Schottky diodes are used. Those work at room temperature and thus are easily set up, operated and replaced. Also for Schottky diodes, the sensitivity range depends mainly on the antenna. At KARA, five wave guide coupled (narrow band) ones from Virginia diodes [VDI17] are used [Ste+17b], ranging from 90 GHz up to 140 GHz, from 140 GHz up to 220 GHz, from 220 GHz up to 325 GHz, from 325 GHz up to 500 GHz, and from 500 GHz up to 770 GHz. Also, a quasi-optical (broad band) diode from Advanced Compound Semiconductor Technologies [ACST17] is available, with a sensitivity ranging from 50 GHz up to beyond 2 THz. All of those Schottky diodes have time constants fast enough to resolve individual bunches. The readout can be done using an oscilloscope or KAPTURE (**K**arlsruhe **p**ulse **t**aking and **u**ltra-fast **r**eadout **e**lectronics) [Cas+14]. The latter allows for continuous acquisition of the data for every single bunch (with a spacing of 2 ns) with up to four channels. Those channels can sample the same detector pulse with delays of up to 100 ps, or alternatively pulses of up to four separate detectors.

The measured CSR radiation power

$$P_{\text{CSR}}(t, I) = \int_{f_1}^{f_2} s(f) \times P_{\text{CSR}}(t, I, f) \, df \quad (4.1)$$

also involves the spectral sensitivity curve $s(f)$ of the detector. Typically, the detected energy is given in arbitrary units. The reason is that $s(f)$ is also dependent on the shape and the duration of the incident pulse but in the data sheet it is only given for DC signals.

When $\partial P_{\text{CSR}}(t, I)/\partial t \neq 0$, the fluctuation is often (e.g. [Kus09; Fei+11; Eva+12; Shi+12]) visualized using a spectrogram³

$$P_{\text{CSR}}(f, I) = \left| \int_{-\infty}^{\infty} P_{\text{CSR}}(t, I) e^{-i2\pi ft} \, dt \right|. \quad (4.2)$$

²Especially the synchronous phase which is harder to validate than the relative phase.

³A good introduction can be found in [Bro+16b].

Notice, that here f is in the kHz range, so it should not to be mixed up with the wavelength of the emitted photons. While the time domain data might look almost chaotic, the periodicity of the burst follows a very incisive pattern.

4.3. Electro-Optical Methods

Electro-optical measurements sample the electric field using the Pockels-effect, either using individual, short pulses or a long, chirped laser pulse. The first method allows to scan large time intervals and is referred to as (classical) electro-optical sampling (EOS), the latter method – called EOSD – allows for single-shot measurements.

EOS was first used to measure repetitive ultrafast electrical transients [VMG83; MV86]. Latter, the principle has been applied for THz radiation, also advancing the technique to EOSD [WZ95]. Electron bunch lengths and profiles were first obtained at linear accelerators [Wil+02; Jam+06]. KARA is the storage ring where first EO measurements were carried out both, using the THz radiation in far field CSR (see Sec. 2.3.2) at a beam port [Ple+09], as well as using a dedicated in-vacuum setup. The latter, since its installation in 2013 [Hil+13], is in regular operation at KARA. The setup is based on an EO crystal that is placed right next to the electron beam and allows for measuring bunch profiles. The principle uses the fact that the coulomb-field has almost the same shape as a charge distribution moving at relativistic speed (see Sec. 2.3.1). Comprehensive studies using the original setup, can be found in [Hil13], some more recent results – also using the optimized design – are in Sec. 5.4.

4.3.1. The Pockels-Effect

The measurement principle is based on a crystal that becomes birefringent when it is exposed to an electric field. This can be explained by the Pockels-effect. The Pockels-effect is an electro-optical effect that imposes a linear change to the refractive index when an external electric field E_{mod} is applied [Poc06; Voi08]. To be precise, the refractive indexes n along two of the crystal axes are proportional to the electric field. They are referred to as the “slow”-axis (refractive index n_s) and the “fast”-axis (n_f). With $E_{\text{mod}} = 0$, $n_s = n_f = n_0$. As described in [Ste07, p. 38] the maximum EO effect is obtained by alignment according to Fig. 4.2.

When a laser pulse is sent through the crystal with the Pockels coefficient r_{41} , this birefringence then turns the linear polarization of a laser pulse into an elliptical one. The modulation Λ is a function of the laser wavelength λ , the unperturbed refractive index n_0 and the crystal thickness d

$$\Lambda = \frac{2\pi d}{\lambda} \Delta n = \frac{2\pi d}{\lambda} n_0^3 r_{41} E_y. \quad (4.3)$$

At KARA, a Gallium Phosphide (GaP) crystal and a near-infrared laser (central wavelength 1030 nm) are used. The crystal thickness originally was $d = 5$ mm, but was increased to $d = 7$ mm to optimize performance (see Sec. 5.3).

4.3.2. The Electric Field Inside the EO Crystal

For EO measurements in the near field, it is used that the Coulomb field of the bunch has almost the same shape as the longitudinal bunch profile (see Sec. 2.3.1). The Pockels-effect, however, does not allow to differentiate between the primary coulomb field and wake fields (cf. Sec. 3.2). In fact, what is really measured is not the Coulomb field itself: As the EO crystal has a refractive index different from 1, the field can no longer co-propagate the

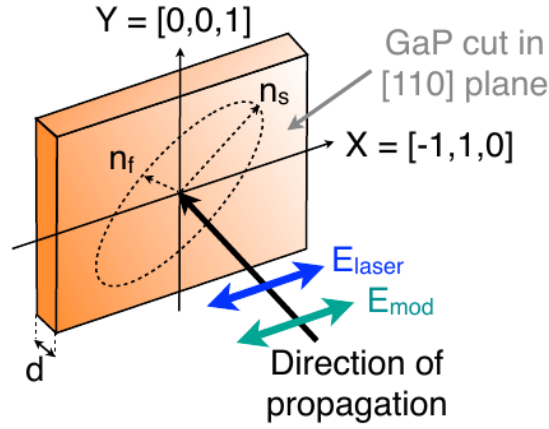


Figure 4.2.: EO Crystal Orientation

The electric field to be measured \vec{E}_{mod} and the laser, both propagate into the same direction. The latter is linearly polarized with $\vec{E}_{\text{laser}} \parallel \vec{E}_{\text{mod}}$. To obtain the strongest modulation, the crystal is placed in a way that $[-1, 1, 0] \parallel \vec{E}_{\text{mod}}$. Here, the electron beam would pass the crystal either on the left or on the right hand side. The direction of propagation for the electric fields is perpendicular to the (110)-plane with their directions of polarization being parallel to the (-110)-axis X. Also, the fast and the slow axis are displayed. [Hil13]

electron bunch and thus forms a primary wake field. Its shape corresponds to the Coulomb field but is clinched in the direction of propagation because of the higher refractive index n of the material. The wake field traverses the crystal from its front side to the back with $v = c/n$. This stripping off of the Coulomb field can be compared to the effect of a collimator for electron beams. From the bottom side, due to the limited resistivity, a second wake field penetrates the crystal. This wake field is referred to as a Cherenkov cone [Jam+06] and can be thought of as a resistive wall wake. As a counter-measure, it has been suggested [Ste16] to apply a metalized surface to the crystal's bottom side. Alternatively, when the height of the crystal is large enough, there are radial positions, where the Cherenkov cone and the Coulomb-like peak do not overlap.

On a longer time-scale, also secondary wake fields play a role. They develop due to a multitude of reflections of the electric fields – not only at the crystal borders but also at the crystal mount, parts of the laser transport system, and the electron beam pipe. To get an impression, the (typically complex) geometry has to be considered in simulations⁴. An example is shown in Fig. 4.3: The black curve marks the electric field that can be probed electro-optically. It has to be noted that the first peak is just a fraction of the signal. Similar effects have already been observed in the very first EOSD experiments, that measured a THz pulse in free space [WZ95]:

“The minor peaks following the main peak are mainly due to multiple reflections of the electrical pulse in the photoconductive emitter (impedance mismatch along the biased electrodes), and the terahertz pulse in the sensor crystal.”

For the measurement in an accelerator beam pipe, the first source translates to primary wake fields that are also present without the crystal. Those fields can be relevant for beam dynamics. When drawing conclusions from such measurements of such a long-ranged wake field, its position dependency has to be considered: The blue curve in Fig. 4.3 marks the vertical component of the electric field at the beam orbit, the red curve the longitudinal one. To influence the beam energy, only the longitudinal wake field plays a role⁵. It can be

⁴Sec. 5.2 will explain how the EO signal is derived from a simulation of the electric fields.

⁵The fields are of course related, see [Ng06, Chap. 1].




Figure 4.3.: Simulated Fields inside the EO Crystal and at the Electron Orbit

The EO setup is capable to measure the component of the wake field E_y caused by the electron bunch that is parallel to the laser modulation (black curve). However, this does not imply the electrons see a field of the same shape: First of all, the EO crystal is next to the electron beam path (here at a distance of $y = 11.5$ mm), but not right at the same position ($y = 0$ mm). In fact, the vertical field component of the Coulomb field is 0 V/m at the center of the electron bunch. Hence, the blue curve only shows wake fields but no Coulomb-like peak at $t = 0$. Secondly, the influence of interest would be in the longitudinal direction (E_z , red curve).

seen, that the the shape of the measurable field $E_y(y \neq 0)$ deviates from the one that can influence the beam energy $E_z(y = 0)$, also, the amplitudes differ by one order of magnitude. So, the component of the measured wake field that is used to study beam dynamics is the peak at $E_y(y \neq 0, t \approx 0)$. As described above, this peak is directly related to the Coulomb field of the electron bunch and – as long as its overlap with the Cherenkov cone is negligible – has a very similar shape. Especially for the bunch tail, however, modulation of the laser will also happen while it is counter-propagating the electron beam. This can lead to a significant signal distortion of the measured longitudinal profile [Hil13, p. 105].

4.3.3. Electro-Optical Spectral Decoding

As usually only the part of the wake field representing the bunch shape is of interest, and as the bunch shape varies with time, typically measurements are performed single-shot using EOSD. The working principle uses a long chirped laser pulse, where $f(\tau) \propto \tau$, so that the time information can be encoded in the laser frequency. An implementation is depicted in Fig. 4.4, the steps are:

1. The used laser [Hil13, Sec. 5.2] is a Ytterbium-doped fiber laser with a central wavelength of 1050 nm. Using a piezo-stretcher, its repetition rate is actively synchronized to the storage ring RF master oscillator (500 MHz) and – to be synchronous with the bunch – a harmonic of 2.7 MHz.
2. The (vertically polarized) laser pulse is sent through a polarization maintaining, dispersive fiber ($\partial n(\lambda)/\partial \lambda \neq 0$) which also introduces a chirp. The total chirp can be controlled (or pre-compensated) using a grating compressor inside the amplifier part of the laser. A second pulse stretcher/ compressor is placed in the free space part in proximity to the EO crystal. Currently, however, this second compressor is not used to decrease reflection losses. In the chirped pulse, a linear relation between the pulse wavelength and the arrival time can be assumed.
3. The EO crystal is placed inside the electron beam pipe. The required optics to couple the laser beam out of the fiber, into the EO crystal and later back into the next




Figure 4.4.: Schematic drawing illustrating the working principle of the EOSD setup at KARA. (Adapted from [Rot18, Fig. 2.6], see text for description.)

fiber, is placed on an optics plate, called the *EO monitor*. The in-vacuum part – the *EO arm* – features the crystal mount and a reflective surface so that the laser beam propagates parallel to the electron orbit (see Sec. 5.3 for details).

4. A polarizing beam splitter is also part of the EO monitor. It is used to couple vertically and horizontally polarized parts into different fibers.
5. Typically, only the vertically polarized pulse is used. Neglecting losses, in absence of the electron bunch, the *unmodulated pulse* contains the full power. In presence of the electron bunch, it is called *modulated pulse*.
6. To detect the modulated pulse, e.g. a grating spectrometer can be used. (For details, see Sec. 5.1.)
7. As the modulation Λ is proportional to the electric field [Eq. (4.3)], the bunch profile can be reconstructed by dividing the modulated pulse by the unmodulated one.

Similar to streak camera measurements, there is a fast and a slow time axis. For the slow time axis, the number of a detected single shot profile can be used. As the laser repetition rate is precisely known, no further calibration is needed here. While the repetition rate and timing of the electrons are fixed, the laser pulse can be shifted using fiber stretchers. The fast time axis can be calibrated based on the beam overlap: When the laser pulse is delayed by a known time, the position of the modulation will move accordingly.

5. Optimization of the EO Setup

“You would make a ship sail against the winds and currents by lighting a bon-fire under her deck? I have no time for such nonsense.”

Napoleon (about steamships)

To be able to resolve the rich dynamics of the electron bunches during bursting CSR, the existing EO setup was significantly enhanced. This chapter describes the steps taken to achieve turn-by-turn recordings of the bunch profile that show significant sub-structures as they are also predicted by the simulations (Chap. 10). In particular, it will be shown that accepting a (known) deformation of the recorded pulse allows to significantly increase the modulation strength. Measurements presented in this chapter have been carried out purely to verify the improvements made to the EO setup. Actual bunch-profile measurements also utilize other detectors introduced in the previous chapter and will be presented in Chap. 6.

5.1. Turn-By-Turn Detector Readout

The nature of the bursting CSR (Sec. 3.3.2) creates the requirement of the EO setup to resolve both, fast and slow processes. Synchrotron motion [Eq. (2.39)] defines the time scale of the fast ones. Structures on the bunch form and disappear within times $\Delta t < T_s$. The relaxation of an excited beam happens on the slow time scale $\Delta t > \tau_d$, defined by the damping time τ_d (Sec. 2.3.3).

These requirements are hard to meet: Commercially available spectrometers have repetition rates up to [cite something]. Compared to the tens of kHz of the synchrotron motion, this is not enough. Alternatively, a time-stretch approach [Rou+15] can be used. It uses a long,

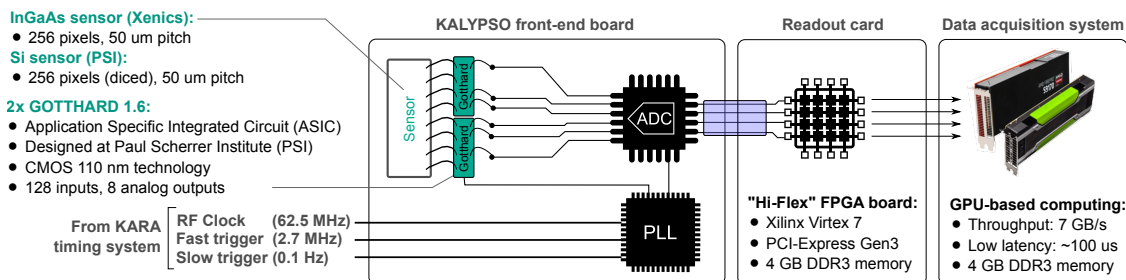


Figure 5.1.: Schematic Drawing of KALYPSO

The detector consists of a line detector with 256 pixels. The pixels are individually connected to one of two GOTTHARD (gain optimizing microstrip system with analog readout) frontend chips, so they can be read out all in parallel. Digitization is performed using a total of sixteen parallel channels. (Courtesy of Lorenzo Rota.)




Figure 5.2.: Histogram of KALYPSO Background Measurement

Displayed is the number of ADC counts for a measurement of the KALYPSO detector with no illumination by the laser. Sources for the noise include ambient light, fluctuations of the voltage supply and ground levels, and thermal noise. Notice that most pixels (6 %) show 0 ADC counts, but due to a slight asymmetry of the distribution, the fitted Gaussian distribution peaks at -0.83.

dispersive fiber to elongate the modulated, chirped light pulse after the EOSD measurement. Afterwards, it can be detected using a fast photo diode which is attached to an oscilloscope. This method offers an unmet repetition rate, however the recording time is limited by the amount of needed memory. In [Rou+15], the readout rate of 88 MHz was limited by the laser rather than the readout, the total time recorded was $420/(847 \text{ kHz}) = 0.5 \text{ ms}$.

To overcome these limitations, in cooperation between DESY (Deutsches Elektronen-Synchrotron), Łódź University, PSI (Paul Scherrer Institute), and KIT (Karlsruhe Institute of Technology), KALYPSO (**K**arlsruhe **l**inear **a**rray detector for **M**Hz-**r**epeatition rate **s**pectroscopy) was developed. A brief overview can be found in [Rot+17], a more comprehensive description is in [Rot18]. KALYPSO consists of a linear pixel array detector that can be made of either Si or InGaAs. Each of the 256 pixels is individually connected to one of two GOTTHARD frontend chips, allowing a readout at 2.7 MHz. So for KARA, it allows the continuous acquisition of one bunch at every revolution. The data is digitized using a 14-bit eight-channel ADC. (A schematic is shown in Fig. 5.1.) To have a basic data consistency check and properly aligned data (cf. Chap. 7), the two padding bits are used as a counter, so the total data rate is $1\,328.4 \text{ MB/s}$. Using a PCI-Express interface, the raw data can be directly streamed to a computer – either to the main memory or to the graphics memory.

For an error estimation, measurement results for a constant signal are histogrammed. Doing so, not only the standard deviation is known, but the complete probability density function can be estimated. Figure 5.2 shows such a histogram for the KALYPSO detector when the EO laser is turned off. As no signal is present, the measurement allows to investigate the background noise. The fact that the distribution is almost perfectly Gaussian might be seen as an indication that multiple independent sources contribute to the total noise of the detector (σ_{KALYPSO}). Sources for the fluctuation include ambient light, limited stability of the voltage supply and ground levels, and thermal noise. The first two sources (σ_{FEC}) are correlated for all pixels or for all pixels connected to the same GOTTHARD frontend chip. Thus it is possible to separate them from uncorrelated noise σ_{thermal} . This separation is shown in Fig. 5.3.

Notice that this separation has only been possible due to the knowledge of the correct result. If there are for example unknown laser modulations, since the instantaneous error of the baseline is not known, it is not possible to compensate for baseline fluctuation of the front end electronic. It however, helps to identify possible measures to reduce the noise level of the electronics.

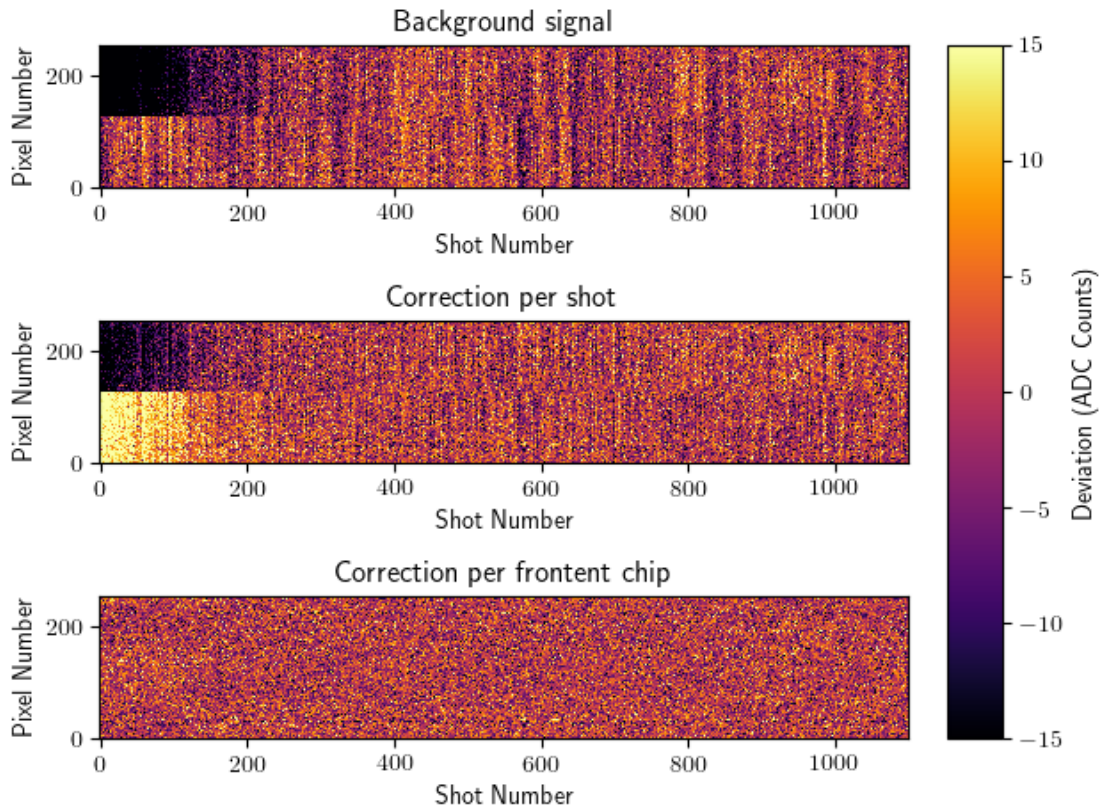


Figure 5.3.: Electronic Contributions to the EOSD Background Signal

The EOSD background signal (top) for the first number of shots shows a systematic drift of the background on all pixels attached to one GOTTHARD front-end chip as well as a fluctuation baseline afterwards. To investigate the source, average values are subtracted. If the same average is used for all pixels in every shot, noise contributions by e.g. ambient light or a shared ground level are corrected for. The resulting shot-to-shot intensity fluctuation is lowered but still visible for the two chips (center). When background level fluctuation treated individually on a per-chip basis, to compensate for e.g. per-chip voltage supply fluctuations, what remains is only uncorrelated noise (bottom). For comparison: The modulation due to the EO effect is up to 350 ADC counts. The amplitude of the substructures that are to be resolved, however, is only in the per-cent range of that. It has to be noticed that in the real experiment changing illumination from the laser cannot be distinguished from these background effects, so they have to be treated as an error.

The total statistical uncertainty for a measurement of the laser profile σ_{total} also has a contribution from the laser intensity fluctuation and from a timing jitter. The latter might lead to a varying charge loss at the moment the pixel is sampled.

$$\sigma_{\text{thermal}} = 5.94 \text{ ADC Counts}, \quad (5.1a)$$

$$\sigma_{\text{ADC}} = 6.78 \text{ ADC Counts}, \quad (5.1b)$$

$$\sigma_{\text{total}} = 12.81 \text{ ADC Counts}. \quad (5.1c)$$

As the contributions are independent, the errors due to the laser fluctuation σ_{laser} , and the




Figure 5.4.: Design of the EOS v1 EO Arm

In the sketch, metal is displayed in gray, light gray indicates parts that lay behind the cutting plane. The laser path (red) is from the top side, first to a glass prism with a metalized surface (the blue triangle). It serves as a mirror and is held by two metal fingers. It is reflected back from the back of the EO crystal (purple) with a small angle in the image plane and follows the same path on the way back. (Photo courtesy of Nicole Hiller [Hil13].)

frontend chip σ_{FEC} can be calculated as

$$\begin{aligned}\sigma_{\text{laser}} &= \sqrt{\sigma_{\text{total}}^2 - \sigma_{\text{ADC}}^2} \\ &= 10.87 \text{ ADC Counts},\end{aligned}\tag{5.1d}$$

$$\begin{aligned}\sigma_{\text{FEC}} &= \sqrt{\sigma_{\text{ADC}}^2 - \sigma_{\text{thermal}}^2} \\ &= 3.27 \text{ ADC Counts}.\end{aligned}\tag{5.1e}$$

In comparison, $\sigma_{\text{laser}} > \sigma_{\text{ADC}}$ and $\sigma_{\text{thermal}} > \sigma_{\text{FEC}}$. So to reduce the noise, first the laser stability should be improved. Afterwards better cooling of the detector might be considered.

5.2. Reconstruction of the Measured Wake Fields

To measure individual bunches in a multi-bunch environment, the wake field inside the EO crystal (see Sec. 4.3.2) has to decay before the next bunch arrives. For KARA, this means the fields may last for ideally less than 2 ns. The original geometry (*EOS v1*, see Fig. 5.4) was designed to be used at linear accelerators [SSM09] with a bunch spacing of 200 ns – and did not allow measurements in multi-bunch operation at KARA [Hil13]. In a first step needed to optimize the in-vacuum part of the EO setup, it is crucial to verify that the signal measured using the original geometry is well understood. To do so, the EO signal is reconstructed based on simulations using *CST Particle Studio* [CST17]. This software can be used to calculate electro-magnetic fields based on a computer aided design (CAD) model.

In the wake field solver [CST17], the electron beam is approximated by line charge with a Gaussian profile exciting the structure. The particle velocity is taken to be $\beta = 1$. When the particles pass discontinuities, the fields are scattered and create wake fields. Also, resistivity and magnetic permeability of the used materials is considered for resistive wall wakes. The field strength is saved at positions defined by so-called “field probes”. The solver also provides results such as the wake impedance, and loss factors.

The maximum frequency of the electro-magnetic field used for calculations is dependent on the set length of the electron bunch. A second parameter, the ‘wake length’, which describes the maximum distance from a source point, at which a wake field is evaluated.




Figure 5.5.: EO sampling signal using EOS v1 normalized to the first peak (Coulomb-like peak at $t = 0$). For *Simulation A* and *B* the bunch length is set to 3.4 mm respectively 4.5 mm, which are both possible values for the (non-constant) bunch length. Furthermore the position of the line where the laser passes the crystal differs by about 2 mm. Note that *Simulation B* matches the measured curve better for small times but has the tendency to overestimate the amplitude of the wake fields at later times, while it is the other way around for *Simulation A*. (First published in [Sch+15].)

Benchmarks have shown that the results converge already for a small wake length and are almost independent from the wake length until they get numerically unstable when it is set too large. The calculated impedance on the other hand becomes inaccurate for the highest sampled frequencies.

There are three things that have been taken into account to reconstruct the EO modulation Λ . First, the field is not constant at all positions inside the crystal. To attack this, a number of field probes recording E_y is placed along the line where the laser passes the crystal. Secondly, the laser pulse has a finite length but only the integrated intensity is measured while sampling. Thirdly, the speed of the laser light is taken into account while integrating over the laser path. Defining the point in time when the laser hits the reflecting surface of the crystal as $t = 0$ one yields an EO signal [Sch+15]

$$\Lambda \propto \int_{-T}^T E_y(z(t), t) dt = \int_{-T}^T E_y\left(\frac{c}{n} \times |t|, t\right) dt, \quad (5.2)$$

where T is the time, the laser needs to pass through the crystal in one direction, and n is the refractive index of GaP.

The result of this method can be seen in Figure 5.5. The first peak corresponds to the Coulomb field of the electron bunch, the ringing is due to secondary wake fields (see also Secs. 3.2 and 4.3.2). The differences between simulation and measurements might be explained in several ways. Of course, there are (and always will be) differences between the CAD model and the real setup. For example the (small) angle between incoming and outgoing direction of the laser has been neglected. Also the absolute position where the laser passes the crystal is not known. Here, it has been slightly varied to find a fair match. As the effects of those differences accumulate, they play a bigger role for later times. Also, the bunch length is constant in the simulation. In the measurement however, as a sampling method is used to span the wide time range, the bunch charge is decreasing with every sampling point. And while the decreasing signal strength has been corrected in the measurement data, this cannot be done for the decreasing bunch length. At last, the bunch in the measurement should not assumed to be Gaussian. So it has a significantly different spectrum that would have to be considered in the simulation to model the wake field. Considering all these points, the overall shape is matched pretty well.

The most prominent ringing has a periodicity of $T \approx 110$ ps, so it corresponds to $\lambda_0 \approx 3.3$ cm in vacuum or in GaP (refractive index $n = 3.3793$) to $\lambda \approx 1$ cm. The latter corresponds to

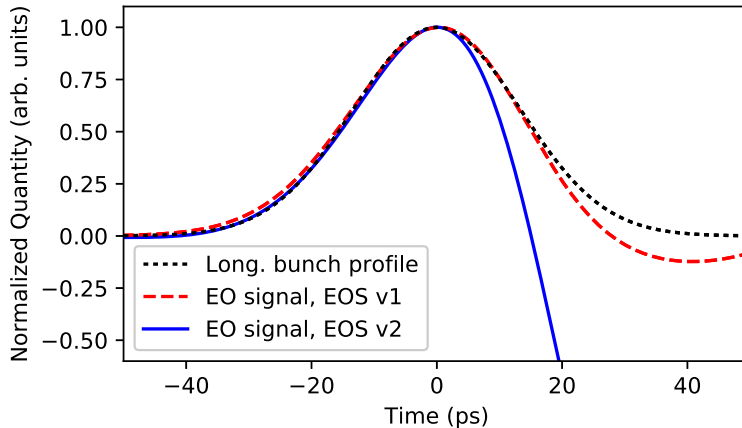


Figure 5.6.: Comparison of Bunch Profile and EO Signal

The Gaussian bunch profile used for the simulation (black), causes an electric field E_y inside the EO crystal. The resulting signal (colored) that would be measured using the two versions of the EO setup was calculated according to Sec. 5.2. In the measurement it seems as if the bunch was leaning to its tail. This is the opposite to what it is doing because of the potential well distortion in the CSR wake. Due to the reduced distance between laser and electron beam, the effect is stronger for *EOS v2* (cf. Sec. 4.3.2). As the deformation is known, it can be considered in the data evaluation.

the height of the crystal which can be identified as the main source of the long ranging, secondary wake field and also¹ to the diameter of the vacuum inside the EO arm. The simulations show that both can be seen as cavities that store energy for a longer period of time, so they should be changed in an updated design.

5.3. Optimization of the EO Arm Geometry

The most simple and efficient measure to avoid wake fields is to increase the distance between the structure and the electron beam. Without any further adjustments, however, some parts of the wake field may stay unchanged [Hil13, Sec. 7.2.2] whereas the change would also decrease the Coulomb-like peak. As the important distance for the detection of this peak is the one between the laser and the electron beam, a solution is to decrease the distance between the laser and the lower edge of the EO arm. This change alternatively allows to move the laser closer to the beam axis. To further increase the height of the Coulomb-peak, the crystal width can be increased: Based on the dispersion of the field entering the crystal at the front surface, the phase retardation first grows approximately linear with the thickness. At a certain point that is dependent on the bunch length, saturation is reached. As an example: For GaP crystals and bunches with a length of 100 fs, effects of the saturation are visible starting from a crystal thickness of 300 μm , but phase retardation still grows until 400 μm . Extrapolating from the results shown in [Ste07, p. 73], for bunch length of 2 ps or more, saturation should not be reached before a crystal thickness of 10 mm. Also a pulse lengthening is predicted for thicker crystals. So, the increased signal comes at the cost of a limited resolution.

For the optimized EO arm (*EOS v2*), both changes are applied: The crystal height (and width) is decreased from 10 mm to 5 mm, and its length is slightly increased from 5 mm to 7 mm. Notice that these changes also increase the effect of the Cherenkov cone (see Sec. 4.3.2). Even for measurements using *EOS v1*, it seems as if the bunch was leaning to its

¹Just for information. As the distance is in vacuum, the inner arm diameter cannot account for the most prominent ringing.




Figure 5.7.: Design of the EOS v2 EO Arm

Metal parts are displayed in gray, the EO crystal (GaP) in violet. The metal surface facing up left replaces the prism used as a mirror in EOS v1. The other main improvements are the conical shape inside the EO arm and the decreased distance between laser path and lower edge of the EO arm.

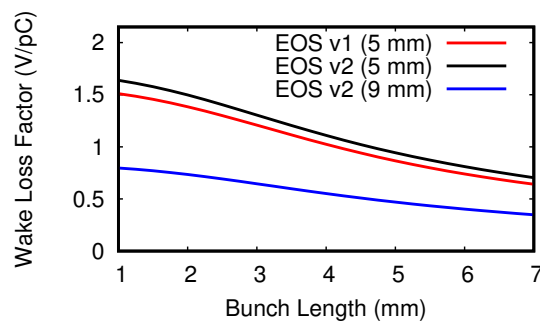


Figure 5.8.: Wake Loss of EOS v1 and EOS v2

Wake losses calculated using *CST Particle Studio*, assuming a Gaussian bunch profile. It can be seen, that the losses can be approximately halved using the 9 mm setting. In the 5 mm setting, on the other hand, the modulation is approximately doubled.

tail, but Fig. 5.6 shows that the increased signal height comes at the cost of a significantly stronger deformation of the measured profile. As the deformation is systematic, it can be considered in the data evaluation. The changed crystal dimensions, on the other hand, increase the signal-to-noise ratio by a factor Λ_2/Λ_1 dependent on the distance Δy between the lower end of the EO arm and the electron beam. Based on Eqs. (2.13) and (5.2), it can be estimated by

$$\begin{aligned} \frac{\Lambda_2}{\Lambda_1} &\approx \frac{l_2}{(\Delta y + h_2/2)^2} \times \frac{(\Delta y + h_1/2)^2}{l_1} \\ &= \frac{(\Delta y + 5 \text{ mm})^2}{(\Delta y + 2.5 \text{ mm})^2} \times \frac{7 \text{ mm}}{5 \text{ mm}}. \end{aligned} \quad (5.3)$$

An exact expression would have to consider not only the velocity of the electrons, but also the exact path of the laser. Here, it has been taken to be centered in the crystal, neglecting also small angles that have to be present independently from the actual alignment. In total, the new geometry realistically allows to increase the modulation strength by up to a factor of four. In this case, however, it will probably still be limited to single-bunch operation.

The typical distance between the bottom of the EO arm and the electron beam Δy was estimated to be approximately 5 mm. To have a modulation as strong as before, using EOS v2, the new design distance is 9 mm.




Figure 5.9.: Wake impedance of EOS v1 and EOS v2

Real part of the wake impedance for EOS v1 (red) and EOS v2. Δy is the distance between the lowest part of the EO arm and the nominal beam orbit. For EOS v2, the optimized distance between laser and lower edge of the arm leads to the possibility to increase the distance to $\Delta y = 9$ mm at the original signal level. For EOS v2 and $\Delta y = 5$ mm the EO signal will be increased. (First published in [Sch+17a].)

To disallow resonant modes, parallel surfaces are avoided in the metal part of EOS v2. For the mirror, replacing the prism holder (the two metal fingers) by an integrated mirror surface also allows to keep small size so that the advantages of the smaller crystal can be fully exploited. The inner side has a conical shape (see Fig. 5.7) and was inspired by a pickup beam position monitor that was also designed to avoid after-ringing [Ang+12; Ang+15].

Figure 5.8 shows the wake losses calculated for EOS v1 ($\Delta y = 5$ mm) and EOS v2 ($\Delta y = 5$ mm and $\Delta y = 9$ mm) and a range of bunch lengths between 1 mm to 7 mm. It can be seen that the losses slightly increase for the closer setting – which gives approximately doubled modulation when compared to EOS v1, but the alternative setting without increased modulation approximately halves the losses.

To have a more general impression, that is also valid for non-Gaussian bunch profiles, the real part of impedance $\Re[Z(f)]$ has to be considered. Figure 5.9 shows that it has decreased almost over the whole simulated frequency range, only a contribution around 5 GHz is significantly stronger than before. The predicted wake can be found in Appendix C.1, a measurement result is shown in the next section. To investigate the region around $30 \text{ GHz} = c/1 \text{ cm}$ where there was the most prominent ringing in EOS v1 (cf. Fig. 5.5), the number of grid points would have to be increased above the limit that was allowed with the used CST license.

In total, the new EO arm design allows to increase the distance to the electron beam, leading to a significantly lower wake loss. Alternatively, the modulation (and thus the signal-to-noise ratio) can be increased by a factor of up to four. These improvements come at the cost of a deformation of the signal that has to be considered in the data evaluation. Furthermore, a conical shape of the inner side as well as the avoidance of parallel surfaces reduce the impedance and also counter-act the persistence of secondary wake fields.

5.4. Test Measurements

To evaluate the optimized EO arm, a series of test measurements was carried out. At first, the absolute signal was recorded using a photo diode while scanning the time delay between laser pulse and the revolution signal from the accelerator. The result yields the needed delay for measurements of the electron bunch’s electric field as well as a signature of the trailing wake fields inside the crystal. As it was the standard procedure before the upgrade,





Figure 5.10.: EOS wake field measurements using EOS v1 and EOS v2

The overall improvement is clearly visible: For EOS v2, the wake field decays much faster in comparison to the measurement with EOS v1, where the wake field persists over the entire measurement range of $t = 2.5$ ns. (First published in [Sch+17a].)

for this test, the EO arm was moved to the electron beam as close as possible without loosing electrons. Figure 5.10 displays the measured wake fields for EOS v1 and EOS v2. In the measurement with EOS v1 the wake field persists over the entire measurement range of $t = 2.5$ ns. Using EOS v2, the overall improvement is clearly visible: the wake field decays much faster. At $t = 2$ ns, where, in a multi-bunch environment, the subsequent bunch would be located, the wake field has already dropped roughly to background level. Notice that this major goal of the redesign of the EO arm could be reached even without increasing the distance to the beam.

Two single-shot electron bunch profile measurements – recorded before and after the upgrade to EOS v2 – are shown in Fig. 5.11. The improved data quality is clearly visible: As predicted by simulation, the modulation is increased approximately by a factor of two². Due to continued firmware optimizations, the signal-to-noise is even further increased. The electronic noise contribution was reduced from $\sigma_{ADC} = 11.7$ to $\sigma_{ADC} = 6.8$. [See Eq. (5.1) for individual noise contributions.] The sensitivity has significantly improved, allowing for clearer observation of substructures on the bunch profiles (see Sec. 6.2).



(a) EOS v1 modulation profile

(b) EOS v2 modulation profile

Figure 5.11.: Single bunch profile recorded using EOSD

The modulation profile recorded using EOS v1 (a) is dominated by noise. In this single profile it is hard to see substructure. For EOS v2 (b), the modulation is stronger – approximately by the factor of two that has been predicted in the simulations. As expected, the adjacent wake is influencing the overall shape of the measured profile. Due to the deformation of the measured pulse but also of the electron bunch, structures are still hard to see in a single profile. The darker, inner error bands are due to the electronic noise ($\sigma_{ADC} = 11.7$, $\sigma_{ADC} = 6.8$), the lighter, outer ones include all sources of noise – especially laser power fluctuation. The influence of detector is clearly visible, there are pixel defects (e.g. pixel 104). (Data first published in [Sch+17a]).

²A second option to influence the modulation is alternating the separation in the beam splitter. Changing the laser focusing on the other hand can only influence the total intensity.

6. Bunch-Profile Measurements

“We are never in a position to say what really is or what really happens, but we can only say what will be observed in any concrete individual case.”

Erwin Schrödinger

This chapter shows possibilities and provides experimental data obtained using the detectors introduced in Chap. 4. The data can serve as benchmark scenarios to show that simulation results describe real beam dynamics. It also points out the complementary advantages of using a streak camera (Sec. 6.1) or the near field EO setup (Sec. 6.2), and compares bunch profiles obtained using the two methods (Sec. 6.3). At last (Sec. 6.4), an experimental example for the connection between the evolution of the bunch profile and the emission of CSR is shown. Interpretation is done later (Chap. 10), based on numerical studies.

6.1. Streak Camera Measurements

To obtain the average bunch profile including a precise measurements of the average bunch length, streak cameras (see Sec. 4.1) are the weapon of choice. While they in principle allow for single turn measurements, due to the strong influence of shot noise, they lack the ability to resolve short substructures on highly compressed, low-charge bunches.

The result of a typical single streak camera measurement is shown in Fig. 6.1. The horizontal and vertical time axis are variable and can be calibrated by the temporal gradient of the sweep voltage. In the shown example, the sweep voltages are set so that the horizontal axis shows 500 μs , and the vertical axis 191 ps. For a bunch length measurement, the time span of the horizontal (slow) sweep should be chosen so that the synchrotron oscillation can be resolved. However, a very short timescale introduces high statistical uncertainty due to the photon emission's Poisson noise. Notice that single photons can be identified in Fig. 6.1 and also that the individual pulses of light arriving at the streak camera have a certain width. So the horizontal axis has a uncertainty based on the spot size. Even though single-bunch resolution is possible also in a multi-bunch environment, due to the small photon count, a slower horizontal sweep combined with single-bunch operation is preferable. This way, resolution on the slow time axis is sacrificed to an extent that it is no longer possible to resolve individual turns, to get the best resolution for the averaged bunch profile. Without single-bunch operation the signals of all bunches overlap and it is not be possible to distinguish between them or (even worse) between pulse shape and synchrotron oscillation. So the strong optimization for the averaged profile requires single bunch operation.

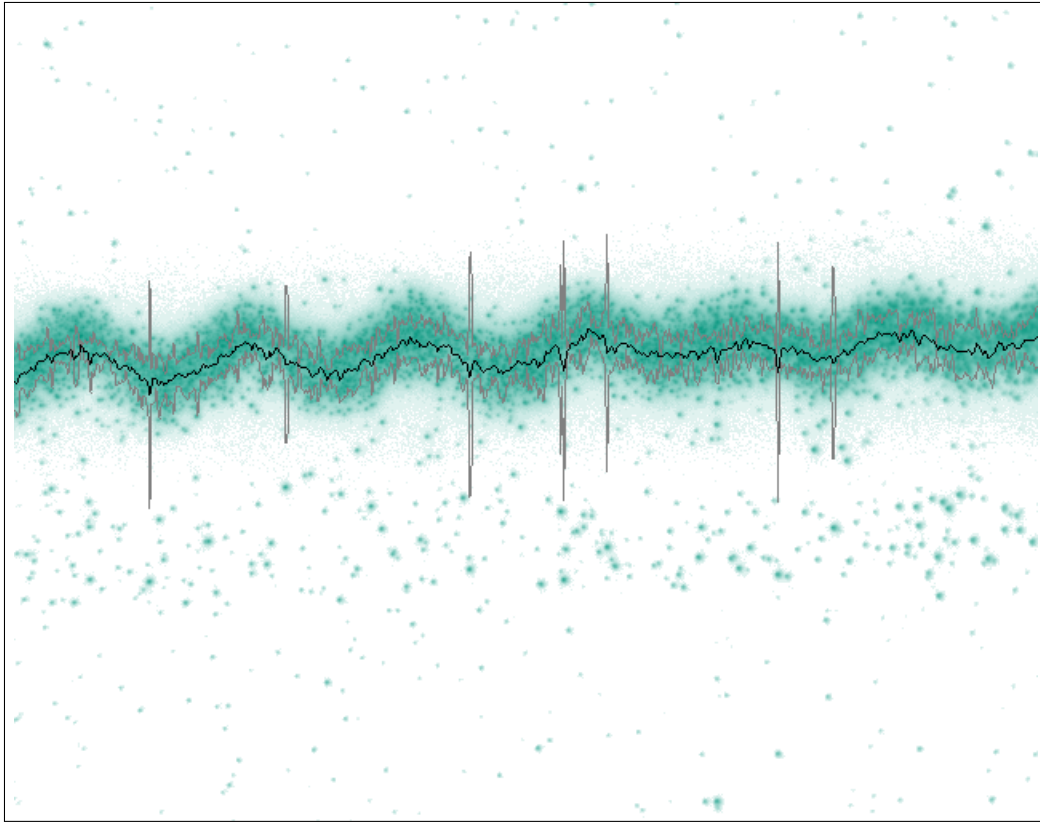


Figure 6.1.: Raw Image Recorded using a Streak Camera.

The position the photo electrons (green) hit the screen correlates to its arrival time. In the shown case the sweep voltages are set so that the horizontal axis shows 500 μs , and the vertical axis 191 ps. However, even without sweep, there is non-zero spot size coming from the electron beam diameter. To minimize the effect of the transverse shape on the bunch length measurement result, the spot has been chosen to be wide and flat. There are overlays added for the bunch centroid (black) and its duration (gray) based on single vertical slices.

The bunch length is obtained from a streak camera image using the projected bunch profile to the vertical axis

$$\rho(y) = \sum_{x=0}^{N_x} I(x, y), \quad (6.1)$$

where $I(x, y)$ is the recorded intensity at the pixel indexed x, y . As it is more robust against shoulders in the signal and remaining outliers, the FWHM is the measure of choice to describe the bunch length from experimental data. To remove the influence of the coherent synchrotron oscillation, first identifiable background photons are suppressed by discriminating outliers

$$I'(x, y) = \begin{cases} I(x, y) & \text{if } y_{\min} < y < y_{\max}, \\ 0 & \text{else,} \end{cases} \quad (6.2)$$

where $y_{\min/\max}$ are the positions closest to $\max(\rho(y))$ where $\rho(y) = 0$. Afterwards, the synchrotron oscillation is deducted, moving the signal to the vertical center

$$I'(x, y) = I(x, y + N_y/2 - \text{COM}(y(x \pm \Delta x))). \quad (6.3)$$

$\text{COM}(y(x \pm \Delta x))$ marks the center of mass y position in the surrounding of x . The two steps described by Eqs. (6.2) and (6.3) are repeated iteratively until the projection has

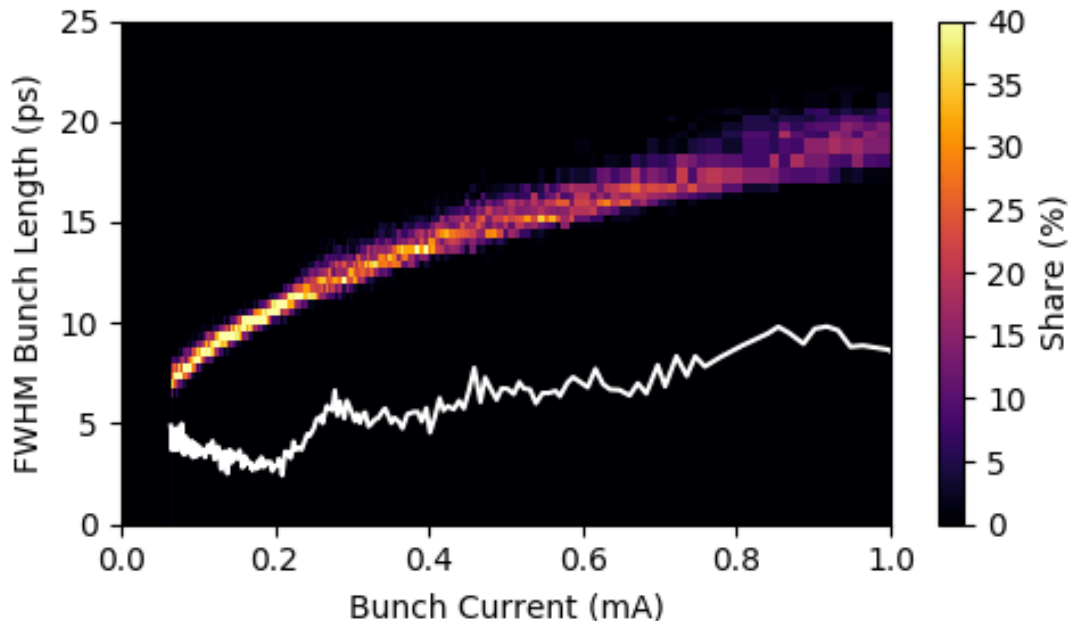


Figure 6.2.: Bunch Length Histogram (Using Streak Camera Measurements)

For constant currents, there were series of 100 bunch length measurements using a streak camera. The results are shown in a histogram where each vertical line stands for one of these measurement series. The color indicates the share of these 100 measurements that was evaluated with a given lengths. To visualize the spread of the obtained distribution of bunch lengths, 10 times the standard deviation is added as a white line. Its increase below currents of 0.2 mA can be attributed to Poisson noise, but for higher currents the influence of noise decreases. Instead, the distribution becomes broader because of fluctuating length.

converged. More details to the data evaluation including benchmarks based on simulated data can be found in [Sch13].

Using this method, static bunch deformation (see Sec. 3.3.1) and lengthening as well as fluctuating bunch length (see Sec. 3.3.2) can be observed. The latter, however not as a dynamical process: Obtaining the bunch length – e.g. using the algorithm described above – requires to set the horizontal axis to a time range that is able to resolve the coherent synchrotron oscillation. Otherwise, for a changing amplitude of the oscillation will cover the changing length. Thus, a statistical approach is chosen. Figure 6.2 shows this fluctuation: For constant currents, series of 100 bunch length measurements were done. The results are then represented in a color-coded histogram. It has to be mentioned that the uncertainty is much lower than the width of the distribution: The increasing spread below currents of 0.2 mA can be attributed to Poisson noise, but for higher currents the influence of noise decreases. Above 0.2 mA, the fluctuation is not a measurement artefact, the actual values, however, are influenced by the averaging done to obtain the single bunch lengths. So high-frequency fluctuation cannot be resolved using this method.

As the FWHM pretty much has to be used for bunch length measurements but theory [e.g. Eq. (3.24)] focuses on the standard deviation, the bunch lengthening is quantified based on simulation results (see Sec. 10.1.3).

6.2. Electro-Optical Measurements

Electro-optical spectral decoding offers single-shot measurements of the bunch profile. In contrast to streak camera measurements, every single revolution can be separately recorded




Figure 6.3.: EOSD measurements using EOS v1 and EOS v2, both using the KALYPSO 2 detector system. In both cases, the bunch head points to the bottom side (lower pixel numbers). The time-range has been chosen to that substructures emerge roughly in the center of the horizontal axis. The different signal to noise ratio is due to the increased modulation – the EO arm has been moved as close to the beam as the one before. (Upper dataset already published in [Rot+17], lower plot already published in [Rot18].)

with a signal-to-noise ratio that allows to see dynamic changes. A further advantage is that continuous recordings are possible, only data storage sets the limit. As displayed in Fig. 5.6, the measurement principle may introduce significant deformations to the measured profile. The effect is systematic, so it does not influence the observation of temporal deformations, significant structures and their turn-to-turn evolution. However, when doing quantitative analysis, it should be considered. In Sec. 6.3, it will be discussed in comparison to streak camera recordings taken at the same time.

Figure 6.3 shows two example recordings featuring emerging substructures. As for the streak camera data (Fig. 6.1), the synchrotron oscillation is clearly visible. In contrast to streak camera recordings, the vertical slices feature individual bunch profile measurements. For the EOS v1 data, the influence of the substructures is visible e.g. in a $6f_s$ modulation at the bunch tail (upper side of the bunch profile after turn 4500). Due to the signal-to-noise ratio, the traces are, however, easy to oversee. For the EOS v2 data, thanks to the improved modulation strength and the decreased electronic noise, there is a clear intensity modulation visible also in the bunch center. Also notice that the bunch recorded using EOS v2 seems to lean backward. This is not real but due to the aforementioned measurement artifact.

As continuous recordings cannot be evaluated by hand, efforts have been undertaken to allow automatic categorization of the bunch profiles using unsupervised machine learning. Details on the method (investigated based on simulation data) can be found in [Boll17].

6.3. Comparison of Bunch Profile Measurement Methods

In the previous sections, the advantages and disadvantages of the two measurement principles, streak camera and EOSD, have been pointed out individually. To directly compare results, Figure 6.4 shows one averaged bunch profile measurement, simultaneously recorded using both detectors.

For streak camera measurements, it can clearly be seen that the bunch is leaning forward, due to the effect of the potential well distortion (cf. Sec. 3.3). For the EO measurements




Figure 6.4.: Longitudinal Bunch Profiles measured using SC and EOSD


Comparison of averaged longitudinal bunch profiles [Hil13, p. 107], simultaneously measured using a streak camera (black) and EOSD (turquoise). The head is to the left. For the streak camera measurements, it can be seen that the bunch is slightly leaning forward. This is expected due to the potential well distortion in the CSR wake (cf. Sec. 3.3). For the averaged EOSD measurements this effect is not visible. Using EOS v2, it is even stronger (see Fig. 5.11). However, as the effect is known, it can be considered in the data analysis.

(in this case using *EOS v1*), on the other hand, the effect becomes invisible, the maximum is centered. As introduced in Sec. 4.3.2, this is because EO measurements are sensitive to all electric fields that have the correct polarization E_y . So the measurement result relies heavily on the wake field inside the crystal. For short time ranges, mostly the primary wake fields, namely the Coulomb like wake field and the Cherenkov cone, are of importance. The additional laser modulation caused by the Cherenkov cone is mostly visible at later times, so it completely shadows the fact that the bunch is leaning forward. In contrast, for a Gaussian bunch profile the impression would be the bunch was leaning backwards. Another possible problem is that the field profiles recorded using EOSD show adjacent wake fields, even when the effects of the Cherenkov cone is minimized. This makes it hard to determine the point where the bunch profile has dropped back to background level (see right hand side in Fig. 6.4).


Summarizing, the sensitivity to arbitrary wake fields introduces two problems to the near field EO measurements that both effect absolute results. For a qualitative analysis, they are not important. Streak cameras, on the other hand, typically do not yield single turn recordings that allow to properly identify the shape of the bunch profile at all (cf. Fig. 6.1). If averaged, the error due to dispersion in the beam line and other sources of point spread is significantly lower than the error for EO. The remaining challenge for streak camera recording is to correctly cope with the synchrotron oscillation, which is complicated by the possibly low number of recorded photons for single revolutions.

6.4. Correlation with THz-Bursts

Both detector systems, the streak camera and the EO setup allow the use of measurement triggers for synchronized data recordings. The successful synchronization was verified by applying steps to the phase of the accelerating voltage. Using this well understood mechanism (see Sec. 2.4) for a beam based method, all detectors are simultaneously exposed to changing beam conditions that are easy to interpret. Details on the synchronisation and its evaluation can be found in [Keh+17b].



(a) 18 ms of simultaneous measurement



(b) Zoom into the first burst

Figure 6.5.: Synchronous EOSD and THz Measurements

(a) shows 18 ms of a simultaneous measurement of the longitudinal bunch profile (top) and the CSR (bottom). In this time, here, four sawtooth-like “bursts” are observable. These bursts are easily identified in the detected radiation. In the displayed time frame, the amplitude of the coherent synchrotron oscillation changes several times, as can be seen best at the borders of the recorded bunch profile. The zoom-in to the first burst (b) shows (barely visible) structures that are present on the bunch profile during the burst. (THz detector data courtesy of Miriam Brosi.)

Figure 6.5 shows synchronized turn-by-turn recordings of the longitudinal bunch profile (using EOS v1) and the emitted CSR detected using a broadband Schottky diode attached to KAPTURE. Unlike previous measurements (e.g. [Rou+15]), they are not dependent on a trigger that just happens to be right before the burst. Instead, for the first time, a time span containing multiple bursts can be recorded, allowing to zoom into any region of interest afterwards.

The THz radiation shows a characteristic pattern that consists of a fast rise marking the beginning of the burst and a significantly slower decrease. In between, the power oscillates with a much lower amplitude. In the bunch profile measurement, there is a clear coherent synchrotron oscillation (Fig. 6.5(b)). In the long-term measurement (Fig. 6.5(a)), it becomes visible that the amplitude of the synchrotron oscillation is not constant but more or less abruptly changing. It should be noted that there is no obvious coincidence with the bursts. On the contrary, all possible combinations can be observed in this dataset: At the beginning of the first burst, there is no change, at the beginning of the second burst, synchrotron oscillation decreases, at the beginning of the third burst, it increases, and for the last one, again, there is no change. Another feature visible in the zoom-in is the emerge of substructures in the burst. They are slightly visible as a periodic pattern with five to six times the synchrotron frequency. Those correlate to the fluctuation of the measured THz intensity that show the same frequency. (A comprehensive description of the dynamics, based on simulation data, is presented in Sec. 10.2.2.)

7. Parallelized Computing

“There are only two hard things
in Computer Science:
cache invalidation, naming things,
and off-by-1 errors.”

attributed to Phil Karlton

To perform systematic numerical studies, a fast simulation tool is needed. One cornerstone in programming such a code is parallelized computing. Here, some fundamentals are given, aiming to help understanding design choices in the implementation.

Although most today’s computers work differently, *Von Neumann* computers are still broadly taught in computer lectures (e.g. [BU10; GS17]). When one neglects multiple cores and the presence of different buses for instructions (e.g. the program) and data, the description can still serve as a model for a single *Compute Core*¹ or processing unit (PU). One reason is that additions like *pre-fetching*, *caching*, *non-executable-bits*, etc. are mostly transparent – they are literary invisible to the usual programmer. For a superficial overview, even the aforementioned splitting of buses for instructions and data – the main memory still is shared – can be ignored. In this work, the most important features of this concept are just named briefly.

A simple processor will consist of at least² a control unit (CU) that controls the instruction stream, and an arithmetic logic unit (ALU) that does the actual computation. The *Von Neumann instruction cycle* consists of five steps that are processed in an infinite loop. They are:

1. **Fetch:** The CU fetches the next command from the command queue.
2. **Decode:** The CU interprets the command and directs it to an ALU.
3. **Fetch Operands:** The parameters for the command are loaded to the ALU.
4. **Execute:** The actual computation is performed (at the ALU).
5. **Write Back:** The result is written back to the memory.

For parallelized computations, the obvious approach is to have some of the components (only ALUs, or whole PUs) be present more than once; but as Sec. 7.1.1 shows, this does not have to be the case. Before going into detail, the next section will structure the possibilities in an abstract way.

¹OpenCL terminology

²At this high level of abstraction, registers do not need to be explicitly named. Also, memory may be considered to be managed outside the processor.

7.1. Parallelization Categories

A widely used scheme to categorize the different types of parallelization is called *Flynn's taxonomy* [Fly72]. It differentiates whether data or instruction streams are parallel, by naming four different kinds of parallel computing (see Fig. 7.1).

SISD (single instruction stream, single data stream) : Historically, this was the most common case for desktop computers. It will be discussed in Section 7.1.1.

SIMD (single instruction stream, multiple data streams) : Many ALUs share one CU. This data parallel approach offers the best hardware-scalability. However, algorithms have to be fitted to this model (see Section 7.1.3).

MISD (multiple instruction streams, single data stream) : This parallelization model can be implemented using different computers, 'voting' on the result of a computation to cope with bit-flips, e.g. in a high-radiation area. As it is (only) important for fault tolerance, it will not be further discussed in this work.

MIMD (multiple instruction streams, multiple data streams) : Every piece of data can be processed using its own instruction, so this model is more flexible than the SIMD approach but also needs a more complex hardware. It will be discussed in Section 7.1.2.

To do efficient parallel programming, one has to be familiar with the pros and cons of the different models. For example, it might be worth to rewrite an algorithm in a SIMD way, but it is also very easy to oversight a detail that will need MISD instead.

7.1.1. Single Instruction, Single Data

A *Von Neumann* computer, as briefly (re-) introduced in the beginning of the chapter is an example for a SISD design: It has one stream for both, data and instructions. However, even this architectural approach with a single processor (core) allows a limited amount of parallelization. This approach is called instruction-level parallelism and may be achieved using e.g. *pipelining*, *superscalar execution*, or *very long instruction words* (VLIW). The method will be explained using the following simple program (example taken from [Wiki18]):

```
1 e = a + b
2 f = c + d
3 g = e + f
```

The first two operations are independent and could be computed in parallel. The third operation, however, depends on the result of the first two, and thus has to wait until the results are saved to **e** and **f**, respectively.

In general, operations need more than one clock cycle [Gas+12] and different operations may take different times. Traditionally calculation was computationally more expensive than memory access. Nowadays, in some cases this relation does no longer hold true: A simple computation on data already present will be done even before an electric signal can travel the multiple centimeters to the memory modules. To simplify the explanation, this fact will be ignored in the following. Furthermore, to keep the example short, fetching and decoding the instruction will also be neglected. Strictly speaking, this implies the use of different buses for instructions and data, but as stated before, for current hardware this is the case anyways. For a very simple computer architecture, the above example may consist of the following subcommands:

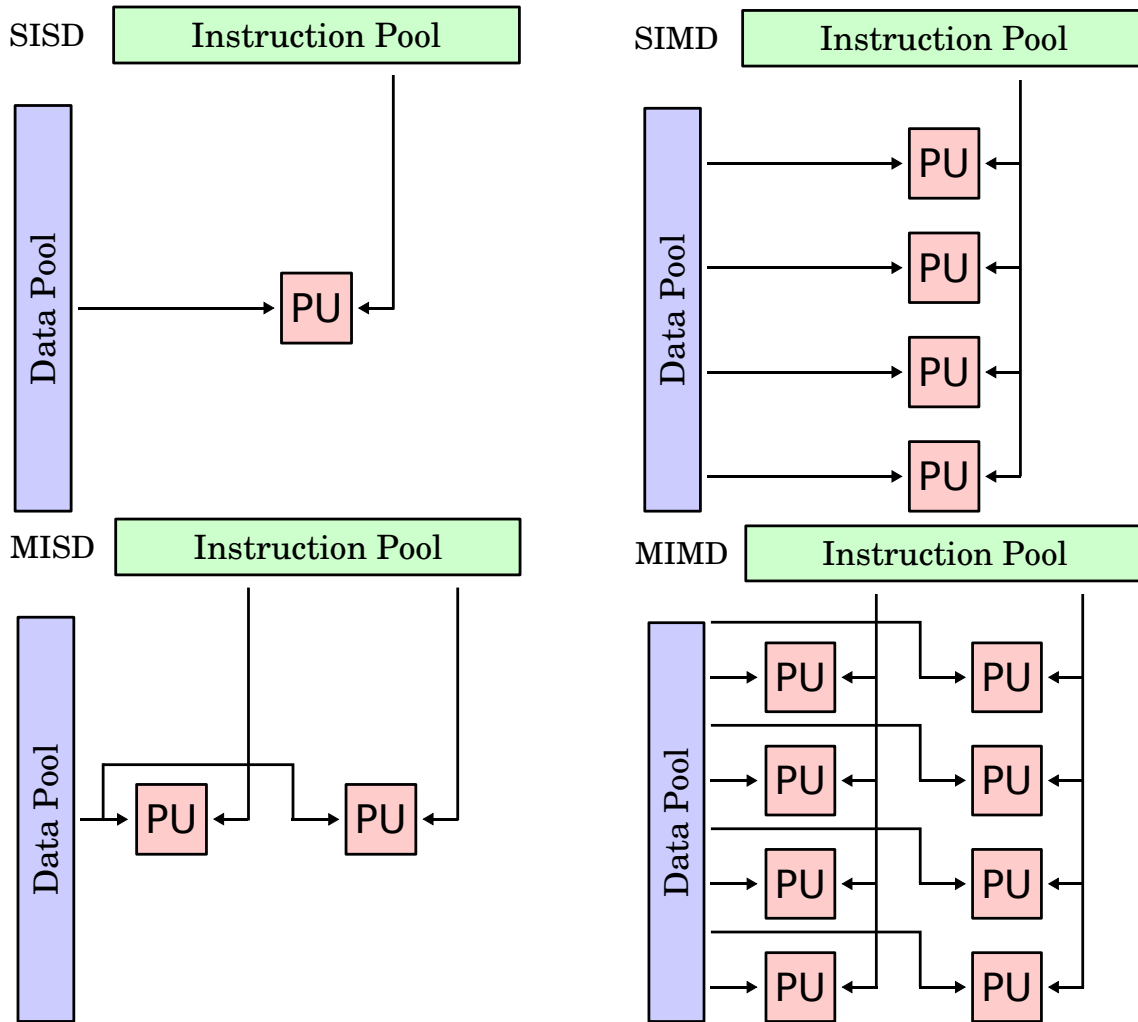


Figure 7.1.: Flynn's taxonomy of computer architectures

Flynn's taxonomy [Fly72] classifies computer architectures by the number of instruction streams and data streams. Traditionally, most (Von-Neumann-inspired) computers were SISD. With just a single PU, they use a single instruction stream (SI...) to work on a single data stream (...SD). A typical example for SIMD is found in graphic cards: They can apply the same instruction to multiple data streams (...MD) at the same time, e.g. when blending two overlapping, semi-transparent images. MISD is mostly relevant for fault tolerance: If the PUs yield different results, an error must have occurred. MIMD includes clusters and super-computers, but as of today, most PCs have multi-core processors and those also do clearly belong into this category. (Illustrations taken from [Wiki17], similar also in [Fly11].)

1. Read a
2. Read b
3. Add a and b
4. Write result to e
5. Read c
6. Read d
7. Add c and d
8. Write result to f
9. Read e
10. Read f
11. Add e and f
12. Write result to g

Pipelining can mean that parts of hardware serving different purposes (here: reading, calculating and writing) are used independently from each other. This is visualized in Fig. 7.2. It is also possible to further divide operations into smaller chunks. So for example a new computation can be started before the previous one is finished. A superscalar processor [Gas+12] will automatically check data dependencies and issue commands from a sequential instruction stream. This can be combined with out-of-order execution to allow reshuffling independent commands to be more effective in keeping the full pipeline busy. Of course, this needs logic in the processor, so as an alternative, using VLIW the complexity of decision is delegated to the compiler. To do so, very specific knowledge on the processor is needed at compile time.

Clock Cycle	Read	Add	Write
1	a		
2	b		
3	c	a,b	
4	d		e
5	e	c,d	
6	f		f
7		e,f	
8			
9			g

(a) Scalar, pipelined computation

Clock Cycle	Read	Add	Write
1	a,b		
2	c,d	a,b	
3		c,d	e
4			f
5	e,f		
6		e,f	
7			g

(b) Pipelined computation with vector reads

Figure 7.2.: Visualization of an instruction pipeline processing the aforementioned example. In a pipelined, scalar execution model (a), the execution of the 12 subcommands takes 9 clock cycles. Slowly moving away from a purely scalar design, vector reads (b) are an option to allow reading adjacent chunks of memory in one step. Even with sequential processors further optimization is possible. *Superscalar execution* will allow to check dependencies in hardware and *out-of-order execution* will then fill the pipeline with commands that do not need current ones to be completed. Another approach is to use *very long instruction words* that allow the compiler to explicitly state which (scalar) operations are done in parallel. This, however, needs very specific knowledge on the processor at compile time.

7.1.2. Multiple Instruction, Multiple Data

MIMD is the simplest approach to increasing the amount of work per clock cycle: Instead of a single processor (or processor core), multiple of them are used. *Symmetric multiprocessing* (SMP), where a number of equal processors is connected to the same memory through one shared bus, has been used for decades. Multi-core architectures simply scale this model down to a single CPU, plus they allow for more sharing of resources. On the other extreme,

computing clusters do not share any local resources on a single computer but are built of multiple computers connected by a network.

Notice that in a multi-threaded environment, the subdivision of the operations becomes an important (anti-) feature: Consider `int i = 0;` to be incremented using `i++;`. If there are two threads doing this there are two possible results: It will be the intuitive result (`i==2`) if the second thread reads after the first one has written its result. Otherwise, it might also be `i==1`. Notice that even plain assignments or reads cannot be assumed to be done in one step. For example, when accessing a 64 bit variable on a 32 bit computer, two separate operations are performed: one for the upper, one for the lower 32 bits [Pres17]. To handle this, there are so called *atomic operations* that may not be divided.

In contrast to the explicitly multi-threaded approach, task-based programming offers a higher level of abstraction [Mey14]. Logically parallel tasks do not have to be executed in parallel. Dependent on hardware resources it might be useful to run them sequentially to avoid context switches. When there is only one CU, doing one thing after the other is faster than continuously changing between multiple task.

In MIMD, a further differentiation can be made by taking the memory into account. There is *uniform memory access* (UMA) and *non-uniform memory access* (NUMA). For the latter, there is local memory that can be accessed faster but might not always be synchronous. For the first, all CUs share the same memory³. Examples are SMP systems (UMA) and computing clusters (typically NUMA).

7.1.3. Single Instruction, Multiple Data

At a high level, SIMD can be seen like operating with vectors. For example, multiplying a vector of size n by r ,

$$\vec{b} = r \times \vec{a}$$

means

$$b_i = r \times a_i \quad \forall i \in \mathbb{N}^*, i \leq n.$$

Here, the “ \times ” is the *single instruction*, while the a_i can all be loaded at once, using *multiple data streams*. But this is not where SIMD ends. Consider particle tracking: A total of m particles at time t_0 with their speeds $\vec{v}_j(t_0)$ is located at positions $\vec{r}_j(t_0)$. Assuming that $\partial v / \partial t|_{t_0} = 0$, the positions at a later time can be easily computed as

$$\vec{r}_j(t_0 + \Delta t) = \vec{r}_j(t_0) + \Delta t \times \vec{v}_j(t_0).$$

Writing the vectors with their 3D-coordinates component-wise yields

$$r_{i,j}(t_0 + \Delta t) = r_{i,j}(t_0) + \Delta t \times v_{i,j}(t_0) \quad \forall i, j \in \mathbb{N}^*, i \leq 3, j \leq m.$$

Using SIMD, this means just two sequential instructions: the multiplication followed by the addition. Again, using multiple data streams all $r_{i,j}, v_{i,j} - 3m$ pieces of data – can be loaded at once. However, branching should be avoided: If parts of the data take different paths in the code, more than one instruction is needed, so SIMD can no longer be used.

If the particles should interact according to a force $\vec{F} \propto (\vec{r}_j - \vec{r}_k) / |\vec{r}_j - \vec{r}_k|^3$, division by zero could happen. To check for it during the computation would include a branch for every pair of particles. When the force is repulsing, it is better to just rely on this than to check. For an attractive force on the other hand, the algorithm might be changed at a different point: Implementing macro-particles of different “weight” and merging those that

³This has to be seen on the logical level of the computation. Hardware may provide *coherent* local caches.

are very close would avoid branching (at this point) and also gradually reduce the total number of particles.

Notice that also loops include a branch: After every iteration it has to be checked whether to return to the beginning of the loop or to continue. If loops are of the same size globally, this is not a fundamental problem: All pieces of data are still operated on by the same instruction. The reader should also bare in mind, that hardware designed for MIMD is always capable to do SIMD. Therefore, the different instruction streams just have to be fed with the same instruction.

7.2. Example Hardware Implementations

Before OpenCL, the hardware and vendor independent industry standard for parallel computation, is being introduced in the next section, here a brief overview summarizing the range of currently available options is given. It will show the various flavors of parallel computation and the rapid change of products that called for an abstract solution that allows program independently from the individual piece of hardware.

Even for single-CPU systems, MIMD can come in different variants. Next to multi-core architectures there is *simultaneous multithreading* (SMT) [IBM17b]. Intel introduced it (as *hyper threading*) [Intel17] in 2002 as an extension to *out-of-order execution* (for SISD). It interleaves instructions from multiple threads to make (even) more effective use of the existing resources [Gas+12]. As an alternative, AMD introduced *clustered multithreading* (CMT) [AMD05; Anand12], which is more on the multi-core side. Instead of keeping resources busy by efficient scheduling, multiple cores share some ALUs that are (in a statistical mean) not used that often.⁴

Also worth mentioning is the trend to integrate classical CPU and graphic-processing cores into one chip. For this, AMD coined the term *heterogeneous system architecture* (HSA) [HSA16]. It has the advantage that it allows *heterogeneous uniform memory access* (hUMA) also including the graphics processor. For workloads, where synchronization between the central processor and the graphic cards is a bottleneck, this comes as a great advantage. Although being advertised especially by AMD and the HSA Foundation (other members are ARM and Samsung), Intel processors do support a similar feature. The first systems to support hUMA were Sony's PlayStation 4 gaming consoles [Sony13].

Since 1997, desktop CPUs started to support SIMD commands. Intel's *MMX* that allows vector operations (128 bit) for integers was first extended for floating points by AMD (*3DNow!* in 1998) and separately by Intel (*Streaming SIMD Extension*, SSE, 1999). Today, using *Advanced Vector Extension* (AVX) 256-bit vector operations allow up to 8 `bin32` operations per clock cycle [Gas+12]. These extensions are best combined with vector reads. When properly aligned multiple data points can be read as one bigger piece of data. As a consequence it might be faster to fetch four `bin32` in one 128 bit read than three `bin32` which make up to two (128 bit) reads. To restore memory alignment, padding bytes can be used. In the example, it will add unused 32 bits after the 96 bits of the three `bin32`, to again have usable chunks of 128 bits. As a result, cubic interpolation can be as fast as quadratic interpolation or (without padding) can even outperform the latter. (See Section 8.1.3 for details on interpolation.)

The theoretical peak performance of a chip can be calculated by its clock, number of parallel execution units, vector length and special instructions. [DLP01] One that is available for all the examples mentioned here, is fused multiply-add (FMA or `fmadd`) [IBM17a; IEEE08]. It

⁴With its new "Zen" architecture, [AMD17d] AMD is now also favoring SMT over CMT.

doubles the theoretical peak performance⁵ which is usually measured in billions of floating point operations per second (GFLOPS). An Intel Core i7-5820 with six cores at 3.3 GHz and AVX, for example, offers 317 GFLOPS at a power consumption of 140 W [Intel14].

The layout of graphics cards is a bit different. To emphasize the importance of pipelining, their hardware is even referred to as *stream processors*. Program counters are managed per SIMD vector (typically 512 bit) so hardware threads (propagating as *wavefronts*) are vectors [Gas+12]. This approach offers a large (theoretical) computing power using smaller and thus more power and cost efficient chips. In the mid-range a Radeon R7 360 using 100 W offers 1536 GFLOPS [AMD17b] – more than twice the performance of the the aforementioned i7 at just 60 % of the electrical power. When it comes to highest performance on a single piece of hardware, both Nvidia (Titan X, 10 974 GFLOPS, 250 W [NV17]) and AMD (Radeon Pro Duo, 16 384 GFLOPS, 350 W [AMD17a]) dramatically increase both computational performance and power consumption. It has to be kept in mind that graphic cards combine many ALUs with just a few CUs. As an extreme example, the Titan X has a total number of just 28 CUs for 3584 ALUs⁶ to run different instructions. This means that branches are really bad: They will in fact disallow parallel execution, which is heavily based on a SIMD approach.

Although not a fundamental problem, for real hardware tight computation loops might come with a significant slow down – even if they have a fixed size. One reason is that of course also the control logic takes its time. Secondly – even if all data takes the same branch, the *CU* needs to be active. As discussed above, this is disadvantageous especially for stream processors. If the number of iterations is static (i.e., at compile time) the loops can be unrolled to increase the ratio of arithmetic to control instructions [DLP01; Gas+12].

7.3. OpenCL

To make use of the various features that can accelerate program execution the (binary) program code has to be fitted to the actual hardware. To get optimal computational performance, this can be done by hand. *OpenCL* [OCL16] offers a more general approach: It is a hardware-independent solution for parallel computation and offers high-level access to hardware-extensions like vector-units by an abstract, C-style language. Using a specific compiler allows for efficient mapping onto the hardware. The result is parallel in a SIMD fashion, but OpenCL also comes with management-functions like command queues that can be used to implement task-based parallelization. In fact, there will always be at least two (logical) threads, one for the *host* and one for the *compute device*. Between those, memory does not have to be in sync, which is called *relaxed consistency model*. It requires explicit synchronization and can be seen as a variant of NUMA, but with the addition that consistency is not even guaranteed when host and compute device are the same hardware.

To make the usage of OpenCL more concrete, consider adding vectors of size N , $\vec{C} = \vec{A} + \vec{B}$. Implemented in *C*, a function to perform this computation in a serial way might read:

```

1 //element-wise addition of A and B, stored in C
2 void add(float* C, float* A, float* B, int N) {
3     for (int i=0; i<N; i++) {
4         C[i] = A[i] + B[i];
5     }
6 }
```

⁵Another advantage is increased precision: “The operation fusedMultiplyAdd(x,y,z) computes $x \times y + z$ as if with unbounded range and precision, rounding only once to the destination format.” [IEEE08, p. 4]

⁶The 3584 “CUDA Cores” [NV17] are arranged in 128-wide vectors.

In OpenCL, the central element are so-called *kernels* that are executed on a device. The unit of concurrent execution is called a *work-item*; each of those work-items will execute the kernel's function body. With OpenCL, the goal is often to represent parallelism programmatically in the finest granularity possible [Gas+12]. The driver (or *OpenCL Platform*) will then create appropriate chunks of work-items based on the used hardware. (It is also possible to specify this manually, but without cross-dependencies the automatic estimation works fine.) Following the above design-goal, the kernel for the vector-addition will just add two elements and store the result. Which elements to take can be found out using the built-in OpenCL-C-function `get_global_id(int dimension)`.

```

1  __kernel void add(__global float* C, __global float* A,
    __global float* B) {
2    int i = get_global_id(0);
3    C[i] = A[i] + B[i];
4  }
```

The kernel has to be compiled for the desired hardware. This will usually be done at runtime. For that, first, the platform and the device have to be set up. After compilation, the arguments have to be copied from the host memory to the memory of the computation device and the parameters have to be set. However, what has to be done in the host program to perform these steps will be omitted here. To understand the working principle, it is enough to know that the kernel is executed using the function

```

1  cl_int cl::CommandQueue::enqueueNDRangeKernel(
2      const Kernel& kernel,
3      const NDRange& offset,
4      const NDRange& global,
5      ... [optional parameters]);
```

where `kernel` names the OpenCL kernel to use, `offset` allows to start from arbitrary `i` (would be set to 0 here), `global` gives the size of the work (would be set to N this example). Work can also be N-dimensional: The type `NDRange` stands for N-dimensional range, which is an N-dimensional array containing the size of the work in N dimensions. For example, a kernel working on a 2D-array of size 256x256 might be called with `global = {256,256}`.

To circumvent branches, it is very useful to know predefined functions. For example, there is a clamp function [cf. Eq. (9.4)] that acts like

```

1  float clamp ( float x, float minval, float maxval) {
2    if (x < minval)
3      return minval;
4    if (x > maxval)
5      return maxval;
6    return x;
7  }
```

but is implemented in hardware – without branching. For automatic loop unrolling, the `#pragma unroll` can be used in front of the loop. Even when the number of iterations is not known at compile time, (partial) automatic unrolling can be performed if the number of iterations is guaranteed to be a multiple of a $N \in \mathbb{N}$. This can be indicated by `#pragma unroll N`.

As stated in Sec. 7.1.3, for all SIMD techniques data-parallel thinking is very important. Once started, it can become quite intuitive and OpenCL offers help in avoidance of branches or unnecessary control structures. For further information the reader is referred to the OpenCL website [OCL16]. It includes a good and up-to-date list of books and tutorials.

7.4. Applicability

The aforementioned Intel Core i7-5960X [Intel14] for example multiplies its computational power gained from the clock by a total factor of 128 by using eight cores (MIMD) and vector units (SIMD). Graphic cards even have up to 8192 parallel ALUs (AMD Radeon Pro Duo, [AMD17a]). As modern processors and compilers offer some level of automated optimization (see Section 7.1.1), comparison to purely sequential computing is not really fair, but it helps to get an impression.

In the framework of this thesis, parallelized computing in general, and SIMD in particular are useful at two different points: The detector used in the EOSD spectrometer (KALYPSO, see Sec. 5.1), can continuously stream 256 pixels of 14 bit data at 2.7 MHz. For each of these pixels, first the background has to be subtracted before the modulation strength can be computed by relating the measured signal for this particular pixel to the one of a reference pulse. These tasks are easily parallelized in a SIMD manner. Secondly, also simulation of beam dynamics can profit from this approach. The example of particle tracking has already been mentioned, but here the numerical solution of the VFPE will be used. When the charge density function $\psi(q, p)$ is discretized on a grid, the needed number of grid points is typically much lower than the number of needed particles needed for a simulation. However, before parallelization was applied, “only a limited number of synchrotron periods [could] be simulated in reasonable computation time, which means that statistics (...) [was] not optimal” [Sch+14]. Besides overcoming this limitation, the almost instant availability of simulation results also opens the door to more exploratory work flows. It is an interesting side-note, that many needed computation steps are also common in image processing – another field where it is very common to parallelize computation.

8. Numerical Solution of the VFPE

“The best programs are the ones written when the programmer is supposed to be working on something else.”

Melinda Varian

This chapter discusses the steps that are performed to simulate the dynamics of the longitudinal phase space. Some choices are fixed in the simulation code *Inovesa*, created in the framework of this thesis [Sch+17b]. Besides the algorithms, there are two kinds of parameters, that influence the result: *numerical* and *physical* ones. For example when an electron beam is accelerated by a voltage $V = \sin(2\pi \text{ 500 MHz} \times t) \times 1 \text{ MV}$ for a time $T = 2 \text{ ns}$ (both physical settings), there has to be a numerical parameter Δt which describes how often the equations is evaluated. Wherever reasonable, both kinds of parameters can be given at runtime (see lists in Appendix D).

8.1. Discretization

The discretization described here assumes that fields can be assumed to form instantaneously. If their speed of propagation has to be considered, this requires a full backlog of the bunch configurations. There exists a multi-threaded code implementing this [Ter+15], however, as the bunch shape does not significantly change in the time scales defined by Eq. (2.19), the approximation is in order. The general algorithm, this work is based on, was derived in plasma physics [CK76], and first applied to particle accelerators by [WE00].

8.1.1. Time Steps

Single particle motion under the Hamiltonian H can be described by a nonlinear map M with

$$r' = M(t_2|t_1)(r), \quad (8.1)$$

meaning that a trajectory passing through r at time t_1 passes through r' at time t_2 . To solve this equation, a naive implementation would use the finite difference method. This approach, however, is very challenging when it comes to the numerical stability. Because of this understanding, [WE00] followed another approach that explicitly exploits conservation of probability for their VFPS. For a trajectory starting at r at time 0 evolving to r' at time t and volume element dr evolving to dr' , conservation of probability requires that

$$\psi(r', t)dr' = \psi(r, 0)dr \quad (8.2)$$

is fulfilled. Additionally, time evolution in Hamiltonian systems is volume-preserving (or *symplectic*), meaning that $dr' = dr$. It follows that

$$\psi(r, t) = \psi(M(0|t)(r), 0). \quad (8.3)$$

Assuming that the collective force keeps constant for small time intervals Δt , it can be considered as an external force [WE00], so that

$$\psi(r, t + \Delta t) \approx \psi(M(t|t + \Delta t)_{\psi(t)}(r), t), \quad (8.4)$$

where $M(t|t + \Delta t)_{\psi(t)}$ is a map that varies as ψ evolves in time. Furthermore,

$$M_{\psi(t)} = K_{\psi(t)} \circ S, \quad (8.5)$$

where the dependency on $\psi(t)$ is in K , and the S is constant. For better readability, the subscript $\psi(t)$ will be omitted in the following. In the solution of the Vlasov Equation [Eq. (3.3)], S is a rotation map describing the synchrotron motion and K describes the kick due to the collective effects [Eq. (3.5)]. The remaining Fokker-Planck terms can be approximated, e.g. using a finite difference method, yielding the map D . Details on the actual implementation of D are discussed in Sec. 9.2.2.

Following the assumptions, the rotation map S has the largest effect. Consequently, the numerical parameter describing the discretization of time is given as the number steps per synchrotron period N_θ . To describe K , the wake potential (see Sec. sec: calculating the wake potential) has to be considered. When $Z(s) \equiv 0 \forall s$ with $B_y(s) \neq B_0$, one can approximate an accelerator by an iso-magnetic ring with bending radius R_0 and scale time and frequencies accordingly¹. This approach, however, comes to problems when straight sections play a role – e.g. because of the resistive wall impedance. In that case, an averaged wake potential is used.

8.1.2. The Grid Coordinates

When solving the equations of motion for systems of a multitude of particles, in principle there are two possibilities to reduce the amount of needed computing time to a manageable amount: The first one is to reduce the number of particles, the second one is to approximate the particle distribution as a continuum ψ which is then expressed at a number of sampling points. Notice that also the other quantity, e.g. $\psi(q, p)$, is (implicitly) discretized by storing it digitally (see Sec. 9.1).

Following the approach of [WE00], ψ it is discretized on a grid. For the boundary conditions, a minimum and a maximum coordinate are defined, where the charge distribution may be assumed to be $\psi(q, p) \approx 0$. For practical reasons the new coordinate system's origin is no longer $q = p = 0$ but $q = q_{\min}$, $p = p_{\min}$. With Δp , the granularity of the grid in energy direction, the transformation between the coordinate systems is defined to be

$$y_r := (p - p_{\min})/\Delta p, \quad (8.6)$$

and accordingly for x_r . So the generalized coordinates $q, p \in \mathbb{R}$ become the grid coordinates $x_r, y_r \in \mathbb{R}_{\geq 0}$. The grid, however, only stores function values $\psi(x_r, y_r)$ for the integer parts of the grid coordinates. To emphasize that fact

$$x := \lfloor x_r \rfloor, \quad y := \lfloor y_r \rfloor \quad (x, y \in \mathbb{N}) \quad (8.7)$$

will be used, where a point is guaranteed to be covered by the grid. For complete mapping and especially for interpolation also the fractional part is needed. It will be denoted $\{x\} := x_r - x$, and accordingly for $\{y\}$.

¹In *Inovesa* the parameter R can be set to $R < 0$ to explicitly state that an iso-magnetic ring is simulated.

-1,1	4,0	4	-0.5,1	4,1	9	0,1	4,2	14	0.5,1	4,3	19	1,1	4,4	24
-1,0.5	3,0	3	-0.5,0.5	3,1	8	0,0.5	3,2	13	0.5,0.5	3,3	18	1,0.5	3,4	23
-1,0	2,0	2	-0.5,0	2,1	7	0,0	2,2	12	0.5,0	2,3	17	1,0	2,4	22
-1,-0.5	1,0	1	-0.5,-0.5	1,1	6	0,-0.5	1,2	11	0.5,-0.5	1,3	16	1,-0.5	1,4	21
-1,-1	0,0	0	-0.5,-1	0,1	5	0,-1	0,2	10	0.5,-1	0,3	15	1,-1	0,4	20

Figure 8.1.: An example grid of $N_x = N_y = 5$ cells in each dimension spanning from $(q, p) = (-1, -1)$ to $(q, p) = (1, 1)$. The numbers inside the cell give the position in the different coordinate systems. The first pair is q, p , the second pair is x, y , the last number is n .

In total, the charge density in phase space is sampled at $N_x \times N_y$ points, where N_x, N_y is the number of grid cells in direction of location and energy, respectively. To speed up computation, it is valuable to also define

$$n := N_y \times x + y. \quad (8.8)$$

Doing so allows to rewrite maps $M : \psi_t(\mathbb{N}^2) \rightarrow \psi_{t+\Delta t}(\mathbb{N}^2)$ in a one dimensional form $M_{SM} : \psi_t(\mathbb{N}^1) \rightarrow \psi_{t+\Delta t}(\mathbb{N}^1)$, which depends only on n (see Fig. 8.1 for an example and Sec. 8.2 for details).

When choosing a grid size (N_x and N_y), of course the size of the structures that have to be simulated, plays a big role. Furthermore, to solve the VFPE, it has been assumed that $\varrho(t) \approx \varrho(t + \Delta t)$. To guarantee this, one can choose a grid size so that the maximum displacement during one computation step is less than the distance of two grid points. Assuming $(q, p) = (0, 0)$ to be centered in the grid, the relations

$$N_{x,y} < \frac{T_s}{\pi \Delta t} = \frac{N_\theta}{\pi} \quad (8.9)$$

have to be fulfilled². As at the edges of the grid $\psi(q, p) \approx 0$, this equation does not have to be strictly fulfilled (see Sec. 9.3.1 for benchmarks).

8.1.3. Interpolation

As the algorithm to solve the VFPE assumes small changes – often smaller than one grid cell – interpolation is obviously needed. In fact, a straightforward rotation causes artifacts even for big changes (an intuitive example is at [Tan08]). For interpolation we take into account the surrounding data points of the point (x_r, y_r) . In one dimension, interpolation can be expressed as

$$\psi(x_r) \approx \psi_{\text{pi}}(x_r) = \vec{\psi}(x)^T \vec{p}(\{x\}), \quad (8.10)$$

where $\vec{\psi}(x)$ is a vector containing the particle density at the grid points surrounding the coordinate of interest, and $\vec{p}(\{x\}) = (p_{0,n}, p_{1,n}, \dots, p_{n,n})^T$ is a vector containing $n + 1$ interpolation coefficients, each a function of $\{x\}$.

In two dimension Eq. (8.10) becomes

$$\psi(x_r, y_r) \approx \vec{p}^T(x_r) I_\psi(x, y) \vec{p}(y_r), \quad (8.11)$$

where

$$I_\psi(x, y) := \begin{pmatrix} \ddots & & & & & \ddots \\ \cdots & \psi(x-1, y-1) & \psi(x-1, y) & \psi(x-1, y+1) & \cdots & \\ \cdots & \psi(x, y-1) & \psi(x, y) & \psi(x, y+1) & \cdots & \\ \cdots & \psi(x+1, y-1) & \psi(x+1, y) & \psi(x+1, y+1) & \cdots & \\ \ddots & & & & & \ddots \end{pmatrix}, \quad (8.12)$$

²For the outermost grid points, it is $1 > (\Delta x)_{\text{max}} = \tan\left(\frac{2\pi\Delta t}{T_s}\right) \frac{N_y}{2}$.

is a matrix containing the particle density at the grid points surrounding the coordinate of interest.

For the interpolation coefficients $p_{\nu,n}(\{m\})$, the Lagrange basis polynomials [Pre+07]

$$l_{\nu,n}(\{m\}) := \prod_{k=0 \neq \nu}^n \frac{\{m\} - m_k}{m_\nu - m_k} \quad (8.13)$$

have been chosen. The interpolation method is then called spline interpolation. For $n = 1$ the polynomial forms a linear interpolation. For $n \geq 2$, not only the next neighbors are taken into account. In 2D this method then interpolates using a total number of 9 or 16 points. This is then called biquadratic ($n = 2$) respectively bicubic ($n = 3$) interpolation.

Important properties are that $|\vec{p}| = 1$ and that the curves directly pass the values at the grid points³. Lagrange basis polynomials are defined so that $l_i(x_j) = \delta_{ij}$, which exactly is the desired property. If it was different, even without a real change ($\Delta q = \Delta p = 0$) the interpolation would change the grid content. So Lagrange polynomials are used for Eq. (8.10) for interpolation in *Inovesa*.

Another, option would be sinc interpolation which also fulfills the above mentioned requirement but also guarantees the cardinal theorem of interpolation [Kot33]⁴ is adhered to. This algorithm however, includes algorithms of at least $\mathcal{O}(N \log N)$, whereas local spline interpolation is only $\mathcal{O}(N)$. The increased performance of course comes at a cost: As the Lagrange basis polynomials $l_{\nu,n}$ can be negative, interpolating may lead to non-conservation of monotony or – in the picture of the Vlasov solver – even transform electrons to positrons. Details of this problem, called overshooting, and possible solutions are explained in Sec. 9.2.1.

8.1.4. Integration

Integration is especially needed to calculate the wake potential [see Eq. (3.9)]: The bunch profile is obtained by the projection of the charge density in phase to the position axis $\varrho(x) = \int \psi(x, y_r) dy_r$.

If the function $f(x)$ is a polynomial, the Newton-Cotes formula [Eq. (8.14a)] gives the exact integral [Pre+07, p. 179]. This is not the case here, but the assumption that the charge distribution may be locally approximated by Lagrange basis polynomials $l_{i,n}$, gives a natural choice for the weights ω_i [Eq. (8.13)]

$$\int_a^b f(x) dx \approx \sum_{i=0}^n f(x_i) \times \omega_i \quad (8.14a)$$

$$= \sum_{i=0}^n f(x_i) \times \int_a^b l_{i,n} dx, \quad (8.14b)$$

that is consistent with the chosen interpolation scheme. Applied to N equally sized sub-intervals, and choosing $n = 3$ yields the composite Simpson's rule [Pre+07, Eq. 4.1.13]

$$\int_a^b f(x) dx = h \left[f(x_0) + \frac{2}{3} \sum_{k=1}^{N/2-1} f(x_{2k}) + \frac{4}{3} \sum_{k=1}^{N/2} f(x_{2k-1}) + f(x_N) \right] + \mathcal{O}\left(\frac{1}{N^4}\right), \quad (8.15)$$

with $h = (b - a)/N$ and $x_k = a + k \times h$. It is noteworthy that choosing $n = 4$ does not add to the accuracy and that the alternating occurrence of $2/3$ and $4/3$ has no deeper meaning but is an artifact of the building blocks using spline interpolation [Pre+07, p. 160].

³In contrast Bézier curves, another interpolation method not used here, do not reproduce the values at all grid points.

⁴It is often referred to as the *Nyquist-Shannon sampling theorem*, but according to [Lük17] the first theoretically exact formulation is found here.

8.2. Source-Mapping

As discussed in Sec. 8.1.1, the map for each time step $M : \psi(q, p, t) \rightarrow \psi(q, p, t + \Delta t)$ is composed of a number maps for “rotation” (S), “kick” (K), “damping and diffusion” (D)

$$M = D \circ K \circ S.$$

that are solved in individual simulation steps. The rotation comes from the the solution of the two-dimensional harmonic oscillator part [Eq. (2.40)] of the VFPE [Eq. (3.4)]. The kick models the effect of the wake treated as a perturbation to the harmonic oscillator (cf. Chap. 3). The last step implements the RHS of Eq. (3.4). To reduce the number of needed interpolations, implementations often combine rotation and kick into one step. However, since every rotation step always does the same, it is worth putting some initial effort into speeding up the rotation steps. As will be shown below, the gain of the optimization of the rotation overweighs the additional step – leading to a performance gain.

To do so, the direct implementation of each of these simulation steps is split into two sub-steps. The information on the actual coordinates (q, p) is used by the first half simulation step. So the original map

$$M : \psi(q, p, t) \rightarrow \psi(q, p, t + \Delta t), \quad q, p \in \mathbb{R} \quad (8.16a)$$

becomes

$$M_{\text{SM2D}} : \psi(x, y, t) \rightarrow \psi(x, y, t + \Delta t), \quad x, y \in \mathbb{N} \quad (8.16b)$$

which will be called a ‘source map’ (SM). The name reflects the fact that the implementation uses backward mapping⁵ to first evaluate which $(x_r(t), y_r(t))$ evolves to $(x(t + \Delta t), y(t + \Delta t))$ and then assembles all x_i, y_i that are required for the interpolation to approximate x_r, y_r . This approach makes sure simplicity is conserved even when only grid coordinates x, y are used. In contrast to previous implementations, the SM avoids re-computation by storing all (constant) coefficients needed for the implementation. Further, using the definition in Eq. (8.8), the original map can be rewritten in a one-dimensional form

$$M_{\text{SM1D}} : \psi(n, t) \rightarrow \psi(n, t + \Delta t), \quad n \in \mathbb{N}. \quad (8.16c)$$

This one-dimensional map – by construction – produces the same results as the single-step implementation. Practically speaking, the source map expresses the information which data of the current simulation step contributes to a grid point for the next simulation step directly in terms of position in the computer’s memory. For many functions – such as rotation – the SM looks the same for the whole runtime of the program. For that reason, it only has to be computed once during a simulation run and can be kept for multiple usages.

In total, there are three implementation details that improve the computational performance:

- Caching of intermediate results (using SM),
- the reduction of the problem’s dimension (SM1D), and
- parallelization using OpenCL.

The implementation of the latter can be based on a previously created SM for the problem: The SM formalism does not only allow to keep intermediate results, it also gives a handy interface to implement arbitrary manipulators on the phase space. OpenCL has been chosen because it can utilize not only multi-core CPUs but also graphic processors. OpenCL advocates a data parallel approach that processes multiple data with one single instruction.

⁵Source can also be read as an abbreviation for ‘speicherorientiert und reduzierte Computerlast ermöglicht’ – German for ‘memory-oriented and allowing reduced computational load’.

Table 8.1.: Runtime of different implementations of the rotation on an Intel Core i5-4288U with integrated GPU (used for OpenCL). “Rotation” stands for a direct implementation of the rotation, SM2D implements a source map that still has two dimensions, SM1D marks a source map that uses only one-dimensional indexes [see Eq. (8.8)]. For this test case a `binary32` grid of 512px squared has been chosen. The time it takes to create the SM with/without C++ compiler optimization (-O3/-O0) is in the order of 0.1s/1s. (OpenCL code is always being optimized by the OpenCL compiler.)

	-O0	-O3
Rotation	5m52s	33s
SM2D (seq.)	3m50s	32s
SM1D (seq.)	3m00s	25s
SM1D (par.)	0m03s	03s

So in contrast to other forms of parallelization, there is no need for parallel tasks (see Chapter 7 for more details).

The results of a benchmark of the computational performance for the particular case of rotation using cubic interpolation is shown in Table 8.1. It can be seen, that SM brings a significant speedup for non compiler optimized binary code, but since it only does simple tasks, it does not profit as much from compiler optimization as traditional rotation does. In fact, a decent combination of compiler and processor does a lot of automatic optimizations. However even in compiler optimized sequential operation, SM clearly outperforms traditional rotation. Notice that parallelization harmonizes well with the used quadratic interpolation: All four coefficients can be read in one 128 bit vector-read, multiplication can be performed in four-wide vector units. In fact, a straightforward implementation of quadratic interpolation would be slower. To achieve competitive speed, padding bits would have to be introduced, to regain memory-alignment.

Results of a benchmark of the computational performance including the full three-step map M are shown in Fig. 8.2. To allow vector-reads when applying the maps K and S , the definition of n [Eq. (8.8)] is chosen in a way that n and $n + 1$ in most cases are adjective in energy direction. In total, a non-optimized program takes days for a typical simulation run. Using the method described above, *Inovesa* can reduce this to 20 minutes when running on a customer grade graphics card. Comparing this to the 24 minutes the i7 CPU takes, shows that there is a lot more potential for optimization: Theoretically, the R9-290 has 4850 GFlops [AMD17c] and should outperform the i7-5820 (317 GFlops) at least one order of magnitude. Measuring the timings of the individual maps⁶ shows that indeed the calculation is much faster on the GPU (e.g. 1 μ s vs. 62 μ s for K). The typical latency between the time the calculation is scheduled and the actual calculation on the other hand can be in the order of a millisecond for the GPU (0.1 ms for the CPU). In another example (Late 2013 MacBook Pro featuring an Intel Core i5-4288U running on MacOS Sierra), verifies these results: Here, when the CPU is used, 50% of the between enqueueing the command to the computation pipeline and the end of the calculation is spent for computation, but on the integrated GPU, this number drops to 10%. Notice that the inverse of this time does not give the frequency commands are processed: Commands do overlap in the pipeline. To completely evaluate this issue, full profiling would have to be applied. However the timings combined with the total computational time already hint that the pipeline is a lot less populated on the GPU (see Chap. 7).

⁶The actual values are very dependent on the driver.

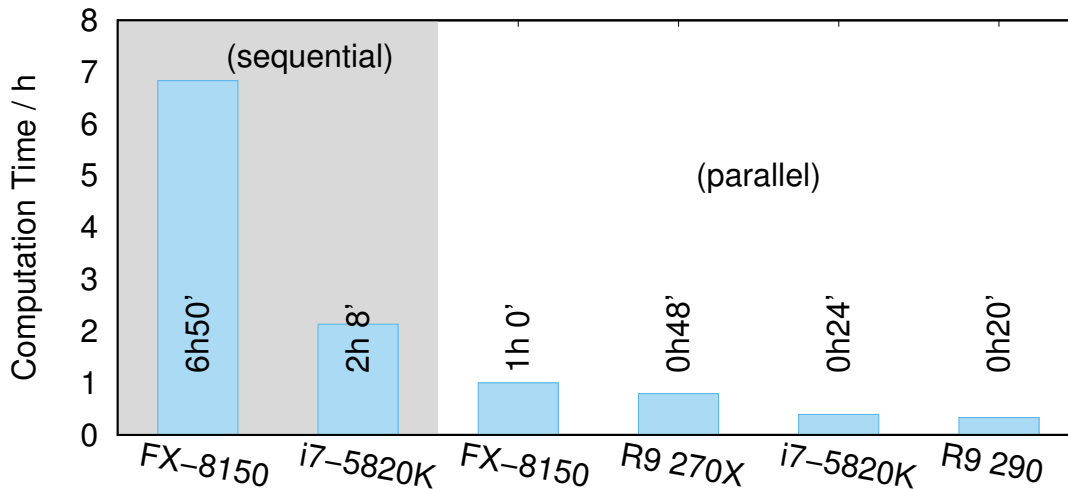


Figure 8.2.: Computational Performance Using Source-Mapping

Computational time needed by different hardware to perform one simulation run of 500 synchrotron periods, each consisting off 1000 steps, using a grid of 512 points per axis. When running on the FX processor, a non-optimized program would take more that two days.

8.3. Manhattan-Rotation

While rotation mathematically is a symplectic operation, the direct implementation (as in [WE00], ‘standard rotation’) easily loses this property when an infinitesimal rotation is performed. For example, the needed two-dimensional interpolation which involves $N \times N$ data points leads to interpolation coefficients $l_{\nu,N}^2 \propto \{x\}^{N-1} \times \{y\}^{N-1} \ll 1$. So the algorithm is very vulnerable to numerical absorption. As an alternative, a symplectic integrator [Wol14a] can be used that is area preserving also for truncated power series. The symplectic map R_S splits the rotation into a energy-dependent drift R_D followed by a location dependent RF kick R_K

$$R_S(x, y) = R_K(x) \circ R_D(y). \quad (8.17)$$

Each requires a one-dimensional interpolation, involving N data points. The resulting coefficients are $l_{\nu,N}^1 \propto \{m\}^{N-1} \gg l_{\nu,N}^2$. Since a point is rotated by moving on straight lines in perpendicular directions, the term ‘Manhattan rotation’ was coined for this approach [Sch+17c].

Splitting the rotation into two maps also allows to model additional physical properties. Recall (see Sec. 2.4) that the rotation is the solution to the linearized Hamiltonian with $\eta_c \approx \alpha_{c,0}$ and $V = (\hat{V}^2 - U_0^2)^{1/2}$. Now, also higher orders of α_c [see Eq. (2.28)] as well as the shape of the RF bucket can be considered (cf. [Sch+18a]): The first are part of R_D , the latter of R_K . As translations are always symplectic, the implementation can be directly based on Eq. (8.1), for example

$$\psi(x_r, y, t) \rightarrow \psi(x + \eta'_c(y), y, t + \Delta t), \quad (8.18a)$$

where $\eta'_c(y)$ marks the full polynomial slip factor scaled to the grid coordinates. As

$$x + \eta'_c(y) \stackrel{!}{\in} \mathbb{N},$$

the source point is $x_r \in \mathbb{R}$ and has to be evaluated by interpolation. The complete map can be expressed as

$$R_D : \sum_{i=0}^{N_x-1} \omega_i(\eta'_c(y)) \psi(x_i, y, t) \rightarrow \psi(x, y, t + \Delta t), \quad (8.18b)$$

where the the interpolation coefficients $\omega_i(\eta'_c(y))$ must satisfy

$$\sum_{i=0}^{N_x-1} \omega_i(\eta'_c(y)) = 1. \quad (8.19)$$

8.4. Calculating the Wake Potential

To calculate the wake potential, needed for the kick map K , most VFPS codes [WE00; Ven+05; Kle12; Kus12; Rou14] perform convolution in time domain [see Eq. (3.9a)]. Of course, expressions for the impedance were derived based on the wake function, and there are still cases where only a wake function is analytically known. In such cases it might be favorable to work with the wake function. However, the computation of the convolution “requires special care since it [the wake function] is sharply peaked in a region compared to typical mesh cells in the Vlasov integration.” [Ven+05] This gives a first strong reason to work in frequency domain [see Eq. (3.9b)]: The impedances presented in Section 3.2 are all continuous and well defined for the whole range of interest. Furthermore, as $\varrho(q) \stackrel{!}{=} 0$ at the borders of the used grid, working in frequency domain allows for zero-padding the bunch profile. This increases the number of sampling points for the impedance and – for the bunch profile – gives a computational cheap option equivalent to sinc interpolation of the bunch spectrum. So the numerical accuracy of this step is also increased.

When it comes to the computational performance, fast Fourier transforms behave like $\mathcal{O}(n \log n)$, while straight-forward convolution is a more computational expensive [$\mathcal{O}(n^2)$] algorithm. Further notice that common implementations of discrete real-to-complex Fourier transforms like [FFTW16; cFFT16] only compute $\tilde{f}(k < k_{max}/2)$ while $\tilde{f}(k < 0)$ is ignored. The real result would be Hermitian, meaning $\tilde{f}(-k) = -\tilde{f}^*(k)$, but for the (Hermitian-to-real) backward Fourier transform, this optimization is considered. This means the impedance for negative frequencies does not have to be explicitly set and – more important – speeds up the computation by a factor of two. Also, libraries like the ones mentioned above already implement effective SIMD algorithms. In total, multiplication followed by a Fourier transform very optimized. In fact, a fast implementation of the convolution approach would also be implemented as a multiplication in Fourier space [Pre+07].

To obtain the kick map K from the wake potential, it has to be scaled after computation – especially, when using a normalized charge with $\int \varrho dq = 1$. To obtain the desired offset in grid cells Δy for one time step Δt , the fraction of one revolution $s/C = f_0 \Delta t$ and the bunch charge $Q_b = I_b/f_0$ are needed. The factor $y(1 \text{ eV})$ transforms eV to grid cells, and $dk/d\tilde{q} = c/z(q=1)$ takes into account that the Fourier transform is done in normalized space. The combined scaling factor is

$$S_{E \rightarrow y} = \frac{s}{C} \times Q_b \times \frac{1}{y(1 \text{ eV})} \times \frac{c}{z(q=1)} = \frac{I_b \Delta t c}{y(1 \text{ eV}) z(q=1)}. \quad (8.20)$$

From $\Delta y(x) = S_{E \rightarrow y} \times \Delta E(x)$, it can be continued analogue to Eq. (8.18).

8.5. Single-particle dynamics⁷

In the previous sections, the maps are seen in the spirit of Eq. (8.3) because this picture offers the combination of computational efficiency and numerical stability. Performing the calculations on the particle density $\psi(x, y)$, however, comes with the drawback that it does not allow to draw conclusions on the single particle dynamics. But there is a cost-efficient solution: Recall that the original definition of a map – Eq. (8.1) – focused on exactly this:

⁷A more detailed description, also considering stochastic emission, can be found in [Sch+18a].

A map describes the displacement of a particle between two points in time. So, when the map is known, it is possible to also use it for particle tracking. As the maps are calculated based on the charge density, it is even possible to use just a small number of representative particles – without any loss of accuracy. In this context, it is useful to express r in terms of x_r and y_r to be able to directly use the same source map as for the calculation of the phase space [see Eq. (8.16b)].

9. Influences of Numerical Effects

“Einstein argued that there must be simplified explanations of nature, because God is not capricious or arbitrary. No such faith comforts the software engineer.”

Frederick P. Brooks

This chapter starts with a discussion of the relevance of the data type chosen for the simulation and how it affects the accuracy of the result. But numerical inaccuracies cannot only be caused by data types but also by the used algorithms. Section 9.2 shows how artifacts can emerge from discretization, interpolation and other algorithms. First, the principles are explained using extreme examples, later, their relevance is proved in simulations of physical systems.

In Sec. 9.3, the numerical accuracy is benchmarked, testing how well it reproduces analytically known properties. The findings of these tests also have implications how starting conditions for a simulation should be chosen if no static solution is available (Sec. 9.4). At last, Sec. 9.5 gives an idea of how numerical effects can influence the found instability threshold.

9.1. Explicit and Implicit Discretization

9.1.1. General Considerations

When doing numerical calculations there is always an implicit discretization when rational numbers are expressed by digital data types. Section 8.1 described the (explicit) discretization of space in a mathematical way, neglecting that function values of $\psi(x, y)$ at the sampling points are also (implicitly) discretized. When integer data types are used, this implicit discretization is obvious: There are 2^N possible numbers that can be represented, where N is the number of bits the integer number is composed of. For 32 bits, this creates a linear scale up to $2^{32} \approx 4 \times 10^9$.

Arguing from this point of view, the approach of particle tracking can be seen to be very similar: While for VFPS, the particle density at discrete points is stored, particle tracking has to save the spacial position of every particle – which results in an implicit discretization of space. In total, both representations have N_P macro-particles (arguing with a fixed point representation) jumping on a grid of size N_G . The difference is, that $N_P \gg N_G$ for VFPS, and $N_G \gg N_P$ for particle tracking¹.

¹ Based on this understanding, *Inovesa* uses the same maps that were implemented for charge density tracking in the VFPS to track individual particles. As now N_G, N_P are both considerably small, the results are not considered in the simulation but meant for visualization purposes only.

Usually the explicit discretization is much harsher than the implicit one, thus only this one is considered as an approximation. However, being aware of implicit discretization, allows to influence and make use of this step. Also be aware of the fact that explicit and implicit discretization are independently chosen – thus the error contributions are not correlated. So the overall error will be most likely dominated by the explicit discretization and will not be significantly reduced by increasing the precision. Both discretization types will be analyzed in the following.

9.1.2. Data Types

For the implicit discretization, the choice of the data type plays a crucial role. The size of the data type is important as it gives the number of values that can be expressed.

Additionally, for floating points, the implementation becomes important. In modern programming languages, the IEEE 754 [IEEE08] representation

$$w_{\text{flt}} = (-1)^s \times b^q \times c \quad \text{with } s, b, q, c \in \mathbb{N}, \quad (9.1)$$

is commonly used. The `binary32` type (often referred to as single precision) has $s \in \{0, 1\}$, $b = 2$, $-149 \leq q \leq +104$, and $0 \leq c \leq b^{24} = 16777216$. This type allows a total of about 2.5×10^9 numbers between 0 and 1 to be represented: There are 2^{24} possible coefficients to $2^{-149} < b^q < 1$. So these numbers are not equally distributed, but there are exponential distributed clusters, each consisting of linearly distributed points.

In contrast, fixed point numbers are equally spaced. They can be described as [FPML09]

$$w_{\text{fixp}} = \sum_{i=0}^{N-2} b_i \times 2^{i-f} - b_{N-1} \times 2^{N-1-f}, \quad (9.2)$$

where $b_i \in \{0, 1\}$, f is the bit-count of the fractional part and N is the total width of the data type (e.g. 32). For example, with $N = 8$ and $f = 2$, the range of this type is from -32 to $+31.75$, and the fractional part can encode multiples of $2^{-2} = 0.25$. There are three advantages of using fixed point numbers. First, for almost any² w_i there is a w_j , so that $w_i + w_j = 1$. This property of course helps to keep normalization. In principle, fixed point arithmetics in combination with stochastic rounding (see Appendix B.3) could completely avoid rounding errors. Second, the implementation is flexible, so f can be freely chosen to allow more accuracy. For example $N = 8$ will also allow to have a range from -2 to $(2 - 2^{-6})$. This flexibility can even be used to temporarily increase the accuracy for critical steps with low computational overhead. Third, fixed point math typically is computationally cheaper than floating point maths.

9.1.3. Simple Example

The Gaussian distribution will serve as an example for explicit and implicit discretization. For a grid, (explicit) sampling points have been chosen to be at multiples of $\Delta_{\text{exp}} = 0.1$. The (usually) implicit sampling is set to 101 values \hat{y} in the range $[0 : 1]$, following

$$(\hat{y})_{\text{lin}} = 0.01 \times n, \quad \text{or} \quad (9.3a)$$

$$(\hat{y})_{\text{exp}} = \exp(-0.01 \times n), \quad (9.3b)$$

where $n \in \mathbb{N}$, $n \leq 100$ in both cases. This configuration leads to $\Delta_{\text{exp}} > \Delta_{\text{imp}}$. Usually one would have $\Delta_{\text{exp}} \gg \Delta_{\text{imp}}$. To get an idea of the differences between expression of values in

²The only exception is $b_i = 0 \forall i \neq (N - 1)$, $b_{N-1} = 1$.




Figure 9.1.: Gaussian distribution sampled at multiples of 0.1

To give a feeling for implicit discretization, original values (black) are rounded down to the next of 101 values, which are exponentially (red) respectively linearly (blue) distributed [see Eq. (9.3)]. The points (“+”) are connected by lines obtained by local cubic spline interpolation. On the left hand side, the actual curves are displayed. The right hand side shows the mismatch of the sampled data to the analytic curve. The linearly distributed values are advantageous near the maximum and have a smaller maximum absolute error, whilst the exponentially distributed ones can still model the small changes at $y \approx 0$. In total, the maximum relative error is smaller for the exponentially distributed values.

fixed point and in floating point, every function value y is rounded down³ to the nearest \hat{y} . The result is shown in Fig. 9.1. It also features local cubic interpolation between the sampling points. Even though this simplified example ignores the variable coefficients, which floating point data types usually have, it still gives an idea of the advantages and disadvantages of the different kinds of data types: Fixed point values have a limited absolute mismatch, so they show less relative mismatches for higher values. When the values become small, on the other hand, of course accuracy is lost.

For both cases, linearly and exponentially distributed values, the relative mismatch $\Delta y(x) = [y(x) - \hat{y}(x)]/y(x)$ grows to 100% as the Gaussian distribution vanishes. As a rule of thumb, for values distributed over several magnitudes, expression as fixed point minimizes absolute mismatches, floating points minimize relative mismatches.

9.2. Numerical Artifacts

9.2.1. Interpolation

Interpolation using negative weights [cf. Eq. (8.13)] has possible overshooting, which results in “interpolated” values larger (or smaller) than all neighboring ones (see Fig. 9.2). While this problem does not exist for sufficiently smooth surfaces, especially when it is desired to simulate small structures the impact should be checked. The problem is because the combination of small changes and local polynomial interpolation do not consider the cardinal theorem of interpolation (cf. comment in Sec. 8.1.3).

As an example a test pattern consisting of four different quadrants is used: One is kept empty, one is at the maximum value but contains a diagonal spoke at background level, the third sector decreases linearly with one of the dimensions and the last one decreases

³This is the default for integer numbers and fixed point number implementations implemented using them. IEEE 754 floating point numbers have a rounding-direction attribute. The different options are presented in Appendix B.3.

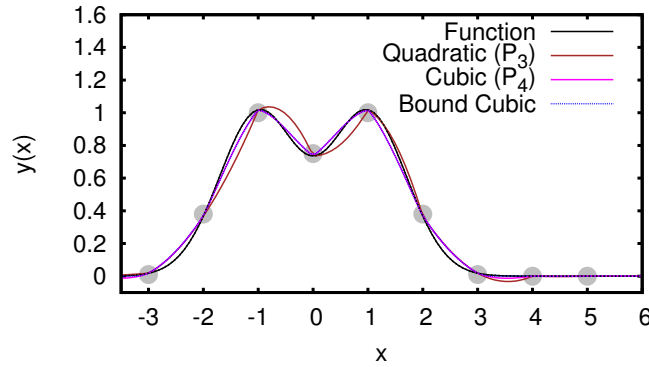


Figure 9.2.: An example function (black curve) has been sampled at the points marked in gray. The colored curves give the result of the approximation using different interpolation schemes. In the region $3 < x < 4$ quadratic and cubic interpolation show an overshoot: the interpolation result exceeds the value of the neighboring data points. This can be cured using clamped interpolation as described Eq. (9.4).

in both dimensions when moving off the center. When this pattern undergoes a series of infinitesimal rotations, the aforementioned artifacts evolve in a wave-like structure, as can be seen in Fig. 9.3. A possible countermeasure could be a finer grid (more grid points to describe a structure) combined with (a limited amount of) intentional blurring to wash out sharp edges. The blurring can be seen as a cheap low pass filter and helps to avoid the overshoots that also higher order polynomials show. In total, the minimum resolvable structure would still be the same but the cardinal theorem of interpolation would be adhered to. However, even though still relying on the computationally cheap $\mathcal{O}(n)$ local spline interpolation, due to the increased number of grid points in all simulation steps, in total this approach is time-consuming $\mathcal{O}(n^2)$. But there is an alternative to sinc interpolation $\mathcal{O}(n \log n)$ that can also be $\mathcal{O}(n)$: An interpolation function that preserves monotony⁴. One solution of this type uses clamped interpolation. This simple technique is used e.g. in real-time computer graphics and is for example part of OpenGL since its initial release [OGL17] in 1991. Applying it, Eq. (8.10) becomes [Sch+17c]

$$\psi_{\text{clamp}}(x_r) = \begin{cases} \max[\psi(x), \psi(x+1)] & \text{if } \psi_{\text{pi}}(x_r) > \max[\psi(x), \psi(x+1)], \\ \min[\psi(x), \psi(x+1)] & \text{if } \psi_{\text{pi}}(x_r) < \min[\psi(x), \psi(x+1)], \\ \psi_{\text{pi}}(x_r) & \text{else.} \end{cases} \quad (9.4)$$

This approach does not cure the principle problem of local spline interpolation but it addresses the problem of overshoots at a relatively low cost.

To summarize, interpolation can introduce artifacts that form unphysical substructures⁵. Possible counter-measures include using computationally more expensive algorithms. However, using clamping allows to suppress the shown type of artifacts still using an $\mathcal{O}(n)$ algorithm.

⁴It may of course not act globally because non-monotonic substructures (the micro-bunches) shall be simulated.

⁵The relevance of these self-amplifying and propagating structures for realistic physical cases will be shown in Sec. 9.2.3.

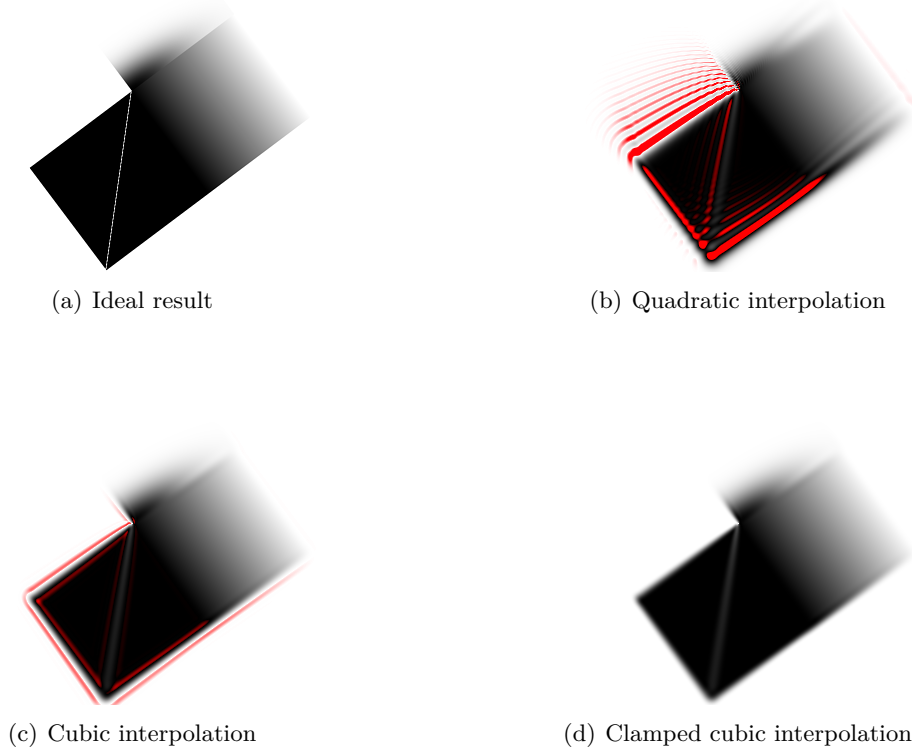


Figure 9.3.: Comparison of different interpolation types applied while rotating a test pattern by 127° in steps of $360^\circ/4000 = 0.09^\circ$ (using the standard rotation algorithm). Ideal result for the rotated test pattern (a), result obtained using quadratic interpolation (b), result obtained using cubic interpolation (c), and result obtained using clamped cubic interpolation (d). Areas displayed in red are above or below the limits of the initial test patterns – these values are clear interpolation artifacts.

9.2.2. Numerical Differentiation

To solve the Vlasov-Fokker-Planck-Equation's damping and diffusion part [map D , right-hand side of Eq. (3.4)], the first and second derivative of ψ are needed. The first numerical derivative often is carried out as a difference quotient [Pre+07], either single-sided

$$\left. \frac{\partial f(x)}{\partial x} \right|_{x=x_0} \approx \frac{f(x_0 + \Delta x) - f(x_0)}{\Delta x}, \quad (9.5a)$$

or two-sided

$$\left. \frac{\partial f(x)}{\partial x} \right|_{x=x_0} \approx \frac{f(x_0 + \Delta x) - f(x_0 - \Delta x)}{2\Delta x}. \quad (9.5b)$$

As Eq. (9.5b) is the average of Eq. (9.5a) evaluated in both directions, it is intuitive that the latter method yields the more stable results. However, even with the more advanced method, significant numerical artifacts come up, when it is applied a huge number of times (see Fig. 9.4).

To solve this problem, one idea is to first use spline interpolation for a larger number of points and then use the (known) analytical derivative at the position of interest. Starting with the Lagrange polynomials [see Eqs. (8.10) and (8.13)] we find two well known expressions




Figure 9.4.: Damping of a test pattern carried out in multiple time steps

All pictures have the same scale, white corresponds to 0, black to 1. The total integral (“amount of dark”) should be is conserved. Between two images, damping using numerical differentiation as shown in Eq. (9.5b) has been applied a 1000 times. On the right hand side of each picture, the originally white areas have been colored red, to make numerical artifacts more visible. They can be seen as a periodic white pattern emerging from the gray area in the center.

[identical to Eq. (9.5)]

$$\left. \frac{\partial P_1(x)}{\partial x} \right|_{x=x_0} = \frac{f(x_0 + \Delta x) - f(x_0)}{\Delta x}, \quad (9.6a)$$

$$\left. \frac{\partial P_2(x)}{\partial x} \right|_{x=x_0} = \frac{f(x_0 + \Delta x) - f(x_0 - \Delta x)}{2\Delta x}, \quad (9.6b)$$

as well as higher orders

$$\left. \frac{\partial P_3(x)}{\partial x} \right|_{x=x_0} = \frac{-2f(x_0 - \Delta x) - 3f(x_0) + 6f(x_0 + \Delta x) - f(x_0 + 2\Delta x)}{6\Delta x}. \quad (9.6c)$$

Notice that this differentiation scheme uses four points and thus is also practical for SIMD parallelization. Furthermore, there is no additional error added, as the method is equivalent to local cubic spline interpolation which is applied anyways.

For the cases studied within the framework of this thesis, the approach discussed here was sufficient – also when it comes to the avoidance of artificial structures. Examples that have lower energy and thus slower damping (c.f. Sec 9.3.1), however, show bad convergence or – when starting with theoretical prediction – even slow (numeric) drifts. An alternative scheme to solve the Fokker-Planck term – using operator splitting – can be found in [ZMM99]. It features unconditional numerical stability, so it does not require the change between time steps to be within certain limits. Using a such a scheme for D should allow to work with any damping time.

9.2.3. Example Studies featuring Physical Cases

The formation of substructures due to numerical artifacts voids any possible knowledge that can be gained from the simulation. In principle, those artifacts can be circumvented by the choice of a different combination of numerical parameters, e.g. a different grid size or a different time step. However, this would require manual checks for artifacts to make sure a combination of parameters remains stable for the specific physical conditions, e.g. an increased grid size would in principle be prone to the same artifact with a (in physical units) different period length. As an alternative, those numerical artifacts can be completely suppressed by clamping [Eq. (9.4)].

The first study presented here was first published in [Sch+17c]. It features two simulation runs starting with the same current distribution, one run using cubic interpolation, the second using quadratic interpolation. The complete set of simulation parameters is listed in Table 9.1. For the simulation run that uses cubic interpolation, the distribution stays in a relatively calm state with just little oscillation⁶. Note that this is not equilibrium: As expected for the unshielded CSR case and a normalized current $\xi > 0.5$, it is $\sigma_E > \sigma_{E,0}$ [SH02]. Also, there is an oscillation with $f \approx 2f_s$.

If there was no or little influence of the different interpolation schemes, one would expect no difference between the two runs. However, as Fig. 9.5 depicts, this is not the case: The energy spread simulated using quadratic interpolation rapidly increases to a higher value after one synchrotron period ($T = 1 T_s$). On a longer time scale, it will damp down again, and after that a new numerical instability might rise. If the aim is to find out whether the simulated conditions are above the micro-bunching threshold these artifacts might not matter – below the threshold the initial, numerical modulation should not be amplified. However, we want to track the evolution of the charge distribution, and this numerical artifact might be interpreted as an unphysical slow bursting frequency. So we have to avoid numerical artifacts as they occur in the run using quadratic interpolation (cf. Fig. 9.5(c)): The period length of the ripples is exactly two grid cells, and they continue to exist even to a position where they create negative charge densities. In this particular case, the instability is clearly triggered by numerical artifacts (overshoots) of the interpolation.

Another example can be found in Fig 9.6. It shows a charge modulation in the phase space similar to the one expected in physical cases (cf. Sec. 10.2.1) but with an amplitude in the order of magnitude of the numerical accuracy. As they do not grow to the same extent they would during the instability, in this work, they are referred to as *sub-critical substructures*. The set current in the example is $I_b = 186.66 \mu\text{A}$ which is slightly below the micro-bunching instability threshold current for the used simulation settings $I_{th} = 190 \mu\text{A}$ (see Fig. 10.10(e)). The structures are presumably triggered by the occurrence of the numerical noise best visible near $q = 0$ in Fig. 9.6(d). These small spikes have the amplitude of the numerical accuracy of the `binary32` data type (see Sec. 9.1.2) that is approximately 7 decimal digits, and thus (in absolute numbers) lower near the maximum of the charge distribution. When clamping is used, the bigger structures (at $q \approx \pm 2$) sometimes do not appear. In these cases, it seems to suppress the growth of the numeric noise from the implicit discretization.

It has to be noted, that although clamping often suppresses the formation of probably unphysical substructures, it is not a universal remedy but also introduces new problems.

Table 9.1.: Parameters for an example run to check for numerical artifacts.

For the given set no artifacts were observed. Changing to quadratic interpolation however, triggered the occurrence of (non-physical) structures with a period length of two grid cells (see Fig. 9.5).

Parameter	Value
Grid points per axis	256
Steps per T_s	4000
Rotation method	Manhattan
Interpolation method	cubic
Impedance model	free space
Normalized current (ξ)	0.516
Damping time	$200 T_s$

⁶That state was used as the starting distribution.

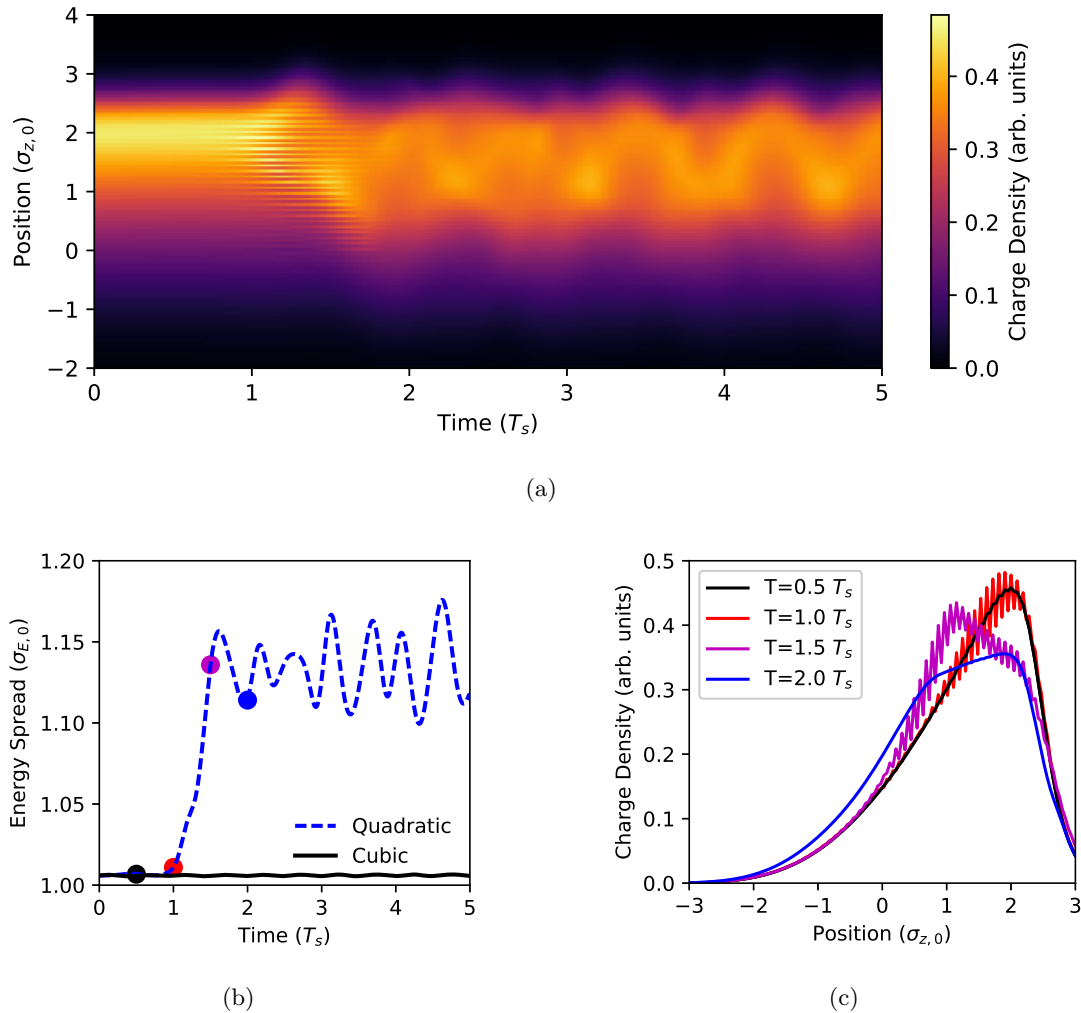


Figure 9.5.: Self-amplifying interpolation artifacts

Evolution of the bunch profile (a), and energy spread (b) over time, as well as example bunch profiles (c) at the points in time marked by the disks. For the case where cubic interpolation is used [(b), solid black line], the energy spread stays at $\sigma_\delta \approx 1.01\sigma_{\delta,0}$. When using quadratic interpolation [(b), dashed blue line], an increase in energy spread can be observed at $T = 1 T_s$. The bunch profiles in (a) and (c) reveal that this increase is a numerical artifact: The bunch profiles computed during the initial increase of the energy spread show large ripples with a period length of two grid cells. The earlier and later profiles do not show such structures. This implies that the increase of energy spread is driven by a numerical instability. (The bottom row was first published in [Sch+16].)

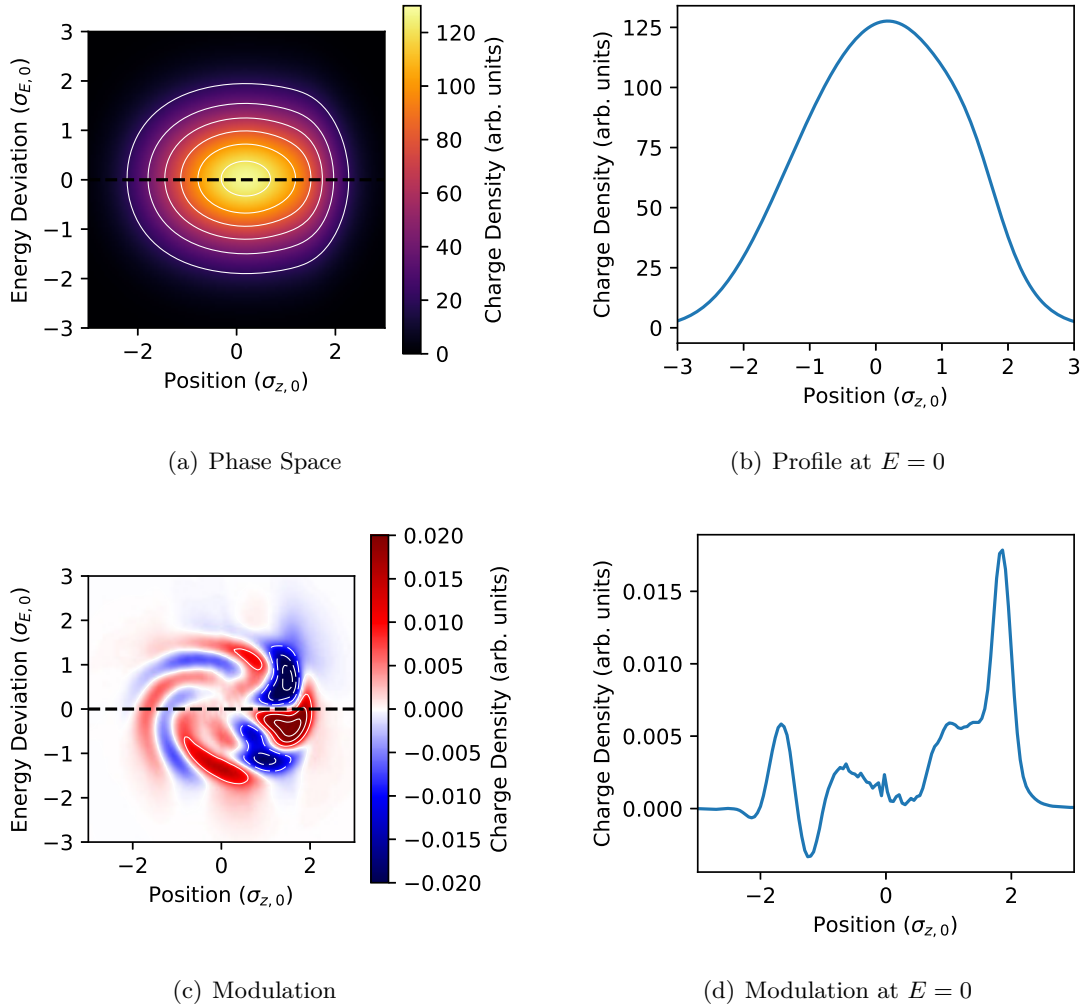


Figure 9.6.: Suspected amplification of quantization noise

On top of the distribution in phase space (a), there is a small modulation (c). For both, contour lines are added to the color-code to guide the eye. How the volatile part forms on the otherwise static distribution will be discussed in Chapter 10, here the focus is on a numerical problem. To get a better impression, (b) and (d) show slices at $E = E_0$. In (d) at $z \approx 0$ there is discretization noise arising from sampling the largest values of the peak region (cf. Sec 9.1.2). The shape of these *sub-critical substructures* at $z = \pm 2\sigma_{z,0}$ is the same as for physical substructures (cf. Sec. 10.2.1), in this particular case, however, they are presumably triggered by the occurrence of the numerical noise. This view is supported by the fact that the set current $I_b = 186.66 \mu\text{A} < I_{th} = 190 \mu\text{A}$, where I_{th} is the micro-bunching instability threshold current for the used simulation settings, as deduced in Sec. 10.3.

Clamping adds a switch to be sure no numerical substructures of the a length of $2\Delta q$ are formed due to interpolation. This comes at two possible costs: As will be shown in Sec. 9.4.2, also physical structures might be suppressed. Second and probably worse, there are also examples where clamping introduces unphysical substructures below the instability threshold (e.g. in Sec. 9.5).

9.3. Benchmark Scenarios

In this section, it will be checked how well analytically known behavior is reproduced by the simulation. It will be shown that the unperturbed problem converges almost independently from the chosen data types but that there are differences dependent on the chosen explicit discretization and the used algorithms (Sec. 9.3.1). Also, the long-term stability will be checked (Sec. 9.3.2) by testing the conservation of probability.

9.3.1. Convergence of Unperturbed VFPE

For stable beam conditions, the energy is distributed Gaussian with $\sigma_p(t) \equiv 1$. This equilibrium of damping and diffusion is independent from the starting conditions, thus every simulation of a stable beam should converge to this case. To be able to also check the temporal evolution, the collective impedance can be set to $Z = 0$, so that the damping time τ_d is reproduced without any change due to collective effects⁷.

Figure 9.7 shows how an initial disturbance is damped in an unperturbed case with $\tau_d = 10 T_s$. An initial displacement leads to a coherent synchrotron oscillation following

$$\mu_z(t) = \cos(z(t=0) \times \exp(-t/\tau_d)). \quad (9.7)$$

An initially increased energy spread⁸ decreases asymptotically to $\sigma_{E,0}$. Notice that the time constant is halved [Wol14b, Eq. (120)]

$$\sigma_E(t) = \sigma_E(t=0) \times \exp(-2t/\tau_d) + \sigma_{E,0} \times [1 - \exp(-2t/\tau_d)] \quad (9.8)$$

$$= \sigma_{E,0} + \Delta\sigma_E \times \exp(-2t/\tau_d) \quad (9.9)$$

for this second order moment: It reduces with the incoherent synchrotron motion of particles that also have radiation excitation [San69, p. 120]. As the energy spread allows to obtain both, the damping time as well as the asymptotic behavior (here $\sigma_{E,0}$) in just one fit, there is no need for an initial displacement.

In a first study, the data type is varied. The types include `binary64` (bin64), `binary32` (bin32) as floating point [see Eq. (9.1)] data types⁹ as well as various fixed type data types that are named using the number of bits of the fractional part [e.g. `fix29`, see Eq. (9.2)]. As stated in Sec. 9.1.2, the explicit discretization overweights the implicit one: Figure 9.8 confirms that there is no noticeable difference coming from the used data type. For example, the mismatch between `binary32` and `binary64` is in the order of 10^{-7} , which is the difference in the numerical accuracy. Also, for the other data types, increasing the accuracy does not bring the results closer to the analytical one but to another (see also Appendix C.2). Only, when decreasing the accuracy to a 29 bit fixed point data type, a significant change can be observed at $t \approx 20 T_s$: In this case, the numerical accuracy does not seem to be sufficient to model the physical system properly and the trend starts to visibly deviate from the theoretical curve. As libraries for mathematical function, such as [clFFT16; FFTW16; Bst17a] are only available for floating point data types, and as the choice of the data type does only play a secondary role for the result, `binary32` is chosen as the default.

⁷For the study varying the data type, it is also important that K is implemented based on a DFT and the used libraries only work with floating point data types.

⁸Because $Z = 0$, the bunch length would follow the same curve as the energy spread.

⁹The type `binary16` was not evaluated because at the time of writing it is not supported by many platforms.




Figure 9.7.: In equilibrium, the amplitude of the coherent synchrotron motion (bunch position, oscillating blue line) as well as the energy spread σ_E damp exponentially. The blue envelope follows $z(t=0) \times \exp(-t/\tau_d)$, while it is $\Delta\sigma_E(t=0) \exp(-2t/\tau_d) + \sigma_{E,0}$ for the energy spread.

In the second series of benchmark runs, the grid size N_x, N_y , the size of the time step Δt , and the damping time β_d are varied. Furthermore, standard and Manhattan rotation are compared. The results are displayed in Fig. 9.9. The quality of the rotation and the interpolation used in the map R is dependent mostly on the rotation angle. If the number of steps for one full rotation becomes too large, the influence of numerical absorption [Lop+18] rises: For example with linear interpolation [Eq. (8.13) with $n = 1$] the factor $1 - \{m\}$ is susceptible to absorption. This means $f = (1 - \{m\}) + \{m\}$ will be computed to be $f < 1$. In that case the content of $f(\lfloor m \rfloor)$ will be underrepresented while the content of $f(\lfloor m \rfloor + 1)$ will be overrepresented. For higher order interpolation techniques the problem is even more severe, since $\{m\}^2 \ll \{m\}$ appears. This is one reason why the Manhattan rotation performs better compared to standard rotation.

For the stability of the map D (see Sec. 9.2.2) the change in energy in one time step is the key parameter. It is directly proportional to β_d , and it can be clearly seen that the slowest damping ($\beta_d = 0.001$) is the most difficult one. The grid size – as long as it is sufficiently large – does not seem to have much of an influence. It seems as if fast damping can compensate for the inaccuracies introduced by a coarse grid, but there is no gain visible due to finer grids.

For the examples studied in Chap. 10, it is always $\beta_d < 0.002$, so convergent settings can be found for those cases. To achieve short run times combined with good convergence, typically a combination of $N_x = N_y = 256$, Manhattan rotation, and $N_\theta = 500$ was chosen.

Of course, the numerically reproduced bunch shape does also have an effect on the simulated bunch spectrum and thus on the instability threshold. Such influences will be discussed in Sec. 9.5.

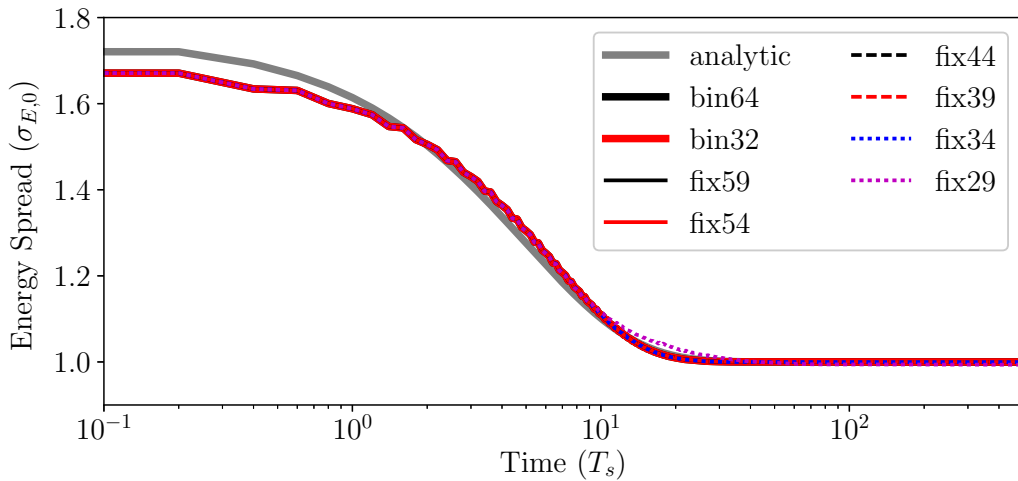
9.3.2. Conservation of Probability

The algorithm discussed in Chap. 8 was designed to explicitly fulfill Eq. (8.2) and thus conservation of probability¹⁰. In other words, the integral over the grid should be a constant in time


$$\frac{\partial [\iint \psi(q, p, t) dq dp]}{\partial t} \equiv 0. \quad (9.10)$$

If all outer grid points held $\psi(x, y) = 0$ there is no problem due to rotations and translations but in fact it is just $\psi(x, y) \approx 0$. If the grid models a sufficiently large area of the phase

¹⁰In this context, one may as well call it conservation of charge.



(a)



(b)

Figure 9.8.: Exponential damping of the energy spread σ_E as a function of time for different data types, where bin marks (binary) floating point types and fix marks fixed point types (cf. Sec. 9.1.2). There is an initial offset as well as a jitter due to quantization noise in the test pattern that has been read in from a 16 bit *PNG* file. Even with this problematic starting conditions, after $T = 60 T_s$, all data types have converged to $\sigma_q = \sigma_p \approx 1$. Note that all simulation curves are almost perfectly overlapping. Except for fix29 (at $T \approx 20 T_s$), there is no noticeable difference coming from the used data type visible in (a). The difference from the analytic result is emphasized in (b). Here also small differences in the asymptotic behavior of $\sigma_{E,s}$ become visible. However, in most examples the offset from the analytic value is not significantly dependent the chosen data type, so it must be due to the explicit discretization and the used algorithms.

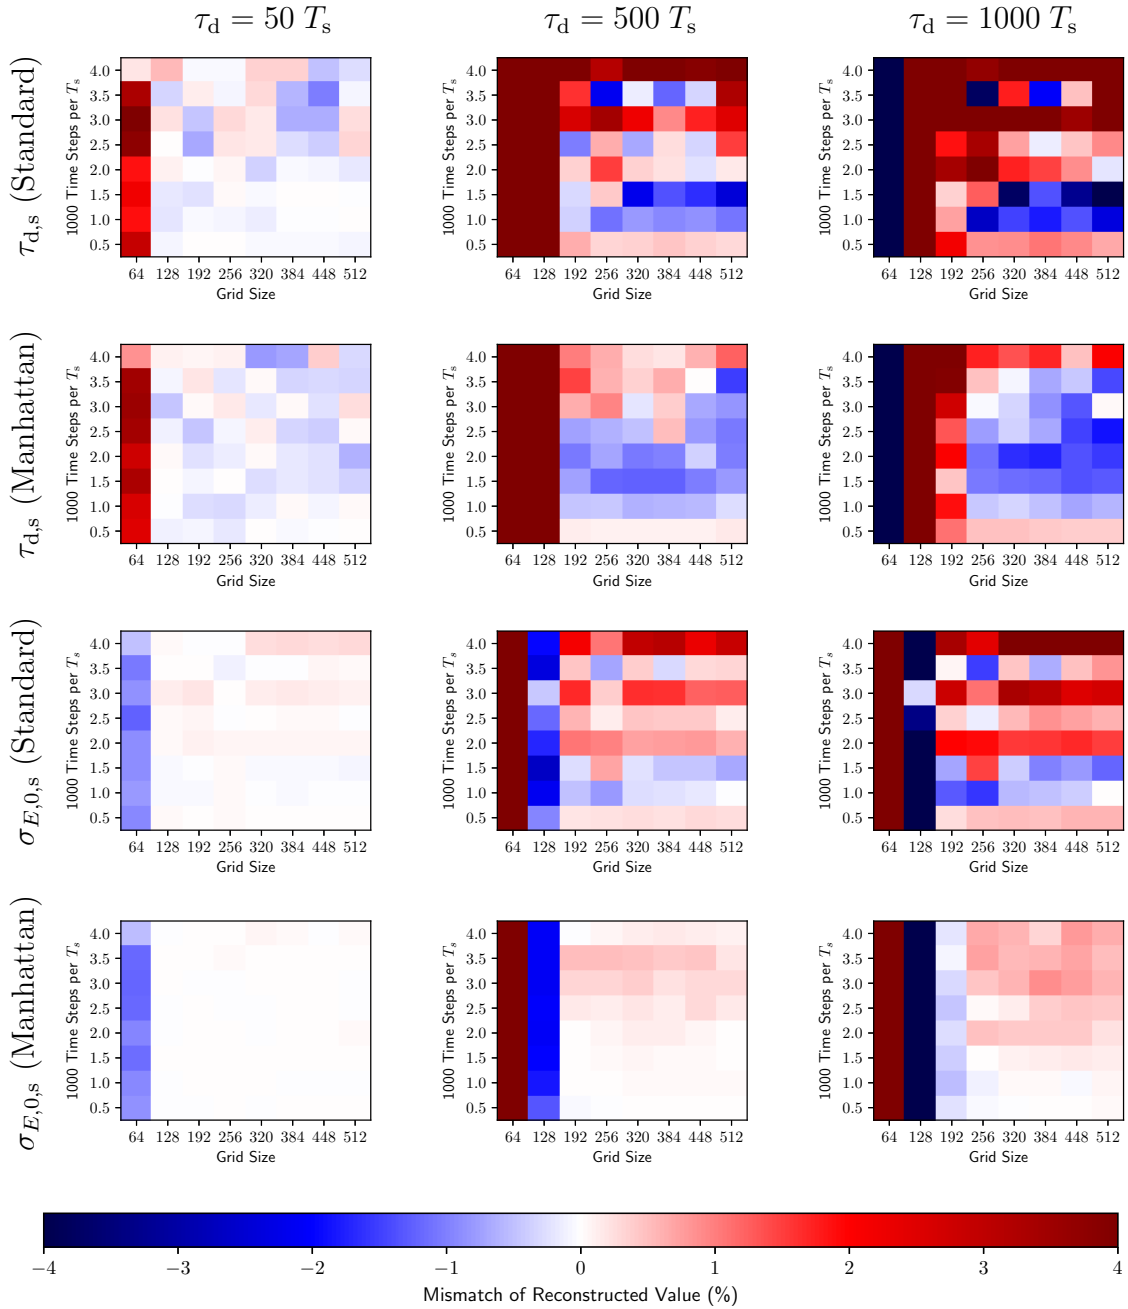


Figure 9.9.: Reconstructed damping time ($\tau_{d,s}$, upper two rows) and natural energy spread ($\sigma_{E,0,s}$, lower two rows) for simulation runs with different grid and time step sizes. The set damping time is $\tau_d = 50 T_s$ ($\beta_d = 0.02$) for the left column, $\tau_d = 500 T_s$ ($\beta_d = 0.002$) for the middle one (results for both first published in [Sch+17c]), and $\tau_d = 1000 T_s$ ($\beta_d = 0.001$) for the right one. It can be seen that slow damping is numerically more challenging, and that – independently from the numerical settings – Manhattan rotation is more robust in reconstructing the set values.

space, for $\iint \psi(q, p) dq dp$ the relative error is very small. Typically a quadratic grid of $12 \sigma_{z,0}$ and $12 \sigma_{E,0}$ is a good choice, so

$$q_{\max} - q_{\min} = p_{\max} - p_{\min} = 12. \quad (9.11)$$

When bunch lengthening [see Eq. (3.24)] becomes too strong, of course bigger maximums and minimums have to be chosen. It is common practice to choose the coordinates so that $q = p = 0$ has an explicit grid point (nearly) at the center of the grid, meaning

$$q(x = \lceil N_x/2 \rceil) = p(y = \lceil N_y/2 \rceil) = 0. \quad (9.12)$$

But as the bunch can become asymmetric due to $\Re(Z)$, it can be a good idea to reflect this in the grid, by choosing

$$|q_{\max}| \geq |q_{\min}| \text{ for } \alpha_c \geq 0. \quad (9.13)$$

Typically Eq. (9.10) gives a good indicator¹¹ when the coordinate system has to be adjusted.

It has to be kept in mind that absorption errors are most severe in the center of the phase space, where usually the maximum of $\psi(q, p)$ is located: Here, $\{m\}$ in Eq. (8.13) is smallest for the rotation map R . Also notice that the order of magnitude of the absorption error in the worst case can be many orders of magnitude bigger than the accuracy of the chosen data type. For example for `binary64` the accuracy is $< 10^{-15}$ but the absorption error can be $> 10^{-8}$ (cf. Sec. 9.1.2). Over- and undershoots (as well as clamping) do also lead to numerical gains or losses. So even when the grid is chosen large enough, for long simulation runs, numerical changes can accumulate to a significant value – practically ad infinitum. To fix this loss of normalization, a renormalization can be applied from time to time. A simplistic method implements

$$\hat{\psi}(n) = \frac{\psi(n)}{\iint \psi(q, p, t) dq dp} \quad (9.14a)$$

as


$$\hat{\psi}(n) \approx \frac{\psi(n)}{N_x \times N_y \sum_{i=1} \psi(i)}, \quad (9.14b)$$

where $\hat{\psi}$ marks the renormalized charge density. As the normalized bunch profile ϱ is computed in every time step, it is also possible to use it for the computation


$$\hat{\psi}(n) \approx \frac{\psi(n)}{\sum_{i=1}^{N_x} \varrho(i)}. \quad (9.14c)$$

This approach ignores the more elaborated methods of integration presented in Sec. 8.1.4 for $\int dq$, but follows them for $\int dp$. It is used to obtain the results shown in Fig. 9.10. This approach seems to be sufficient here – especially for the dynamics well above the threshold – the numerical effects should be negligible. On the other hand, when nonlinear effects come into play, even small changes and the perturbation due to the renormalization might become significant and influence the instability threshold. As Sec. 9.5 shows, this is typically not the case.

¹¹*Inovesa* prints the value $1 - Q/Q_0$, which consequently should be approximately zero.



(a) Change of average bunch length due to renormalization



(b) Simulated bunch length over time

Figure 9.10.: In a simulation run, the integrated charge is normalized and should always be 1. However, due to numerical effects the charge can drift off this value. As a benchmark 1500 synchrotron periods each modeled in 1000 time steps on a grid of 256×256 points are simulated for $\beta = 0.002$. (To guide the eye, the data points are connected.) In this test case, without renormalization a numerical equilibrium is found with an integrated charge of 1.00234 and a length of 1.0012 $\sigma_{z,0}$. This can get more severe, when improper starting conditions drive part of the charge out of the space spanned by the grid. To overcome this problem, renormalization may be applied. The difference between set and reconstructed bunch length for different frequencies of renormalization is shown in (a). It can be seen that the simulated bunch size becomes smaller and the difference to the analytical value bigger, the more often it is renormalized. The time domain signal for four example numbers of renormalizations per synchrotron period (b) shows a slight numerical fluctuation due to the renormalization. It is worth noting that the numerical equilibrium is disturbed by the renormalization, and thus even though the curves look smoother for more frequent renormalization, it is just a trade-off between amplitude and frequency. Notice that the deviation is relatively small for all shown settings. In fact, loss of normalization becomes severe for long simulation times (this will be visible in Sec. 9.5).

9.4. Starting Conditions to Simulate Unstable Beams

Under stable beam conditions, there is an equilibrium and the energy distribution can be described by a Gaussian function with $\sigma_p = 1$. To obtain the bunch profile, the Haïssinski equation [Eq. (3.19)] can be solved¹². To describe the dynamics of the micro-bunching instability above the threshold current, often (e.g. [WE00; BCS10; Kle12]) the Haïssinski solution right at the threshold is taken for the initialization of $\psi(q, p)$. This, however, is not always a good choice.

9.4.1. Well above Instability Threshold

If the dynamics of the micro-bunching instability well above the threshold current is of interest, the non-equilibrium charge distributions have little similarities to the ones in equilibrium. Furthermore, there is no simple function describing any physical state of a possible charge distribution $\psi(t_0)$. One simply has to start with a distribution that might be physically unreachable and give it time to converge to a possible $\psi(t)$. In fact, well above the instability threshold, any distribution will reach the physical state of the instability. So, alternatively, to the Haïssinski distribution, one can – for example – start the simulation with a Gaussian distribution that is broader than the expected charge distribution. It will damp down until the (unstable) physical conditions are reached. This initial conditions show good convergence compared to narrower distributions: The latter can lead to the formation of unphysical structures that might remain for a long time. If available, it is a good idea to start with a charge distribution obtained by a simulation run with similar




Figure 9.11.: Evolution of the RMS energy spread over a time of 500 synchrotron periods for $I = 1.5$ mA. The width of the line is caused by high frequency oscillations. For the starting distributions the Haïssinski solution (equilibrium at $I = 675$ μ A), a Gaussian distribution with $\sigma_p = \sigma_q = 1.5$, and distribution that happens to be the last of the previous run (“adiabatic”, $I = 1.54$ mA) were used. Note that the Haïssinski distribution is immediately blown up and becomes larger than the Gaussian distribution, and that the high frequency oscillation is systematically higher until $T \approx 300 T_s$. (First published in [Sch+17c].)

¹²No VFPS is needed in such cases.




Figure 9.12.: Even when numerical and physical settings are the same, depending on the starting-conditions two simulation runs may converge to different states. In the above example, the red curve represents a simulation run that clearly ends up in a state that describes “regular bursting”: The width is due to rapid oscillation (see Sec. 10.2.1 for details). The blue curve on the other hand is in a meta-stable result state with only small oscillation.

physical settings. Such distributions are especially given when a whole range of parameters is scanned, e.g. to investigate the current-dependent changes of the dynamics.

In Fig. 9.11 the evolution of the energy spread over 500 synchrotron periods is shown for three different starting distributions at $I = 1.5$ mA. When using the final distribution of a previous run with slightly higher current (here $I = 1.54$ mA), the simulation converges quasi instantaneously. As shown, the Haïssinski distribution is immediately blown up and becomes larger than the Gaussian distribution that has been set to be larger than the expected distribution. Also in the beginning ($t < 300 T_s$) the oscillation is systematically enlarged.¹³

9.4.2. Near the Instability Threshold

Near the instability threshold, the situation is different: The closer the physical state comes to stable conditions the longer it will take an arbitrary charge distribution to converge. Furthermore, small numerical inaccuracies gain more weight. Figure 9.12 shows two simulation runs for the same physical and numerical settings – both initialized with different starting distributions $\psi_{1,2}(t_0)$ with $\sigma_z(\psi_1(t_0)) < \sigma_z(\psi_2(t_0))$. For ψ_1 , the bunch length first grows, then shows some transitional jitter and finally settles in a state with strictly periodic oscillations (“regular bursting”, see Sec. 10.2.1). For ψ_2 , the bunch length shrinks with decreasing length fluctuation and settles in a state of constant length. As Fig. 9.13 shows, this state of constant length is an artifact.

¹³Paragraph first published in [Sch+17c].




Figure 9.13.: A simulation run that has converged to a meta-stable state can be triggered to switch the state by introducing a (temporary) change. Here, the same starting conditions are used as for the blue curve in Fig. 9.12. This time, after 500 synchrotron periods, clamping is tuned off, and after another 500 synchrotron periods again back on. In the end, the simulation is in the same state as the red curve in Fig. 9.12.

In the given example, the substructures that are characteristic to the physical state are not formed when starting with ψ_2 . The reason is, that clamping (see Sec. 9.2.1) may not only suppresses the growth of (unphysical) numerical substructures but also the formation of ones that are physically correct. In the simulation run displayed in Fig. 9.13, the substructures are generated by temporarily turning the clamping off. As soon as the substructures are present, they persist even with activated clamping.

In total, convergence is more delicate under these – almost stable – conditions. How this can effect the determination of the instability threshold will be discussed in the next section.


9.5. Error Estimation for Bursting Thresholds

The example shown in Sec. 9.4.2, would lead to an overestimation of the bursting threshold. Figure 9.6, on the other hand, shows substructures in the order of magnitude of the numerical accuracy even below the instability threshold. So, when the threshold estimation is based on the presence of substructures, the bursting instability threshold current can as well be underestimated.

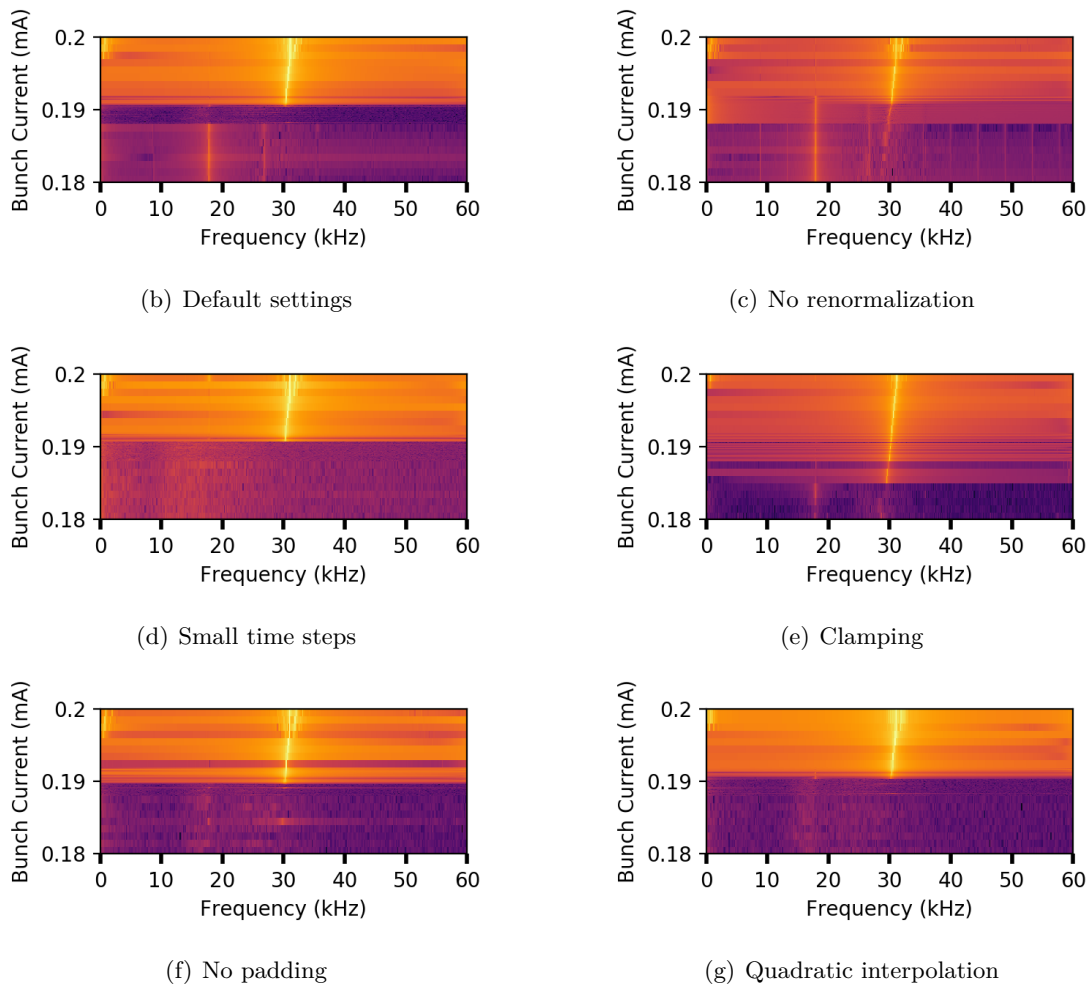
Typically, the threshold is determined based on the fluctuation of the CSR σ_{CSR} , on a spectrogram of the same quantity, or on the average energy spread $\langle\sigma_E\rangle$. All three are displayed in Fig. 9.14 for different simulation runs with the same physical but different combinations of numerical settings. The default settings are $N_x = N_y = 256$, $\theta = 1/500$, 50 renormalizations per synchrotron period, no clamping, cubic interpolation and 8-fold padding which results in $8 \times N_x = 2048$ sampling points for the impedance¹⁴.

On a linear scale, the sub-critical substructures mentioned in Sec. 9.2.3 are too small to be visible when compared to an unstable beam. The results using the frequencies of the fluctuation of the CSR power do not differ from the ones obtained using σ_{CSR} or $\langle\sigma_E\rangle$. Note that – because of their logarithmic power axis – some of these plots show track of the sub-critical substructures. In these cases, the determined threshold will be dependent on the chosen scale when sloppily arguing based on the spectrograms. In the example shown here, the most severe differences are created by clamping: In the current range of the instability threshold and below, it introduces significant sub-critical substructures that

¹⁴The padding is equivalent to sinc interpolation of the form factor, so these points yield 8-fold frequency resolution.



(a) Comparison of simulation runs

**Figure 9.14.: Bursting Threshold obtained with Different Numerical Settings**

In (a) values are computed at multiples of $0.1 \mu\text{A}$, lines are to guide the eye. The curves obtained using cubic and quadratic interpolation overlap almost perfectly. Smaller energy spread as well as the “fluctuation” baseline visible without renormalization, are artifacts due to the loss of charge (which is 0.1%). What remains is a slightly higher current when the curves start to rise. The default settings are a grid size of 256, a step size of $\theta = 1/500$, renormalization, no clamping, cubic interpolation (IP=4), and zero-padding of the bunch profile to be able to sample the impedance with 2048. Changing the step size to $\theta = 1/4000$, introduces numerical jitter and effects σ_E but not the thresholds. Here, clamping, obviously introduces severe artifacts, whereas a decreased accuracy when sampling the impedance (256 instead of 2048 points) shifts the dynamics to lower current values but does not introduce other differences. For the power fluctuation spectrograms, the results are as well very comparable. The $f_s/2$ lines in (c) seem to be artifacts. Also, it shows track of the sub-critical substructures. In this cases, the threshold might be underestimated – especially when not considering the logarithmic power axis.

express in all three methods. These structures are probably caused by the magnitude of the needed renormalization (0.1 % in $25 T_s$). Doing no renormalization leads to a small (0.1 % in $150 T_s$) loss of charge which is visible in two artifacts: a smaller energy spread as well as a the “fluctuation” baseline. The drift does not influence the spectrogram by the same extent, however, the current axis becomes erroneous and a slightly higher threshold is determined. At last, the decreased accuracy when sampling the impedance shifts the dynamics to lower current values but does not introduce other differences.

In the study presented here, the spread of most results is in the range of $0.1 \mu\text{A}$. Situation is worse, for example when charge normalization is lost or numeric fluctuation is introduced, e.g. by a combination of clamping and renormalization. In such cases results can deviate by up to $4 \mu\text{A}$, but in total, the instability threshold obtained by simulation is pretty robust against numerical influences. When considering the difficulty of the representation of damping and diffusion, the method might be less accurate in simulating dynamics with higher β_d or where the exact damping time becomes more important.

10. Simulation of the Longitudinal Phase-Space

“A good simulation, be it a religious myth or scientific theory, gives us a sense of mastery over experience.”

Heinz R. Pagels

In this chapter, first, quasi-static or averaged influences of the self-interaction due to CSR emission are described (Sec. 10.1). These effects are present for stable beam conditions but also when the micro-bunching instability rises. Section 10.2 shows how bursting CSR and the underlying beam dynamics look like in ideal-typical cases. Afterwards (Sec. 10.3) gives some brief examples how the dynamics can be influenced by adding additional impedances. Section 10.4 gives a deeper look into the frequency dependence of the CSR spectrum and its fluctuation. The chapter is closed (Sec. 10.5) with an analysis of the microscopic dynamics (e.g. single-particle movement) and tries to build an intuitively accessible explanation how these lead to the formation of the micro-structures.

10.1. Current-Dependent, Static Phenomena

The phenomena presented in this section already start in the regime of the potential well distortion (see Sec. 3.3.1). The displayed results for bunch deformation (Sec. 10.1.1) and the tune spread (Sec. 10.1.2), are chosen so that only small time-dependent changes are visible. For the examples for bunch lengthening (Sec. 10.1.3) and synchrotron radiation (Sec. 10.1.4), however, the instability threshold is crossed to an extent, that this is no longer true. Here, only an example with significant fluctuation is shown, the dynamics will be discussed in Sec. 10.2.

10.1.1. Bunch Deformation

The longitudinal bunch profile changes with current – even when the beam does not show any time-dependent fluctuation or instability. As introduced in Sec. 3.2, for example a purely resistive impedance makes the bunch lean forward, while an inductive impedance lengthens it in a symmetrical manner.

For an accelerator dominated by the CSR impedance, the effect is biggest when shielding is low (cf. Sec. 3.2.4). As the vacuum chamber has a fixed height, this means that short bunches will be effected more. For them, the potential well becomes very asymmetric – leading to a shark fin shape bunch profile. An example ($f_{s,0} = 7.7$ kHz, $V_{RF} = 1.5$ MV) for bunch shapes in the deformed potential well is displayed in Fig. 10.1. It is clearly visible that the bunch leans forward for higher currents, slowly deviating it from a nearly Gaussian




Figure 10.1.: Example Deformed Bunch Profiles

The bunch is leaning forward due to the potential well distortion. For $I = 100 \mu\text{A}$, it is already unstable, but the modulation is too small to be seen here. The shape does not only depend on the current, but also on the impedance and the natural bunch length.

distribution at $I = 5 \mu\text{A}$ to a distribution with a quite flat top at $I = 100 \mu\text{A}$. The profile at $I = 100 \mu\text{A}$ already has some almost invisible, time-varying micro-structures. The others are all static cases without time-varying components. So instead of using a full featured VFPS simulation – as done here, they also could have been calculated using the Haïssinski equation [Eq. (3.19)].

The bunch deformation is intuitively understandable when looking at the path of the individual particles in the longitudinal phase space. Before discussing this in detail in the next section, a purely resistive impedance may serve as a first example: It leads to energy losses, and due to synchrotron motion, more particles will aggregate in the front of the bunch.

10.1.2. Shift and Spread of the Synchrotron Tune

The contribution of wake fields adds to the accelerating voltage leading to a deformed potential well. Radiation damping and drifts in dispersive sections on the other hand remain unchanged. In total, the synchrotron motion deviates from its originally symmetric path¹. Fig. 10.2 illustrates these perturbed paths for example particles (cf. Sec. 8.5) in the phase space for two different beam currents ($f_{s,0} = 7.2 \text{ kHz}$, $V_{\text{RF}} = 1.4 \text{ MV}$).

First, it has to be noticed that the position of the synchronous particle (the one that does not move) shifts to the bunch head². This is a result of the coherent phase shift (see Sec. 2.4): The increased energy loss due to the real part of the (CSR) impedance has to be compensated by the higher accelerating voltage that is found at the new synchronous phase. The position of the synchronous particle should not be mixed up with the average position of all particles – the center of mass or bunch position. The bunch position gives the average energy loss but when the energy losses are position dependent, there is a relevant difference. Notice that

$$\left| \frac{\partial E(-|\phi|)}{\partial t} \right| < \left| \frac{\partial E(|\phi|)}{\partial t} \right|, \quad (10.1)$$

where $\phi = 0$ marks the synchronous phase [Eq. (2.32)]. This fact results in an asymmetric

¹In this thesis only α_0 is considered and the coordinate system is normalized so that the unperturbed motion would follow a cycle.

²For $\alpha_c < 0$, the movement is to the bunch tail.

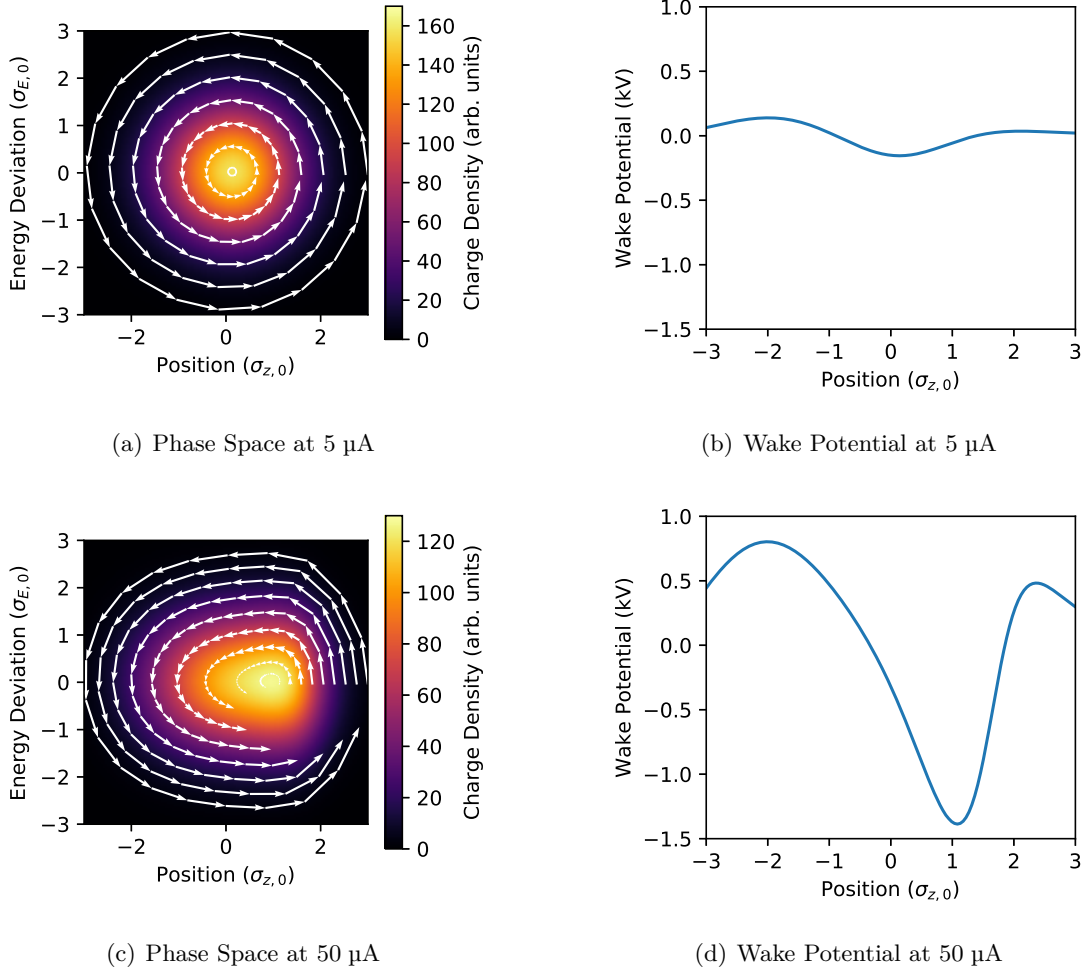


Figure 10.2.: Particle Movement in the Longitudinal Phase Space

(a) and (c) display charge densities in the phase space, the arrows represent the movement of example particles in 20 time intervals of $\Delta t = T_{s,0}/20$, each. (b), (d) show the corresponding wake potential. As both currents are below the bursting threshold ($62 \mu\text{A}$, according to BBT), average quantities are constant in time. Notice that energy changes faster ahead of the synchronous phase, where also less charge can be found. The inverted time difference the particles need for one round-trip is referred to as the *tune spread*, the difference from $f_{s,0}$ is called (incoherent) *tune shift*.

bunch shape where more particles are trailing the synchronous particle than leading it

$$\int_{-\infty}^0 \varrho(\phi) d\phi > \int_0^{\infty} \varrho(\phi) d\phi. \quad (10.2)$$

Center of mass and the synchronous phase are compared in Fig. 10.3. In this plot ($f_{s,0} = 7.2 \text{ kHz}$, $V_{\text{RF}} = 1.4 \text{ MV}$) the situation under unstable conditions is not shown to circumvent situations where $\partial V/\partial t \neq 0$, so also $\partial \phi/\partial t \neq 0$. Furthermore, the intuitive definition of a single synchronous phase becomes ill defined when there is more than one minimum of $V_{\text{Wake}}(t)$. In that case, there will be more than one synchronous phase.

For stable beam conditions it can be used that the average synchrotron motion $\langle J \rangle$ is slower for particles located where the particle density is high. So the synchronous phase can be identified with the peak of the charge distribution. Be $\vec{r}(t) = (q(t), p(t))^T$ the path of a particle, so

$$\psi(\vec{r}_1) > \psi(\vec{r}_2) \Rightarrow \langle J(\vec{r}_1) \rangle < \langle J(\vec{r}_2) \rangle. \quad (10.3)$$




Figure 10.3.: Synchronous Phase Compared to Bunch Position

The bunch position (dashed red) is defined by the center of mass of the electrons, while the synchronous phase (solid blue) was determined as the minimum of the total potential well including both, \hat{V} and V_{Wake} . The band gives the size of the grid cell that contains the synchronous phase position. The quantities follow a similar trend but differ by about a factor of ten. This fact can be explained by the asymmetry of the charge distribution (see Figs. 10.1 and 10.2).

Meaning the charge density is higher where particles with a lower average instantaneous synchrotron motion are passing through. This relation is pretty intuitive as particles will aggregate where $J \approx 0$ but move away from those locations where J is bigger.

The $\langle J \rangle$ can be expressed in terms of (incoherent) synchrotron frequencies $f_{s,i}$, e.g. by looking at the time it takes an electron for one rotation in phase space³. As in most cases $f_{s,i} \neq f_{s,0}$, this phenomenon is called the (incoherent) *tune shift*⁴. The fact that $f_{s,i}$ is not a single frequency but position dependent, is referred to as *tune spread*. Note that the potential well moves alongside with the bunch, so the shift of the coherent synchrotron frequency $f_{s,c}$ which describes the motion of the bunch as a whole is a higher order phenomenon. Besides deformation, the tune shift can also manifest in a changed bunch length. For example, $f_{s,i} < f_{s,0}$ corresponds to bunch lengthening.

10.1.3. Growing Bunch Length and Energy Spread

While the potential well distortion alone deforms the bunch profile and thus also changes the bunch length, the energy spread stays constant as long as there is a stable equilibrium. Only when an instability (here: micro-bunching) rises, also the energy spread increases. In fact, this has been used as *the* indicator for microwave instabilities for both, experimental and simulation data (for example [BOZ05; BCS10]). When the beam dynamics are faster than the synchrotron motion, damping can no longer compensate the coupling between position and energy.

The energy spread as well as the bunch length for two examples ($f_{s,0} = 7.7$ kHz, $V_{RF} = 1.5$ MV and $f_{s,0} = 5.76$ kHz, $V_{RF} = 1.3$ MV) are shown in Fig. 10.4. For both, at lowest currents, the constant energy spread is visible. Also, under unstable conditions, the correlation between bunch length and energy spread can clearly be seen. Especially note that 10.4(b) shows important differences from the behavior in 10.4(a): First, there is an additional instability for $0.02 \text{ mA} < I < 0.035 \text{ mA}$. This **weak micro bunching** instability

³In this thesis, the $f_{s,i}$ are deduced as the inverse of twice the time it takes an electron to cross $p = 0$.

⁴To be exact, the word *tune* refers to the corresponding frequency divided by f_{rev} .

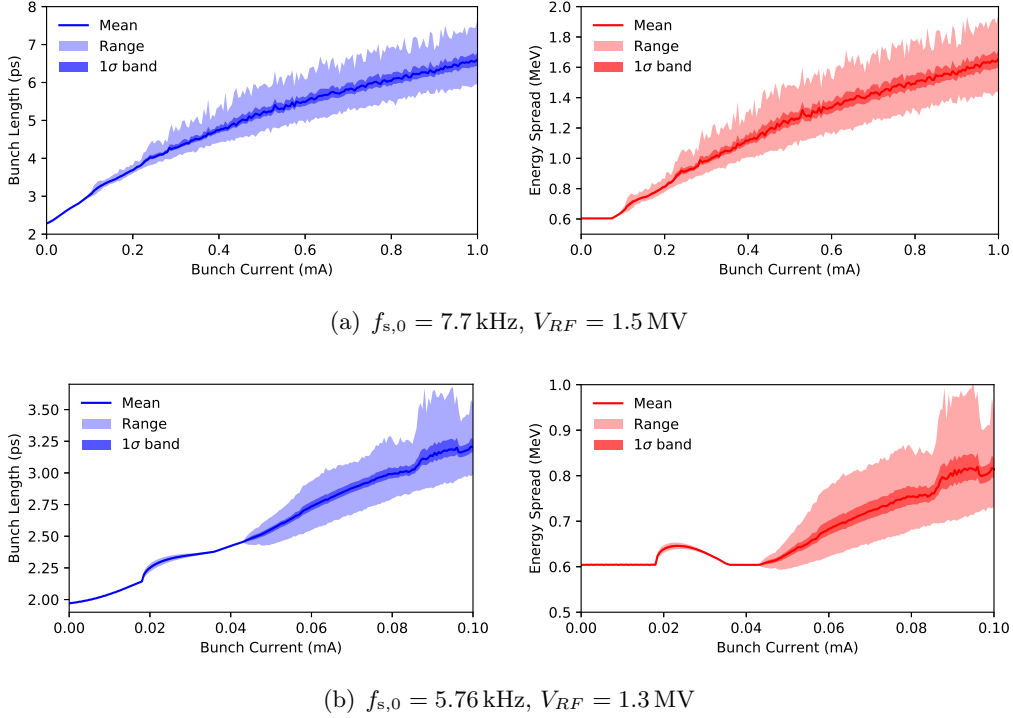


Figure 10.4.: Current-dependency of Bunch Length and Energy Spread

Left is the bunch length (RMS), right the corresponding energy spread. The evolution shown in (a) looks comparable to the one at $0.04 \text{ mA} < I < 0.08 \text{ mA}$ for the case (b). In the latter, the peak at low currents can be attributed to an additional weak instability (see text) that does not occur for (a).

(WMB)⁵ has been predicted in [BCS10], first experimental evidence is in [Mül+13]. More detailed studies have been performed both experimentally [Bro+16a] as well as numerically [Kus17]. This instability does not occur for the other case. For higher currents, first the beam becomes stable again, before another instability rises. Until $I \approx 0.08 \text{ mA}$, again the behavior is similar to the one in 10.4(a). At $I \approx 0.09 \text{ mA}$, there is another clear difference: The average bunch length and energy spread first instantly goes up and down afterwards. These steps are accompanied by a significant rise in the upper limit of both quantities while the lower limit follows a monotonic curve that just has a varying growth rate. This phenomenon is connected to a slower frequency of the sawtooth bursting as it will be explained in Sec. 10.2.2. The changing growth rate at $I \approx 0.25 \text{ mA}$ in 10.4(a) might be a less intense variant of the change.

The growth of the of the average bunch length above the instability threshold can be estimated as a power law. Using Eq. (3.24), also bunch lengthening has been utilized to determine the bursting threshold. (The author did this in [Sch13; Sch+13]). This approach is justified as long as the exponent in the power law is known – for example because it is constant over the whole spectrum. Measurements [Cai11; Mül+12] suggested that the bunch lengthening is dominated by the effect of the unshielded CSR impedance, meaning $\sigma_z \propto I^{3/7}$.

In Fig. 10.5, the RMS and FWHM bunch length are shown for simulations using $f_{s,0} = 7.7 \text{ kHz}$ and $V_{RF} = 1.5 \text{ MV}$. The double-logarithmic representation, suggests that the bunch length follows a power law. However, due to bunch deformation, the exponent is different for FWHM and RMS. Another result is, that in the case displayed here, $\sigma_z \not\propto I^{3/7}$,

⁵Also referred to as **short bunch length bursting** (SBB).

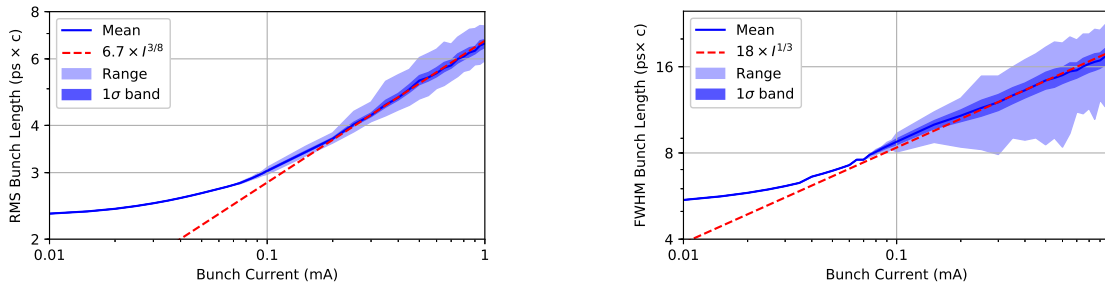


Figure 10.5.: Scaling of the Bunch Length with Current

Lengthening of RMS σ_z (left) and FWHM l (right) bunch length with current. The central dark blue lines mark the arithmetic mean, they are enclosed by a 1σ fluctuation band and the total range of the fluctuation. The error is too small to be visible. The dashed red lines display possible parameterizations for $0.1 \text{ mA} < I < 1 \text{ mA}$.

but rather $\sigma_z \propto I^{3/8}$ for the RMS and $l \propto I^{1/3}$ for the FWHM⁶.

To summarize, both bunch length and energy spread fluctuate due to the micro-bunching instability. At the same time, also the averages of these quantities rise. While the energy spread is constant for stable conditions, this is not the case for the bunch length. So when it comes to the average values, an increased energy spread can be seen as a reliable indicator for the micro-bunching instability while bunch lengthening can only give an estimate.

10.1.4. Synchrotron Radiation

Because of the changing form factor, bunch deformation directly influences the emission of coherent synchrotron radiation. Here, it is studied based on the example using $f_{s,0} = 7.7 \text{ kHz}$ and $V_{RF} = 1.5 \text{ MV}$. The results are displayed in Fig 10.6. It is visible, that the peak of almost full coherence ($\tilde{q}(f) \approx 1$ for the form factor) at very low frequencies becomes narrower (less bandwidth) but higher (more spectral power) due to bunch lengthening with increasing current. The position of the peak also moves to slightly lower frequencies. Already the static deformation due to the potential well also enhances radiation a higher frequency range ($I > 20 \mu\text{A}$, $f > 150 \text{ GHz}$). The rise of emission in the higher frequency range ($f > 800 \text{ GHz}$) happens as a sudden change at $I \approx 75 \mu\text{A}$. Sec. 10.2.1 will show, that it is caused by the appearance of micro-structures sitting on top of the bunch profile.

As introduced in Sec. 3.3.2, another consequence of the micro-bunching instability is the fluctuation of the emitted CSR. From the perspective of a synchrotron light source, this has negative as well as positive consequences: If stable emission is needed on a pulse-to-pulse basis, this fluctuation might be a problem. (As Secs. 10.2, 10.4 will show, the fluctuation amplitude can differ depending on the emitted frequency.) On the other hand, experiments relying on (coherent) emission at high frequencies can make use of the amplification – especially when only the average power is of interest. Average as well as minimum and maximum CSR power and the average fluctuation (1σ fluctuation band) are displayed in Fig 10.7.

From Eq. (2.22) it follows that $P_{CSR} \propto F(I) \times I^2$. To quantify the increase of the emitted power, here the form factor will be taken to be almost constant $F(I) \approx F(I + \Delta I)$ for small current changes ΔI to find the exponent $b(I)$ in the parameterization $P_{CSR} = k(I) \times I^{b(I)}$. It can be computed to be

$$b(I) = \frac{\ln [P(I)/P(I + \Delta I)]}{\ln [I/(I + \Delta I)]}, \quad (10.4)$$

⁶When uncertainty of a measurement is high, it might not be possible to reject $l \propto I^{3/7}$.




Figure 10.6.: Averaged CSR Spectrum

Both spectrograms show the averaged spectral CSR power of an electron bunch with $f_{s,0} = 7.7 \text{ kHz}$ and $V_{RF} = 1.5 \text{ MV}$. The left-hand side is a zoom-in with an adjusted color range, also the choice of the contour lines was changed. In the lower frequency range mainly changes due to bunch deformation in the potential well are visible. The peak position ($f \approx 70 \text{ GHz}$) is caused by the shielding: Emission of lower frequency radiation is suppressed by the height of the beam pipe. There are two effects that define the current-dependent differences: Bunch lengthening reduces the maximum frequency of full coherence, and bunch deformation introduces higher frequencies. At around $I \approx 75 \mu\text{A}$, also in the higher frequency range ($f > 800 \text{ GHz}$), CSR starts to become visible. This is due to the formation of substructures that are shorter than the (now modulated) charge distribution.

and is shown in the right hand side of Fig 10.7. It can be seen that the CSR emission grows more than linear but at a decreasing relative growth rate until the threshold. (The absolute growth rate increases until the threshold.) Above the instability threshold, the average integrated emission grows less than linear. Notice that this is the emission integrated over the whole frequency range, in isolated frequency bands the trends may be different (cf. Sec. 10.4).

10.2. Dynamics of Typical CSR Bursts

In Sec. 10.1.4, current-dependent changes of the CSR have been discussed. It has been shown that due to the micro-bunching instability, there are changes in the average power, its fluctuation as well as in the spectrum. Here, the dynamics of the fluctuations are studied in time-domain. A *burst* is the increase in CSR emission during a time interval where the emission of CSR peaks. It might be seen as time between two minimums of the emission, this however is not a fixed definition, as sometimes two bursts can overlap.

Depending on bunch current and machine settings, different types of burst can be observed. They vary in length, shape and amplitude. Current ranges that always show the same or very similar kinds of bursts, are called *bursting regimes*. In the following, two examples featuring very different characteristics are shown.

10.2.1. Regular Bursting Regime

In the “regular” regime, the substructures in phase space are quasi-static: They grow while moving from the back to the front of the bunch and shrink while moving into the other direction. The result is that bunch length and energy spread fluctuate a bit when the extrema are at the bunch front or back. The whole dynamics are dominated by one distinct frequency, the regular bursting frequency f_b . Also the CSR power is fluctuating with that frequency, and so are the small changes in the overall shape: As there is no low-frequency component in the fluctuation, the pattern of substructures in phase space has a nearly constant amplitude, so also its coherent emission is fluctuating with a relatively low amplitude (see also Appendix C.3). As the regular bursting regime is the




Figure 10.7.: Integrated CSR Power and its Fluctuation

From Eq. (2.22) it follows that $P_{CSR} \propto F(I) \times I^2$. Above the bursting threshold [(a), $I \approx 70 \mu\text{A}$], the growth rate of the CSR Power drops. This is because the changes in the form factor overweight the current change. But also in the stable region, the changing form factor plays a vital role. The green points in (b) gives the variable exponent $b(I)$ in the parameterization $P_{CSR} = k(I) \times I^{b(I)}$ of the CSR power. Until the threshold $b > 1$, so the slope is increasing up to that point; above the threshold $b < 1$, so the further increase of the power is less than linear. For corresponding bunch profiles, see Fig. 10.1, for bunch length and energy spread, see Fig. 10.4(a).

only one not featuring low-frequency oscillations, it is sometimes even referred to as “stable CSR” [Bar+15].

For a deeper analysis, it is convenient to express the bunch profile ϱ as [Bol17]

$$\varrho(\tau, t) = \langle \varrho(\tau, t) \rangle + \zeta(\tau, t) \quad (10.5a)$$

$$= \langle \varrho \rangle(\tau) + \zeta(\tau, t), \quad (10.5b)$$

where $\langle \varrho \rangle$ marks the average bunch profile and ζ the instantaneous modulation. For the form factor follows

$$F(f, t) = |\tilde{\varrho}(f, t)|^2 \quad (10.6a)$$


$$= \left| \langle \tilde{\varrho} \rangle(f) + \tilde{\zeta}(f, t) \right|^2 \quad (10.6b)$$

$$= \underbrace{\left| \langle \tilde{\varrho} \rangle(f) \right|^2}_{\text{static}} + \underbrace{\left| \tilde{\zeta}(f, t) \right|^2}_{\text{modulation}} + 2 \underbrace{\left[\Re \left(\langle \tilde{\varrho} \rangle(f) \right) \Re \left(\tilde{\zeta}(f, t) \right) + \Im \left(\langle \tilde{\varrho} \rangle(f) \right) \Im \left(\tilde{\zeta}(f, t) \right) \right]}_{\text{mixed}}. \quad (10.6c)$$

The first term, labeled “static”, comes from the average bunch profile, the second term from the modulation. The “mixed” term plays a role where the volatile modulation spectrum and the static spectrum from the average profile overlap. Notice that because $\langle \zeta(\tau) \rangle \equiv 0$, also $\langle \zeta(f) \rangle \equiv 0$. This means that in the average CSR spectrum the mixed term vanishes. But as $\langle |\zeta| \rangle \geq 0$, there will always be a contribution from a non-zero modulation. In total,

$$\langle F \rangle(f) = \left| \langle \tilde{\varrho} \rangle(f) \right|^2 + \left\langle \left| \tilde{\zeta}(f, t) \right|^2 \right\rangle. \quad (10.7)$$

However, as for the regular bursting $\langle \varrho \rangle(\tau) \gg |\zeta(\tau)|$, the average spectrum is (over a wide frequency range) dominated by the static contribution and the integrated power is almost identical to the one of the spectrum of the average bunch profile (cf. Appendix C.3). For the same reason, when there is spectral overlap, the mixed term can easily dominate



(a) “Regular” Bursting Signal

(b) Average Phase Space

(c) Average CSR Spectrum


(d) Modulations (difference from average) at $t = 0.64 T_s$ (e) Modulations (difference from average) at $t = 0.68 T_s$ (f) Modulations (difference from average) at $t = 0.80 T_s$

Figure 10.8.: The “regular” bursting regime is characterized by fast bursts with a constant period that is shorter than T_s [cf. (a)]. In the given example, the modulations are small in comparison to the average charge distribution (b). The average CSR spectrum is displayed in (c). To make the changes visible, (d) to (f) show instantaneous differences from the average values (“modulations”), at the points in time marked by black lines in (a). The situation at the lowest integrated CSR Power is displayed in (d). With respect to (a), in (e) the structures in the phase space are rotated a bit, in (f) the maximums and minimums change positions. The increase in the integrated CSR power can be identified with the changes just above 100 GHz. As the modulation is always present, its inverted length is mostly part of the average spectrum [cf. Eq (10.7)]. It is noteworthy that a narrow-band radiation detector e.g. at 70 GHz would measure a different phase for the CSR bursts: This frequency peaks in (d) and is below average in (f).

the modulation term in the power fluctuation. The longest kind of modulation typically manifests as bunch length fluctuation.

An example series of bursts with $f_{s,0} = 28.130$ kHz, $V_{RF} = 1.048$ MV, $I = 605$ μ A in an iso-magnetic ring⁷ with $f_{rev} = 8.582$ MHz ($R = 5.5597$ m) is displayed in Fig. 10.8. It shows the described behavior, especially the fluctuation of the integrated CSR power is anti-correlated with the bunch length. Narrow band detectors, however, might measure a different phase: In the given example, the emission at 70 GHz and 105 GHz is anti-correlated [see right panels in Figs. 10.8(d), (e), (f)]. Although not yet explicitly mentioned there, the same can be found in experimental data (e.g. in [Ste+17b]). Notice that the actual shape of the CSR spectrum depends on the impedance and especially the position of the peak is determined by shielding. Further notice that phase and paraphase appear in the linear charge density modulation [see Fig. 10.8 (d) and (f)]. These modulations obviously have the same form factor, so the difference in the total CSR spectrum is (only) due to the mixed term in Eq. (10.6). In the case, where the modulation has a fixed amplitude and a time-varying phase, it is even possible that

$$|\tilde{\zeta}(f, t)|^2 = \left\langle |\tilde{\zeta}(f, t)|^2 \right\rangle \quad \forall t, \quad (10.8)$$

which means that there is no time-varying contribution to the CSR spectrum from the modulation term. In the displayed example, this is not the case. The same should be true for all CSR impedances: The real part of the CSR impedance introduces an asymmetry also for the substructures. Figure 10.8 showed that phase and paraphase are present; the charge profiles in between, however, differ not only in phase but also in magnitude. So Eq. (10.8) is not fulfilled, and – even if the mixed term vanished –

$$\frac{\partial F(f, t)}{\partial t} \neq 0. \quad (10.9)$$

An analysis of the frequency-dependence of the fluctuation including an outlook on remaining options on how to get stable CSR from substructures of an electron bunch is given in Sec. 10.4.


10.2.2. Sawtooth Bursting Regime

The “sawtooth” bursting regime was first observed in the SLAC damping rings [Kre+93]. Here at high beam intensities the bunch length started to oscillate in a sawtooth like pattern. With changing current, there might even be different shapes of sawteeth.

An example for a typical bursting cycle is shown in Fig. 10.9. It was simulated for KARA settings with $f_s = 7.7$ kHz, $V_{RF} = 1.5$ MV, and $I = 735$ μ A. The cycle starts when the bunch length is decreasing. As for the regular bursts, this leads to an increasing CSR power. For the sawtooth bursting regime, however, the relation $\langle \varrho \rangle(\tau) \gg |\zeta(\tau)|$ is no longer valid. So, although the fluctuation of the bunch length is very significant, the power fluctuation is not dominated by this process. Instead, the fact that the amplitude of the substructures (in phase space) varies becomes very important. Between the bursts the amplitude decreases. Depending on the time span between the burst, the substructures might almost vanish.

While the bunch length is still decreasing, substructures again start to grow. They do so until their potential well blows up the bunch. In the process, also the substructures become more filamented. Those complex structures then emit more CSR at higher frequencies, leading to the maximum in the integrated CSR power. Now, these structures start to vanish, leading to a relatively smooth, blown-up electron bunch. It will damp until the effect of

⁷The settings correspond to $f_{s,0} = 8.9$ kHz and $I = 191$ μ A for KARA.



(a) Typical “sawtooth” signal

(b) Zoom into first burst shown in (a)




(c) Situation at $t = 12.5 T_s$ (d) Situation at $t = 18.4 T_s$ (e) Situation at $t = 21.5 T_s$

Figure 10.9.: The “sawtooth” bursting regime is dominated by big fluctuations of bunch length and CSR power. A typical time-domain signal is shown in (a). The zoom-in (b) shows black lines at the times where more details are presented. In the beginning (c), the charge distribution is rather smooth. The position axis has the same range but different units are given to ease the comparison to the synchrotron radiation spectrum. With time, the bunch profile is slowly shortened due to synchrotron radiation damping, leading to increased radiation in the low-frequency region [cf. peak at 40 GHz in (c) and (d)]. At the onset of a burst (d), prominent structures form, increasing the spectrum in the mid-frequency range (150 GHz). Afterwards (e), the charge distribution almost explodes, also a multitude of fine structures forms. The increased length leads to less radiation in the low-frequency range, but the fine structure increases coherent amplification for the highest displayed frequencies. However, the wake-field cannot sustain the substructures. As a consequence, the charge density starts to damp again – until it becomes small enough to start a new cycle.

the instability-driving frequencies is high enough to trigger the formation of substructures in a new bursting cycle. For the given example, the damping time is a bit more than half τ_d , since the driving force of the wake potential is always present⁸.

It is noteworthy that there is something like a critical length, at which the substructures grow again. This lengths can be estimated by Eq. (3.24) or Eq. (3.27). As it increases with current, it effectively leads to a lengthening of the average bunch length. To be specific, it is not a fixed length but rather has to do with the bunch spectrum which might still feature remains of the substructures from the last burst. The seed of the new cycle is hard to predict even in this classical model neglecting the random nature of quantum particles. Thus, the process has components that can be considered deterministic chaos. Also, as the potential well acts in energy direction but is caused by a bunch profile along the position axis, a lag is introduced due to the synchrotron motion: The energy spread might still shrink due to radiation damping although the charge density being rotated by a fraction of a synchrotron period is already compressed enough to trigger a new blow-up due to the potential well. An example of experimental data showing this temporal mismatch between the rise of the energy spread and the CSR can be found in [Keh+17a].

10.3. Contribution of Additional Impedances

The effect of additional impedances is studied just briefly by adding the additional impedances (collimator and resistive wall) as introduced in Sec. 3.2 directly to the CSR impedance Z_{PP} . As before, shielding is modeled by parallel plates with a distance of $g = 32$ mm. This treatment of course introduces qualitative inaccuracies, e.g. the resistive wall for straight sections is not handled separately without CSR, and even the collimator is not localized but handled as an average contribution once around the ring. This approach is justified, as one synchrotron period is long compared to one revolution.

Since the focus of this thesis is on bursting CSR, in the cases relevant here, CSR impedance will always be most important. For any additional impedances, they typically are almost negligible compared to Z_{PP} , so that over a wide frequency range $\sum Z_i \approx Z_{PP}$, where the Z_i are the impedances including Z_{PP} . This relation implies that the additional impedances might influence the beam conditions but will not totally change it. Hence, a qualitative study is done to get a general idea of the influences of the different additional impedances⁹.

The resistive wall impedance describes a round ($r = 16$ mm) pipe, made of aluminum ($\sigma_c = 3.77 \times 10^7$ S and $\mu_r = 1 + 2.1 \times 10^{-5}$). The collimator is taken to be infinitesimally short with an inner radius of $r_{\text{Coll}} = 2$ mm. The outer radius is the same as for the beam pipe. Even though this collimator radius might be unrealistically small, it is $Z_{\text{Coll}}(f) \gtrsim Z_{PP}(f)$ only where the shielding suppresses a large fraction of the CSR ($f \lesssim 74$ GHz).

As the collimator is modeled to have a purely resistive impedance, its effect is that the bunch is leaning forward a bit further. This lengthens the bunch but also enhances the higher frequency components in the form factor. The resistive wall impedance has a resistive and a capacitive (so: shortening) component. Applying only Z_{RW} , makes the bunch lean forward and, for the chosen parameter set, in total also slightly lengthens it. One example for inductive impedances would be ferrite rings. They are discussed e.g. in [Ng06, p. 71] as a possibility to lengthen a bunch and thus suppress microwave instabilities. As this approach would currently counteract the goal to produce CSR, it is not considered here. Notice that for any possible bunch deformation, enhancing one frequency typically means suppression of other frequencies. The reason for this is that the form factor is normalized $\int \tilde{\rho}(f) df = \sqrt{2\pi}$.

⁸The length damps twice as fast as the single particle position.

⁹There exist studies including more elaborated impedance models. Those studies focus on the instability threshold (e.g. [Ban16]).

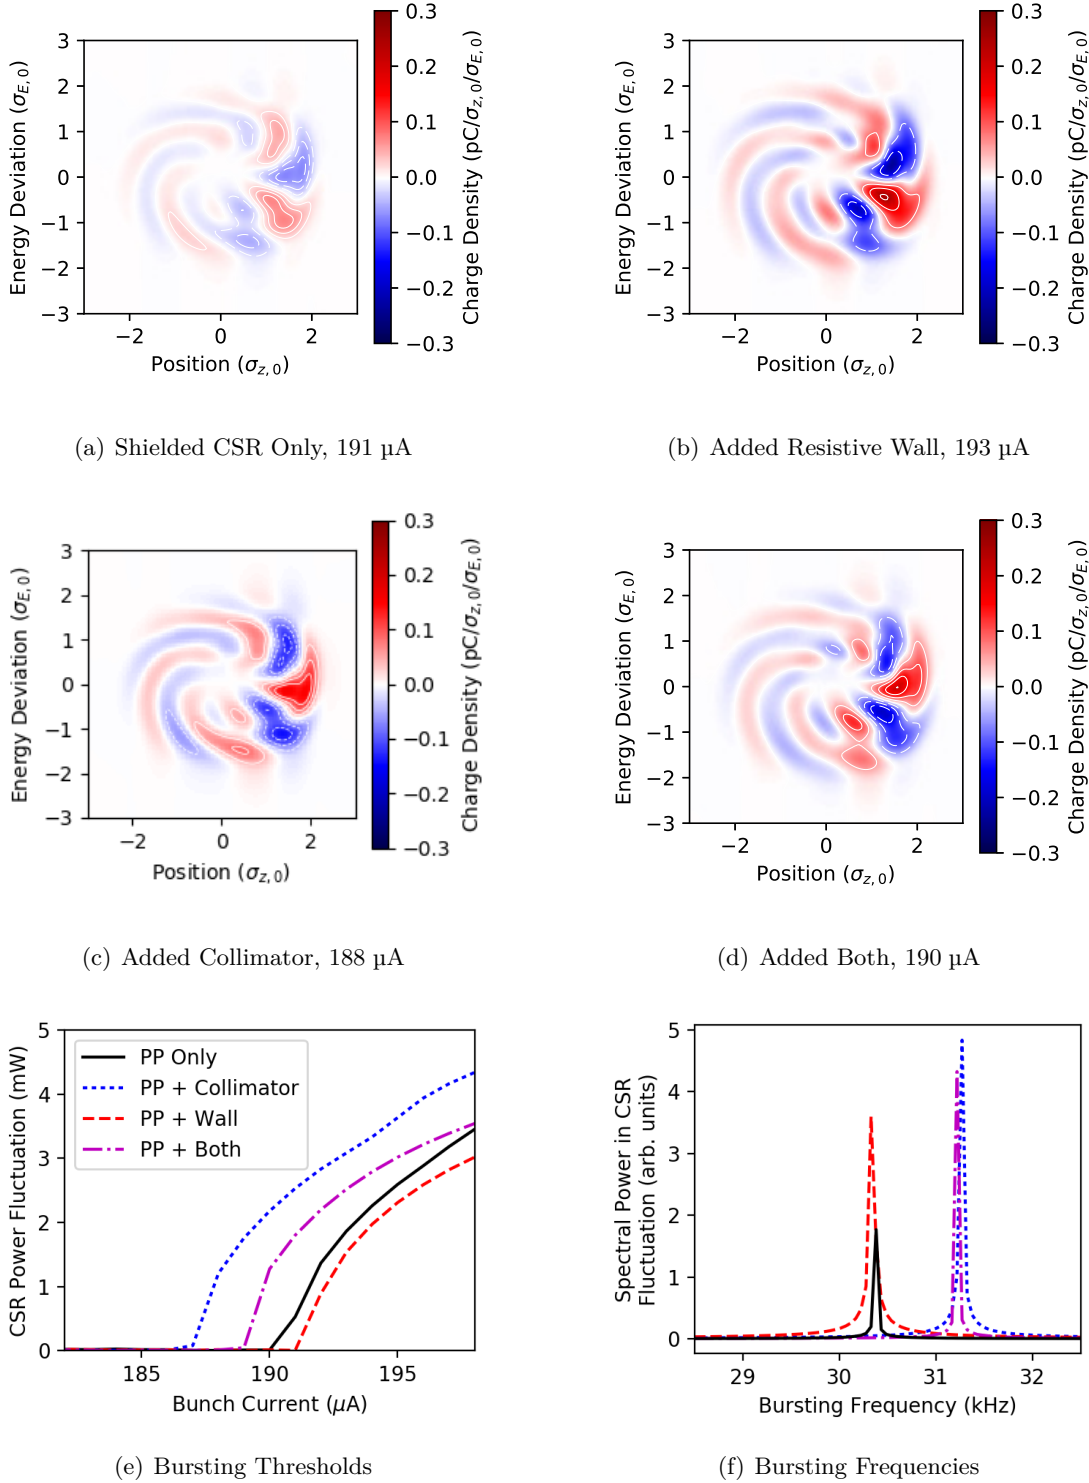


Figure 10.10.: (a) to (d) show the substructures on the charge distribution slightly above the micro-bunching threshold for different impedances added to the parallel plates CSR model. The amplitude of the structures is the most significant difference. To keep it at a level that still allows to see the similarities, individual currents but similar phases have been chosen for the different impedances. This different modulation is linked to the power fluctuation (e) at the chosen current. The lines connect the CSR power fluctuation for simulation runs performed for integer value μA currents. The collimator impedance lowers the micro-bunching threshold, whereas the resistive wall has the opposite effect. In the spectrum of the bursts (f), the collimator influences the first occurring bursting frequency to a higher value, the resistive wall does not seem to have a significant effect. The frequencies can differ despite identically shaped substructures, as the impedances also introduce a synchrotron tune shift.

Results of the four simulation runs (CSR impedance only, CSR impedance plus resistive wall, CSR impedance plus collimator, and CSR impedance plus resistive wall and collimator) are displayed in Fig. 10.10. As expected, the overall structure stays almost unchanged for all cases, only the amount of charge forming the substructure is changed¹⁰. This modified modulation-strength correlates with a shift in the instability threshold. For the collimator impedance, for example, it is an increase in modulation-strength and the instability threshold moved to lower currents. Although the effect of only the resistive wall impedance would deform the bunch in a way that the form factor would be enhanced at high frequencies – something that should benefit the rise of the instability – it accomplishes the opposite.

A possible explanation is found in the shape of the impedance. While the collimator impedance is just constant, the resistive wall impedance is $Z_{RW} \propto \sqrt{f}$. This shifts prominent features of the total impedance $Z = Z_{PP} + Z_{RW}$ like the zero-crossing of the imaginary part or the peak of the real part to higher frequencies. Thus the resonance frequency (band) of the total impedance should also move to higher frequencies. To quantify and in the end possibly verify this explanation, more detailed studies of the impedances have to be undertaken.

Experimental studies using a scraper show a comparable result to the simulation including the collimator for the shift of the bursting frequency (see Sec. 11.2). For a full comparison, Appendix C.4 gives spectrograms for the cases studied here, and also the average CSR power.

10.4. Wavelength Dependency of CSR Power Fluctuation

In Sec. 10.2.1, it has been shown that it is mathematically possible to have non time-varying emission from a beam that has substructures. The required conditions are a fixed amplitude of the form factor of the substructures and a vanishing mixed term in Eq. (10.6). The latter can be the case when substructures are very short when compared to the shape of the average bunch profile. As only the regular bursting has persistent substructures that do not vary even on millisecond timescales, it is conditions closest to the described CSR emission from short substructures. However, as the modulation is asymmetric¹¹, they do not only change in phase but also in amplitude. Thus the emitted power for single frequencies is not constant, but shows a small oscillation.

It might be possible to counter-act the asymmetry caused by the real part of the impedance, and by that to minimize the power fluctuation. Possibilities to achieve this include modification of the higher orders of the momentum compaction factor α_c , modulation of the accelerating voltage V_{RF} , and usage of additional impedances. This step, on the other hand, will also counter-act the formation of the substructures. Still, there might be an optimal point for the desired conditions. A deeper, more systematic analysis of impedances and how they influence the bunch shape will be in [Sch18b]. For this thesis, the search for such conditions is beyond scope, for experiments using synchrotron radiation, however, it might be sufficient to have limited fluctuation in the emitted CSR only for a specific range of wavelengths.

Figure 10.11 shows example bunch profiles in the regular bursting regime, as well as emitted CSR spectra¹². The comparison shows a similar trend for all for all three examples. For example, it is possible to distinguish areas where static, modulation and the mixed term dominate the form factor. Another common feature is the existence of a side-peak at $f \approx 100$ GHz. There is a link to the PP CSR impedance which is $\Im[Z_{PP}(f <$

¹⁰To be precise, there are also slight differences that can be explained by a different phase of the rotation.

¹¹This lies in the very nature of the mechanism that creates them, see Sec. 10.5.2.

¹²For temporal evolution see Sec. 10.2.1.




(a) $f_{s,0} = 7.7$ kHz, $V_{RF} = 1.5$ MV, $\sigma_{z,0} = 2.28$ ps(b) $f_{s,0} = 8.9$ kHz, $V_{RF} = 1.048$ MV, $\sigma_{z,0} = 3.78$ ps(c) $f_{s,0} = 14$ kHz, $V_{RF} = 1.396$ MV, $\sigma_{z,0} = 4.46$ ps

Figure 10.11.: Average bunch profile right at the micro-bunching threshold and an example modulation in the regular bursting (left), as well as CSR spectrum (right) for different possible machine settings. CSR below the bursting threshold is displayed in black and in the regular bursting regime in blue and red. As the regular bursting regime spans a smaller current range for those settings, the current values picked for (b) are closer to each other. For bursting CSR, bands are drawn between the minimum and maximum emission. The average CSR increases with current; but as fluctuation also increases, the minimum instantaneous values can decrease. The amplitude of the fluctuation of CSR emission is also dependent on the emitted wavelength. In (a), for constant current bursting CSR emission above 800 GHz stays in the same order of magnitude but still oscillates by up to a factor of four. Also note that the shorter natural bunch length seems to increase not only the static CSR emission but also the substructure induced emission at high frequencies. The (almost static) side-peak at $f \approx 100$ GHz might be linked to the structure PP CSR impedance (it is also present in the form factor, cf. Appendix C.3). Values below 10^{-10} mW/GHz can be considered a numerical noise floor.




Figure 10.12.: CSR Fluctuation Averaging Out

The two adjacent frequency bands (colored lines) have been chosen in a way so that their emitted CSR power shows similarly shaped bursts with almost identical fluctuation amplitude with a relative phase of π . So for their sum (black curve), the amplitude of the oscillation is much smaller. ($f_{s,0} = 7.7$ kHz, $V_{RF} = 1.5$ MV, $I = 80$ μ A)

100 GHz)] $< 0 \Omega < \Im[Z_{PP}(f > 100 \text{ GHz})]$, and seems to cause a peak in the form-factor (cf. Appendix C.3). When it comes to the emitted spectrum, the local maximum at $\Re[Z_{PP}(f \approx 140 \text{ GHz})]$ slightly shifts this peak in the form factor. In all three examples, the emission at low frequencies is relatively stable. Notice that – in absolute values – the power fluctuation in this frequency range still dominates the overall power fluctuation [cf. Fig. 10.8, which corresponds to profile and blue curve in 10.11(b), and Appendix C.5]. For higher frequencies, the relative power fluctuation increases. The average bunch profile does not change to an extent that can explain the minimum power – which is orders of magnitude below the value for stable emission. So, these fluctuations can be attributed to the mixed term in Eq. (10.6). This term can have a negative sign and decrease the emitted CSR. In total, the amplitude of the power fluctuation is dependent on the frequency but there is no constant emission due to the substructures observed at any frequency.

Alternatively, a spectral range can be chosen so that the integrated power is still (almost) constant

$$\partial \left[\int_{f_1}^{f_2} P(f, t) df \right] / \partial t \approx 0 \quad (10.10)$$

for well chosen f_1, f_2 . An example is shown in Fig. 10.12. Here, the emitted CSR power in two adjacent frequency bands shows similarly shaped bursts with almost identical amplitude that are phase shifted by π . The emitted power in the combined frequency range is still oscillating – but with a much smaller amplitude and a doubled frequency. Using a detector with a sensitivity that is different for different wavelength, it is probably possible to find settings where the total measured power is constant.

In an experimental setup, a radiation detector usually has a fixed range of sensitivity. It is almost certain that this specific range is not perfectly matched to the range of minimized power fluctuation just by chance. So the minimal fluctuation in the measured power will not be zero. Still, local minimums with $\sigma_P(I) > 0$ are to be expected because of averaging out power oscillations with different phasis. This fact has to be kept in mind when analyzing data. Such a behavior for the standard deviation detected CSR signal has already been seen, e.g. in [Bro14, Fig. 5.6]. When just seeing the data, one tempts to explain these findings by the shift spectral range that is fluctuating most. As seen here, situation is not that simple. Without further information, none of the two reasons can be excluded.

For the total power signal in a given frequency range, averaging out of different signs

of course cannot happen. Consequently, one should also consider this information to understand what is the reason for oscillation minimums. An example showing a curve comparable¹³ to the ones for the standard deviation referenced here, can be found in [Jud13, Fig. 5.16 (a)]. Local minimums of the average power showing just a slight decrease – found in measurement and simulation – are presented in [Ste+17a]. In [Fei+11], there are no local extrema.

The difference between a total oscillation minimum and a minimum just of the integrated signal can also be identified in the bursting spectrogram: When the effect is due to the frequency-shift of the mixed term, the main bursting frequency only undergoes its default current-dependent frequency-shift. In the example given in Fig. 10.12, for a small current range the dominant frequency is a higher harmonic of that main bursting frequency. Also notice that the 1σ band in Fig. 10.4(a) ($f_1 = 0, f_2 = \infty$) does not show clear local minimums and maximums at all. This implies that the choice of the frequency range is crucial for the shape of $\sigma_P(I)$.

10.5. Further Interpretation of the Results

10.5.1. Fast Bursting Frequencies

The fast bursting frequency which is visible during regular bursting is tightly connected to the substructures on the bunch profile. As they rotate in phase space, they can trigger modulations of the emitted light intensity with a multiple of f_s . This link is pretty intuitive and [Pod99] already attributed substructures to the eigenmodes of synchrotron motion. In this thesis, the link was already mentioned shortly after introduction of the eigenmodes in Sec. 3.3.2.

When there is no significant potential well distortion, [Rou+14a] predict the relation

$$\varphi = \frac{2\pi f_{s,0}}{f_b} \quad (10.11)$$

for the angle between two substructures. When there are substructures all around the center of the rotation, this is equivalent to


$$f_b = m \times f_{s,0}, \quad m \in \mathbb{N}. \quad (10.12)$$

They attribute deviating observations to a changing shape of the structure during the rotation. In this view, the structures are material: They (and the particles they are made of) follow the synchrotron motion. When potential well distortion plays a role, the changes could then be expressed by the incoherent tune spread. The tune spread also is of importance for the WMB, which is dependent on a small incoherent tune spread and thus can be Landau-damped easily [BCS10].


Figure 10.13 images the charge modulation at different points in time and paths of example particles. Unlike to the case of the static potential well distortion in Sec. 10.1.2, $\psi(q, p)$ is no longer constant in time. For 10.13(a), and 10.13(b), the settings describing the accelerator are the same ($f_{s,0} = 5756$ Hz, $V_{RF} = 1.3$ MV), but the beam current differs ($I = 30$ μ A vs. $I = 45$ μ A). For the first, two substructures can be observed that do one full rotation in $t_m \approx 1.15 T_s$.

In the same time, the fastest example particle also did one full revolution, the slowest particle, however, only did 90% ($0.77 f_{s,0} \lesssim f_{s,i} \lesssim 0.87 f_{s,0}$). The bursting frequency is $f_b = 9670 \pm 10$ Hz, which makes $f_b/f_{s,0} = 1.680$. For the second, one full rotation of the


¹³The author, however, attributes the local minimums to detector saturation.



(a) f05864, WMB



(b) f05864, Regular



(c) f05493, Regular

Figure 10.13.: Movement of substructures in phase space here is compared to the path of example particles (black). The point in time picked for the right column is chosen so that the sub-structures are identical to their initial configuration. In (a) structures and all particles have almost the same frequency, $1.15 T_s$ makes approximately one revolution for both, structures and particles. In (b) the outer structures move with about $f_{s,0}$, the inner ones with a lower frequency: At $1.05 T_s$, the modulation is identical to the initial one. However, the outer structures did one full revolution but the inner structures just did two thirds. The latter in between also formed a structure similar to the small, central isle at $0 T_s$. The inner particles stay in their minimum for the whole displayed time, while the outermost particles lag behind their structure. Also in (c) structures and particles move with different frequencies. At $1.35 T_s$, all particles are inside the minimum which previously hosted the second particle from the left.

$m = 2$ structures takes $t_m \approx 1 T_s$, but in the center there are $m = 3$ structures which again are at their starting position at $t_m \approx 1.5 T_s$. Here, it is $f_b/f_{s,0} = 1.895$ ($f_b = 10910 \pm 10$ Hz). For the example particles, it is $0.65 f_{s,0} \lesssim f_{s,i} \lesssim 0.85 f_{s,0}$. So, the outer structures rotate faster than the individual particles.

The example displayed in Fig. 10.13(c) was simulated using $f_{s,0} = 7.2$ kHz, $V_{RF} = 1.396$ MV, and $I = 70$ mA, here $m = 3$ structures are visible with a rotation period of $t \approx 1.3 T_s$. With a fast bursting frequency of $f_b = 16150 \pm 50$ Hz, $f_b/f_{s,0} = 2.243$, and $0.6 f_{s,0} \lesssim f_{s,i} \lesssim 0.85 f_{s,0}$. Another example, corresponding to $f_{s,0} = 8.9$ kHz, $V_{RF} = 1.048$ MV was already shown in Fig. 10.8. There, it is $f_b/f_{s,0} = 3.395$, with $m = 4$ structures at a rotation period of $t_m \approx 1.2 T_s$, and in the center $m = 5$ structures with a rotation period of $t_m \approx 1.46 T_s$.

Table 10.1.: Overview of frequencies (all in kHz) of unperturbed synchrotron motion $f_{s,0}$, of spread of example perturbed incoherent synchrotron motion $f_{s,i}$, of bursting CSR f_b , and rotation of the substructures $f_m = 1/t_m$. (The accuracy of the frequencies is $0.01 \times f_{s,0}$.)

	m	$f_{s,0}$	$f_{s,i}$	f_b	f_m
f05864, WMB	2	5.8	4.4 : 5.0	9.7	5.0
f05864, reg	2		3.7 : 4.9	10.9	5.8
	3				3.8
f05993, reg	3	7.2	4.3 : 6.1	16.2	5.5
f05135, reg	4	8.9	6.0 : 8.2	30.0	7.4
	5				6.1

These examples (put together in Table 10.1) suggest that in regular bursting and WMB¹⁴, the relation expressed in Eq. (10.12) should be modified to

$$f_b = m \times f_m \quad (10.13)$$

with $f_m = 1/t_m \leq f_{s,0}$. The difference cannot be explained solely by the incoherent tune shift, as in many cases there is not just one angle but the number of substructures in the center of the phase space may be different from the number further outside. One should also note that the substructures move with distinct frequencies, whereas the tune spread of the individual particles results in a continuous frequency band. However, even if not strictly proportional, with increasing number of substructures also the ratio $f_b/f_{s,0}$ increases.

An explanation could be that the impedance drives a constant modulation length (or a range of lengths) that can resonate with the synchrotron motion eigenmodes. Key to this understanding is the zero-crossing of the imaginary part or the CSR impedance Z_{PP} (cf. Sec 3.2.4): Short structures sample the inductive part of the impedance, and thus become longer, while the opposite is true for long structures. When the bunch is longer (e.g. because of $f_{s,0} \propto \sigma_{z,0}$), higher modes are excited. Due to tune shift and tune spread, also the single frequency $f_{s,0}$ becomes a frequency band. So there are no discrete steps between the particle's oscillation modes. Also note that the example particles do not co-propagate with the substructures but can move with a different $\langle J(\vec{r}) \rangle$. So the structures are not material but are formed by the equilibrium of a dynamic process, as described in the next section.

¹⁴Brief checks showed Eq. (10.11) works best during the bunch-shrinking phase of the sawtooth bursting cycle.

10.5.2. Growth and Decay of Substructures

One of the results of the previous section is, that there is not only a tune-spread (cf. Sec. 10.1.2) but particles and substructures take different times for one full revolution in phase space. Here, an approach is undertaken to link the local minimums of the charge distribution to the areas where J is small, and vice versa. This description does not exclude a mixture of azimuthal and radial eigenmodes of the synchrotron oscillation (see Sec. 3.3.2) but might be more intuitive than a purely mathematical description.

To study this processes including the mode mixing, a possibility would be to transform the coordinates using a non-linear scale that shadows the average potential well distortion in the position coordinate. Doing so would make the substructures as well as the bunch shape more symmetric, so that in equilibrium cases the paths of the particles will again be round. Another option is to go into a rotating, non-linear reference frame where $\langle J \rangle(q, p) \equiv 0$, so that only the time-varying part of the local instantaneous synchrotron motion is non-zero. The substructures in phase space and their formation also show noteworthy similarities to the formation of the spiral arms of galaxies that can be explained by a density wave theory [LS64]. As the potential well only acts along the energy axis of the phase space but gravitational forces in a disk act in two (spacial) directions, the theory is not directly applicable. The similarity, however, hints to the possibility to find a suitable explanation for the micro-bunching in the field of fluid dynamics, where also the density wave theory originates.

Here, the coordinate system is still normalized using the unperturbed case. Still, it can be made plausible that the substructures are reproducing themselves when moving to the bunch head, and that the substructures form because of the accumulation of particles where J is smaller. For the static case and the overall charge distribution, Eq. (10.3) already stated that particles accumulate in regions with small J . Here, the situation is no longer static. $\partial J(q, p, t)/\partial t \neq 0$, and the forming substructures are constantly moving. Figure 10.14 shows substructures in the phase space, simulated for $f_{s,0} = 7.2$ kHz, $V_{RF} = 1.396$ MV, and $I = 70$ mA at two different points in time. For this example, as well as for all other studied substructures, it can be observed that the substructures have greater local charge density when they are at the bunch head. This implies that they are growing when moving from the bunch tail to the head and shrinking when moving into the other direction¹⁵. To simplify the discussion, deviating from the original definition in Eq. (3.1) the synchronous phase *including CSR* is now defined to be $q = 0$, so the motion is again centered

$$\langle J \rangle(0, 0) \stackrel{!}{=} 0, \quad (10.14a)$$

and particles with negative energy offset move to the bunch front and vice versa

$$\text{sign}[\partial q(p)/\partial t] = -\text{sign}(p). \quad (10.14b)$$

First, it can be seen that the modulation of the longitudinal profile is mostly due to the substructures at $\text{sign}(q) = \text{sign}(p)$. The substructures in the quadrants $\text{sign}(q) = -\text{sign}(p)$ are almost parallel to the position axis and thus do not contribute much. Further, Figs. 10.14(b) and 10.14(d), show that the modulation of the wake potential due to the modulation of the bunch profile typically have the same sign. For $q < 0$, this means that areas of low charge density lose less energy per time (lower J), areas of high charge density loose more (higher J). In $q < 0, p < 0$ the modulation of the wake and of the charge density co-propagate, meaning that particles will accumulate. For $q < 0, p > 0$ the movement of

¹⁵Notice that all examples have $\alpha_c > 0$, so particles with negative energy deviation move to the bunch front and vice versa, $\text{sign}(\partial z(E)/\partial t) = -\text{sign}(E)$. For $\alpha_c < 0$, many of the following arguments reverse.

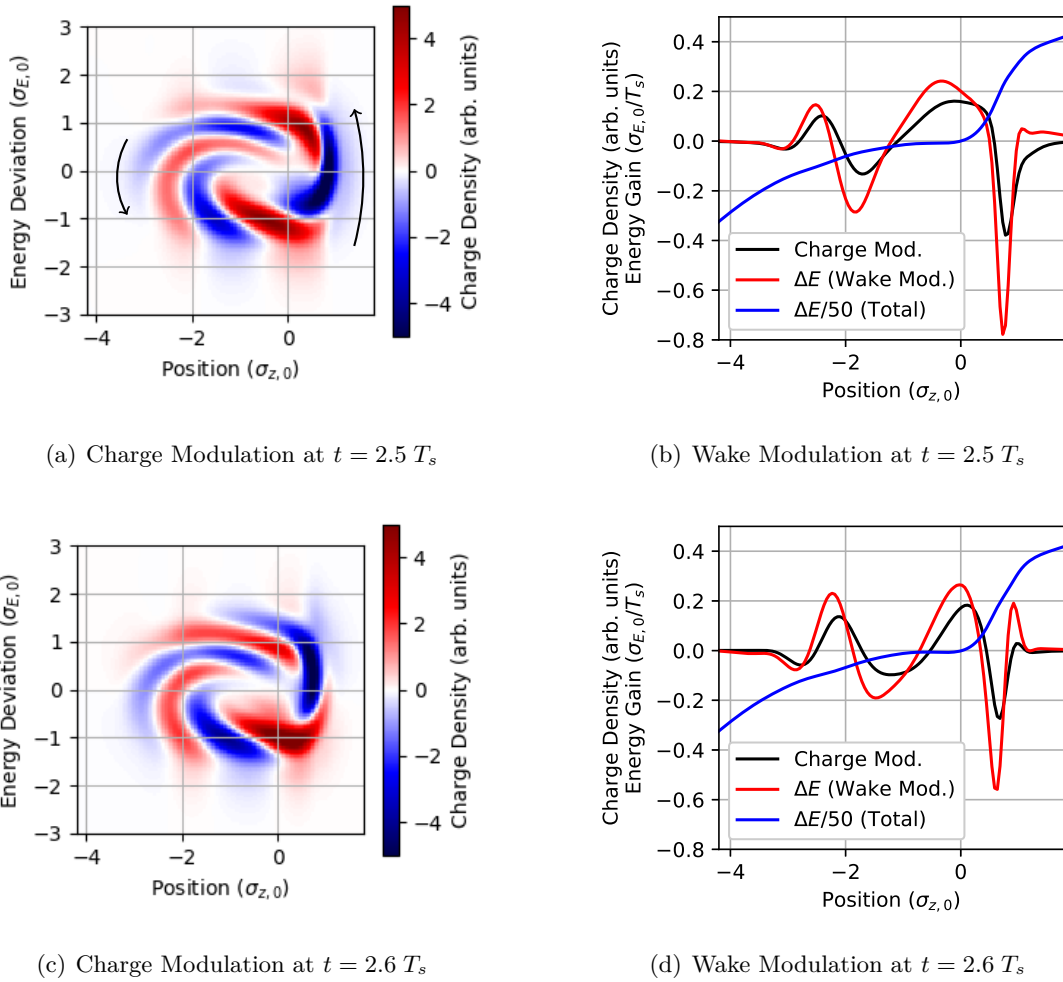


Figure 10.14.: Formation of Substructures in the Longitudinal Phase Space

In contrast to the original definition in Eq. (3.1), here, the coordinate system has been shifted so that $p = q = 0$ for the synchronous particle (center of rotation) *in the perturbed case*. Substructures in the lower energy half of the phase-space are parallel to the energy axis, while in the upper half they are parallel to the spatial axis. So the lower energy substructures contribute dominantly to the charge modulation [black in (b), (d)]. This results in the wake potential (blue, divided by 50 to match the scale) being modulated the same way (red). When wake modulation and wake point into opposing directions, particles will accumulate because of the slow-down. When this wake modulation is co-propagated by a charge modulation, the substructure self-amplifies.

the substructures is opposing the movement of the wake-modulation. As they are oriented along the energy axis, the structures wash out. In the bunch front ($q > 0$), the total energy gain due to the static wake modulation results in a higher instantaneous synchrotron motion which is very directed ($\partial p/\partial t > \partial q/\partial t$). So although the structures are principally damped, they can survive these quadrants. In fact,

$$\left| \frac{\partial p(-|q|)}{\partial t} \right| < \left| \frac{\partial p(|q|)}{\partial t} \right|, \quad (10.15)$$

to an extent that electrons reside at the bunch tail 50 % longer than in the bunch head, so the density of the substructures decreases (and afterwards increases) mostly in the quadrants $q < 0$. Also note that the wake modulation slightly lags behind the charge modulation, thus the amplifying effect is also (partly) behind the structure. This configuration seems to counter-act the movement of the substructures, making them slower than the synchrotron motion.

The given qualitative investigation can already make plausible how the substructures in phase space reproduce. The explanation also solves the argumentative problem that the substructures rotate at discrete frequencies but the individual electrons have a continuous tune-spread. It does so by drawing a self-consistent picture that uses and the fact that the structures propagate with a frequency that is different from the synchrotron frequency of the individual electrons.

11. Explicit Comparison of Simulation and Measurement

“The dream has no way at all of expressing the alternative ‘either ... or’. It usually takes up the two options into one context as if they had equal rights.”

Sigmund Freud

To verify the credibility of the model used in the previous chapter, this chapter shortly compares example measurement results to simulation data. As the simulation method is robust against changing numerical parameters (see Chap. 9), the most error-prone part is to find the correct physical parameters for the simulation. The comparison will start with detailed views and gradually broaden the scope to more expressible quantities: First (Sec. 11.1), the bursting threshold is considered. From an abstract point of view, it boils down all accelerator properties like the beam optics, the particle energy, and the total impedance to just one parameter, the instability threshold current I_{th} . Section 11.2 investigates how well a modification of the geometric impedance is matched using a simple model. When only the change is considered, this benchmark can be seen as a differential method, reducing the influence of the other accelerator properties. The third comparison (Sec. 11.3) looks at bursting spectrograms. Those again contain information on the whole beam dynamics, but in contrast to the bursting threshold alone, they are very feature-rich, so they also allow to argue whether there is a qualitative agreement. At last (Sec. 11.4), the current-dependent bunch lengthening is looked at. This quantity is less complex but very well studied and easily measured. It allows to argue qualitatively as well as quantitatively.

11.1. Bursting Thresholds

As shown in Sec. 9.5, the numerical method is very robust when it comes to the prediction of the instability threshold. However, obtaining the physical parameters for a comparison with an experiment is not a simple task. Figure 11.1 shows measurement data taken at constant machine settings compared to two simulation runs that were both performed based on two unperturbed synchrotron frequencies $f_{s,0}$ measured in the same time period as the CSR data was taken: 8.9 kHz and 9.4 kHz (see also Appendix C.6). It is evident that this significant difference also has a big impact on the threshold determined by the simulation. The simulated instability thresholds of $I_{\text{th},8.9} \approx 190 \mu\text{A}$ and $I_{\text{th},9.5} \approx 215 \mu\text{A}$ frame the experimental value of $I_{\text{th,exp}} \approx 205 \mu\text{A}$. This calls for a systematic study that is not bound to results of single simulations or experiments.

To be more robust against this kind of statistical influences, the comparison should include a model that allows to average the high-uncertainty measurement values. As introduced in Sec. 3.3.2, [BCS10] introduced a scaling law for the instability threshold [Eq. (3.27)]. It was




Figure 11.1.: CSR Power Fluctuation: Measurement and Simulation

Simulation A was performed using a value for $f_{s,0}$ that was obtained at 2 mA using beam position monitors (8.9 kHz), for Simulation B the value was obtained from a CSR power spectrogram (9.4 kHz at 0.19 mA). The curves have been scaled vertically to match the measurement data. Except for the baseline in the measurement, the shapes agree very well. However, the different synchrotron frequency of course leads to different bursting thresholds. (Measurement data courtesy of Miriam Brosi.)

derived based on simulation data using the same impedance model utilized in this thesis. When comparing experimentally obtained bursting thresholds to this prediction, an excellent agreement is found [Bro+16b]. The differences especially include that measurements show a slightly lower instability threshold current. This might be explained by contributions of additional impedances, as studied in Sec. 10.3. The influence of noise on the accelerating voltage, on the other hand, can be considered as rather low [Sch18a]. A test of an example impedance will be in the next section. In general, the simple model already does a good job.

11.2. Influence of a Geometrical Impedance

In Sec. 10.3, the influence of impedances has been studied that were added to Z_{PP} , the latter being the impedance describing the self-interaction due to CSR. To benchmark the according predictions of the simulation, it would be handy if also the impedance of the accelerator could be varied. However, most impedances are naturally constant and cannot be changed. For example, the discussed resistivity of the vacuum pipe (resistive wall impedance), might be influenced by changing the material's temperature, but this is typically not foreseen to do so. There is one relevant exception: The *scraper*, a pair of metal rods that can be moved close to the beam. Historically, the scraper has been introduced to reduce the beam current, but also influences on the beam dynamics have been investigated (e.g. [KST03; Mül+04; Jud13]). The same authors already suggested to use the collimator impedance for an estimated description of the scraper.

As the scraper can be moved in and out, a series of measurements giving sufficient statistics is experimentally easier to obtain for the bursting spectrum (cf. Fig. 10.10(f)) than the bursting threshold (cf. e.g. Fig. 10.10(e)). Figure 11.2 shows such a measurement. As in the simulation, the influence of the scraper shifts the fast bursting frequency to higher frequencies¹. When the scraper is removed, the frequency is restored.

In Sec. 10.3, the fast bursting frequency was changed by $f_{b, \text{Coll}}/f_{b,0} \approx 31.2 \text{ kHz}/30.4 \text{ kHz} = 2.6\%$, in the measurement presented here, the effect is much smaller ($f_{b, \text{Coll}}/f_{b,0} \approx$

¹It should be noted that in his experiment, [Jud13, Sec. 8.1.4] sees a shift in the opposite direction.




Figure 11.2.: The spectrogram shows the fast bursting frequency f_b for a number of data sets with decreasing current (lower current for higher dataset number). Datasets with the scraper moved in are highlighted in lighter color. Under the influence of a scraper f_b moves to higher frequencies. This is expected when approximating the scraper by a collimator (see Sec. 10.3), as suggested by e.g. [KST03; Mül+04; Jud13]. (Plot courtesy of Miriam Brosi.)

14.4 kHz/14.3 kHz = 0.6%). One factor probably is a different distance from the beam, with a simulated collimator opening of a radius of $r_{\text{Coll}} = 2$ mm, whereas the scraper opening was $r_{\text{scraper}} \approx 2.75$ mm. However, already [KST03] found the measured wake losses to be smaller than predicted by the used collimator model. In total, the model of the geometric impedance can give a qualitative estimate. For quantitative studies, a more refined model is needed.

11.3. Bursting Spectrogram

Due to the almost random nature of the seed of a burst² no burst looks exactly like the other in time domain (see Sec. 10.2.2). This makes it hard to directly compare the actual dynamics. The situation is different in the spectral domain: Here, structures become clear and easy to distinguish. Making use of this, bursting spectrograms are a standard tool when analyzing CSR data (see Sec. 4.2). As such spectrograms include information of the whole dynamics seen in the CSR emission, they are sometimes seen as a “fingerprint” [Jud13] for settings and impedance of the accelerator. Of course, as for the instability threshold, which can also be seen as a feature of the spectrogram, it is not trivial to correctly determine all physical parameters to have a simulation with exactly the same conditions as present in the experiment. However, the rich set of multiple features allows for better qualitative comparisons.

In Fig. 11.3 for the simulation the same parameters were used as for Simulation B in Sec. 11.1. There are slight differences in the resulting spectrogram, but the general structure matches quite well: The instability threshold can be identified with the point the isolated single frequency at $f \approx 35$ kHz starts to rise. The second harmonic of this feature is stronger in the

²Notice that the model does not include random processes. Considering the huge number of possible states, it may be considered deterministic chaos.




Figure 11.3.: Example for a measured (left) and a simulated (right) bursting spectrogram. Considering the simplicity of the model, the differences e.g. in the threshold current and in the frequencies are rather small. The general structure of the spectrograms matches well: There is an isolated finger pointing down (at $f \approx 35$ kHz) with the fingertip marking the instability threshold at ($I \approx 0.21$ mA). Fluctuations in the frequency range $f < 10$ kHz start for slightly higher currents. Because of them, also the finger broadens into a frequency comb. For the highest currents displayed here, there is a regime showing broad lines at multiples of approximately $f_{s,0}$. (First published in [Sch+17c].)


simulation. At about 0.24 mA fluctuations in the frequency range $f < 10$ kHz start to occur. They also cause a broadening of the “finger”. Where the low frequency is very narrow, they become visible as a frequency comb in the whole spectrum. For the highest currents displayed here, there is a regime showing broad lines at multiples of $f_{s,0}$. Considering the simplicity of the model, there are surprisingly detailed similarities. The slight differences between the spectrograms will be mostly due to a mismatch of physical parameters such as $f_{s,0}$ and τ_d between simulation and measurement. For example, notice the differences in the lower frequencies that do also result in different occurrences of frequency combs. This hints that probably the damping time was assumed to low.

11.4. Bunch Lengthening


At storage rings, the bunch length is a well studied quantity (e.g. [CG90; Qin+01; Hil+11]) with pretty accessible diagnostics [Cha+13]. If $\sigma_z(I, t) \equiv \langle \sigma_z \rangle(I)$, the scaling law in Eq. (3.24) might hold. In the micro-bunching instability studied here, however, it can start to oscillate to an extent that governs the whole dynamics. Moreover, experimentally, the FWHM is advantageous over the RMS (see e.g. Sec. 4.1). To reflect this, instantaneous values for the FWHM bunch length $l(I, t)$ are put into a histogram (see Fig. 11.4) instead of looking at $\langle \sigma_z \rangle(I)$. This way also allows to compare more information, like the amplitude of the length oscillation.

It can be seen that the development of bunch lengths with current matches quite well between measurement and simulation. In the simulation, the bunches at highest currents (1 mA) are about 2 ps shorter, but at 0.2 mA, there is a slight discrepancy (0.5 ps) into the opposite direction. Also, the level of the fluctuation shows just small deviations: Due to noise on the measurement setup, the length does not reach one constant value even for smallest currents. For higher currents, on the measurement data, there is a small asymmetry in the distribution. This is because in the streak camera measurement the bunch profile is averaged over typically 500 μ s. So, faster fluctuation (and possible short-time changes) cannot be resolved. Another error source again lies in the choice of simulation parameters. One that was not named before is the calibration of the accelerating voltage, so that \hat{V} is over-estimated in the simulation and the simulated bunch lengths are thus shorter than the measured ones.

In total, the differences can be considered to be small. The model, which describes the electrons as a rigid line charge circulating between infinitely large parallel plates with no resistivity, reproduces bunch lengthening at KARA very well.



(a) Streak Camera Measurement



(b) Simulation using Inovesa

Figure 11.4.: Distributions of Measured and Simulated Bunch Lengths

(a) shows a histogram of bunch length measurement results, the data is identical to the one in Fig. 6.2, (b) displays simulation results in a comparable fashion. It is visible that the development of the bunch length matches quite well between measurement and simulation. The main difference is the level of the fluctuation at low currents: In the simulation, the length reaches a constant value, whereas this is not the case for the measurement. Besides noise on the measurement setup, there is also a data-evaluation artifact that influences the distribution especially for higher currents. In the streak camera measurement the bunch profile is averaged over typically $500 \mu\text{s}$. This means, that faster fluctuation cannot be resolved. As a result, the measured distribution of bunch lengths has an asymmetry which is not expected from the simulation. The mismatch in the bunch length at highest currents displayed here is probably due to a slight mismatch between the real accelerating voltage and the one used for the simulation.

12. Implications and Outlook

“There is more to reality than meets a normal eye.”

Willem Rebergen

This chapter reviews the insights gained in Chapter 10 from a perspective that is more the one of an experimentalist. It makes explicit how some of the findings have to be considered in measurements of the phase shift (Sec. 12.1) and the bursting threshold (Sec. 12.2). Further, it summarizes the possibilities to influence the bursting and hints how they can probably be evolved (Sec 12.3). At last, it suggests methods to verify the numerical result that the substructures in the phase space are not material but rather created as a dynamic equilibrium (Sec. 12.4).

12.1. Radiation Losses and Phase Shift¹

Measuring the energy in pulsed light is a challenge because of the lacking availability of properly calibrated detectors [Ste+17a]. This becomes even more severe when the spectrum has a wide frequency range spanning several orders of magnitude. However, at a synchrotron light source, there is potential to indirectly measure the energy loss by determining the energy that is regained due to \hat{V} .

An absolute phase information is hard to obtain, but phase differences are easily measured. In fact, it is common practice to track the shift of the coherent phase (e.g. [PS98; RBT11; Che+17]). This way, the so called *loss factor*

$$k(N) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \Re[Z(\omega)] |\tilde{\rho}(\omega)|^2 d\omega \quad (12.1)$$

is determined, a measure how increasing charge increases the total losses. There are two common approximations: For small phase variations $\phi \ll 1$, it is $\sin(\Psi) - \sin(\Psi_0) \approx \phi \cos \Psi_0$, yielding [RBT11]

$$k(N) = \frac{\Delta \langle \phi \rangle (N)}{eN} \hat{V}. \quad (12.2)$$

Also, when $\partial \rho / \partial Q = 0$, $k(N)$ is constant and $\phi \propto N$, as in [PS98]. Notice that dependent on the shape of Z , $k(N)$ can also be constant with changing $\rho(Q)$.

Recall that the synchronous phase $\Psi_{z,0}$, unlike $\langle \Psi \rangle$ as used here, is defined to be the phase of the particle that gains exactly the energy it loses [Eq. (2.32)]. Based on the definition, in Sec. 10.1.2 it was shown that the synchronous phase is not equivalent to the average particles position of a deformed bunch. In fact, phase shift and deformation are

¹The thoughts presented in this section were a motivation for [Sch+18b], which contains some of them.

coupled: Only when $\Re(Z) \neq 0$ – which also introduces an asymmetry – a changing bunch population can modify W_0 . Following this particle-centered definition of the synchronous phase, however, introduces two counter-intuitive features: As the distortion of the potential well – which is equivalent to the position-dependent losses – moves along with the electron bunch, under synchrotron oscillation the synchronous phase is not at a fixed location. Also, there might be more than one synchronous phase in a micro-bunched beam.

To summarize, changes in the average energy emitted per electron can be measured by tracking the shift of the center of mass position of an electron bunch. As the emission typically depends on the particles position in the bunch, this is not the synchronous phase as defined for a single particle. If it was possible to get a reference point the phase measurement would further allow absolute power measurements. Experimentally, this might be challenging, however, it is probably easier than calibrating detectors for the whole range of emitted frequencies.

12.2. Methods to Determine the Micro-Bunching Threshold

There are various strategies to experimentally determine the bursting threshold. The methods include the observation of practically all phenomenons discussed in Chap. 10. This section reviews those methods, also commenting on the possible influence of available detectors².

A first method, that does not require fast detectors, is to measure the average synchrotron radiation level. In [Jud13], the description reads that there is a quadratic increase with the current below the micro-bunching threshold, with a further increase at higher current due to the substructures. This observation, however, highly depends on the spectral sensitivity of the used detector. With a flat detector response, the integrated radiation level should see great influences of bunch deformation in the potential well distortion also before the instability threshold. The example in Sec. 10.1.4, showed a growth rate below the threshold that was not quadratic but still more than linear. Above the threshold, however, the synchrotron radiation power per electron even decreased. Thus it may be disputed whether a measurement of the average radiation level is a good indicator for the micro-bunching instability threshold: It might happen that there is no observable change at all.

The situation is different when explicitly using a narrow band (or short-wavelength) detector. Such a detector can be chosen to be sensitive at wavelengths of the size of the substructures but might not be influenced by the change of the average bunch shape (cf. Fig. 10.6). If the sensitivity curve is not flat the results might be biased by the occurrence of different substructures or by remaining visibility of the bunch deformation, but in total, the emerge of short-wavelengths components in the emitted CSR should be a very reliable indicator.

Alternatively, the evolution of the total loss rate might be determined by observing the coherent phase shift. As $\langle \Psi \rangle \propto P/I$, Fig. 10.7 suggests, there should be a turning-point at the micro-bunching threshold. This feature, however, has not been identified in any measurement, yet. One reason is that the total impedance budget can be dominated by other sources than CSR [Sch+18b].

When fast detectors are available, also the fluctuation of the CSR signal can be considered. In early experiments, this has been done by looking at time-domain traces [Byr+02], but it is also possible to use the standard deviation (examples in [Sch+13; Bro+15]). The idea behind is that a stable beam should also emit a stable synchrotron radiation. In a real accelerator, the detected intensity is influenced by motion of the radiation source point.

²This section might be seen as a revisit to [Sch+13; Sch13], now with a broader knowledge also on the theoretical background.

For example, a slight RF phase jitter introduces a coherent synchrotron oscillation which leads to a $1 \times f_{s,c}$ or $2 \times f_{s,c}$ modulation of the detected signal. The latter is the case if for example the amplitude but not the sign of the sinusoidal motion is of importance. As this method allows to select the frequencies of interest after the measurement, this problem can be overcome. Choosing a narrow band of CSR signal fluctuation frequencies can also reduce the noise. This elaborated version has been first used by [Bro14; Bro+16b]. Recalling that the relative power fluctuation in the regular bursting regime is very dependent on the wavelength of the emitted light, the shape of the fluctuation curve is of course very dependent on the detector: As shown in Sec. 10.4, narrow-band detectors possibly show a different evolution of the power fluctuation. Also notice that it is possible that the power fluctuation for the chosen range of photon wavelengths is insignificant. However, the situation with an averaged-out power fluctuation (see Fig. 10.12) can be considered an edge-case.

Often, the spectrum of the CSR signal fluctuation (*bursting spectrogram*) is taken as an indicator. To produce this, requires fast detector to be sensitive to fluctuations with time constants $t < 1/(N \times f_{s,0})$ (see Sec. 4.2). A slower detector would not see the regular bursting regime. In the bursting spectrogram, the “finger pointing down” marking the regular bursting with the bursting threshold at the fingertip is a very prominent feature that is easily identified. However, in the spectrogram, it is hard to name an exact criterion where the instability begins. So, decision based on the integrated power fluctuation in a certain frequency range – as discussed above – seems to be more accurate.

Instead of the coherent radiation, also the charge distribution can be investigated. Sec. 6.2 has shown that a direct measurement of substructures is possible. But for a threshold determination, it should be considered that – as for fluctuation of bunch length and energy spread – those structures are just starting to emerge and thus are very small at the threshold. It sounds to be more promising to quantify the size of these structures to be able to extrapolate from larger modulations. To handle the amounts of acquired data when continuously recording, T. Boltz started exploring the possibilities of unsupervised machine learning to allow automatic classification of the substructures. The results gained from *Inovesa* simulations show that this is possible if the signal to noise ratio is good enough [Bol17].

Another approach pins down the increase of energy spread (see Sec. 10.1.3). This method is again well tested. Measurements for example can be performed indirectly using the transverse bunch size in a dispersive section [Kre+93]. There might also be other effects that lead to a transverse beam-blowup, but in principle this method should be quite robust. An example measurement for KARA is found in [Keh+17a]. Especially for simulations (e.g. in [BCS10]), the increase of energy spread at the bursting threshold is a good indicator.

Analogue to the increase of the energy spread, it sounds possible to view the bunch lengthening. This approach has been followed, e.g. in [Jud13; Sch+13]. It is less reliable because of the potential well distortion, so that there is no constant level before the threshold. Also, the lengthening during the instability depends on the part of the impedance that is sampled. This can be simple, if the impedance is e.g. governed by unshielded emission of CSR, but this assumption usually cannot be done. In total, bunch lengthening is not a precise indicator but it can give an estimate.

Fluctuation of the bunch length was the first sign of the micro-bunching instability [Kre+93; Pod99]. For example in [Sch+13], it was also tried to use it to determine the micro-bunching threshold. However, the fluctuation is typically very small at the threshold, but strong in the sawtooth regime (cf. Sec. 10.2.2). Section 10.2.1 showed, that the fluctuating bunch length is strongly connected to the fluctuation of the emitted CSR. So it (and analogously the fluctuation of the energy spread) should in principle be usable as an indicator. The

method still will be less sensitive than using the fluctuating emission of CSR, as the latter is a non-linear effect that enhances the visibility of small length-fluctuations.

12.3. Influencing the Bursting CSR

As shown in Section 10.3 the strength of the structures as well as the bursting threshold can be influenced by additional impedances. For short bunches, however, it is hard not to be dominated by the effect of the shielded CSR impedance. It would be worth to investigate, whether a narrow-band impedance could form substructures that then self-amplify due to their own CSR. This approach is currently being explored by P. Schreiber using *Inovesa* [Sch18b]. Another option would be to influence the formation of the substructures by the incoherent synchrotron tunes (see Section 10.1.2). Especially the WMB is very sensitive in this regard [BCS10]. As it does not introduce bunch lengthening to the same amount as other bursting does (see Section 10.1.3), it still allows full coherence for lower frequency emission than other types of bursting. To influence the tune, the options include higher orders of α_c , a stronger accelerating voltage non-linearities, and introduction of additional impedances. The first option would need tuning of the beam optics using sextupole magnets, the second option might be realized using an additional cavity, the last one would require manipulating the beam pipe but can be evaluated using a scraper.

12.4. Verification of Non-Material Charge Modulation

To experimentally verify the finding that the charge-modulation in the longitudinal phase space is not material but formed in a dynamic process, one possibility would be to measure the tune spread. This information then has to be combined with other data showing discrete frequencies of the substructures.

The direct way is to reconstruct the phase space. A method, called *phase space tomography*, aims to record the changing bunch profile for a period of time. From these projections, the two-dimensional phase space might be reconstructed. However, at this point knowledge of the rotation speed of the structures in phase space is required – which should be a result rather than prior knowledge. Also, the substructures should not deform too much during the measurement which is actually not given. As a last point, different structures can have rotation speeds even if they are present at the same point in time (cf. Tab. 10.1). Still, under certain circumstances, it might be valid to assume the forming substructures to be static. The first example is regular bursting. Because the high degree of (temporal) symmetry a short period of time already holds complete information. Secondly, in WMB, it takes the substructures many synchrotron periods to form, so they are almost static. However, for WMB, the low bunch charge adds another experimental challenge. The optimized EO setup (Chap. 5) can be seen as a first step to realize this method.

Another option would be to image the full phase space at a certain time, using a transverse deflecting structure [EFK00]. The advantage of this approach is that it does not require static substructures. The method, however, is destructive. So two bunches need to be under exactly the same conditions to record two phase spaces that only differ by a fraction of a synchrotron period.

13. Summary

At synchrotron light sources, electron bunches are compressed so that the electrons irradiate coherently for frequencies up to the THz range. For low beam currents and thus low electron densities, the bunch shape is constant in time. In the lower limit it is Gaussian, but typically there is a deformation, called Haïssinski distribution. Above a threshold current, the CSR emission can heavily influence the beam dynamics: It leads to the formation of self-amplifying substructures on the electron bunches. On the one hand, these micro-bunches allow for coherent emission at even shorter wavelength, increasing the radiation level in the region of interest. On the other hand, the self-amplification can lead to unsteady conditions, where the synchrotron radiation is emitted in so-called “bursts”. Also, the overall bunch-shape is lengthened due to the micro-bunching. In order to overcome the limitations implied by this phenomenon, to fully exploit its potential, and to possibly control it, a deeper understanding of the processes is needed. Until now, the dynamics cannot be described analytically. To directly observe the evolution of the bunch profile, a sensitive single-shot turn-by-turn method is needed. At the moment, near-field EOSD is the technique that can best match these requirements. The only setup of this kind that is installed at a storage ring can be found at KARA. To simulate the micro-bunching, it is a well established method to numerically solve the Vlasov-Fokker-Planck equation.

This thesis attacks the topic from two sides: by optimizing the measurement setup and the algorithms used for simulation. In the framework of this thesis, I designed and installed a new in-vacuum component for near-field EOSD measurements. It minimizes the effect of wake-fields that previously limited the usage of the system while at the same time doubling the signal-to-noise ratio. Furthermore, the setup has been equipped with a turn-by-turn readout system called KALYPSO. It was developed in joint collaboration between KIT, PSI, DESY and the University of Łódź, in order to build a spectrometer with MHz repetition rates.

As there is no analytic description nor a closed formulation, one has to rely on simulations to predict the measurement results. In the second half of my thesis, I focused on this part. I did systematic numerical studies of the micro-bunching instability. Before, mostly short periods of time were modeled – often with only a single set of parameters. These systematic studies became possible because I adapted (Chap. 8) the well tested algorithms to make use of general-purpose multi-parallel computing devices (such as graphic processors) for desktop computers. These optimizations lead to a speedup of about a factor of 150 when compared to a non-optimized implementation. Additionally, I added a user interface so that the tool can be used by a wider community.

In particular, I performed studies of the current-dependent changes (Sec. 10.1) as well as an in-depth analysis of two different “bursting regimes” (Sec. 10.2), explaining the underlying mechanisms. In the “regular” bursting regime, there are quasi-static substructures that grow when moving to the bunch head and shrink on their way back. In the “sawtooth” regime, structures form, self-amplify and then lead to a blow-up of the bunch. This

mechanism leads to a strong fluctuation in the emitted radiation, the “bursts”. Also, I present the first systematic numerical studies of the influence of additional impedances on the spatio-temporal dynamics (Sec. 10.3) – previously only the change in the instability threshold was determined. Furthermore, I investigated which frequency ranges contribute to the overall fluctuation of the emitted power. I found that it is possible that these fluctuations have opposing signs for different frequency ranges (Sec. 10.4). When choosing a fitting frequency range for detection, this could lead to a method to obtain stable coherent synchrotron radiation from an unstable beam.

To understand the nature of the substructures, I introduced a novel simulation scheme that combines the Vlasov-Fokker-Planck solver with particle tracking. It allows efficient computation of the whole phase space, yet yields information on single particle dynamics. This method led me to the insight that modulations in the phase space are not material but form as an equilibrium in a dynamic process (Sec. 10.5). For the new findings, I made suggestions how future experimental work can use or verify them (Chap. 12). On this road, my design of an optimized in-vacuum part for electro-optical measurements (Chap. 5) will probably play a significant role.

The view on the “sawtooth bursting regime” gained by the simulation algorithm could already be verified: Using KALYPSO, for the first time the temporal evolution of substructures in near-field could be measured not only for the beginning burst but for several bursting cycles (Sec. 6.4). My work was already continued and expanded by several complementary Bachelor and Master thesis. Based on simulation data, T. Boltz used unsupervised machine learning to automatically find different categories of bunch profiles [Bol17] and also checked the requirements to do the same for measurement data recorded using KALYPSO. Following the long-term goal to design impedances for an explicit purpose, P. Schreiber started systematic studies of the longitudinal beam dynamics for multiple impedances [Sch18b]. The influence of physical noise in the form of phase jitter and a varying amplitude of the accelerating voltage was investigated by J. Schestag [Sch18a]. As expected, he triggered a coherent synchrotron oscillation, and found out that the noise does not significantly effect the instability threshold.

To summarize, significant improvements were achieved for both, measurement and simulation of longitudinal electron bunch profiles. On the measurement side, now turn-by-turn single-shot measurements are possible with increased signal-to-noise ratio and strongly reduced influence of interfering wake-fields. Especially after 2 ns, where a consecutive bunch in a multi-bunch environment would be located, the remaining wake field could be reduced to background level. On the simulation side, the speedup and the improved numerical stability helped to perform extensive studies. Also thanks to the combination of the Vlasov-Fokker-Planck solver with particle tracking, these studies gave new insights into the dynamics of the micro-bunching.

Appendix

A. Complete Lists of Used Simulation Settings

To maintain a balance between feature richness and comparability of scenarios, simulation parameters have been chosen to correspond to a total number of four different machine settings. These settings are named using a so-called fill number and listed in Tab. A.1.

f4623: Figs. 10.1, 10.5, 10.4(a), 10.6, 10.7, 10.9, 10.11(a), 10.12

f5135: Figs. 10.8, 10.10, 10.11(b)

f5493: Figs. 10.2, 10.13(c), 10.14, (with modified $f_{s,0}$) 10.11(c)

f5864: Figs. 10.4(b), 10.3 10.13(a) 10.13(b)

The fill number can also be used to identify the machine shift at KARA where measurements were carried out with the same parameters. The simulation results are available at [f4623; f5135; f5493; f5864].

Table A.1.: Simulation settings used for the results obtained in this thesis.

Quantity	Unit	f4623	f5135	f5493	f5864
V_{RF}	MV	1.5	1.048	1.396	1.3
$f_{s,0}$ (KARA)	kHz	7.7	(8.9)	7.2	(≈ 5.76)
$f_{s,0}$ (set)	kHz	7.7	28.130	7.2	(n/A)
E_0	GeV	1.285	1.285	1.285	1.285
σ_δ	10^{-6}	470	470	470	470
R	m	5.559		5.559	5.559
f_c	GHz	23	23	23	23
τ_d	ms	10.4	3.353	10.4	10.4
f_{rev}	MHz	2.71579	8.582	2.71579	2.71579
g	mm	32	32	32	32
α_0	10^{-6}				151.6
h	1	184	58.21	184	184

B. Mathematical Definitions and Discussions

B.1. Airy and Bessel Functions

Mathematical functions were used according to the definition in the ‘‘Boost’’ library for C++. The Bessel functions read [Bst17c]

$$J_\nu(z) := \left(\frac{z}{2}\right)^\nu \sum_{k=0}^{\infty} \frac{(-z^2/4)^k}{k! \Gamma(\nu + k + 1)}, \quad (\text{B.1})$$

and [Bst17d]

$$I_\nu(z) := \left(\frac{z}{2}\right)^\nu \sum_{k=0}^{\infty} \frac{(z^2/4)^k}{k! \Gamma(\nu + k + 1)} \quad (\text{B.2})$$

$$K_\nu(z) := \frac{\pi I_{-\nu}(z) - I_\nu(z)}{2 \sin(\nu\pi)} \quad (\text{B.3})$$

for the modified Bessel functions. The Airy functions are [Bst17a]

$$\text{Ai}(z) := \begin{cases} \frac{\sqrt{z}}{3} \left[\text{J}_{-1/3} \left(\frac{2}{3} z^{2/3} \right) + \text{J}_{1/3} \left(\frac{2}{3} z^{2/3} \right) \right], & \text{for } z < 0 \\ \frac{1}{3^{2/3} \Gamma(2/3)}, & \text{for } z = 0, \\ \frac{1}{\pi} \sqrt{\frac{z}{3}} \times \text{K}_{1/3} \left(\frac{2}{3} z^{2/3} \right), & \text{for } z > 0 \end{cases} \quad (\text{B.4})$$

and [Bst17b]

$$\text{Bi}(z) := \begin{cases} \sqrt{\frac{z}{3}} \left[\text{J}_{-1/3} \left(\frac{2}{3} z^{2/3} \right) - \text{J}_{1/3} \left(\frac{2}{3} z^{2/3} \right) \right], & \text{for } z < 0 \\ \frac{1}{3^{1/6} \Gamma(2/3)}, & \text{for } z = 0. \\ \sqrt{\frac{z}{3}} \left[\text{I}_{-1/3} \left(\frac{2}{3} z^{2/3} \right) + \text{I}_{1/3} \left(\frac{2}{3} z^{2/3} \right) \right], & \text{for } z > 0 \end{cases} \quad (\text{B.5})$$

B.2. Fourier Transforms

As the aforementioned mathematical functions, also the definition of the Fourier transform was chosen to be consistent with the used software implementation. In [FFTW16; cFFFT16] the “forward” Fourier transform is defined to be

$$\tilde{f}(\omega) = \int f(t) e^{-i\omega t} dt. \quad (\text{B.6})$$

The direction is important, as also the real-to-complex Fourier transform is of the forward type. As unnormalized transforms are performed in the implementation, the sign in the exponent is the only difference to a backward transform. The transform from a (complex) Hermitian function to a real function is of the backwards type. Of course, the Fourier transform has to be normalized at some point. To do so, here, the factor of $1/(2\pi)$ is applied to the backward transform

$$f(t) = \frac{1}{2\pi} \int \tilde{f}(\omega) e^{i\omega t} d\omega. \quad (\text{B.7})$$

This is definition is in agreement to the choice in [Kha06].

It should be noted that this is not the only possible definition. In [Ng06] the direction with the negative exponent is called the “inverse” Fourier transform. Arguing from [Ng06, p. 16] the naming is in agreement, as the inverse FT is used for complex to real transformations. However, the sign in the exponent differs. For example in [Ng02, p. 1-13]

$$W'(z) = \frac{1}{2\pi} \int Z^{\parallel}(\omega) e^{-i\omega z/v} d\omega, \quad (\text{B.8})$$

a Hermitian to real (“backward” or “inverse”) transform has a negative sign in the exponent. Effectively, the only difference is a flipped sign in the imaginary part of the result, so this does not result in a physical difference.

B.3. Rounding

As they are working with discrete numbers (defined by the number of bits), computers have to implement some sort of rounding. For example, division often has a remainder. A commonly known case is that the remainder is discarded in integer division

$$\begin{aligned} 5/4 &= 1 \\ 7/4 &= 1. \end{aligned}$$

This behavior adds a bias to the average result. To avoid this, one usually would round, so $7/4 = 2$.) As the remainder 0.5 is as close to 1 as to 0, often special attention is taken for these kinds of values, called *ties*. (Otherwise there still would be a small bias.) IEEE 754 floating point representations offer *unbiased rounding*: It is guaranteed that rounding random numbers do not bias the average. To achieve this, there are different possibilities that can be chosen using a so called rounding-direction attribute [IEEE08, p. 16]. The default (*roundTiesToEven*) is to round ties to the next number with an even least significant digit. Other options for **binary** formats are to always round upward, downwards or towards zero. These three options are not unbiased anymore. Another possibility to implement unbiased rounding would be to randomly round upwards or downwards if there is a tie. However, this approach is not foreseen in the IEEE standard.

To have no bias even when the same rounding is performed a multitude of times, the above treatments offer no solution. In such cases, *stochastic rounding* [Cai+18] could be performed: The decision to round up or down has to be random not only for ties. Introducing the random number $0 \leq p < 1$, stochastic rounding can be described as

$$A/B = \lfloor A/B \rfloor + \begin{cases} 0 & \text{if } p \geq (A \bmod B)/B \\ 1 & \text{if } p < (A \bmod B)/B. \end{cases} \quad (\text{B.9})$$

This leads to a situation where – on average – the result is correct.

C. Complementary Data

C.1. Prediction of EOS v2 wake

As stated in Sec. 5.2, the result of an EO measurement is dependent on the (changing) bunch current, and the exact laser path through the EO crystal. The path is dependent on the alignment of and cannot be determined by measurement in the installed setup. So, for the simulation of the EOS v1 wake, the path set in the simulation has been varied until a satisfactory agreement to the measurement was obtained. For EOS v2, this manual fitting was not performed. Instead, the shape of the wake field was only predicted for a rather pessimistic alignment. Figure C.1 adds this prediction to the data already displayed in Fig. 5.10. Notice that all of the points that can explain mismatches that have been named in Sec. 5.2 are also valid here:

- The CAD model of the EO arm is simplified with respect to the original geometry.
- The position where the laser passes the crystal is not known. Further, for EOS v2, it has not been fitted.
- The small angle between incoming and outgoing direction of the laser has been neglected.
- For the simulation, a Gaussian bunch shape and a constant bunch length have been assumed.

Considering all these, the agreement between the simulation result for EOS v2 and the measurement still can be considered to be reasonably good.




Figure C.1.: Simulated EOS v2 wake-field compared to measurements

The result is dependent on the (changing) bunch current, and the exact laser path through the EO crystal (c.f. Sec. 5.2). For the prediction of the EOS v2 wake field, simulation parameters were used that provided reasonably good agreement for EOS v1. This might explain why in some parts the simulation is more similar to EOS v1 than to EOS v2.

C.2. Implicit Discretization: Differences in Convergence

One topic of Chapter 9 was the effect of the choice of the data type. It has been stated that it can be seen as an implicit discretization (see Sec. 9.1). In most cases, it is overweighted by the explicit discretization. This statement was verified by benchmarking it simulating the convergence for the case of the unperturbed VFPE in Sec. 9.3.1: When the numeric accuracy of the chosen data type is high enough, the simulation converges to an almost common value that is dominated by other parameters. Still, there are small differences between the possible data types. They are displayed in Fig. C.2.

The mismatch between `binary32` and `binary64` is in the order of 10^{-7} , which is the difference in the numerical accuracy. For the other data types, it is similar: With increased accuracy the results become more similar (up to 10^{-11} for `fix59`) – but do not converge closer to the analytic result. When decreasing the accuracy to a 29 bit fixed point data type, a significant change can be observed: The difference has a maximum at $t \approx 20 T_s$. At $t \approx 42 T_s$, the curves cross. All simulations have converged at latest at $t = 90 T_s$. From that point on the values do not change anymore.

C.3. Contributions to the Form Factor

The form-factor can be expressed as the sum of the form-factor of the average bunch profile (static term), the form-factor of the modulation (modulation term), and a mixed term [c.f. Eq. (10.6)]. An example is visualized in Fig. C.3. The static term always starts at $F(f = 0) = 1$, which means that there is full coherence for low frequencies. The peak at $f = 100$ GHz is not due to substructures but comes from a static bunch deformation. It can be linked to the PP CSR impedance which is $\Im[Z_{PP}(f < 100 \text{ GHz})] < 0 \Omega < \Im[Z_{PP}(f > 100 \text{ GHz})]$. So – neglecting the real part – structures emitting at 100 GHz remain unchanged, but longer structures are lengthened and shorter structures are shortened. In the frequency range $f > 400$ GHz, the contribution from the modulation term dominates the total form factor. The mixed term shows several magnitudes of fluctuation in the whole frequency range. So, where its average contribution is larger than the one of the average modulation term, also the fluctuation is dominated by the mixed term. This is especially the case at $f \approx 100$ GHz, where also the main modulations of the spectra were visible in Fig. 10.8. Also notice that the mixed term can be negative. This is especially visible because of the two dips in the minimum form factor at $f \approx 400$ GHz: They are several orders of magnitude lower than the static and the modulation term.




Figure C.2.: Convergence based on Data Type

For a better comparison of the influence of the data type, here the result of the `binary64` simulation run has been subtracted as the reference dataset. When the simulation has converged after $t = 90 T_s$, most displayed data types except for the 34 and the 29 bit point ones (`fix34` and `fix29`) have a result closer to the one obtained using `binary64` than to the analytic result. Also notice that the fixed point results in this logarithmic representations move mostly parallel, the analytic and the floating point results on the other hand show a slight qualitative difference.




Figure C.3.: The form-factor (black) is the sum of the “static term” (dotted red), the “modulation term” (blue) and the mixed term (magenta). Lines mark average values, the bands give the full range from minimum to maximum. Notice that the static term is constant, so it has no band. The total form-factor can be below the static and the modulation term, as the mixed term can be negative, here only the magnitude is displayed. (See text for more explanation.)

C.4. Contribution of Additional Impedances

In Sec. 10.3, the influence of additional impedances that are added to Z_{PP} (Eq. (3.15)) was discussed. Figure C.4 shows complementary spectrograms to the data.

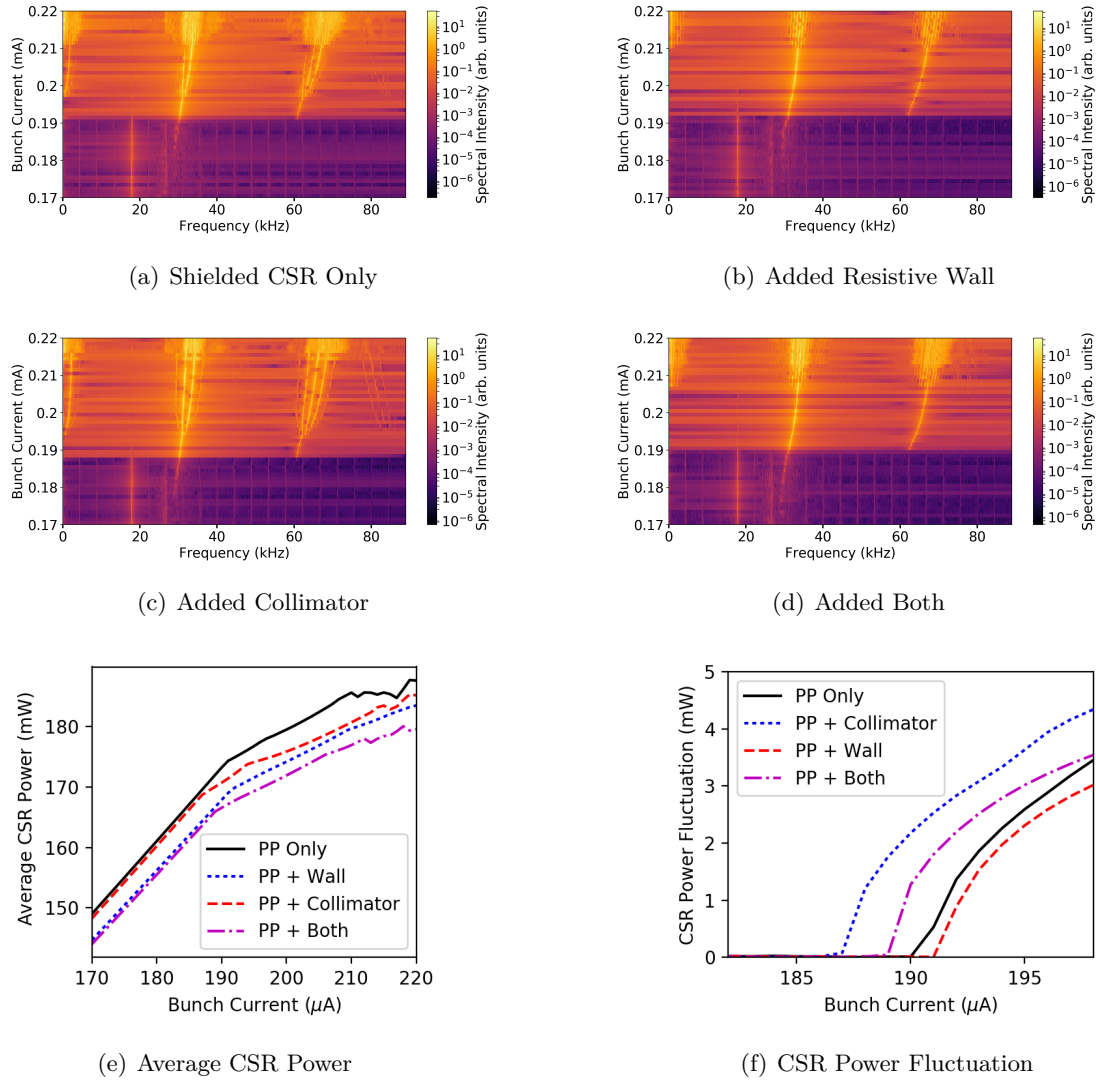


Figure C.4.: The plot of the CSR power function (f) was already shown in Fig. 10.10. The rise from 0 mW provides the easiest and most reliable measure for the micro-bunching instability threshold current. In (e) the average CSR power is shown. It is reduced due to the additional impedances as it adds a further contribution to the potential well distortion. At higher currents, the low frequency (sawtooth) bursting shows significant differences. Those also result in different frequency combs. Notice that also below the bursting threshold, a frequency component close to $f_r = 30$ kHz can be observed – but with a very low intensity. These can be linked to structures in the order of magnitude of the numerical resolution (see Fig 9.6). They do also add a non-zero CSR power fluctuation below the instability threshold.

C.5. Fluctuation of the CSR Spectrum

In Sec. 10.4, the wavelength dependence of the CSR power fluctuation was discussed. It included an example case where the total fluctuation was minimized by choosing a specific sensitivity range for the radiation (see Fig. 10.12). Figure C.5 shows the same data in time domain and including the information over the full spectral range.

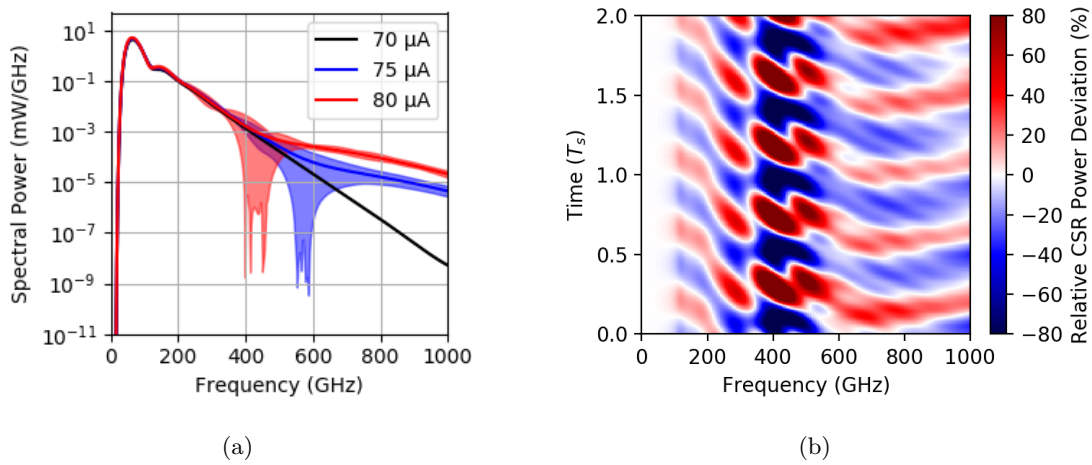


Figure C.5.: (a) is identical to Fig. 10.11(a). As the average bunch profile does only change slightly in the displayed current range, the most of the differences can be attributed to the substructures, with the spectrum for $I = 70 \mu\text{A}$ being the one without. (b) shows the relative deviation from the average CSR spectrum for the example of 80 mA over time. Maximums and minimums are cropped, the extreme values are +730% and -100%.

C.6. Difficulty to Obtain the Synchrotron Frequency

In Chap. 11, it has been stated that one issue when comparing simulation results to measurement data is to choose the correct unperturbed synchrotron frequency for the simulation that corresponds to the measurement. Figure C.6 gives example measurement results to confirm this claim.

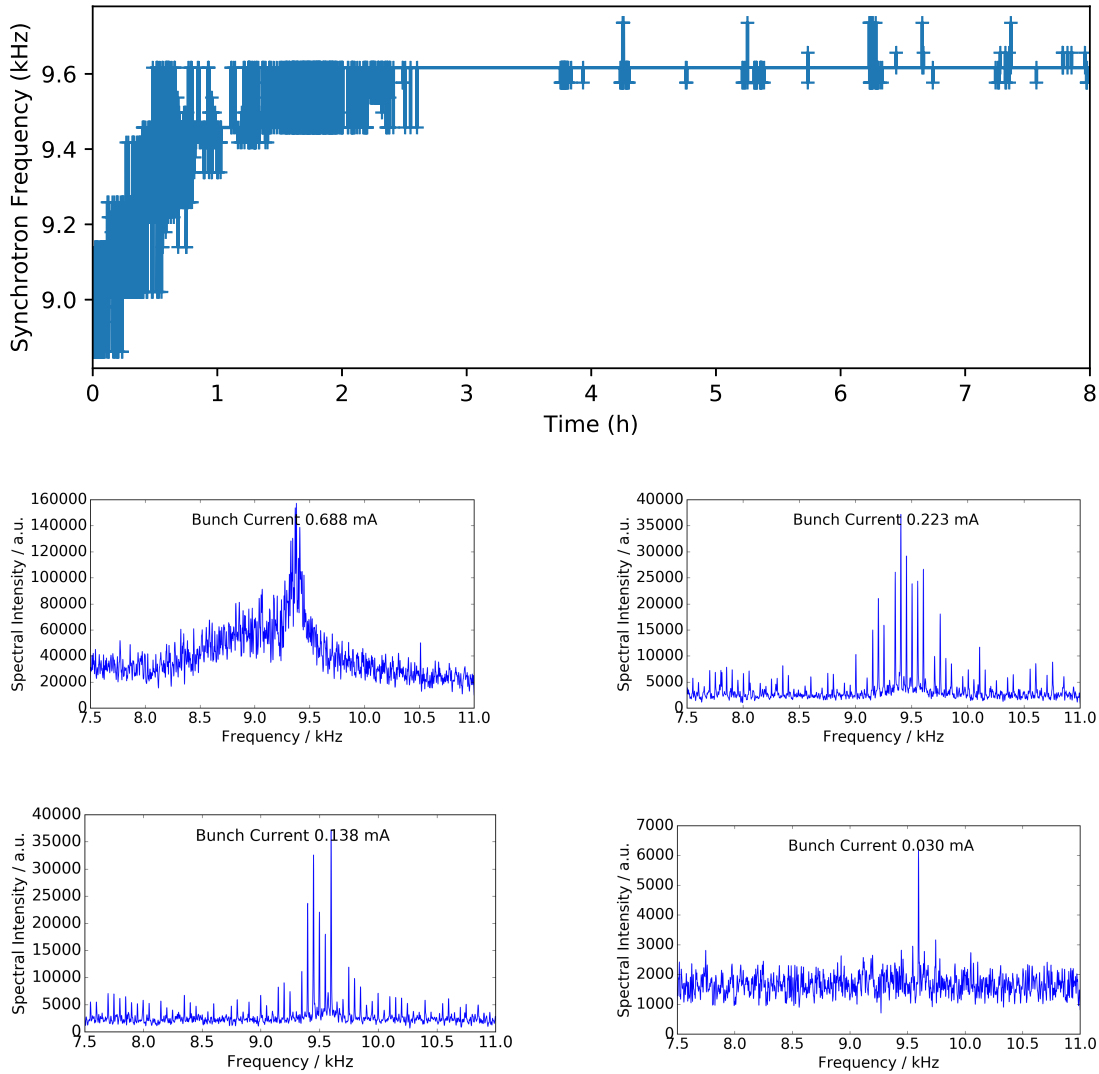


Figure C.6.: Difficulty to Obtain the Synchrotron Frequency

The upper panel shows the synchrotron frequency as it is automatically obtained by the accelerator controls over time. Values are saved to the database every time a change is detected. At early times, the bunch current is highest, later it decreases. It can be seen that in the first hour, the value drifts by as much as 500 Hz and that it fluctuates by with an amplitude of about 100 Hz. This fluctuation is fastest for the highest currents. The lower four panels show CSR spectrograms for four example currents. Two of them show very distinct 50 Hz lines, which hints that the reason for the fluctuation might be found in unstable power supplies. The shift of the coherent tune, however, cannot be explained within the used model. (Plots courtesy of M. Brosi.)

D. Inovesa Quick Reference

Inovesa is implemented as a non-interactive command line tool. For every simulation parameter there is a preset value that can be overwritten using a key value style config file (text format). Parameters that are given using command line parameters have priority over the two other values (if set).

The simulation results are written to a file that can be defined using the option `--output`. This can either be a file name ending on `.h5` – to save the computed information to a *HDF5* [HDF516] container – or a file name ending on `.png` – to save the final phase space of the simulation to a greyscale *PNG* [PNG16] file. When no output file is set, a *HDF5* file is created that is named using the current *UNIX* timestamp, to not save the results, use `--output /dev/null`. Additionally there is an optional graphical live-view (parameter `--gui`) that shows the phase space currently being simulated as well as the corresponding bunch profile, wake potential and emitted CSR.

D.1. Program Information

help print help message

copyright print copyright information

version print version string

D.2. General Program Parameters

cldev OpenCL device to use
 -1: lists available devices
 0: do not use OpenCL

config name of a configuration file

ForceOpenGLVersion use specific version of OpenGL for displaying data

gui show phase space view

output name of file to save results, might be:
 HDF5 file (‘.hdf5’, ‘.h5’)
 greyscale PNG (‘.png’)
 ‘/dev/null’ to explicitly state no read-in

SavePhaseSpace save full phase space to HDF5

SaveSourceMap save every outstep’s source map to HDF5 file

tracking file containing starting positions (grid points) of particles to be tracked

verbose print information more detailed

D.3. Numerical Parameters for Simulation

steps simulation steps for one synchrotron period

outstep save results every n simulation steps

padding factor for zero padding of bunch profile when FFT is performed

RoundPadding Always do zero padding up to 2 to the power of N

PhaseSpaceSize size of phase space

PhaseSpaceShiftX shift grid center by X grid points

PhaseSpaceShiftY shift grid center by Y grid points

RenormalizeCharge might be:

- > 0: renormalize charge every n-th simulation step
- = 0: do just one initial renormalization
- < 0: no renormalization

RotationType used implementation for rotation

- 0: standard rotation without source map
- 1: standard rotation with source map
- 2: Manhattan rotation

GridSize number of grid points per dimension

rotations simulated time in number of synchrotron periods

derivation number of grid points to be used to numerically find derivative

InterpolationPoints number of grid points to be used for interpolation

InterpolateClamped restrict result of interpolation to the values of the neighboring grid points

D.4. Physical Parameters for Simulation

alpha0 linear momentum compaction factor (1)

alpha1 quadratic momentum compaction factor (1)

alpha2 cubic momentum compaction factor (1)

SynchrotronFrequency Synchrotron frequency (Hz),
will overwrite alpha0

RevolutionFrequency revolution frequency (Hz)

DampingTime damping time (s)

HarmonicNumber harmonic number (1)

InitialDistFile might be:

- Inovesa result file (‘.hdf5’, ‘.h5’)
- quadratic greyscale PNG (‘.png’)
- text file (‘.txt’) containing particle coordinates
- ‘/dev/null’ to explicitly state no read-in

InitialDistStep Select step of HDF5 file for initial distribution

InitialDistZoom Magnification for generation of initial distribution

BunchCurrent ring current due to a single bunch (A)

BendingRadius Bending radius of accelerator (m)
negative: calculate from RevolutionFrequency

BeamEnergy beam energy (eV)

BeamEnergySpread natural energy spread (relative)

VacuumGap full height of vacuum chamber (m)

- < 0: free space CSR impedance
- = 0: no CSR impedance
- > 0: parallel plates CSR impedance

UseCSR Switch to turn off CSR for VacuumGap $\neq 0$

CollimatorRadius Radius of collimator opening (m)
 ≤ 0 : no collimator

WallConductivity Conductivity of the vacuum pipe (S/m)
 ≤ 0 : perfect conductor

WallSusceptibility Magnetic susceptibility of the vacuum pipe (1)
 $< -1 : -1$

CutoffFreq beamline cutoff frequency (Hz)

AcceleratingVoltage accelerating voltage (V)

E. List of Symbols

D	dispersion (magnet optics function) 5
E_0	energy of the reference particle 4
E	energy 3
J	instantaneous synchrotron motion 17
K	map describing the kick due to a collective force ('kick map') 52
L_0	circumference of a circular accelerator 4
N_θ	number of time steps per $T_{s,0}$ 52
Q	charge 3
R	bending radius 5
S	map describing the synchrotron motion ('rotation map') 52
T_0	revolution time 4
W_0	per-turn radiation losses of the synchronous particle 8
Z_0	vacuum impedance 15
α_c	momentum compaction factor 9
β_d	inverted damping time (in units of $T_{s,0}$) 13
β	ratio of the velocity to the speed of light 3
δ	relative momentum deviation 4
η_c	slip factor 9
γ	Lorentz factor 4
\hat{V}	amplitude of accelerating voltage 10
μ_0	magnetic constant 15
ψ	charge density in longitudinal phase space (normalized charge) 12
τ_d	longitudinal damping time 8
τ	relative arrival time ($\tau = z/c$) 4
θ	dimensionless time in multiples of the synchrotron period length 12
ε_0	electric constant 7
ϱ	longitudinal bunch profile (normalized charge) 7
\vec{B}	magnetic field 3
\vec{E}	electric field 3
\vec{F}_L	Lorentz force 3
\vec{p}	momentum (of a particle) 4
\vec{v}	velocity 3
$\{x\}$	grid coordinate of space (fractional part) 52
$\{y\}$	grid coordinate of momentum (fractional part) 52
c	speed of light 3
f_{RF}	frequency of accelerating voltage (typically in the RF range) 4
f_b	(dominant) bursting frequency 86
f_c	cutoff frequency (beam line cutoff) 9
f_{rev}	revolution frequency 4
$f_{s,0}$	synchrotron frequency 11
g	gap (full height) in PP CSR impedance 16
h	harmonic number 4
i	imaginary unit, $i = \sqrt{-1}$ 7
m	rest mass of a particle 4

n	grid coordinate represented as relative position in memory 53
p	dimensionless coordinate of momentum space 12
q	dimensionless coordinate of spacial space 12
x_r	grid coordinate of space 52
x	grid coordinate of space (integer part) 52
y_r	grid coordinate of momentum 52
y	grid coordinate of momentum (integer part) 52
z	physical coordinate of longitudinal space 4

F. Acronyms

ALU	a rithmetic l ogic u nit
ANKA	Ä ngströmquelle K arlsruhe
BBT	b unched b eam t heory
CAD	c omputer a ided d esign
CSR	c oherent s ynchrotron r adiation
CU	c ontrol u nit
DC	originally d irect c urrent, but also used for signals with no time-varying component
DESY	D eutsches E lektronen- S ynchrotron
EO	e lectro- o ptical
EOS	e lectro- o ptical s ampling (using multiple laser pulses)
EOSD	e lectro- o ptical s pectral d ecoding (using one chirped laser pulse)
FEL	f ree- e lectron l aser
FWHM	f ull w idth h alf m aximum
GOTTHARD	g ain o ptimizing microstrip system w ith a nalog r eadout
HEB	h ot e lectron b olometer
ISR	i ncoherent s ynchrotron r adiation
KALYPSO	K arlsruhe l inear a rray detector for MHz- r epetition rate s pectroscopy
KAPTURE	K arlsruhe p ulse t aking and u ltra- f ast r eadout e lectronics
KARA	K arlsruhe R esearch A ccelerator
KIT	K arlsruhe I nstitute of T echnology
MIMD	m ultiple i nstruction streams, m ultiple d ata streams
MISD	m ultiple i nstruction streams, s ingle d ata stream
NSLS	N ational S ynchrotron L ight S ource V acuum U ltraviolet ring

PSI	P aul S cherrer I nstitute
PU	p rocessing u nit
RMS	r oot m ean s quare; often used synonym to standard deviation, when the position of the average is (defined to be) zero
SBB	s hort b unch length b ursting
SIMD	s ingle i nstruction stream, m ultiple d ata streams
SISD	s ingle instruction stream, s ingle d ata stream
SLC	S tanford L inear C ollider
VFPE	V lasov- F okker- P lanck e quation
VFPS	V lasov- F okker- P lanck s olver
WMB	w eak m icro b unching instability

Bibliography

- [Abo+02] M. Abo-Bakr, J. Feikes, K. Holldack, G. Wüstefeld, and H.-W. Hübers. “Steady-State Far-Infrared Coherent Synchrotron Radiation detected at BESSY II”. *Phys. Rev. Lett.* 88 (25 June 2002), p. 254801. DOI: 10.1103/PhysRevLett.88.254801.
- [Abo+03] M. Abo-Bakr, J. Feikes, K. Holldack, P. Kuske, and G. Wustefeld. “Coherent emission of synchrotron radiation and longitudinal instabilities”. *Proceedings of the Particle Accelerator Conference*. Vol. 5. IEEE. 2003, pp. 3023–3025. DOI: 10.1109/PAC.2003.1289801.
- [ACST17] ACST GmbH. *Quasi-optical detectors*. 2017. URL: http://www.acst.de/?page=4.5.&view=thz_detector (visited on 08/30/2017).
- [AMD05] *Presentation at AMD Analyst Day*. 2005.
- [AMD17a] *Radeon Pro Duo*. 2017. URL: <http://www.amd.com/en-us/products/graphics/desktop/radeon-pro-duo> (visited on 04/13/2017).
- [AMD17b] *Radeon R7*. 2017. URL: <http://www.amd.com/en-us/products/graphics/desktop/r7> (visited on 04/13/2017).
- [AMD17c] *Radeon R9*. 2017. URL: <http://www.amd.com/en-us/products/graphics/desktop/r9> (visited on 10/18/2017).
- [AMD17d] *The "Zen" Core Architecture*. 2017. URL: <http://www.amd.com/en/technologies/zen-core> (visited on 04/13/2017).
- [Anand12] *The Opteron 6276: a closer look*. 2012. URL: <http://www.anandtech.com/show/5279/the-opteron-6276-a-closer-look/6> (visited on 04/13/2017).
- [Ang+12] A. Angelovski, A. Kuhl, M. Hansli, et al. “High bandwidth pickup design for bunch arrival-time monitors for free-electron laser”. *Phys. Rev. ST Accel. Beams* 15 (11 Nov. 2012), p. 112803. DOI: 10.1103/PhysRevSTAB.15.112803.
- [Ang+15] A. Angelovski, M. Kuntzsch, M. K. Czwalińska, et al. “Evaluation of the cone-shaped pickup performance for low charge sub-10 fs arrival-time measurements at free electron laser facilities”. *Phys. Rev. ST Accel. Beams* 18 (1 Jan. 2015), p. 012801. DOI: 10.1103/PhysRevSTAB.18.012801.
- [Axi17] Axis Photonique Inc. *Ultrafast Streak Cameras*. 2017. URL: <http://axis-photon.com/ultrafast-streak-cameras/> (visited on 08/30/2017).
- [Ban16] K. Bane. “Longitudinal Stability Study for the FACET-II e+ Damping Ring”. *ArXiv e-prints* (Nov. 2016). URL: <https://arxiv.org/abs/1611.08042>. 08042arXiv: 1611.08042.
- [Bar+13] J. Barros, C. Evain, L. Manceron, et al. “Coherent synchrotron radiation for broadband terahertz spectroscopy”. *Review of Scientific Instruments* 84.3 (2013), p. 033102. DOI: 10.1063/1.4793558.

- [Bar+15] J. Barros, C. Evain, E. Roussel, et al. “Characteristics and development of the coherent synchrotron radiation sources for THz spectroscopy”. *Journal of Molecular Spectroscopy* 315 (2015). Spectroscopy with Synchrotron Radiation, pp. 3–9. ISSN: 0022-2852. DOI: 10.1016/j.jms.2015.03.012.
- [BCS10] K. L. F. Bane, Y. Cai, and G. Stupakov. “Threshold studies of the microwave instability in electron storage rings”. *Phys. Rev. ST Accel. Beams* 13 (10 Oct. 2010), p. 104402. DOI: 10.1103/PhysRevSTAB.13.104402.
- [BIF17] BIFO company. *Streak camera and frame cameras*. 2017. URL: <http://bifocompany.com/eng/p-cameras.php.htm> (visited on 08/30/2017).
- [Bol17] T. Boltz. “Comprehensive Analysis of Micro-Structure Dynamics in Longitudinal Electron Bunch Profiles”. MA thesis. Karlsruhe Institute of Technology, 2017. DOI: 10.5445/IR/1000068253.
- [Bon+80] T. D. Bonifield, F. H. K. Rambow, G. K. Walters, M. V. McCusker, D. C. Lorents, and R. A. Gutcheck. “Time resolved spectroscopy of xenon excimers excited by synchrotron radiation”. *The Journal of Chemical Physics* 72.5 (1980), pp. 2914–2924. DOI: 10.1063/1.439490.
- [Bou75] D. Boussard. *Observation of microwave longitudinal instabilities in the CPS*. Tech. rep. CERN-LabII-RF-INT-75-2. LabII-RF-INT-75-2. Geneva: CERN, Apr. 1975. URL: <https://cds.cern.ch/record/872559>.
- [BOZ05] K. L. Bane, K. Oide, and M. Zobo. *Impedance Calculation and Verification in Storage Rings*. Tech. rep. SLAC-PUB-11007. SLAC, 2005. URL: <http://www.slac.stanford.edu/pubs/slacpubs/11000/slac-pub-11007.pdf>.
- [Bro+15] M. Brosi, M. Caselle, E. Hertle, et al. “Online Studies of THz-radiation in the Bursting Regime at ANKA”. *Proc. IPAC*. (Richmond, Virginia, USA). International Particle Accelerator Conference 6. Geneva, Switzerland: JACoW, May 2015, pp. 82–84. ISBN: 978-3-95450-168-7. URL: <http://jacow.org/ipac2015/papers/mopha042.pdf>.
- [Bro+16a] M. Brosi, E. Blomley, E. Bründermann, et al. “Systematic Studies of Short Bunch-Length Bursting at ANKA”. *Proc. IPAC*. (Busan, Korea). International Particle Accelerator Conference 7. Geneva, Switzerland: JACoW, June 2016, pp. 1662–1665. ISBN: 978-3-95450-147-2. DOI: 10.18429/JACoW-IPAC2016-TUPOR006.
- [Bro+16b] M. Brosi, J. L. Steinmann, E. Blomley, et al. “Fast Mapping of Terahertz Bursting Thresholds and Characteristics at Synchrotron Light Sources”. *Phys. Rev. Accel. Beams* 19 (11 Nov. 2016), p. 110701. DOI: 10.1103/PhysRevAccelBeams.19.110701.
- [Bro14] M. Brosi. “Untersuchung des Burstingverhaltens von Synchrotronstrahlung im THz-Bereich”. MA thesis. Karlsruhe Institute of Technology, 2014.
- [BS12] G. Bellisola and C. Sorio. “Infrared spectroscopy and microscopy in cancer research and diagnosis”. *American Journal of Cancer Research* 2.1 (Nov. 2012), pp. 1–21. ISSN: 2156-6976. URL: <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC3236568/>.
- [Bst17a] *Airy Ai Function*. 2017. URL: http://www.boost.org/doc/libs/1_52_0/libs/math/doc/sf_and_dist/html/math_toolkit/special/airy/ai.html (visited on 09/15/2017).
- [Bst17b] *Airy Bi Function*. 2017. URL: http://www.boost.org/doc/libs/1_52_0/libs/math/doc/sf_and_dist/html/math_toolkit/special/airy/bi.html (visited on 09/15/2017).

- [Bst17c] *Bessel Functions of the First and Second Kinds*. 2017. URL: http://www.boost.org/doc/libs/1_52_0/libs/math/doc/sf_and_dist/html/math_toolkit/special/bessel/bessel.html (visited on 09/15/2017).
- [Bst17d] *Modified Bessel Functions of the First and Second Kinds*. 2017. URL: http://www.boost.org/doc/libs/1_52_0/libs/math/doc/sf_and_dist/html/math_toolkit/special/bessel/mbessel.html (visited on 09/15/2017).
- [BU10] U. Brinkschulte and T. Ungerer. *Mikrocontroller und Mikroprozessoren*. 3. Aufl. eXamen.press. Berlin: Springer, 2010. ISBN: 978-3-642-05397-9. DOI: 10.1007/978-3-642-05398-6.
- [Byr+02] J. M. Byrd, W. P. Leemans, A. Loftsdottir, et al. “Observation of Broadband Self-Amplified Spontaneous Coherent Terahertz Synchrotron Radiation in a Storage Ring”. *Phys. Rev. Lett.* 89 (22 Nov. 2002), p. 224801. DOI: 10.1103/PhysRevLett.89.224801.
- [Cai+18] Y. Cai, C. Liang, Z. Tang, H. Li, and S. Gong. “Deep Neural Network with Limited Numerical Precision”. *International Conference on Applications and Techniques in Cyber Security and Intelligence*. Ed. by J. Abawajy, K.-K. R. Choo, and R. Islam. Cham: Springer International Publishing, 2018, pp. 42–50. ISBN: 978-3-319-67071-3.
- [Cai11] Y. Cai. “Theory of Microwave Instability and Coherent Synchrotron Radiation in Electron Storage Rings”. *Proc. IPAC*. (San Sebastián, Spain). International Particle Accelerator Conference 2. Geneva, Switzerland: JACoW, Sept. 2011, pp. 3774–3778. URL: <http://jacow.org/ipac2011/papers/frxaa01.pdf>.
- [Car+01] G. Carr, S. Kramer, J. Murphy, R. Lobo, and D. Tanner. “Observation of coherent synchrotron radiation from the NSLS VUV ring”. *Nuclear Instruments and Methods in Physics Research A* 463 (2001), pp. 387–392. DOI: 10.1016/S0168-9002(01)00521-6.
- [Car+99] G. Carr, S. Kramer, J. Murphy, J. LaVeign, R. Lobo, D. Reitze, and D. Tanner. “Investigation of coherent emission from the NSLS VUV ring”. *Particle Accelerator Conference, 1999. Proceedings of the 1999*. Vol. 1. IEEE. 1999, pp. 134–136. DOI: 10.1109/PAC.1999.795644.
- [Cas+14] C. Caselle, M. Brosi, S. Chilingaryan, et al. “A Picosecond Sampling Electronic “KAPTURE” for Terahertz Synchrotron Radiation”. *Proc. 3rd International Beam Instrumentation Conference*. (Monterey, CA, USA). Geneva, Switzerland: JACoW, 2014, pp. 24–28. ISBN: 978-3-95450-141-0. URL: <http://jacow.org/IBIC2014/papers/mocz1.pdf>.
- [CG90] A. W. Chao and J. Gareyte. “Scaling Law for Bunch Lengthening in SPEAR II”. *Particle Accelerators* 25 (1990), pp. 229–234. URL: <https://cds.cern.ch/record/210836/files/p229.pdf>.
- [Cha+13] A. W. Chao, K. H. Mess, M. Tigner, and F. Zimmermann. *Handbook of accelerator physics and engineering*. World Scientific, 2013. ISBN: 978-981-4417-17-4.
- [Che+17] W. Cheng, B. Bacha, K. Ha, and O. Singh. “Precise Synchronous Phase Measurements”. *Proc. of International Particle Accelerator Conference (IPAC’17), Copenhagen, Denmark, 14–19 May, 2017*. (Copenhagen, Denmark). International Particle Accelerator Conference 8. Geneva, Switzerland: JACoW, May 2017, pp. 487–490. ISBN: 978-3-95450-182-3. DOI: 10.18429/JACoW-IPAC2017-MOPAB152.

- [CK76] C. Z. Cheng and G. Knorr. “The integration of the Vlasov equation in configuration space”. *Journal of Computational Physics* 22.3 (1976), pp. 330–351. ISSN: 0021-9991. DOI: 10.1016/0021-9991(76)90053-X.
- [clFFT16] *clFTW Manual*. 2016. URL: <http://clmathlibraries.github.io/clFFT/> (visited on 04/21/2016).
- [CM65] K. Codling and R. P. Madden. “Characteristics of the “Synchrotron Light” from the NBS 180-MeV Machine”. *Journal of Applied Physics* 36.2 (1965), pp. 380–387. DOI: 10.1063/1.1713998.
- [CST17] *CST PARTICLE STUDIO – Wakefield Solver*. <http://www.cst.com/products/cstps/Solvers/WakefieldSolver>. 2017. (Visited on 02/22/2017).
- [Dea+77] D. A. G. Deacon, L. R. Elias, J. M. J. Madey, G. J. Ramian, H. A. Schwettman, and T. I. Smith. “First Operation of a Free-Electron Laser”. *Phys. Rev. Lett.* 38 (16 Apr. 1977), pp. 892–894. DOI: 10.1103/PhysRevLett.38.892.
- [DLP01] J. Dongarra, P. Luszczeky, and A. Petitet. *The LINPACK Benchmark: Past, Present, and Future*. Tech. rep. University of Tennessee, 2001. URL: <http://www.netlib.org/utk/people/JackDongarra/PAPERS/hpl.pdf>.
- [DRS95] Y. S. Derbenev, J. Rossbach, and E. L. Saldin. *Microbunch Radiative Tail-Head Interaction*. Tech. rep. TESLA FEL-Report 1995-05. DESY, 1995. URL: <http://cds.cern.ch/record/291102/files/SCAN-9511114.pdf?version=1>.
- [EFK00] P. Emma, J. Frisch, and P. Krejcik. *A Transverse RF deflecting structure for bunch length and phase space diagnostics*. Tech. rep. SLAC Linac Coherent Light Source, 2000. URL: <http://www-ssrl.slac.stanford.edu/lcls/technotes/lcls-tn-00-12.pdf>.
- [Eld+47] F. R. Elder, A. M. Gurewitsch, R. V. Langmuir, and H. C. Pollock. “Radiation from Electrons in a Synchrotron”. *Phys. Rev.* 71 (11 June 1947), pp. 829–830. DOI: 10.1103/PhysRev.71.829.5.
- [Eva+12] C. Evain, J. Barros, A. Loulergue, et al. “Spatio-temporal dynamics of relativistic electron bunches during the micro-bunching instability in storage rings”. *EPL (Europhysics Letters)* 98.4 (2012), p. 40006. DOI: 10.1209/0295-5075/98/40006.
- [f4623] P. Schönfeldt. *Inovesa result f04623*. 2018. DOI: 10.5445/IR/1000084475.
- [f5135] P. Schönfeldt. *Inovesa result f05135*. 2018. DOI: 10.5445/IR/1000084485.
- [f5493] P. Schönfeldt. *Inovesa result f05493*. 2018. DOI: 10.5445/IR/1000084490.
- [f5864] P. Schönfeldt. *Inovesa result f05864*. 2018. DOI: 10.5445/IR/1000084506.
- [Fei+11] J. Feikes, M. von Hartrott, M. Ries, et al. “Metrology Light Source: The first electron storage ring optimized for generating coherent THz radiation”. *Phys. Rev. ST Accel. Beams* 14 (3 Mar. 2011), p. 030705. DOI: 10.1103/PhysRevSTAB.14.030705.
- [FFTW16] *FFTW Manual*. 2016. URL: http://www.fftw.org/fftw2_doc/fftw_2.html (visited on 04/21/2016).
- [Fly11] M. Flynn. “Flynn’s Taxonomy”. *Encyclopedia of Parallel Computing*. Ed. by D. Padua. Boston, MA: Springer US, 2011, pp. 689–697. ISBN: 978-0-387-09766-4. DOI: 10.1007/978-0-387-09766-4_2.
- [Fly72] M. J. Flynn. “Some Computer Organizations and Their Effectiveness”. *IEEE Transactions on Computers* C-21.9 (Sept. 1972), pp. 948–960. ISSN: 0018-9340. DOI: 10.1109/TC.1972.5009071.

- [FPML09] *Fixed Point Class Documentation*. 2009. URL: <http://codeproject.com/Articles/37636/Fixed-Point-Class> (visited on 06/17/2016).
- [Gas+12] B. R. Gaster, L. Howes, D. R. Kaeli, P. Mistry, and D. Schaa. *Heterogeneous Computing with OpenCL*. Waltham, MA, USA: Morgan Kaufmann, 2012. ISBN: 9780124058941.
- [GS17] H.-P. Gumm and M. Sommer. *Informatik*. Vol. Band 2: Rechnerarchitektur, Betriebssysteme, Rechnernetze. De Gruyter Studium. Berlin: De Gruyter, 2017. ISBN: 978-3-11-044236-6. DOI: 10.1515/9783110442366.
- [Hai73] J. Haissinski. “Exact Longitudinal Equilibrium Distribution of Stored Electrons in the Presence of Self-Fields”. *Il Nuovo Cimento B* 18 (1973), pp. 72–82. ISSN: 1826-9877.
- [Ham17] Hamamatsu. *Streak cameras*. 2017. URL: <http://www.hamamatsu.com/jp/en/product/category/5001/5011/index.html> (visited on 08/30/2017).
- [HDF516] *HDF5 home page*. 2016. URL: <http://www.hdfgroup.org/HDF5/> (visited on 06/22/2016).
- [Hil+11] N. Hiller, A. Hofmann, E. Huttel, V. Judin, B. Kehrer, M. Klein, S. Marsching, and A. Müller. “Status of Bunch Deformation and Lengthening Studies at the ANKA Storage Ring”. *Proc. IPAC*. (San Sebastián, Spain). International Particle Accelerator Conference 2. Geneva, Switzerland: JACoW, 2011, pp. 2951–2953. URL: <http://jacow.org/ipac2011/papers/thpc021.pdf>.
- [Hil+13] N. Hiller, A. Borysenko, E. Hertle, et al. “Electro-Optical Bunch-Length Measurements at the ANKA Storage Ring”. *Proc. IPAC*. (Shanghai, China). International Particle Accelerator Conference 4. Geneva, Switzerland: JACoW, May 2013, pp. 500–503. ISBN: 978-3-95450-122-9. URL: <http://jacow.org/ipac2013/papers/mopme014.pdf>.
- [Hil13] N. Hiller. “Electro-Optical Bunch Length Measurements at the ANKA Storage Ring”. PhD thesis. Karlsruhe Institute of Technology, 2013. URL: <https://publikationen.bibliothek.kit.edu/1000041159>.
- [Hin08] F. Hinterberger. *Physik der Teilchenbeschleuniger und Ionenoptik*. Springer Berlin Heidelberg, 2008. ISBN: 9783540752820.
- [HSA16] *HSA Foundation*. 2016. URL: <http://www.hsafoundation.com/standards/> (visited on 04/13/2017).
- [IBM17a] *fmadd or fma (Floating Multiply-Add) instruction*. 2017. URL: http://www.ibm.com/support/knowledgecenter/ssw_aix_61/com.ibm.aix.alangref/idalangref_fmadd_instrs.htm (visited on 04/13/2017).
- [IBM17b] *Simultaneous multithreading at IBM Knowledge Center*. 2017. URL: http://www.ibm.com/support/knowledgecenter/ssw_aix_61/com.ibm.aix.genprog/smt.htm (visited on 04/13/2017).
- [IEEE08] *IEEE Standard for Floating-Point Arithmetic*. 2008. DOI: 10.1109/IEEESTD.2008.4610935.
- [Intel14] *Intel Core i7-5820K Processor*. 2014. URL: <https://ark.intel.com/products/82932/> (visited on 10/18/2017).
- [Intel17] *Intel Hyper-Threading Technology*. 2017. URL: <http://www.intel.eu/content/www/eu/en/architecture-and-technology/hyper-threading/hyper-threading-technology.html> (visited on 04/13/2017).
- [Jac98] J. D. Jackson. *Classical Electrodynamics*. 3rd. Wiley, 1998. ISBN: 978-0471309321.

- [Jam+06] S. Jamison, G. Berden, A. MacLeod, D. Jaroszynski, B. Redlich, A. van der Meer, and W. Gillespie. “Electro-optic techniques for temporal profile characterisation of relativistic Coulomb fields and coherent synchrotron radiation”. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 557.1 (2006), pp. 305–308. ISSN: 0168-9002. DOI: 10.1016/j.nima.2005.10.090.
- [Jud13] V. Judin. “Untersuchung von Bunch-Bunch-Wechselwirkungen und des Einflusses der geometrischen Impedanz bei der Erzeugung kohärenter THz-Strahlung”. PhD thesis. Karlsruhe Institute of Technology, 2013. URL: <http://digbib.ubka.uni-karlsruhe.de/volltexte/1000039528>.
- [Kam+00] A. Kaminski, J. Mesot, H. Fretwell, et al. “Quasiparticles in the Superconducting State of $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$ ”. *Phys. Rev. Lett.* 84 (8 Feb. 2000), pp. 1788–1791. DOI: 10.1103/PhysRevLett.84.1788.
- [Kar+10] E. Karantzoulis, G. Penco, A. Perucchi, and S. Lupi. “Characterization of coherent THz radiation bursting regime at Elettra”. *Infrared Physics & Technology* 53.4 (2010), pp. 300–303. ISSN: 1350-4495. DOI: 10.1016/j.infrared.2010.04.006.
- [Keh+17a] B. Kehrer, E. Blomley, M. Brosi, et al. “Time-Resolved Energy Spread Studies at the ANKA Storage Ring”. *Proc. IPAC*. (Copenhagen, Denmark). International Particle Accelerator Conference 8. Geneva, Switzerland: JACoW, May 2017, pp. 53–56. ISBN: 978-3-95450-182-3. DOI: 10.18429/JACoW-IPAC2017-M00CB1.
- [Keh+17b] B. Kehrer, M. Brosi, J. L. Steinmann, et al. “Simultaneous Detection of Longitudinal and Transverse Bunch Signals at the KARA Storage Ring”. *ArXiv e-prints* (Sept. 2017). submitted to PRAB. arXiv: 1709.08973 [physics.acc-ph].
- [Kha06] S. Khan. *Collective Phenomena in Synchrotron Radiation Sources*. Springer, 2006. ISBN: 978-3-642-07068-6.
- [Kle12] M. Klein. “Optics Calculations and Simulations of Longitudinal Beam Dynamics for the Low-Alpha Mode at ANKA”. PhD thesis. Karlsruhe Institute of Technology, 2012. ISBN: 978-3-8439-0647-0.
- [Kot33] V. A. Kotelnikov. “On the transmission capacity of the ‘ether’ and of cables in electrical communications”. *Proceedings of the first All-Union Conference on the technological reconstruction of the communications sector and the development of low-current engineering*. (Moscow, USSR). Jan. 1933. URL: <http://ict.open.ac.uk/classics/1.pdf>.
- [Kre+93] P. Krejcik, K. Bane, P. Corredoura, et al. “High Intensity Bunch Length Instabilities in the SLC Damping Rings”. *Proceedings of PAC’93*. 1993, p. 3240. URL: http://accelconf.web.cern.ch/accelconf/p93/PDF/PAC1993_3240.PDF.
- [KS69] E. Keil and W. Schnell. *Concerning longitudinal stability in the ISR*. Tech. rep. CERN-ISR-TH-RF-69-48. ISR-TH-RF-69-48. Geneva: CERN, July 1969. URL: <https://cds.cern.ch/record/1229157>.
- [KST03] E. Karantzoulis, V. Smaluk, and L. Tosi. “Broad band impedance measurements on the electron storage ring ELETTRA”. *Phys. Rev. ST Accel. Beams* 6 (3 Mar. 2003), p. 030703. DOI: 10.1103/PhysRevSTAB.6.030703.

- [Kus09] P. Kuske. “Investigation of the Temporal Structure of CSR-Bursts at BESSY II”. *Proc. of Particle Accelerator Conference 09*. (Vancouver, BC, Canada). Particle Accelerator Conference. Geneva, Switzerland: JACoW, 2009, pp. 4682–4684. URL: <https://accelconf.web.cern.ch/accelconf/pac2009/papers/fr5rfp063.pdf>.
- [Kus12] P. Kuske. “Calculation of Longitudinal Instability Threshold Currents for Single Bunches”. *Proceedings of the 11th International Computational Accelerator Physics Conference*. (Rostock-Warnemünde, Germany). International Computational Accelerator Physics Conference 11. Geneva, Switzerland: JACoW, 2012, pp. 267–269. ISBN: 978-3-95450-116-8. URL: <http://accelconf.web.cern.ch/AccelConf/ICAP2012/papers/thxdc3.pdf>.
- [Kus17] P. Kuske. “Short Bunches at the Transition From Strong to Weak Longitudinal Instability”. *Proc. IPAC*. (Copenhagen, Denmark). International Particle Accelerator Conference 8. Geneva, Switzerland: JACoW, May 2017, pp. 3696–3699. ISBN: 978-3-95450-182-3. DOI: 10.18429/JACoW-IPAC2017-THPAB007.
- [KW85] S. Krinsky and J. M. Wang. “Longitudinal Instabilities of Bunched Beams Subject to a Nonharmonic RF Potential”. *Part. Accel.* 17 (1985), pp. 109–139.
- [Lar97] J. Larmor. “Philosophical Transactions of the Royal Society of London: Containing papers of a mathematical or physical character.” A. Royal Society (Great Britain), 1897. Chap. On a Dynamical Theory of the Electric and Luminiferous Medium, Part 3, Relations with material media.
- [Lié98] A. Liénard. *La Théorie de Lorentz et celle de Larmor*. Paris, France: G. Carré et C. Naud, 1898.
- [Liu+09] J. Liu, D. J. Stevens, L. F. Haire, P. A. Walker, P. J. Coombs, R. J. Russell, S. J. Gamblin, and J. J. Skehel. “Structures of receptor complexes formed by hemagglutinins from the Asian Influenza pandemic of 1957”. *Proceedings of the National Academy of Sciences* 106.40 (2009), pp. 17175–17180. DOI: 10.1073/pnas.0906849106.
- [Lop+18] M. K. Lopes, L. B. França, L. A. V. Dias, and A. M. da Cunha. “Enhancing Range Analysis in Software Design Models by Detecting FloatingPoint Absorption and Cancellation”. *Information Technology – New Generations, 14th International Conference on Information Technology*. Advances in Intelligent Systems and Computing. Springer International Publishing, 2018, pp. 453–458. ISBN: 978-3-319-54977-4. DOI: 10.1007/978-3-319-54978-1.
- [LS64] C. C. Lin and F. H. Shu. “On the Spiral Structure of Disk Galaxies.” *The Astrophysical Journal* 140 (Aug. 1964), p. 646. DOI: 10.1086/147955.
- [Lük17] H. D. Lüke. *For the first theoretically exact formulation of the sampling theorem*. 2017. URL: <http://www.eduard-rhein-stiftung.de/en/fur-die-erstmalige-theoretisch-exakte-formulierung-des-abtasttheorems/> (visited on 11/27/2017).
- [Mey14] S. Meyers. *Effective Modern C++: 42 Specific Ways to Improve Your Use of C++11 and C++14*. 1st. O’Reilly Media, Inc., 2014. ISBN: 9781491903995.
- [MKG97] J. Murphy, S. Krinsky, and R. Gluckstern. “Longitudinal Wakefield for an Electron Moving on a Circular Orbit”. *Particle Accelerators* 57 (1997), pp. 9–64. URL: <http://cds.cern.ch/record/1120287/files/p9.pdf>.
- [Moo95] M. Moore. “Synchrotron X-ray topography”. *Radiation Physics and Chemistry* 45.3 (1995), pp. 427–444. ISSN: 0969-806X. DOI: 10.1016/0969-806X(94)E0061-M.

- [MPG04] Max-Planck Gesellschaft, ed. *The fastest stopwatch in the world*. 2004. URL: <https://www.mpg.de/495195/pressRelease200402241>.
- [Mül+04] A.-S. Müller, M. Pont, F. Pérez, F. Zimmermann, E. Huttel, and I. Birkel. “Investigation of Scraper Induced Wake Fields at ANKA”. *Proc. of European Particle Accelerator Conference (EPAC 2004)*. European Particle Accelerator Conference. Geneva, Switzerland: JACoW, 2004. URL: <http://accelconf.web.cern.ch/AccelConf/e04/PAPERS/WEPLT069.PDF>.
- [Mül+05] A. S. Müller, I. Birkel, B. Gasharova, et al. “Far infrared coherent synchrotron edge radiation at ANKA”. *Proceedings of the Particle Accelerator Conference*. (Knoxville, USA). Geneva, Switzerland: JACoW, May 2005, pp. 2518–2520. URL: <https://accelconf.web.cern.ch/accelconf/p05/PAPERS/RPAE038.PDF>.
- [Mül+12] A.-S. Müller, N. Hiller, A. Hofmann, et al. “Experimental aspects of CSR in the ANKA storage ring”. *ICFA Beam Dynamics Newsletter* 57 (2012), pp. 154–165. URL: <http://www-bd.fnal.gov/icfabd/Newsletter57.pdf>.
- [Mül+13] A.-S. Müller, V. Judin, M. Balzer, et al. “Studies of Bunch-Bunch Interactions in the ANKA storage ring with coherent synchrotron radiation using an ultra-fast terahertz detection system”. *Proc. IPAC*. (Shanghai, China). International Particle Accelerator Conference 4. Geneva, Switzerland: JACoW, May 2013, pp. 109–111. ISBN: 978-3-95450-122-9. URL: <http://jacow.org/ipac2013/papers/mopea019.pdf>.
- [Mül10a] A.-S. Müller. “Accelerator-Based Sources of Infrared and Terahertz Radiation”. *Reviews of Accelerator Science and Technology* 03.01 (2010), pp. 165–183. DOI: 10.1142/S1793626810000427.
- [Mül10b] A.-S. Müller. “Accelerator-Based Sources of Infrared and Terahertz Radiation”. *Reviews of Accelerator Science and Technology* 03.01 (2010), pp. 165–183. DOI: 10.1142/S1793626810000427.
- [MV86] G. Mourou and J. Valdmanis. *Measurement of electrical signals with subpicosecond resolution*. US Patent 4,618,819. Oct. 1986.
- [Ng02] K. Y. Ng. *Physics of Intensity Dependent Beam Instabilities*. Tech. rep. Fermilab-FN-0713. Fermilab, 2002. URL: http://lss.fnal.gov/cgi-bin/find_paper.pl?fn-0713.
- [Ng06] K. Y. Ng. *Physics of Intensity Dependent Beam Instabilities*. World Scientific Publishing, 2006. ISBN: 978-981-270-339-2. DOI: 10.1142/5835.
- [Nor+98] M. R. Norman, H. Ding, M. Randeria, et al. “Destruction of the Fermi surface in underdoped high-Tc superconductors”. *Nature* 392 (Mar. 1998), p. 157. URL: <http://dx.doi.org/10.1038/32366>.
- [NS54] J. S. Nodvick and D. S. Saxon. “Suppression of Coherent Radiation by Electrons in a Synchrotron”. *Phys. Rev.* 96 (1 Oct. 1954), pp. 180–184. DOI: 10.1103/PhysRev.96.180.
- [NV17] *NVidia Titan X*. 2017. URL: <http://www.nvidia.com/en-us/geforce/products/10series/titan-x-pascal/> (visited on 04/13/2017).
- [OCL16] *OpenCL home page*. 2016. URL: <http://www.khronos.org/opencv/> (visited on 02/28/2017).
- [OGL17] *OpenGL Wiki*. 2017. URL: <https://www.khronos.org/opengl/wiki/GLAPI/glDepthRange> (visited on 12/04/2017).
- [Opt17] Optronis GmbH. *Streak-Camera*. 2017. URL: <http://optronis.com/en/productcategories/streak-camera/> (visited on 08/30/2017).

- [PB16] J.-P. Provost and C. Bracco. “The 1895 Lorentz transformations: historical issues and present teaching”. *European Journal of Physics* 37.4 (2016), p. 045601. URL: <http://stacks.iop.org/0143-0807/37/i=4/a=045601>.
- [Ple+09] A. Plech, S. Casalbuoni, B. Gasharova, et al. “Electro-Optical Sampling of Terahertz Radiation Emitted by Short Bunches in the ANKA Synchrotron”. *Proc. of Particle Accelerator Conference 09*. (Vancouver, BC, Canada). Particle Accelerator Conference. Geneva, Switzerland: JACoW, 2009, pp. 1150–1152. URL: <http://accelconf.web.cern.ch/AccelConf/PAC2009/papers/tu5rfp026.pdf>.
- [PNG16] *PNG specification*. 2016. URL: <http://www.w3.org/TR/PNG/> (visited on 06/22/2016).
- [Poc06] F. C. A. Pockels. *Lehrbuch der Kristallogoptik*. Sammlung von Lehrbüchern auf dem Gebiete der mathematischen Wissenschaften mit Einschluss ihrer Anwendungen XIX. Leipzig, Berlin: Teubner, 1906. ISBN: 978-9333661065.
- [Pod99] B. V. Podobedov. “Saw-tooth instability studies at the Stanford Linear Collider damping rings”. PhD thesis. SLAC, 1999. URL: <http://www.slac.stanford.edu/spires/find/books/www?cl=QCD191:P6:1999>.
- [Pre+07] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. 3rd ed. New York, NY, USA: Cambridge University Press, 2007. ISBN: 9780521880688.
- [Pres17] *Atomic vs. Non-Atomic Operations (Preshing on Programming)*. 2017. URL: <http://preshing.com/20130618/atomic-vs-non-atomic-operations/> (visited on 02/01/2017).
- [PS98] B. Podobedov and R. Siemann. “New apparatus for precise synchronous phase shift measurements in storage rings”. *Phys. Rev. ST Accel. Beams* 1 (7 Nov. 1998), p. 072801. DOI: 10.1103/PhysRevSTAB.1.072801.
- [Qin+01] Q. Qin, Z. Guo, G. Xu, L. Ma, and Z. Zhao. “Studies of bunch lengthening at the BEPC”. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 463.1 (2001), pp. 77–85. ISSN: 0168-9002. DOI: 10.1016/S0168-9002(01)00463-6.
- [RBT11] G. Rehm, C. Bloomer, and C. Thomas. “Loss Factor Measurement Using Time Correlated Single Photon Counting of Synchrotron Radiation”. *Proceedings of the 10th European Workshop on Beam Diagnostics and Instrumentation for Particle Accelerators*. (Hamburg, Germany). 5. Geneva, Switzerland: JACoW, May 2011, pp. 110–112. URL: <http://accelconf.web.cern.ch/AccelConf/DIPAC2011/papers/mopd28.pdf>.
- [RCH11] D. Ratner, A. Chao, and Z. Huang. “Two-chicane compressed harmonic generation of soft x rays”. *Phys. Rev. ST Accel. Beams* 14 (2 Feb. 2011), p. 020701. DOI: 10.1103/PhysRevSTAB.14.020701.
- [RM73] E. Rowe and F. Mills. “Tantalus I: a dedicated storage ring synchrotron radiation source”. *Particle Accelerators* 4 (4 1973), pp. 211–227. URL: cds.cern.ch/record/1107919/files/p211.pdf.
- [Rot+17] L. Rota, M. Balzer, M. Caselle, et al. “KALYPSO: A Mfps Linear Array Detector for Visible to NIR Radiation”. *Proc. of International Beam Instrumentation Conference (IBIC'16), Barcelona, Spain, Sept. 13-18, 2016*. (Barcelona, Spain). International Beam Instrumentation Conference 5. Geneva, Switzerland: JACoW, Sept. 2017, pp. 741–744. ISBN: 978-3-95450-177-9. DOI: 10.18429/JACoW-IBIC2016-WEPEG46.

- [Rot18] L. Rota. “KALYPSO, a novel detector system for high-repetition rate and real-time beam diagnostics”. PhD thesis. Karlsruhe Institute of Technology, 2018. DOI: 10.5445/IR/1000082349.
- [Rou+14a] E. Roussel, C. Evain, C. Sz waj, and S. Bielawski. “Microbunching instability in storage rings: Link between phase-space structure and terahertz coherent synchrotron radiation radio-frequency spectra”. *Phys. Rev. ST Accel. Beams* 17 (1 Jan. 2014), p. 010701. DOI: 10.1103/PhysRevSTAB.17.010701.
- [Rou+14b] E. Roussel, C. Evain, C. Sz waj, et al. “Microbunching Instability in Relativistic Electron Bunches: Direct Observations of the Microstructures Using Ultrafast YBCO Detectors”. *Phys. Rev. Lett.* 113 (9 Aug. 2014), p. 094801. DOI: 10.1103/PhysRevLett.113.094801.
- [Rou+15] E. Roussel, C. Evain, M. Le Parquier, et al. “Observing microscopic structures of a relativistic object using a time-stretch strategy”. *Scientific Reports* 5 (May 2015). DOI: 10.1038/srep10330.
- [Rou14] E. Roussel. “Spatio-temporal dynamics of relativistic electron bunches during the microbunching instability : study of the Synchrotron SOLEIL and UVSOR storage rings.” PhD thesis. Université Lille1 - Sciences et Technologies, 2014. URL: <https://tel.archives-ouvertes.fr/tel-01112910>.
- [San69] M. Sands. “The Physics of Electron Storage Rings: An Introduction”. *International School of Physics, Enrico Fermi, Course XLVI: Physics with Intersecting Storage Rings*. Vol. C6906161. June 1969, pp. 257–411. URL: <http://slac.stanford.edu/pubs/slacreports/reports02/slac-r-121.pdf>.
- [SBZ07] G. Stupakov, K. L. F. Bane, and I. Zagorodnov. *Optical approximation in the theory of geometric impedance*. Tech. rep. SLAC, DESY, 2007. URL: <http://www.slac.stanford.edu/pubs/slacpubs/12250/slac-pub-12369.pdf>.
- [Sch+13] P. Schönfeldt, N. Hiller, V. Judin, A. Müller, M. Schwarz, and J. Steinmann. “Comparison of different approaches to determine the bursting threshold at ANKA”. *Proc. IPAC*. (Shanghai, China). International Particle Accelerator Conference 4. Geneva, Switzerland: JACoW, May 2013. ISBN: 978-3-95450-122-9. URL: <http://jacow.org/ipac2013/papers/mopea020.pdf>.
- [Sch+14] P. Schönfeldt, A. Borysenko, E. Hertle, et al. “Fluctuation of Bunch Length in Bursting CSR: Measurement and Simulation”. *Proc. IPAC*. (Dresden, Germany). International Particle Accelerator Conference 5. Geneva, Switzerland: JACoW, June 2014. ISBN: 978-3-95450-132-8. URL: <http://jacow.org/IPAC2014/papers/mopro068.pdf>.
- [Sch+15] P. Schönfeldt, N. Hiller, B. Kehrer, and A.-S. Müller. “Studies for a Wake-Field-Optimized Near-Field EO Setup at the ANKA Storage Ring”. *Proc. IPAC*. (Richmond, VA, USA). International Particle Accelerator Conference 6. Geneva, Switzerland: JACoW, June 2015, pp. 869–872. ISBN: 978-3-95450-168-7. DOI: 10.18429/JACoW-IPAC2015-MOPHA038.
- [Sch+16] P. Schönfeldt, M. Brosi, A.-S. Müller, and J. Steinmann. “A Parallelized Vlasov-Fokker-Planck-Solver for Desktop PCs”. *Proc. IPAC*. (Busan, Korea). International Particle Accelerator Conference 7. Geneva, Switzerland: JACoW, June 2016, pp. 1658–1661. ISBN: 978-3-95450-147-2. DOI: doi:10.18429/JACoW-IPAC2016-TUPOR005.

- [Sch+17a] P. Schönfeldt, E. Blomley, E. Bründermann, et al. “Towards Near-Field Electro-Optical Bunch Profile Monitoring in a Multi-Bunch Environment”. *Proc. IPAC*. (Copenhagen, Denmark). International Particle Accelerator Conference 8. Geneva, Switzerland: JACoW, May 2017, pp. 227–230. ISBN: 978-3-95450-182-3. DOI: 10.18429/JACoW-IPAC2017-MOPAB055.
- [Sch+17b] P. Schönfeldt, M. Blaicher, T. Boltz, P. Schreiber, and J. L. Schestag. *Inovesa*. 2017. DOI: 10.5281/zenodo.597356.
- [Sch+17c] P. Schönfeldt, M. Brosi, M. Schwarz, J. L. Steinmann, and A.-S. Müller. “Parallelized Vlasov-Fokker-Planck solver for desktop personal computers”. *Phys. Rev. Accel. Beams* 20 (3 Mar. 2017), p. 030704. DOI: 10.1103/PhysRevAccelBeams.20.030704.
- [Sch+18a] P. Schönfeldt, T. Boltz, A. Mochihashi, J. L. Steinmann, and A.-S. Müller. “Elaborated Modeling of Synchrotron Motion in Vlasov-Fokker-Planck Solvers”. *Proc. IPAC*. (Vancouver, BC, Canada). International Particle Accelerator Conference 9. Accepted for *J. Phys.: Conf. Ser.* Geneva, Switzerland: JACoW, May 2018.
- [Sch+18b] P. Schönfeldt, B. Kehrer, E. Blomley, et al. “Study of the influence of the CSR impedance on the synchronous phase shift at KARA”. *Proc. IPAC*. (Vancouver, BC, Canada). International Particle Accelerator Conference 9. Geneva, Switzerland: JACoW, May 2018.
- [Sch07] G. A. Schott. “Über die Strahlung von Elektronengruppen”. *Annalen der Physik* 329 (1907), pp. 635–660. DOI: 10.1002/andp.19073291403.
- [Sch13] P. Schönfeldt. “Systematische Untersuchungen der Bestimmung von Burstingschwellen”. MA thesis. Karlsruhe Institute of Technology, 2013.
- [Sch18a] J. Schestag. “Study of Noise in the Simulation of Collective Effects of Relativistic Electron Bunches”. Bachelor’s Thesis. Karlsruhe Institute of Technology, Mar. 2018. DOI: 10.5445/IR/1000083507.
- [Sch18b] P. Schreiber. “Systematische Untersuchung des Einflusses der geometrischen Impedanz bei der Erzeugung kohärenter THz-Strahlung”. MA thesis. Karlsruhe Institute of Technology, 2018.
- [Sch49] J. Schwinger. “On the Classical Radiation of Accelerated Electrons”. *Phys. Rev.* 75 (12 June 1949), pp. 1912–1925. DOI: 10.1103/PhysRev.75.1912.
- [Sch98] J. Schwinger. *On Radiation by Electrons in a Betatron*. transcribed by Miguel A. Furman. 1998. URL: http://www.nicadd.niu.edu/~piot/phys_671/Papers/schwinger.pdf.
- [SH02] G. Stupakov and S. Heifets. “Beam instability and microbunching due to coherent synchrotron radiation”. *Phys. Rev. ST Accel. Beams* 5 (5 May 2002), p. 054402. DOI: 10.1103/PhysRevSTAB.5.054402.
- [She+11] Y. Shen, X. Yang, G. L. Carr, Y. Hidaka, J. B. Murphy, and X. Wang. “Tunable Few-Cycle and Multicycle Coherent Terahertz Radiation from Relativistic Electrons”. *Phys. Rev. Lett.* 107 (20 Nov. 2011), p. 204801. DOI: 10.1103/PhysRevLett.107.204801.
- [Shi+12] W. Shields, R. Bartolini, G. Boorman, P. Karataev, A. Lyapin, J. Puntree, and G. Rehm. “Microbunch Instability Observations from a THz Detector at Diamond Light Source”. *Journal of Physics: Conference Series* 357.1 (2012), p. 012037. URL: <http://stacks.iop.org/1742-6596/357/i=1/a=012037>.
- [Sony13] *PlayStation 4 includes hUMA technology*. 2013. URL: <http://vgleaks.com/playstation-4-includes-huma-technology/> (visited on 04/13/2017).

- [SSM09] B. Steffen, V. Schlott, and F. Müller. “A Compact Single Shot Electro-Optical Bunch Length Monitor for the SwissFEL”. *Proceedings of DIPAC09*. (Basel, Switzerland). DIPAC. 2009. URL: https://diagnostics.web.psi.ch/publications/DIPAC09_TUPB42.pdf.
- [Ste+17a] J. L. Steinmann, M. Brosi, E. Bründermann, et al. “Continuous bunch-by-bunch spectroscopic investigation of the micro-bunching instability”. *ArXiv e-prints* (Oct. 2017). Accepted for PRAB. arXiv: 1710.09568 [physics.acc-ph].
- [Ste+17b] J. Steinmann, E. Blomley, M. Brosi, et al. “4-Channel Single Shot and Turn-by-Turn Spectral Measurements of Bursting CSR”. *Proc. IPAC*. (Copenhagen, Denmark). International Particle Accelerator Conference 8. Geneva, Switzerland: JACoW, May 2017, pp. 231–234. ISBN: 978-3-95450-182-3. DOI: 10.18429/JACoW-IPAC2017-MOPAB056.
- [Ste07] B. Steffen. “Electro-Optic Methods for Longitudinal Bunch Diagnostics at FLASH”. PhD thesis. University of Hamburg, 2007.
- [Ste16] B. Steffen. private communication. ARD ST3 workshop. 2016.
- [Syd17] Sydor Technologies. *Streak Cameras*. 2017. URL: <http://sydortechnologies.com/imaging-detectors/streak-cameras/> (visited on 08/30/2017).
- [Tak+05] Y. Takashima, M. Katoh, M. Hosaka, A. Mochihashi, S.-i. Kimura, and T. Takahashi. “Observation of Intense Bursts of Terahertz Synchrotron Radiation at UVSOR-II”. *Japanese Journal of Applied Physics* 44.8L (2005), p. L1131. URL: <http://stacks.iop.org/1347-4065/44/i=8L/a=L1131>.
- [Tan08] V. Tan. *Image rotation with bilinear interpolation*. 2008. URL: <http://polymathprogrammer.com/2008/10/06/image-rotation-with-bilinear-interpolation/> (visited on 10/13/2017).
- [Ter+15] B. Terzić, K. Arumugam, A. Godunov, D. Ranjan, and M. Zubair. “High-Performance Simulations of Coherent Synchrotron Radiation on Multicore GPU and CPU Platforms”. *Proceedings, 6th International Particle Accelerator Conference (IPAC 2015): Richmond, Virginia, USA, May 3-8, 2015*. May 2015, MOBC2. URL: <http://accelconf.web.cern.ch/AccelConf/IPAC2015/papers/mobc2.pdf>.
- [TH56] D. H. Tomboulion and P. L. Hartman. “Spectral and Angular Distribution of Ultraviolet Radiation from the 300-Mev Cornell Synchrotron”. *Phys. Rev.* 102 (6 June 1956), pp. 1423–1447. DOI: 10.1103/PhysRev.102.1423.
- [Ung+17] P. Ungelenk, M. Höner, H. Huck, et al. “Continuously tunable narrowband pulses in the THz gap from laser-modulated electron bunches in a storage ring”. *Phys. Rev. Accel. Beams* 20 (2 Feb. 2017), p. 020706. DOI: 10.1103/PhysRevAccelBeams.20.020706.
- [VDI17] Virginia Diodes Inc. *Detectors*. 2017. URL: <http://vadiodes.com/en/products/detectors> (visited on 08/30/2017).
- [Ven+05] M. Venturini, R. Warnock, R. Ruth, and J. A. Ellison. “Coherent synchrotron radiation and bunch stability in a compact storage ring”. *Phys. Rev. ST Accel. Beams* 8 (1 Jan. 2005), p. 014202. DOI: 10.1103/PhysRevSTAB.8.014202.
- [Vla45] A. A. Vlasov. “Theory of Vibrational Properties of an Electron Gas and Its Applications”. *J. Phys. USSR* 9 (25 1945).
- [VMG83] J. Valdmanis, G. Mourou, and C. Gobel. “Subpicosecond electrical sampling”. *IEEE Journal of Quantum Electronics* 19.4 (1983), pp. 664–667.
- [Voi08] W. Voigt. *Magneto- und Elektrooptik*. Mathematische Vorlesungen an der Universität Göttingen III. Leipzig: Teubner, 1908.

- [VW02] M. Venturini and R. Warnock. “Bursts of Coherent Synchrotron Radiation in Electron Storage Rings: a Dynamical Model”. *Physical Review Letters* 89 (Dec. 2002), p. 224802. DOI: 10.1103/PhysRevLett.89.224802.
- [WE00] R. L. Warnock and J. A. Ellison. *A General Method for Propagation of the Phase Space Distribution, with Application to the Sawtooth Instability*. Tech. rep. SLAC-PUB-8404. SLAC, 2000. URL: <https://www.slac.stanford.edu/pubs/slacpubs/8250/slac-pub-8404.pdf>.
- [WebE18] *Aluminium (WebElements)*. 2018. URL: <https://www.webelements.com/aluminium/> (visited on 06/08/2018).
- [Wie15] H. Wiedemann. *Particle Accelerator Physics*. Graduate Texts in Physics. Springer International Publishing, 2015. ISBN: 978-3-540-49045-6.
- [Wiki17] *Flynn’s Taxonomy (Wikipedia)*. 2017. URL: http://en.wikipedia.org/wiki/Flynn’s_taxonomy (visited on 02/03/2017).
- [Wiki18] *Instruction-Level Parallelism (Wikipedia)*. 2018. URL: http://en.wikipedia.org/wiki/Instruction-level_parallelism (visited on 02/01/2017).
- [Wil+02] I. Wilke, A. M. MacLeod, W. A. Gillespie, G. Berden, G. M. H. Knippels, and A. F. G. van der Meer. “Single-Shot Electron-Beam Bunch Length Measurements”. *Phys. Rev. Lett.* 88 (12 Mar. 2002), p. 124801. DOI: 10.1103/PhysRevLett.88.124801.
- [Wil01] K. Wille. *The Physics of Particle Accelerators: An Introduction*. Trans. by J. McFall. Clarendon Press, 2001. ISBN: 978-0198505495.
- [Wil96] K. Wille. *Physik der Teilchenbeschleuniger und Synchrotronstrahlungsquellen*. Teubner, 1996. ISBN: 978-3519130871.
- [Wol14a] A. Wolski. *Beam Dynamics in High Energy Particle Accelerators*. London, United Kingdom: Imperial College Press, 2014. ISBN: 978-1-78326-277-9.
- [Wol14b] A. Wolski. “Low-emittance storage rings” (2014), pp. 245–294. DOI: 10.5170/CERN-2014-009.245.
- [Wüs+10] G. Wüstefeld, J. Feikes, M. Hartrott, et al. “Coherent THz Measurements at the Metrology Light Source”. *Proc. IPAC*. (Kyoto, Japan). International Particle Accelerator Conference 1. Geneva, Switzerland: JACoW, May 2010, pp. 2508–2510. URL: <http://jacow.org/ipac10/papers/wepea015.pdf>.
- [WZ95] Q. Wu and X. Zhang. “Free-space electro-optic sampling of terahertz beams”. *Applied Physics Letters* 67.24 (1995), pp. 3523–3525. DOI: 10.1063/1.114909.
- [Yar+84] J. Yarwood, T. Shuttleworth, J. B. Hasted, and T. Nanba. “A new radiation source for the infrared region”. *Nature* 312 (Dec. 1984), pp. 742–744. DOI: 10.1038/312742a0.
- [ZMM99] M. P. Zorzano, H. Mais, and L. V. Martínez. “Numerical solution of two dimensional Fokker - Planck equations”. *Applied Mathematics and Computation* 98 (1999), pp. 109–117. DOI: 10.1016/S0096-3003(97)10161-8.

Acknowledgments

First of all, I want to thank Prof. Anke-Susanne Müller and Prof. Günter Quast for giving me the opportunity to write this thesis. It was very motivating to be part of an ambitious project while at the same time being able and encouraged to follow my personal curiosity.

The project was funded by the BMBF. I feel also obliged for the support of the Helmholtz International Research School for Teratronics (HIRST) that allowed to broaden my horizon also beyond the scope of physics.

All this would not have been possible without the help and support of my colleagues, especially at the IBPT-IPS-[AN(KA)RA]-LAS¹. In particular, Gudrun Niehus and Michael Nasse maintained and repaired the EO Laser and also operated it many times. Sophie Walther also did EO laser operation, assembled the spectrometers used with KALYPSO and programmed some measurement scripts. Stefan Funkner also joined EO measurement shifts. In particular, he helped with the alignment of the optical path. Edmund Blomley implemented the logging of the EO settings and a feedback loop that allowed to let the laser run freely where previously manual adjustments were needed. Igor Kriznar coped with the motors of the EO setup, especially their control using Ethernet. Nicole Hiller introduced EO to me. She did most of the implementation of the original setup that served as a proof-of-concept for near-field EO at a storage ring. When my setup was installed, she found the first signal during the initial alignment. I am deeply grateful to have had her around. Lorenzo Rota implemented, later operated and optimized KALYPSO to meet the needs of us “accelerator guys”. He also came around with the first user interface to simplify the usage of KALYPSO.

Miriam Brosi did loads of tests in the early days of *Inovesa*. One could say she served as my “local user” for usability tests. She also was my personal Python guru and helped me out with reference THz measurement data. My insights profited a lot from our shared discussions and from her feedback. Michael Holz gave me an introduction to CST. We also had discussions on beam dynamics and possible beam diagnostics. Tobias Boltz joined later. While writing his inspiring work on bunch shapes, their classification and sub-structures, he did (involuntary) *Inovesa* beta testing. With him, I had the first discussions about the different contributions to the CSR spectrum. Patrick Schreiber continued our work. He also pushed the usability of *Inovesa* to a level that I feel can be called “1.0”. Johannes Steinmann was my engineer on duty. With him, I had the most solution-focused discussions e.g. on impedances. He is also the key person that brought me to the conclusion to work in frequency domain. Markus Schwarz explained tons of equations, mathematical definitions, etc. to me. He also double-checked some of my calculations.

Benjamin Kehrer did the first wake field simulations for the original EO arm. He also was the first to study the dynamics of the energy spread for KARA. His work triggered much of what I know today. Steffen Schott helped with engineering problems I had with the optimized EO arm. My first draft was close to optimal from an electro-magnetic point

¹The institute’s name and belonging of personal changed several times in the last years.

of view but not even close to be manufacturable. We had to iterate until both design parameters were met. Marcel Schuh took his time to review my early work and returned it with a lot of feedback. He and Erik Bründermann also had an ear when I needed advice with organizational stuff.

From other institutes, I particularly thank Christophe Sz waj, Serge Bielawski, and Clement Evain, who joined our EO measurement shifts. During that time and at a number of workshops, they also shared some of their insights on Vlasov-Fokker-Planck solvers. Gennady Stupakov first hinted me to the method I referred to as “Manhattan rotation”. Bernd Steffen shared his experiences with the X-FEL EO laser system which is identical to the one we used.

At last, I want to thank my family. My parents always encouraged me to follow my interests, but were understanding when I decided to keep it relaxed. Especially my father, Michel, was a great example when it comes to mental toughness. I know you would be proud. Melanie, my beloved wife, I would not want to be without her. Melanie, it is good to know that we can rely on each other. I love you from the bottom of my heart. Jakob, our little son. He let me type parts of my thesis one-handedly while sometimes sleeping in my other arm. These are the first letters you typed: “f pöudsss j jjb bxcv”

To all of you and also to those who I did not mention explicitly: Thank you all!