Know Your Qubits, Know Your Users: Personas for Quantum Software

Lukas Schmidbauer*
OTH Regensburg
Regensburg, Germany
lukas.schmidbauer@othr.de

Joshua Ammermann*
KIT
Karlsruhe, Germany
joshua.ammermann@kit.edu

Laura Schulz
Argonne National Laboratory
Lemont, USA
lschulz@anl.gov

Jose Garcia-Alonso
University of Extremadura
Badajoz, Spain
jgaralo@unex.es

Robert Wille
TUM
Munich, Germany
robert.wille@tum.de

Sebastian Feld
TU Delft
Delft, Netherlands
s.feld@tudelft.nl

Ina Schaefer
KIT
Karlsruhe, Germany
ina.schaefer@kit.edu

Wolfgang Mauerer
OTH Regensburg
Regensburg, Germany
wolfgang.mauerer@othr.de

Abstract

The advancement of quantum hardware and the intricacies of quantum computing make well-designed quantum software increasingly necessary. Due to the interdisciplinarity of the field, it is crucial to understand the perspectives and specific needs of involved stakeholders, for example, to balance the desired level of abstraction with the exposition of (hardware)-specific details. In this work, we conduct a stakeholder-based analysis to identify personas of quantum software as a means of creating meaningful, user-tailored quantum software. We conducted an expert focus group at a Dagstuhl seminar in 2024 and qualitative interviews with practitioners at conference IEEE QCE in 2025, from which we derive eleven personas of potential users and stakeholders for quantum software. We discuss these personas regarding their use cases, interests, constraints and abstraction level.

CCS Concepts

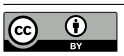
• **Hardware** → **Quantum computation**; • **Software and its engineering**; • **Human-centered computing** → *Interaction design*;

Keywords

Quantum Software, Quantum Software Engineering, Personas

ACM Reference Format:
Lukas Schmidbauer, Joshua Ammermann, Laura Schulz, Jose Garcia-Alonso, Robert Wille, Sebastian Feld, Ina Schaefer, and Wolfgang Mauerer. 2026. Know Your Qubits, Know Your Users: Personas for Quantum Software. In *ACM/IEEE 29th International Conference on Model Driven Engineering Languages and Systems (MODELS Companion 2026), October 04–09, 2026, Málaga, Spain.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3837062.3839073

*Both authors contributed equally to this research.



This work is licensed under a Creative Commons Attribution 4.0 International License.
MODELS Companion 2026, Málaga, Spain
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2903-4/26/10
https://doi.org/10.1145/3837062.3839073

1 Introduction

Quantum computing, as a continuously developing area of research, promises computational or other advantage over classical approaches [1, 15, 46]. Researchers are lowering the bar for achieving practical advantages in quantum agnostic problems with recent advances in mappings to decoding problems [9, 26–28, 41].

Exploiting the advantages of maturing quantum hardware [12] requires the development of practical quantum software, which has become a research focus of its own [8, 33, 55]. One of the key issues in quantum software engineering is designing languages and, in particular, abstractions that can make quantum computing accessible to users from many domains without direct relation to quantum. While languages and frameworks have become available [19], it is common practice to incorporate quantum-specific constructs into existing general-purpose programming languages (*e.g.*, Python). For example, the widely adopted frameworks Qiskit and Cirq employ this approach [50], albeit special-purpose languages have also been proposed [6, 31, 47]. Both offer an abstracted view on hardware and easy deployment of algorithms expressed as quantum circuits, which benefits higher level applications and users with programming experience. Conversely, for some other users, hardware-specific details (*e.g.*, error-rates, topology, qubit type) leverage the (application dependent) optimization potential of available systems (*e.g.*, Co-Design approaches can use exposed details [30, 42, 43]). Thus, balancing the level of abstraction and exposition of details is a non-trivial, but necessary task in the development of quantum software.

Model-Driven Engineering (MDE) applied to quantum software [3, 23] supports the software development process by leveraging models as central artifacts. One major advantage of MDE is that models are created for a specific purpose and stakeholder [51]. Thus, MDE can provide domain experts with tailored Domain-Specific Languages (DSLs) that expose relevant details of quantum software in a familiar way. Researchers in the field of quantum mechanics, for example, could benefit from an interface that allows them to express Hamiltonians in line with their mental models. The work of Murillo et al. [33] gives a comprehensive overview of current quantum MDE research. Ongoing work includes defining domain-specific meta-models for quantum circuits [3, 24] and the high-level Qrisp quantum programming language [7], and extending the widely used Unified Modeling Language (UML) for quantum computing [38].

As an emerging technology spanning a wide range of disciplines, quantum computing involves a variety of domain experts, who could benefit from representations tailored to their needs, such as DSLs. Involved stakeholders range from end users who want to solve problems with quantum computing to engineers who are building the necessary quantum hardware. This interdisciplinarity becomes particularly evident in the field of quantum software, as exemplified by major quantum software events, such as the Munich Quantum Software Forum (MQSF) [52], where the number of attendees with a background in physics still exceeds those with a background in computer science. Therefore, to develop quantum software tailored to stakeholders, it is essential to understand their specific needs and perspectives. Although those working on different parts of quantum software each consider a specific portion of stakeholders (relevant for their project), knowledge about these stakeholders is fragmented across the various workplaces, institutions and smaller specialist communities involved [18].

To address the lack of consolidated stakeholder characterizations, we conduct a *persona-based analysis* to identify and characterise stakeholder groups for quantum software. We use *personas* [11]—a method from interaction design—to define user groups. Personas support practitioners and researchers across domains in the design process and have been widely adopted and adapted [10, 37, 44]. They provide a model for thinking and communicating about user groups, serving as a common language and facilitating communication. The quantum MDE community can benefit from personas, as they provide information on who the tailored models and DSLs are being crafted for. In general, personas are used to design meaningful interfaces that users find satisfactory. Future practitioners may adopt the Goal-Directed Design process [11] (see Sec. 2) in its entirety when designing quantum software products. Therefore, readily available personas for quantum software support the ongoing initiative to create purposeful, engaging quantum software that meets the needs of many different stakeholders. More precisely, to advance towards this goal, in this work we:

- Present results from two studies we conducted at different venues to derive personas for quantum software:
  (a) To incorporate the expertise of the specialists in the field, we conducted a focus group with quantum software engineering experts at the Dagstuhl Seminar 24512 in 2024 (see Sec. 4). The focus group yielded five short-, six mid-, and five long-term personas.
  (b) To incorporate the perspective of practitioners actively working on quantum software, we conducted qualitative interviews with practitioners at IEEE QCE 2025 (see Sec. 5). The interviewees identified nine personas, focusing on the short-term perspective.
- Aggregate 11 personas for quantum software from the insights gathered in the studies (see Sec. 6).

2 Background: Personas in Interaction Design

Interaction Design is concerned with identifying a product's core users and aims to create efficient and appealing design solutions by understanding users and their goals. In this context, Cooper introduced the *Goal-Directed Design* [11] process, which uses *personas* for thinking and communicating about user groups.

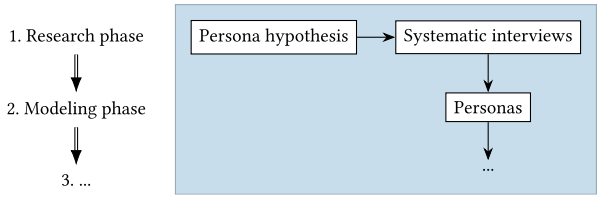


Figure 1: First phases (left) and corresponding activities (right) of the Goal-Directed Design process adapted from [11]. Personas are used in later phases to design user interfaces.

In Goal-Directed Design, three types of models are distinguished: user models (*i.e.*, conceptional or mental models), represented models, and implementation models. Implementation models reflect technology, whereas user models reflect a user's vision, hiding the complex mechanisms they do not need to know or care about. The represented models lie between the designers and the users, capturing how the designers display an application's functionality. Represented models that are closer to the user models are easier for them to understand.

The Goal-Directed Design process has six phases, of which we show the first two in Fig. 1 (left). The initial *research phase* aims to collect qualitative data on potential or actual users of the product. To achieve this, systematic field study techniques are employed, such as user observations and interviews. Interviews are typically conducted one-on-one with broad, open-ended questions, lasting no more than an hour in early interviews, as they focus on gathering domain knowledge from the user's perspective. In this phase, *persona-hypotheses* (see Fig. 1: right) are created based on information from subject matter experts, stakeholders and literature reviews, and serve as an initial attempt at defining the different kinds of users. The outcome of this phase is a set of behaviour patterns that help categorise modes of use of a potential or existing product. In technical domains, they often map to professional roles and are the basis for personas in the modelling phase.

The *modeling phase* uses *personas* as a user model (see Fig. 1), synthesised from previously acquired patterns. Personas (*i.e.*, descriptive models) provide a precise way of thinking and communicating about users. They are based on behaviours and motivations from real persons, and capture users' behaviour, thinking, goals and reasoning. While personas represent groups of users, they are presented as individual humans and thus have a distinct narrative structure. Hence, they provide a unique common language between stakeholders, developers, and designers. In contrast, other user models do not attempt to convey broader human motives and contexts. Personas focus on a specific domain and product, so only personas based on research covering similar contexts should be reused.

In the following phases of the Goal-Directed Design process, the defined personas are used to design meaningful interfaces which are satisfactory for the users. We see immediate benefits in using personas for quantum software, in that they serve as a common language. Though, future practitioners may fully adapt the Goal-Directed Design process to design quantum software products.

3 Research Questions and Methodology

When searching for meaningful quantum software models and DSLs, it is important to consider that they are designed for a particular purpose and user. In the interdisciplinary field of quantum software, it is particularly important to understand the perspectives and specific needs of the stakeholders involved. To better understand these, this work addresses the following research questions:

RQ1 Who are the relevant stakeholders in quantum software, both now and in the future (short-term, mid-term, and long-term)?

RQ2 What are the specific needs of stakeholders in quantum software, and how can abstraction help to reach their goals?

Stakeholders are interested in particular use cases, which are driven by specific interests (*i.e.*, the key objectives they wish to achieve). Constraints or side conditions apply to how objectives can be reached. Different levels of abstraction are required depending on the use case and will vary over time as the capabilities of machines improve. Hence, we derived sub-research questions of **RQ2**:

RQ2.1 What use cases for quantum software do these personas have? Which problems do they want to solve?

RQ2.2 What interests do these personas have in using quantum software?

RQ2.3 Which constraints (side conditions) for quantum software and their use cases do these personas have?

RQ2.4 What are the abstraction levels (*e.g.*, language building blocks and idioms) that these personas want to use?

RQ2.1 regards use cases as (functional) stakeholder goals, which the stakeholders want to achieve using quantum software. In **RQ2.2**, broader motives driving these stakeholders, such as business motives and the pursuit of scientific insight, are investigated. **RQ2.3** considers constraints, meaning limitations or challenges, that stakeholders face when engaging with quantum software. **RQ2.4** aims to identify which abstractions the stakeholders want to interact with quantum software with respect to their mental models. Abstractions in a quantum software stack can range from domain specific descriptions (high-level), over implementation and compilation using programming languages and frameworks, to the low-level physical realisation and integration of (quantum) hardware.

Goal-Directed Design (see Sec. 2) is a widely adopted way to investigate user groups when designing software. It uses the concept of *personas* as descriptive models of users, presented as individual humans. Personas are an established methodology [10, 37, 44] that we propose to use when designing quantum software products. For this work, we do not have a concrete product in mind, but want to give a general overview of possible kinds of users for quantum software. Thus, we do not define concrete personas, but formulate a *persona-hypothesis* as an initial step at defining the different kinds of users according to the Goal-Directed Design process. Future practitioners creating quantum software products may use our persona-hypothesis as a basis for initial interview planning, to then create concrete personas for their product.

To derive meaningful personas, we apply qualitative methods suitable for an exploratory setting, allowing participants to respond more freely [29]. To build personas on the experiences from experts as well as practitioners, we conduct two separate studies. First, we conduct a focus group with the experts of the *Software Abstraction in Quantum Computing* working group [2], as part of a Dagstuhl seminar in 2024 (see Sec. 4). This seminar provides a rare opportunity to interact face-to-face with a range of experts in quantum software engineering. Second, to also incorporate the perspective of practitioners actively working on quantum software, we conduct a qualitative interview study with relevant stakeholders (see Sec. 5) at IEEE QCE 2025, an event attended by academics and industry professionals [34]. To answer **RQ1** and **RQ2**, we then aggregate the insights from both studies as a persona-hypothesis[1] in Sec. 6.

4 Personas by Expert Focus Group at Dagstuhl

We conducted a focus group with the experts from the *Software Abstraction in Quantum Computing* working group [2], in which authors of this work participated, as part of the Dagstuhl Seminar 24512 (link in pdf). In December 2024, the seminar brought together 24 experts in quantum software engineering for five days. Of these, three days were dedicated to discipline-specific working group discussions. The following results[2] were produced by an expert group of eight to ten people, with attendees varying across individual sessions.

According to the expert group, improving quantum hardware is expected to lead to a diversification of the user base over time. Quantum computers featuring only tens to hundreds of physical quantum bits are already available in the *short-term*. This means that execution is fundamentally limited to log-depth circuits [25, 36]. Furthermore, the interconnects necessary to create distributed Quantum Processing Units (QPUs) are lacking. In the *mid-term*, experts project systems with thousands to hundred of thousands of physical qubits and error-correcting codes. These codes may not fully correct all potential errors, but still enable accurate execution of circuits beyond log-depth with a reasonable success rate. Rerunning a calculation may be necessary if too many errors occur, but this limitation is not decisive for practical utility. Initial interconnects to combine tens to hundreds of QPUs are available through hardware vendors aiming to create first of such systems within this decade, according to their roadmaps [48]. A *long-term* goal for quantum hardware is to de-risk physics in the sense that systems can be scaled using multi-QPU setups with reliable physical interconnects that transmit quantum information faithfully across QPUs.

To address **RQ1**, the experts created personas at the *short-*, *mid-* and *long-term* time intervals representing expected stages of technological maturity. The personas involved in quantum software within the time intervals that were identified in the expert group are summarised in Tab. 1. They identified five short-, seven mid-, and six long-term personas[3] covering a variety of domains:

Short-term. Initially **curious** (business) **users** can conduct pilot studies serving as a starting point for scalable quantum operations. Knowledge of concrete pipeline stages (*i.e.*, transformations) and quantum specific details is required [45, 49]. Some quantum learners in the short-term are interested about quantum computing by conducting exploratory studies. Hence, they favour graphical, intuitive, interactive and explaining interfaces not only in circuit level abstraction, but also in IDE's and when applying quantum software

[1] For the sake of brevity, we may refer to the different kinds of users according to our proposed *persona-hypothesis* as *personas* in the remainder of this work.
[2] Take into consideration that we paraphrase. We explicitly do not characterise people.
[3] Note that we expect profiles of stakeholders to change as hardware matures.

Table 1: Personas of quantum software according to the Dagstuhl expert group, categorised by quantum computing hardware availability in the short-, mid- and long-term.

Short-term	Mid-term	Long-term
Quantum curious (business)	Business users	Business application engineer (domain expert)
HPC engineer ("Offloading")	HPC engineer ("QC at scale")	HPC engineer ("QC at large scale")
Embedded quantum developer	Embedded quantum developer	Embedded quantum developer ("power user")
Quantum algorithm researcher	Quantum algorithm developers (CS)	Quantum algorithm developers (CS)
Quantum mechanics researcher	Quantum mechanics researcher	Quantum mechanics researcher / physicist
	Material scientist	Material scientist
	Early adopters	

idioms. Business users operate on a high abstraction level. **High-Performance Computing (HPC) engineers** [16] and **embedded quantum developers** [40] are close to system integration of hardware to (a) develop scalable and distributed quantum computing and (b) develop the necessary low-level interfaces that make system-level integration of quantum computing possible. **Quantum algorithm researchers** also work on interfaces, but at higher levels of abstractions—ranging from purely logical quantum algorithms (not limited to circuit representation) to hardware-aware (*e.g.*, for topology, noise, operations) quantum algorithms. To improve the abilities and efficiency of known algorithms, quantum algorithm researchers in the short-term operate on low-level abstractions (*i.e.*, gates and unitaries), as well as on high-level languages describing the functional behaviour of algorithms. The goal for **quantum mechanics researchers** is scientific knowledge gain, which in the short-term is primarily based on physics. Similar to business users, they are able to conduct (limited) pilot studies for their use cases. They are more specialised in the simulation of various quantum systems, while being constrained by the need to translate to circuit level. Accuracy of simulations and discovery of (new) phenomena are major interests.

Mid-term. The availability of many (mostly) error-corrected qubits enables **business users** to scale their pilot studies to industry relevant problem sizes. While the implementation still necessitates an understanding of quantum-specific intricacies, this requirement is less pronounced compared to pilot studies conducted in the short-term. Although **HPC engineers** and **embedded quantum developers** are still close to system integration, their tasks shift from integrating and optimising single QPUs to multiple (quantum mechanically connected) QPUs [17]. When considering the application of large scale, heterogenous calculations with abstractions in line with, for example, LAMMPS (link in pdf), the ability to interleave communication and coordinate calculations is a major constraint for HPC. As systems scale, quantum algorithms are subject to different (performance) metrics (*e.g.*, the limitation to log depth circuits is gradually relaxed), which entails a broader view on computer science paradigms and complexity theoretic consideration for **quantum algorithm developers**. Also, for **quantum mechanics researchers**, computer science becomes a necessary field of expertise to efficiently execute experiments on multiple connected QPUs. A portion of mid- and long-term quantum mechanics researchers, particularly high energy physicists, are interested in Schwinger models or lattice gauge simulations [22] with attention to accuracy and new phenomena. In line with Lagrangians and Hamiltonians as their main abstractions, they are constrained by the expressibility of low code and maths[4]. Due to the increased availability of qubits, simulations of materials (*e.g.*, multi-atom simulations, macroscopic property simulations) become feasible, which opens the field of quantum computing for **material scientists** to develop new and optimise existing materials and compounds [21, 32, 53]. Concurrently, **early adopters** use quantum computing in production settings.

Long-term. As scalable hardware becomes available, **business users** use quantum computers on industrial scale and become abstract from pipeline and quantum specific details, while focusing on domain expertise. While a subset of long-term **business application engineers** are driven by financial benefits, they are interested in result quality, runtime and customer value in a quantum agnostic abstraction, meaning that the execution on a quantum computer is abstracted away from the user, which enables similar (classical) interfaces as before. **HPC engineers** integrate multiple (distant) quantum computers at large scale. Although there exists a rich environment of (optimising) tools and toolchains in the long-term, **embedded quantum developers** can leverage their system specific expertise to realise systems full potential. In the short- and mid-term, **quantum algorithm developers** are focused on satisfying hardware constraints and optimise resource requirements to be able to execute algorithms. As systems scale in the long-term, additional metrics (*e.g.*, ease of use, customisability, (response) time, cost-effectiveness, political restrictions, security, data privacy, protection of intellectual property) become more pronounced around which new algorithms are designed. Similar to mid-term considerations, **quantum mechanics researchers** are able to execute large scale problems, but need to consider scaling effects, approximation quality and solution guarantees. Also, **material scientists** achieve scalability comparable to quantum systems, enabling new methods (*e.g.*, material interaction simulation on macroscopic scale), which may also make biological simulations possible [35].

5 Personas by Practitioner Interviews at QCE

To gather insights from practitioners, we conducted an interview study with relevant stakeholders. Interviews were conducted as one-on-one interviews with broad open-ended questions, focusing on gathering domain knowledge from the interviewees' perspective. We conducted our study as a workshop session at the IEEE

[4]For instance, Hamiltonian simplifications and their expressibility in frameworks.

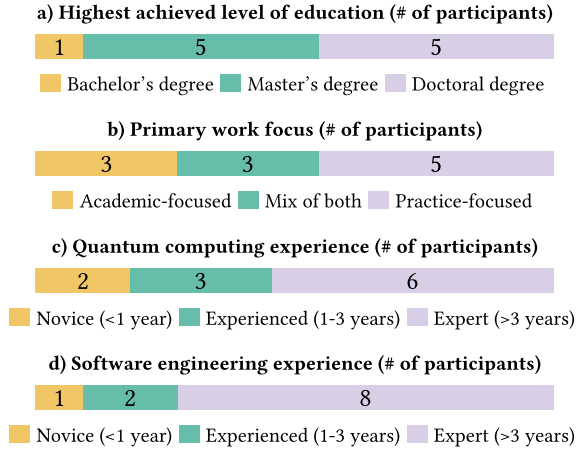


Figure 2: Demographics of QCE 2025 interviewees ($N = 11$).

International Conference on Quantum Computing and Engineering (QCE) in 2025. Before the interviews, the interviewees received a brief presentation (link in pdf) about workshop goals and involved personas, and completed a demographic questionnaire.

Demographic Questionnaire. We asked our participants a series of demographic questions to confirm that they were experienced practitioners from academia and industry. We asked four demographic questions regarding the participants level of education, work focus, and experience in quantum computing and in software engineering. The questions were provided through Google Forms (link in PDF). We obtained demographic information for eleven QCE participants, shown in Fig. 2. Higher education was common: Five participants held a Master's, and another five a doctoral degree (Fig. 2 a). The primary work focus spanned from academic over mixed to practical-focused, leaning towards practical-focused (Fig. 2 b). Quantum computing and software engineering experience of participants (Fig. 2 c, d) are similar, with mostly expert participants.

Practitioner Interviews. To allow for more detailed discussions, we conducted single person interviews (one interviewee, one interviewer, and one note-taker) in multiple sessions of our workshop *Care and Maintenance of Quantum Software Stakeholders* at IEEE QCE 2025. In total, there were three interviewers and three note-takers (all with academic background). The interviews were semi-structured, using guiding questions to allow for flexible conversation and gather domain knowledge from the interviewee's perspective (see Sec. 2). We asked interviewees to share their insights on personas, encouraging them not only to consider those defined in Sec. 4, but also to introduce additional ones when appropriate. To guide the conversation, we use **RQ1**, **RQ2** and **RQ2.1-RQ2.4** as interview guiding questions. Interviews were voluntary — also after completing the demographic questionnaire, and hence we conducted 9 interviews , lasting between 15 and 40 minutes. Due to privacy considerations, we cannot offer raw data of each interview.

Interview Results. The interviewees identified nine personas (**RQ1**). In comparison to Tab. 1 and Sec. 4, business users and researchers were described in depth by the participants and categorized into short-, mid- and long-term. Participants described the business users as a combination of short-, mid- and long-term users, with the key differences in the extent of their application. However, the common factor across all is the profit-driven mindset that motivates decisions and actions. The described researchers (in varying degrees) cover short-, mid- and long-term by (among others) applying quantum algorithms to existing hardware, making pilot studies on industrial use cases for mid-term applications and considering long-term effects on security. They cover a broad spectrum of tasks and hence distinctions to other personas are not sharp. Participants further identified the personas of physicists, chemists, simulation focused engineers, quantum algorithm designer, HPC engineers, platform builders / developers, and government contractors, but did not go into detail about how they will change over time. In the following, we provide paraphrased insights into personas (**RQ2**) described by participants:

Business users are relatively high-level personas, described by three participants. Although they are primarily focused on maximising profit (RQ2.2), they have varying use cases ranging from (a) acquiring new customers, (b) protecting intellectual property, which also considers security aspects of quantum computing, to (c) driving research into industry applications (RQ2.1). Also, one participant mentioned an explicitly longer term business user who has to decide how and to what degree quantum computing benefits a certain problem (*i.e.*, selecting the share of quantum in an (established) process) and thus needs (resource) estimations rather than hardware executable representations (*e.g.*, circuits). In general, the described business users want to abstract from syntax and implementation details (RQ2.4) and use resource estimates to inform decisions, making them dependent on the accuracy and engineering effort for those estimations (RQ2.3). They also favour natural language over rigorous formal languages (RQ2.4).

Two participants described **researchers**. Although – in contrast to business users – their work is academic, (in-)direct connections to industry exist: Industry relevant papers in quantum computing not only lay foundations for industrial utility, but also open the field for research on industry applications. The researcher personas want to show quantum advantage (RQ2.1) and publish findings (RQ2.2), which is usually fast paced and thus requires a fast prototyping environment (RQ2.3). Furthermore, time constraints (of operations) and personnel restrictions (*i.e.*, modular, low-granularity tasks suitable for individual execution rather than large, multi-person work units) are important properties to researchers (RQ2.3).

Physicists were described thrice. Although sometimes associated with lower level details, interviewees describe them as use case oriented. Interests include (a) verifying theories, (b) finding "new stuff", (c) exploring condensed matter or (d) simulating problems (*e.g.*, in the aerospace domain), which make them similar to the business user (RQ2.1, RQ2.2). They show particular awareness of physical limitations (which is true for any persona, but with less importance). Another important constraint is the manual transfer of physical details onto, for instance, (quantum) algorithms — highlighting the importance of automation (RQ2.3). Their levels of abstraction cover programming languages (*e.g.*, Python) over (quantum) instructions to properties of quantum mechanics, resulting in a lower level of abstraction than for business users (RQ2.4).

Analogously, **chemists** are interested in solving hard chemistry problems in varying disciplines (*e.g.*, physical, (in)organic and biochemistry problems; RQ2.1, RQ2.2). They were mentioned by one participant. Similar to the physicist, they are constrained by the laws of physics and chemistry, as well as resource limitations, which results in poor scalability of simulations (RQ2.3). **Simulation focused engineers** were mentioned by the same participant. They are interested in problems representable as Hamiltonians (RQ2.1, RQ2.2) and are constrained by the capabilities of hardware (RQ2.3).

Quantum algorithm designers (potentially hardware-agnostic) were described by two participants. They want to execute an algorithm on multiple backends (RQ2.1) and depend on backend-specific details (to be made available by a software stack; RQ2.3).

HPC engineers focus on quantum-classical integration on multiple (classical) (connected) backends (RQ2.1, RQ2.2). They are thus constrained by communication time (*e.g.*, real time connections; RQ2.3) and were mentioned by one participant.

Two participants describe the **platform builders / developers** that build quantum platforms (*e.g.*, Software Development Kits (SDKs), Application Programming Interfaces (APIs) and abstraction frameworks) to benchmark quantum hardware for use cases (RQ2.1, RQ2.2). They are constrained by modularity and replaceability of (key) components in a toolchain (*e.g.*, transpilers), by scalability of hardware and lack of standardisation in software (RQ2.3). Constraints include (self-)defined interfaces and the balance between sufficient abstraction and expert access to details (RQ2.4).

One participant describes **government contractors**. In distinction to previous personas, they work on strict security requirements and cannot rely on black-box libraries (RQ2.2). They need traceable and deterministic guarantees about software behaviour (RQ2.3).

6 Aggregated Personas for Quantum Software

The preceding sections presented findings primarily from the perspectives of interviewees and subject-matter experts. In this section, we aggregate personas from these findings and complement these perspectives by incorporating our own interpretation of the data. This paper aims to better understand users of quantum software, so interpretations and previously introduced data may be extended.

Regarding **RQ1**, the expert group and QCE interviewees introduced personas for users of quantum software. In particular, the expert group (see Sec. 4) considered business users, early adopters, quantum mechanics researchers, material scientists, quantum algorithm researchers, HPC engineers and embedded quantum developers. Our interview study (see Sec. 5) gives a different perspective: Although interviewees described personas with a significant overlap (but different names), they also considered different users (*i.e.*, government contractor, simulation focused engineer, researcher and platform builder). Fig. 3 compares personas from the Dagstuhl expert group (left) to personas from interviews (right) by showing approximate equivalence ($\approx$) and broader descriptions (*i.e.*, subsets $\subset$). In particular, we found that descriptions of business users, quantum algorithm designers, and HPC engineers are roughly equivalent (see $\approx$-edges). Therefore, we considered them as one persona. Additionally, interviewees described physicists and chemists more generically than the Dagstuhl personas of quantum mechanics researchers and material scientists (see $\subset$-edges). We propose to use

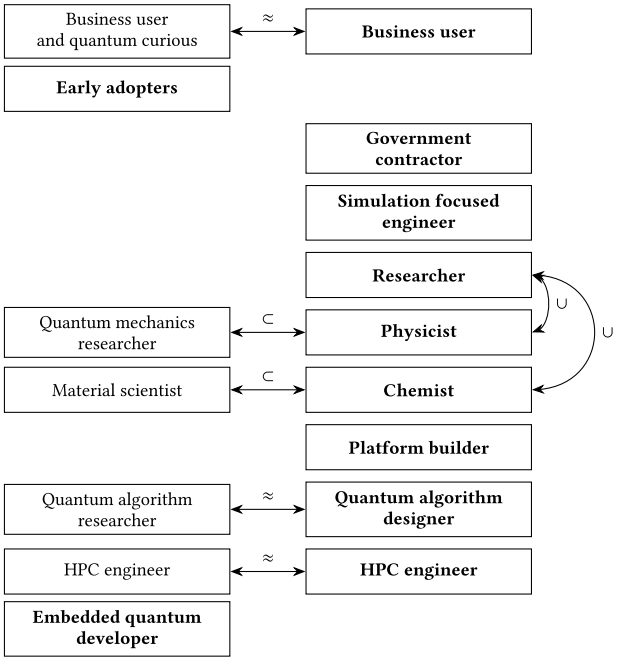


Figure 3: Mapping of Dagstuhl personas (left; Sec. 4) to interview personas (right; Sec. 5). $A \subset B$ means A was described more specialized than B. $A \approx B$ means A and B were described in approximately same scope. Aggregated personas in bold.

the physicist and chemist personas in this broad persona-hypothesis context. Nevertheless, (quantum) mechanics researchers and material scientists could be considered when deriving concrete personas. Furthermore, interviewees mentioned supplementary personas, such as government contractor, simulation focused engineer, and platform builder, which originate from practically relevant fields. In particular, the researcher persona was mentioned as a broad persona (*i.e.*, physicist and chemist can be considered subgroups; see $\subset$-edges). Early adopters and embedded quantum developers were only mentioned by the expert group.

RQ1: Who are the relevant stakeholders in quantum software both now and in the future?

We propose eleven initial personas for quantum software depicted in bold in Fig. 3, which we aggregate based on similarity[5]: business users, early adopters, government contractors, researchers, such as physicists and chemists, platform builders, quantum algorithm designers, HPC engineers, embedded quantum developers and simulation focused engineers.

To address **RQ2**, for each aggregated persona, we describe a representative use case, taking into account their interests and constraints as detailed in Tab. 2. This data is sourced from expert interviews at Dagstuhl, QCE interviews, and our interpretation. Tab. 2 depicts these detailed persona descriptions that are based on the expert interview at Dagstuhl, the interviews at QCE and our own interpretation. These archetypical personas subsume important

[5]Researcher encompasses a broader scope than physicist or chemist.

Table 2: Aggregated personas for quantum software as a generalisation from concrete stakeholders, grouped by application and hardware orientation: application oriented personas (business user, early adopter, government contractor), application and hardware oriented personas (simulation focused engineer, researchers, such as physicist and chemist) and hardware oriented personas (platform builder, quantum algorithm designer, HPC engineer, embedded quantum developer). For each persona we aggregated their use case, interests, constraints and abstraction level according to RQ2.1-RQ2.4 (see Sec. 3).

	Persona	Use case (RQ2.1)	Interests (RQ2.2)	Constraints (RQ2.3)	Abstraction level (RQ2.4)
Application oriented	*Business user*	Scalable (profit-driven) quantum operations	Maximising profit, estimating impact on security	Quantum specific intricacies, resource estimation accuracy	High: focus on domain expertise
	Early adopter	Adopting new quantum technologies in industry applications	Driving business models and revenue growth	Industry relevance, scalability and adaptability	High: focus on market demands
	Government contractor	Government tasks with strict (security) requirements	Data confidentiality, integrity, and availability	Traceable guarantees about software behaviour, no black-box implementations	Mid to high: focus on traceability and observability of software and of algorithms
Application and hardware oriented	*Simulation focused engineer*	Simulating large-scale systems	Accurate simulation of physical, chemical systems or dynamical processes	Scalability, approximation quality, solution guarantees	Mid to high: focus on algorithmic accuracy
	Researchers	Pilot studies, scientific knowledge gain	Scientific discovery, publishing findings	Time constraints, manual transfer of data, resource limitations	Varying: high abstraction for domain specific goal; potentially low for methods
	↪ *Physicist*	Verifying theories, exploring condensed matter, simulation	Solving hard physics problems	physical principles, resource limitations	Mid to low: programming languages over quantum instructions, Hamiltonians for physical processes
	↪ *Chemist*	Solving hard chemistry problems	Scientific / material discovery	chemical principles, resource limitations	Mid: tools / frameworks for chemical simulations, Hamiltonians for chemical processes
Hardware oriented	*Platform builder*	Building quantum platform (SDKs, APIs)	Benchmarking quantum hardware / software	Modularity and replaceability of components, lack of standardization in software	Mid: toolchain development and customization
	Quantum algorithm designer	Designing / improving quantum algorithms	Algorithmic performance (asymptotic & practical)	Quantum-specific model of computation, lack of standardization	Mid to low: on algorithm interface & below (needs access to low-level details)
	HPC engineer	Enabling heterogenous, large scale computation	System integration, interleaving computation	Ability to interleave communication / coordination	Low: hardware and system integration
	Embedded quantum developer	Developing scalable and distributed quantum computing	System-level integration of quantum computing	Interleaving communication/coordination between QPUs, system-specific expertise	Low: focus on hardware / software integration

characteristics of specialized stakeholders, as they are generalisations from concrete people with concrete use cases and characteristics. We group these personas by *application*, and *hardware* orientation with diverse sets of use cases and applications ranging from industry to specialized technical applications. Other personas are working on the technical realization of quantum software — requiring hardware awareness. Similar to the diversity in use cases, interests range from commercial goals and technical objectives to scientific pursuits. Accordingly, personas are subject to a wide range of constraints, that are, limitations or challenges that stakeholders face when engaging with quantum software. The mental models of stakeholders influence the abstractions with which they want to interact when using quantum software. These range from high-level for application oriented personas to low-level for hardware oriented personas. In the following, we provide detailed synthesis of the answers to each sub-research question (see Sec. 3).

RQ2: What are the specific needs of stakeholders in quantum software, and how can abstraction help to reach their goals?

We define use cases, interests, constraints and abstraction level for each aggregated persona in Tab. 2. Furthermore, we provide additional insights into abstractions used by personas in Fig. 4.

RQ2.1. The application-oriented personas (*e.g.*, business user, early adopter, government contractor and some, high-level, researchers) all share the desire to use quantum computing to solve domain-specific problems more effectively than current classical solutions. Use cases of personas that are application- and hardware oriented (*i.e.*, different types of researchers and simulation focused engineer) tend to be more specific and technically involved. For instance, the (physical) simulation of (large-scale) systems requires domain-, as well as simulation-specific details. Mostly hardware-oriented personas require an even deeper understanding of details that are otherwise abstracted to application-oriented personas: Quantum algorithm designers improve algorithms at a theoretical and on a practical level, considering the capabilities of hardware. Platform builders, HPC engineers and embedded quantum developers make a significant contribution to the technical realisation of quantum software and its integration with classical computing.

RQ2.2. Interests of personas fall into three categories: commercial goals (such as maximizing profit and security), technical objectives (such as benchmarking hardware and optimizing runtime), and scientific pursuits (such as making new discoveries and verifying theories). Quantum computing is not only a tool that personas apply, but in itself a field of study for all of the aforementioned categories. For instance, personas involved in building quantum software are interested in its development and improvement.

RQ2.3. Personas face various constraints, meaning limitations or challenges, when engaging with quantum technology. These cover technical limitations imposed by existing machines, as well as the integration of existing interfaces and modularity of software components. Manual data transfer, academic deadlines and guarantees on software behaviour pose operational challenges. Constraints such as scalability, adaptability, modularity and replaceability impact the design of (quantum) programs to suit user needs; there may

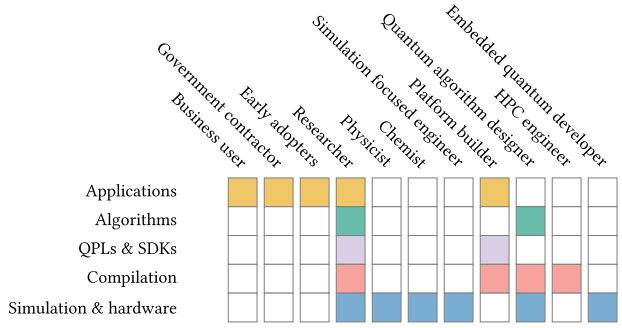


Figure 4: The layers of quantum software (y-axis; additionally coloured) in which each identified persona is mostly active.

be additional use case-dependent constraints. For instance, a simulation use case is fundamentally limited (in precision) by scalability, which depends on – among others – problem structure that might allow for mathematical simplifications or accurate approximations. In physics, some Hamiltonians that describe a physical system that allows for scalable simulations. However, there are a multitude of physical systems requiring different assumptions or modelling techniques to arrive at manageable Hamiltonians, although these use cases are subsumed under the physicist personas.

RQ2.4. Aforementioned constraints and mental model govern which abstractions personas are interested in. For instance, application oriented personas do not consider hardware-specific details, yet, these are relevant to determine use-case feasibility. Similarly, HPC engineers or embedded developers need access to lower-level details to optimise communication and integration; quantum algorithm designers need to consider (future) capabilities of hardware.

To further illustrate the abstractions required by each persona, we categorize them into five (non-exhaustive) abstraction layers in Fig. 4, considering layers similar to those introduced by Bandic et al. [5]. Concretely, from high to low levels of abstraction: end-user *applications*; development of quantum *algorithms*; hardware-agnostic program abstractions such as *Quantum Programming Languages (QPLs)* and *SDKs*; *compilation* of programs into hardware-specific formulations; *simulation* of quantum-mechanical systems and the development of quantum *hardware*. Fig. 4 shows that certain personas (*i.e.*, business users, physicists, chemists, simulation engineers, HPC engineers, and government contractors) mostly work on a single high- or low-level quantum software layer, respectively. Interestingly, there are personas for which multiple layers are relevant, such as the researcher, the quantum algorithm engineer, and the platform builder. Therefore, their requirements also span multiple layers. Since all layers of quantum software[6] are of potential interest to the researcher, each must support the ability for rapid (prototypical) development in the short- and mid-term.

7 Threats to Validity

Despite careful study design, several threats to validity should be considered when interpreting our results.

[6] Specialized personas (see chemist or physicist) may require fewer layers

Participant Selection. For our focus group, we relied on experts who were present at the Dagstuhl seminar (from the authors of this paper, LaS, JG, RW, SF, IS, WM contributed to the discussion). Invitations were limited to pre-selected experts recognized in the field. However, the attendees of the working group on abstractions were chosen by convenience sampling, since experts had to opt in to participate. As the objective of this workshop was to develop an initial hypothesis, it seemed reasonable to base the work on the expertise of specialists in the field. To incorporate the perspective of practitioners we conducted interviews according to the Goal-Directed Design process (see Sec. 5). Interviewer-administered surveys are more social, and therefore at risk of bias due to interviewer variance [14]. As our interviews were conducted by three different interviewers and three different note-takers, we cannot rule out that interviewer variance and bias from the authors' personal opinions influenced the results, yet deem the risk to be negligible, as all interviewers were aware of this problem, and none has commercial or other interest in specific outcomes. Due to the limited time available at the conference and to avoid participants' concerns about confidentiality, we opted for manual note-taking rather than audio recording [20]. To gather a rather diverse set of interviewees, we used convenience sampling at a relevant venue with interest from academia and industry (QCE). Although we deliberately chose QCE for its mix of academic and industry visitors, as well as its diversity of nationalities, our choice of venue may introduce bias. Nevertheless, given the difficulty of sampling from the quantum software community, we believe that our sampling method is reasonable. We were able to obtain nine interviews using this method, limited by the number of workshop attendees. There is some indication of data saturation, as multiple personas were mentioned multiple times: For instance, the last interviewed participant did not provide any new personas, but only added to previously established ones.

Credibility, Resonance, and Transferability . While we carefully considered established empirical standards [39], quantitative criteria (internal, external and construct validity) do not apply to our qualitative survey. We evaluate based on credibility, resonance, transferability. We used a dual approach combining expert knowledge and interviews. The diverse QCE community [34] is reflected in demographics and professional focus (see Fig. 2), indicating multivocality. As Dagstuhl participants are experts in quantum software engineering and most QCE interviewees claim to have significant experience (see Fig. 2), we consider their statements credible. We found agreement between our focus group and interview findings (see Sec. 6), indicating that defined personas resonate with practitioners. Our personas are focused on quantum software, but coarse-grained enough to allow for a broad range of applications within that field. Due to the extent of tailoring, we anticipate limited transferability beyond this quantum software scope.

8 Related Work

Previous works address challenges [33, 54, 55] and application scenarios [8] of quantum software. Personas support practitioners and researchers across domains in the design process [10, 37, 44]. However, no work has, to the best of our knowledge, defined personas or other user models for quantum software. We are convinced that the use of personas will benefit the field, help identify and address users' needs, and enable stakeholder integration [4]. The use of qualitative methods to address human aspects is increasingly recognised in software engineering [29]. Qualitative methods allow participants to respond more freely, making them suitable for use in exploratory studies [29]. While other studies gather insights from active researchers [33] or conduct semi-structured interviews with quantum software practitioners [13], to the best of our knowledge no prior work has conducted systematic, community-level interviews, such as at a quantum software conference.

9 Conclusion and Outlook

In this work, we introduce a persona hypothesis for quantum software based on the results from an expert focus group at the Dagstuhl seminar 24512 and practitioner interviews at the IEEE Quantum Week (QCE) 2025. We define and discuss 11 aggregated personas of quantum software users and stakeholders regarding their use cases, interests, constraints and abstraction level. We view our work as the starting point to consolidate quantum stakeholders at a community level and therefore encourage researchers and practitioners to review and refine user personas introduced in this work. We encourage the quantum software community to gather further insights from stakeholders through user studies and interviews. Future surveys can provide additional relevant insights into the specific needs of stakeholders and challenges they face. We see personas as a useful tool for the design of tailored quantum software products.

Acknowledgments

We thank the organisers of Dagstuhl Seminar 24512, the IEEE QCE and all interviewees for their insightful input. This work was supported by the German Research Foundation, grant 563122858 (MA 9739/1-1 and SCHA 1635/20-1).

References

[1] Scott Aaronson, Andrew M Childs, Edward Farhi, Aram W Harrow, and Barry C Sanders. 2026. Future of quantum computing. *Quantum Mach. Intell.* 8, 1 (2026). https://doi.org/10.1007/s42484-026-00351-2

[2] Shaukat Ali, Johanna Barzen, Andrea Delgado, Hausi A. Müller, and Juan Manuel Murillo. 2025. Quantum Software Engineering. *Dagstuhl Reports* 14, 12 (2025), 63–84. doi:10.4230/DagRep.14.12.63

[3] Shaukat Ali and Tao Yue. 2020. Modeling Quantum Programs: Challenges, Initial Results, and Research Directions. In *Proceedings of the 1st ACM SIGSOFT International Workshop on Architectures and Paradigms for Engineering Quantum Software*. ACM, Virtual USA, 14–21. doi:10.1145/3412451.3428499

[4] Joshua Ammermann, Wolfgang Mauerer, and Ina Schaefer. 2024. Towards View-based Development of Quantum Software. In *INFORMATIK 2024*. Gesellschaft für Informatik e.V., Bonn, 551–554. doi:10.18420/inf2024_43

[5] Medina Bandic, Sebastian Feld, and Carmen G. Almudever. 2022. Full-stack quantum computing systems in the NISQ era: algorithm-driven and hardware-aware compilation techniques. In *2022 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 1–6. doi:10.23919/date54114.2022.9774643

[6] Benjamin Bichsel, Maximilian Baader, Timon Gehr, and Martin Vechev. 2020. Silq: A High-Level Quantum Language with Safe Uncomputation and Intuitive Semantics. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation*. ACM, 286–300. doi:10.1145/3385412.3386007

[7] Sebastian Bock, Raphael Seidel, Matic Petrič, Nikolay Tcholtchev, Andreas Hoffmann, and Niklas Porges. 2025. Designing a Meta-Model for the Eclipse Qrisp eDSL for High-Level Quantum Programming. In *Proceedings of the 13th International Conference on Model-Based Software and Systems Engineering*. SCITEPRESS, 27–39. doi:10.5220/0013121000003896

[8] Cecilia Carbonelli, Michael Felderer, Matthias Jung, Elisabeth Lobe, Malte Lochau, Sebastian Luber, Wolfgang Mauerer, Rudolf Ramler, Ina Schaefer, and Christoph Schroth. 2024. *Challenges for Quantum Software Engineering: An Industrial Application Scenario Perspective.* Springer Nature Switzerland, 311–335.

doi:10.1007/978-3-031-64136-7_12
[9] André Chailloux and Jean-Pierre Tillich. 2025. Quantum Advantage from Soft Decoders. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing (STOC '25)*. ACM, 738–749. doi:10.1145/3717823.3718319
[10] Yen-ning Chang, Youn-kyung Lim, and Erik Stolterman. 2008. Personas: from theory to practices. In *Proceedings of the 5th Nordic conference on Human-computer interaction: building bridges*. ACM, 439–442. doi:10.1145/1463160.1463214
[11] Alan Cooper, Robert Reimann, Dave Cronin, Christopher Noessel, Jason Csizmadi, and Doug LeMoine. 2014. *About Face: The Essentials of Interaction Design* (fourth ed.). Wiley, Indianapolis, Indiana.
[12] Nathalie P. de Leon, Kohei M. Itoh, Dohun Kim, Karan K. Mehta, Tracy E. Northup, Hanhee Paik, B. S. Palmer, N. Samarth, Sorawis Sangtawesin, and D. W. Steuerman. 2021. Materials challenges and opportunities for quantum computing hardware. *Science* 372, 6539 (April 2021). doi:10.1126/science.abb2823
[13] Manuel De Stefano, Fabiano Pecorelli, Fabio Palomba, Davide Taibi, Dario Di Nucci, and Andrea De Lucia. 2024. Quantum Software Engineering Issues and Challenges: Insights from Practitioners. In *Quantum Software: Aspects of Theory and Sys. Design*. Springer Nature Switzerland. doi:10.1007/978-3-031-64136-7_13
[14] Don A. Dillman, Jolene D. Smyth, and Leah Melani Christian. 2014. *Internet, Phone, Mail, and Mixed-Mode Surveys: The Tailored Design Method*. Wiley Publ.
[15] Jens Eisert and John Preskill. 2025. Mind the gaps: The fraught road to quantum advantage. doi:10.48550/ARXIV.2510.19928
[16] Amr Elsharkawy, Xiaorang Guo, and Martin Schulz. 2024. Integration of Quantum Accelerators into HPC: Toward a Unified Quantum Platform . In *2024 IEEE QCE*. 774–783. doi:10.1109/QCE60285.2024.00097
[17] Pau Escofet, Anabel Ovide, Medina Bandic, Luise Prielinger, Hans Van Someren, Sebastian Feld, Eduard Alarcon, Sergi Abadal, and Carmen Almudever. 2025. Revisiting the mapping of quantum circuits: Entering the multi-core era. *ACM Transactions on Quantum Computing* 6, 1 (2025), 1–26.
[18] Sebastian Feld. 2025. Full-Stack Quantum Computing and Distributed Systems: A Community-Centric Approach. In *International Conference on Innovations for Community Services*. Springer, 3–19.
[19] Felipe Ferreira and José Campos. 2025. An Exploratory Study on the Usage of Quantum Programming Languages. *Science of Computer Programming* 240 (Feb. 2025), 103217. doi:10.1016/j.scico.2024.103217
[20] Uwe Flick. 2023. *An introduction to qualitative research* (7th ed.). SAGE, LA.
[21] Maja Franz, Lukas Schmidbauer, Joshua Ammermann, Ina Schaefer, and Wolfgang Mauerer. 2026. Towards Quantum Software for Quantum Simulation. In *7th IEEE/ACM International Workshop (QSE)*. ACM. doi:10.1145/3786150.3788611
[22] Maja Franz, Pía Zurita, Markus Diefenthaler, and Wolfgang Mauerer. 2024. Co-Design of Quantum Hardware and Algorithms in Nuclear and High Energy Physics. *EPJ Web of Conferences* 295 (2024). doi:10.1051/epjconf/202429512002
[23] Felix Gemeinhardt, Antonio Garmendia, and Manuel Wimmer. 2021. Towards Model-Driven Quantum Software Engineering. In *2nd IEEE/ACM International Workshop (Q-SE)*. IEEE, 13–15. doi:10.1109/q-se52541.2021.00010
[24] Felix Gemeinhardt, Antonio Garmendia, Manuel Wimmer, and Robert Wille. 2024. A Model-Driven Framework for Composition-Based Quantum Circuit Design. *ACM Trans. on Quantum Computing* 5, 4 (Oct. 2024), 1–36. doi:10.1145/3688856
[25] Felix Greiwe, Tom Krüger, and Wolfgang Mauerer. 2023. Effects of Imperfections on Quantum Algorithms: A Software Engineering Perspective. In *IEEE QSW*. 31–42. doi:10.1109/qsw59989.2023.00014
[26] Andi Gu and Stephen P. Jordan. 2025. Algebraic Geometry Codes and Decoded Quantum Interferometry. doi:10.48550/ARXIV.2510.06603
[27] Stephen P. Jordan, Noah Shutty, Mary Wootters, Adam Zalcman, Alexander Schmidhuber, Robbie King, Sergei V. Isakov, Tanuj Khattar, and Ryan Babbush. 2025. Optimization by decoded quantum interferometry. *Nature* 646, 8086 (Oct. 2025), 831–836. doi:10.1038/s41586-025-09527-5
[28] Tanuj Khattar, Noah Shutty, Craig Gidney, Adam Zalcman, Noureldin Yosri, Dmitri Maslov, Ryan Babbush, and Stephen P. Jordan. 2025. Verifiable Quantum Advantage via Optimized DQI Circuits. doi:10.48550/ARXIV.2510.10967
[29] Per Lenberg, Robert Feldt, Lucas Gren, Lars Göran Wallgren Tengberg, Inga Tidefors, and Daniel Graziotin. 2023. Qualitative software engineering research: Reflections and guidelines. *Journal of Software: Evolution and Process* 36, 6 (2023). doi:10.1002/smr.2607
[30] Gushu Li, Anbang Wu, Yunong Shi, Ali Javadi-Abhari, Yufei Ding, and Yuan Xie. 2021. On the Co-Design of Quantum Software and Hardware. In *8th Annual ACM International Conference (NANOCOM)*. ACM, 1–7. doi:10.1145/3477206.3477464
[31] Wolfgang Mauerer. 2005. Semantics and simulation of communication in quantum programming. doi:10.48550/ARXIV.QUANT-PH/0511145
[32] Sam McArdle, Suguru Endo, Alán Aspuru-Guzik, Simon C. Benjamin, and Xiao Yuan. 2020. Quantum computational chemistry. *Rev. Mod. Phys.* 92, 1 (2020). doi:10.1103/revmodphys.92.015003
[33] Juan Manuel Murillo, Jose Garcia-Alonso, Enrique Moguel, Johanna Barzen, Frank Leymann, Shaukat Ali, Tao Yue, Paolo Arcaini, Ricardo Pérez-Castillo, Ignacio García-Rodríguez de Guzmán, Mario Piattini, Antonio Ruiz-Cortés, Antonio Brogi, Jianjun Zhao, Andriy Miranskyy, and Manuel Wimmer. 2025. Quantum Software Engineering: Roadmap and Challenges Ahead. *ACM Trans. Softw. Eng. Methodol.* 34, 5, Article 154 (May 2025), 48 pages. doi:10.1145/3712002
[34] Hausi Müller. 2026. IEEE Quantum Week: Why Should I Attend? https://qce.quantum.ieee.org/2026/attend/why-attend/. Accessed on 16.03.2026.
[35] Daniel Needleman and Zvonimir Dogic. 2017. Active matter at the interface between materials science and cell biology. *Nature Reviews Materials* 2, 9 (July 2017). doi:10.1038/natrevmats.2017.48
[36] John Preskill. 2018. Quantum Computing in the NISQ era and beyond. *Quantum* 2 (Aug. 2018), 79. doi:10.22331/q-2018-08-06-79
[37] John Pruitt and Jonathan Grudin. 2003. Personas: practice and theory *(DUX '03)*. ACM, New York, NY, USA, 1–15. doi:10.1145/997078.997089
[38] Ricardo Pérez-Castillo and Mario Piattini. 2022. Design of classical-quantum systems with UML. *Computing* 104, 11 (2022). doi:10.1007/s00607-022-01091-4
[39] Paul Ralph, Nauman bin Ali, Sebastian Baltes, Domenico Bianculli, Jessica Diaz, Yvonne Dittrich, Neil Ernst, Michael Felderer, Robert Feldt, Antonio Filieri, Breno Bernard Nicolau de França, Carlo Alberto Furia, Greg Gay, Nicolas Gold, Daniel Graziotin, Pinjia He, Rashina Hoda, Natalia Juristo, Barbara Kitchenham, Valentina Lenarduzzi, Jorge Martínez, Jorge Melegati, Daniel Mendez, Tim Menzies, Jefferson Molleri, Dietmar Pfahl, Romain Robbes, Daniel Russo, Nyyti Saarimäki, Federica Sarro, Davide Taibi, Janet Siegmund, Diomidis Spinellis, Miroslaw Staron, Klaas Stol, Margaret-Anne Storey, Davide Taibi, Damian Tamburri, Marco Torchiano, Christoph Treude, Burak Turhan, Xiaofeng Wang, and Sira Vegas. 2020. Empirical Standards for Software Engineering Research. doi:10.48550/ARXIV.2010.03525
[40] Ralf Ramsauer and Wolfgang Mauerer. 2025. Towards System-Level Quantum-Accelerator Integration. In *IEEE QCE*. 15–20. doi:10.1109/qce65121.2025.10286
[41] Francesc Sabater, Ouns El Harzli, Geert-Jan Besjes, Marvin Erdmann, Johannes Klepsch, Jonas Hiltrop, Jean-Francois Bobier, Yudong Cao, and Carlos A. Riofrio. 2025. Towards solving industrial integer linear programs with Decoded Quantum Interferometry. arXiv:2509.08328 [quant-ph] https://arxiv.org/abs/2509.08328
[42] Hila Safi, Medina Bandic, Christoph Niedermeier, Carmen G. Almudever, Sebastian Feld, and Wolfgang Mauerer. 2025. Stacking the odds: full-stack quantum system design space exploration. *EPJ Quantum Technology* 12, 1 (Oct. 2025). doi:10.1140/epjqt/s40507-025-00413-7
[43] Hila Safi, Karen Wintersperger, and Wolfgang Mauerer. 2023. Influence of HW-SW-Co-Design on Quantum Computing Scalability. In *IEEE QSW*. 104–115. doi:10.1109/QSW59989.2023.00022
[44] Joni Salminen, Kathleen Wenyun Guan, Soon-Gyo Jung, and Bernard Jansen. 2022. Use Cases for Design Personas: A Systematic Review and New Frontiers. In *CHI Conf. on Human Factors in Comp. Sys.* ACM. doi:10.1145/3491102.3517589
[45] Lukas Schmidbauer, Karen Wintersperger, Elisabeth Lobe, and Wolfgang Mauerer. 2024. Polynomial Reduction Methods and their Impact on QAOA Circuits. In *IEEE QSW*. 35–45. doi:10.1109/QSW62656.2024.00018
[46] Peter W. Shor. 1997. Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *SIAM J. Comput.* 26, 5 (Oct. 1997), 1484–1509. arXiv:quant-ph/9508027 doi:10.1137/S0097539795293172
[47] Krysta M. Svore, Alan Geller, Matthias Troyer, John Azariah, Christopher E. Granade, Bettina Heim, Vadym Kliuchnikov, Mariia Mykhailova, Andres Paz, and Martin Roetteler. 2018. Q#: Enabling Scalable Quantum Computing and Development with a High-Level DSL. In *Proceedings of the Real World Domain Specific Languages Workshop 2018*. ACM, Article 7. doi:10.1145/3183895.3183901
[48] Matt Swayne. 2025. Quantum Computing Roadmaps: A Look at The Maps And Predictions of Major Players. https://thequantuminsider.com/2025/05/16/quantum-computing-roadmaps-a-look-at-the-maps-and-predictions-of-major-quantum-players/. Accessed on 16.03.2026.
[49] Simon Thelen, Hila Safi, and Wolfgang Mauerer. 2024. Approximating under the Influence of Quantum Noise and Compute Power. In *IEEE QCE*. 274–279. doi:10.1109/qce60285.2024.10291
[50] Krishna Upadhyay, Vinaik Chhetri, A.B. Siddique, and Umar Farooq. 2025. Analyzing the Evolution and Maintenance of Quantum Software Repositories. In *IEEE QSW*. 173–184. doi:10.1109/QSW67625.2025.00029
[51] Markus Völter, Thomas Stahl, Jorn Bettin, Arno Haase, Simon Helsen, Krzysztof Czarnecki, and Bettina Von Stockfleth. 2013. *Model-Driven Software Development: Technology, Engineering, Management*. Wiley.
[52] Robert Wille. 2025. Connecting the "who's who" in quantum computing software. https://www.cda.cit.tum.de/research/quantum/2025_mqsf_summary/. Accessed on 16.03.2026.
[53] Arne Wulff, Boyang Chen, Matthew Steinberg, Yinglu Tang, Matthias Möller, and Sebastian Feld. 2024. Quantum computing and tensor networks for laminate design: A novel approach to stacking sequence retrieval. *Computer Methods in Appl. Mech. and Eng.* 432 (Dec. 2024). doi:10.1016/j.cma.2024.117380
[54] Tao Yue, Wolfgang Mauerer, Shaukat Ali, and Davide Taibi. 2023. *Challenges and Opportunities in Quantum Software Architecture*. Springer Nature Switzerland, 1–23. doi:10.1007/978-3-031-36847-9_1
[55] Jianjun Zhao. 2020. Quantum Software Engineering: Landscapes and Horizons. doi:10.48550/ARXIV.2007.07047