

Challenges for Quantum Software Engineering: An Industrial Application Scenario Perspective

Cecilia Carbonelli, Michael Felderer, Matthias Jung, Elisabeth Lobe, Malte Lochau, Sebastian Luber, Wolfgang Mauerer, Rudolf Ramler, Ina Schaefer, and Christoph Schroth

Abstract Quantum software is becoming a key enabler for applying quantum computing to industrial use cases. This poses challenges to quantum software engineering in providing efficient and effective means to develop such software. Eventually, this must be reliably achieved in time, on budget, and in quality, using sound and well-principled engineering approaches. Given that quantum computers are based on fundamentally different principles than classical machines, this raises the question if, how, and to what extent established techniques for systematically engineering software need to be adapted. In this paper, we analyze three paradigmatic applica-

Cecilia Carbonelli
Infineon Technologies AG, Neubiberg, Germany e-mail: cecilia.carbonelli@infineon.com

Michael Felderer
German Aerospace Center (DLR), Cologne, Germany e-mail: Michael.Felderer@dlr.de

Matthias Jung
University of Würzburg, Würzburg, Germany e-mail: m.jung@uni-wuerzburg.de

Elisabeth Lobe
German Aerospace Center (DLR), Braunschweig, Germany e-mail: Elisabeth.Lobe@dlr.de

Malte Lochau
University of Siegen, Siegen, Germany e-mail: malte.lochau@uni-siegen.de

Sebastian Luber
Infineon Technologies AG, Neubiberg, Germany e-mail: sebastian.luber@infineon.com

Wolfgang Mauerer
Technical University of Applied Sciences Regensburg/Siemens AG, Regensburg, Germany e-mail: wolfgang.mauerer@othr.de

Rudolf Ramler
Software Competence Center Hagenberg, Hagenberg, Austria, e-mail: rudolf.ramler@scch.at,

Ina Schaefer
Karlsruhe Institute of Technology, Karlsruhe, Germany e-mail: ina.schaefer@kit.edu,

Christoph Schroth
Fraunhofer IESE, Kaiserslautern, Germany e-mail: christof.schroth@iese.fraunhofer.de

tion scenarios for quantum software engineering from an industrial perspective. The respective use cases center around (1) optimization & quantum cloud services, (2) quantum simulation, and (3) embedded quantum computing. Our aim is to provide a concise overview of the current and future applications of quantum computing in diverse industrial settings. We derive presumed challenges for quantum software engineering and thus provide research directions for this emerging field.

Key words: Quantum Computing, Software engineering, Quantum Software Engineering, Industrial Use Cases, Software Development

1 Introduction

Quantum computers are a reality today, but quantum software development is in its very infancy. Although many small-/medium-sized quantum programs have been written over the years to demonstrate the potentials of quantum computing, barely any of these examples can be seriously called *quantum software*. In other words, there is no such thing as quantum software to date (Deshpande 2022).

In this regard, *software engineering* (SE) is concerned with supporting and improving the development, application, and maintenance of software-intensive systems (Sommerville 2010). SE employs scientific methods, business principles, structured process models, and predefined quality goals to cope with the complexity of software as a whole. Nowadays’s mainstream SE research for classical (i.e., non-quantum) software comprises design principles (e.g., high-level modeling languages fostering abstraction and modularity), development practices (e.g., tasks, roles, and responsibilities), and tool support (e.g., IDEs, code generation, static analysis, version control, issue tracking, unit testing, debugging, etc.). This perspective of SE research on software development is, however, mismatching the current status of quantum software. Zhao et al. were one of the first to coin the term *quantum software engineering* (QSE) to summarize any effort to adopting established SE principles and practices to make it also work for quantum software (J Zhao 2020). However, in this paper, we take on a contrary perspective: research on QSE should, as a first step, identify, understand, and tackle short-term engineering challenges for better support of, usually fully manually crafted, small-/medium scale quantum programs today (i.e., focusing on the programming and deployment phases). More sophisticated and mature concepts including high-level software abstraction as propagated, for instance, in the context of requirements elicitation, object-oriented design patterns, software maintenance and evolution, and re-engineering, are out of scope for now due to the lack of any accessible examples and use cases. To meet the short-term goals, quantum SE should first of all focus on the following challenges.

- Make quantum computing accessible to developers and users through appropriate processes, methods, and tools.

- Facilitate hybrid quantum computing through a combination of classical SE and QSE concepts based on a generic description of a computational problem and (quantum) platform constraints.
- Provide benchmarks and benchmarking processes, methods, and tools for assessing quantum advantage as well as constraints that arise from the integration of quantum software components in an overall (hybrid) software system.

Our goal is to assess the short-term requirements and challenges of SE in the upcoming era of quantum computing. These requirements and challenges are already relevant to the noisy intermediate-scale quantum (NISQ) era. In contrast to other recent works on this subject (J Zhao 2020; Ali, Yue, Abreu 2022; Weder et al. 2022; Yue et al. 2023), we do not follow a top-down approach, but instead, illustrate the status quo of QSE by considering a selection of industrial application scenarios. For each application scenario, we first provide a short general description and then describe selected recent use cases to characterize the common aspects of the respective scenarios. Based on these descriptions, we derive in a bottom-up manner the key challenges for QSE with respect to these application scenarios. Our goal is to gain a better understanding of the principles and practices that will most likely support the development of software systems that solve problems that, at least partly, involve quantum computations. Our claim is that, from an SE point of view, quantum computation is not a new programming paradigm in the first place, but, first of all, a new *computational architecture*. The novel conceptual thinking required for effectively exploiting the frequently promised *quantum advantage* is crosscutting all classical development phases and hierarchies of software systems. Quantum computing will thus potentially influence SE as a whole as we know it today (Serrano, Perez-Castillo, Piattini 2022). Nevertheless, we argue that established solutions developed in SE research over the past decades will not all suddenly become inappropriate and obsolete due to the advent of quantum computing, but instead require careful rethinking and adjustments to also cope with the key characteristics of quantum software. Many of these characteristics and possible side-effects apparent in quantum computations have been considered before in other contexts, whereas the inherent pervasiveness of these characteristics in a quantum setting is indeed a novel aspect. These characteristics include, for instance, the probabilistic nature of computational outcomes and the lack of reference architectures (although `qiskit` may be seen as a de-facto standard today for the majority of computational approaches).

“While many of quantum computing’s promised capabilities could be revolutionary, the realization of this promise requires breakthroughs in several areas, including improvements in the quality of qubits, error correction, and a demonstrable set of practical applications.” (Deshpande 2022) The inflated expectations may result in a *quantum winter* similar to what we experienced with AI, where it took a long time to turn promising theoretical concepts into reality. Thus, the immediately necessary contributions of the SE community to advancing quantum computing lies in moving from first demonstrable examples to real-world applications with practical impact.

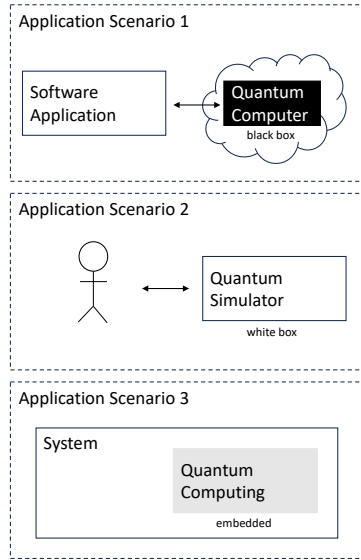


Fig. 1 Application Scenarios for Quantum Computing

2 Paradigmatic Application Scenarios

We next describe potential application domains of QSE by means of paradigmatic application scenarios as illustrated in Figure 1:

- **Application Scenario 1.** Provide quantum computing capabilities as a cloud service to solve optimization problems or machine-learning tasks (quantum-computing-as-a-service).
- **Application Scenario 2.** Perform physical simulations with quantum programs developed by domain experts in a machine-oriented low-level manner.
- **Application Scenario 3.** Embed QPUs as integrated components into hybrid safety- or mission-critical software systems with a special focus on non-functional properties.

The selection of these application scenarios is driven by industrial and academic experiences of the authors and is aligned with the core use cases of the QUTAC consortium ([Bayerstadler et al. 2021](#)). Our aim is to illustrate the diversity of application domains and different perspectives on quantum computing, ranging from recent black-box and white-box views to embedded quantum computing.

2.1 Application Scenario 1: Quantum Cloud Services

2.1.1 Use Cases and Examples

Quantum computing brings new opportunities for **solving optimization problems**, which are among the first industrial applications of the technology (Bayerstadler et al. 2021). One example is the *flight-gate assignment (FGA)* problem in airport and air traffic planning, where the assignment of incoming flights to gates shall be optimized to minimize transfer times (Stollenwerk, Lobe, Jung 2019; Sax et al. 2020). This scheduling problem belongs to a class of NP-hard combinatorial optimization problems. Further examples include *Electronic Design Automation (EDA)* such as placement and routing on design chips and fault detection in electrical power networks (Perdomo-Ortiz et al. 2015), trajectory optimization in air traffic management (Stollenwerk, O’Gorman, et al. 2019), paint-shop scheduling (Yarkoni et al. 2021), and planning problems in highly individualised mass production (Awasthi et al. 2023).

NP-hardness implies that, in practice, real-world instances can only be solved by approximation algorithms or heuristics. Here, quantum computers, taking advantage of entanglement, superposition, and interference, could potentially speed up and improve the optimization. One key property of the mentioned problems is that they can be solved **offline**: one problem instance is solved once, usually without critical time constraints, and the result is used to do something substantial: either conducting further research or going in an operational state, for instance, by applying the optimized flight schedule in an airport scenario.

2.1.2 Approaches and Challenges

A possible solution to bypass insufficient local computational power for effectively solving hard optimization problems is to **pass the work to a quantum cloud system**. For instance, D-Wave’s Leap service (D-Wave Systems Inc. n.d.) provides connections to quantum annealers or IBMs Qiskit interfacing to their quantum machines (Qiskit contributors 2023).

It has been argued that handling the offloading of such computations does not pose any new challenges to SE (Krüger, Mauerer 2020) as quantum computing, essentially appears as a black box with well-defined interfaces. However, an open issue is to properly design such interfaces and to **formulate optimization problems** being tractable by quantum processing units (QPUs). First, an emerging optimization problem may be identified as computational bottlenecks within larger application contexts. These problems are either spotted by mathematical analysis during the design phase or during an optimization stage using a profiler; ideally, they match one of the known quantum primitives (Gemeinhardt, Wille, Wimmer 2021). To this end, developers have to **refactor the overall software system** to isolate and replace the computational component by calls to quantum cloud services. Yet, there may be many such components that are closely tied to specific requirements of the overall

system and which are the result of decades' worth of research and optimization (De Andoin et al. 2022). This makes replacing them a non-trivial endeavor. Examples include sub-tasks of database management systems like join ordering (Schönberger, Scherzinger, Mauerer 2023; Schönberger, Trummer, Mauerer 2023), multi-query optimisation (Trummer, Koch 2016) or transaction scheduling (Bittner, Grope 2020).

Using quantum computing to speed up tasks previously solved by components designed and optimized for classical computers thus requires careful analysis. This includes understanding the underlying problem as well as exploring possible quantum speed-ups under varying workloads, input data characteristics, etc., while simultaneously retaining crucial, yet unrelated functional and non-functional properties of the overall system. Established SE techniques and tools (*e.g.*, for performance analysis and refactoring) may help.

However, it is fair to say that the understanding of **what benefits quantum computers can provide** for what specific problems **is far from being well understood** in comparison to the state-of-the-art in classical algorithms, and *also* in terms of system architecture.

While the potential speedups of seminal approaches (Nielsen, Chuang 2000) like Shor's algorithm (and, more general, quantum phase estimation) or Grover search are rigorously established, the impact of imperfections on these algorithms makes any practical application considerations quickly prohibitive (Greife, Krüger, Mauerer 2023). Likewise, even the actual requirements on the hardware of future machines for comparatively simple co-variables like the number of qubits is subject to discussion, and depends not only on many low-level details of the underlying hardware, but also on the actual payload algorithms (Roffe 2019). A substantial body of the existing literature is dedicated to establishing a comprehensive understanding of the theoretical advantages of quantum sampling approaches (Hangleiter, Eisert 2023); yet, concrete applications of these techniques are thinly spread, and their practical gain especially in comparison to existing heuristics and approximations is still under initial exploration (Deng et al. 2023).

The situation becomes even less straightforward for the more recent class of variational quantum algorithms and QAOA-style approaches (Blekos et al. 2023; Bharti et al. 2022). While it is known that an efficient simulation of specific variants of QAOA would have strikingly unattractive and unexpected consequences on some of the pillars of computational complexity theory, entirely classical replacements for other variants are also known (Medvidović, Carleo 2021). Likewise, the understanding of how to construct efficient classical surrogates for variational algorithms has considerably increased recently (Schuld, Sweke, Meyer 2021; Schreiber, Eisert, Meyer 2022), and restricts potential quantum advantage to increasingly narrower domains. When—unavoidable!—practical constraints are taken into account (Wintersperger, Safi, Mauerer 2022; Safi, Wintersperger, Mauerer 2023), determining a fair basis for comparison is a not yet satisfactorily resolved problem (Becker, Gheorghe-Pop, Tcholtchev 2023; Herrmann et al. 2023), even ignoring the substantial limitations of currently available hardware.

A major challenge is to identify, factor out and transform optimization sub-tasks suitable for quantum computers or annealers and their specific computing architectures (Deb, Dueck, Wille 2020; Kole et al. 2019). Several transformation steps usually reformat the optimization problem. In addition, “glue” code to connect classic and quantum parts is required. Tools like `quark` (Lobe, Stollenwerk 2022)¹ enable users to easily formulate and transform optimization problems, and to handle experiment dispatch and analysis. Likewise, approaches for recommending solution strategies for optimization problems using quantum approaches have been suggested (Poggel et al. 2023). However, such tools to support interacting and experimenting with quantum computers are still subject to research (Wille, Hillmich, Burgholzer 2022).

2.1.3 Conclusions

We are in the phase of evaluating the potentials of quantum computing in solving optimization problems. Providing such capabilities as reliable (black-box) services, however, requires an improved understanding of machine properties obtained from experiments and benchmarking. This necessitates many iterations of interactions with the quantum hardware for parameter tuning. Software development efforts, therefore, increase significantly when dealing with quantum hardware in contrast to well-established classical approaches as this fine-tuning requires not only software skills but also deeper knowledge in fundamental quantum physics. We assume this up-front investment will eventually pay off: If a fast heuristic solution is available and easy-to-access, a user will simply call quantum optimizers as a black-box cloud service over a well-defined interface, hiding transformation complexity and specific hardware requirements.

Summary. Quantum cloud services will allow for accelerating mathematical optimization problems. Automatic means of transforming existing formulations into quantum descriptions have become available; yet, it remains a software architecture and engineering challenge to identify appropriate problems. Integrating quantum solvers into applications from a black-box perspective, including interface design, remains a minor SE challenge. However, the underlying quantum computing software stack including the compilation and hybrid computing process, requires new quantum software engineering approaches.

2.2 Application Scenario 2: Quantum Simulation

2.2.1 Use Cases and Examples

Quantum simulation is one of the most promising application scenarios for quantum computing. It can help understanding real-world chemistry and physics phenomena,

¹ See also the [list of contributors](#) (link in PDF).

improving design methodologies and making experiments much more effective. Simulating quantum mechanics on classical computers is a hard computational problem², and determining relevant properties of quantum systems, *e.g.*, finding their minimal energy, are even harder. To efficiently simulate a quantum system, the simulator might rely on quantum-mechanical dynamics. The basic idea of a quantum simulator is to use a controllable quantum platform to replicate dynamic or static properties of another, usually less controllable, quantum system (Lamata et al. 2018). This is similar to using wind tunnels for testing aerodynamic properties of reduced-scale models in a controlled environment, and then to transfer gained information to full-scale objects in the (uncontrolled) real world. With the rapid growth of quantum computing capabilities, the interest in **(quantum) material science** has also risen significantly. This field targets a large variety of applications ranging from the *design of more efficient batteries and catalysts* to the study of *innovative sensing materials for consumer and automotive applications* (Paudel et al. 2022).

For the latter task, many candidate compounds have to be efficiently screened and evaluated to select or design the best materials with respect to the desired properties. This implies large effort and costs in terms of material procurement, measurement equipment and setup. Direct simulations of the material properties could drastically reduce the required resources and significantly accelerate the discovery process reducing time-to-market. Here, quantum systems promise a fast and more precise simulation tool of real-world mechanisms than their conventional counterparts.

Likewise, the study of new storage materials and the development of innovative battery technology is being pushed by several emerging and established applications, ranging from electric and light electric vehicles to solar energy storage systems and robotics. Researchers aim at understanding the mechanisms impacting efficiency, stability and faster charging of battery operations to predict real world performance. Yet the first fundamental step, again, remains the selection of apt chemical compounds. New families of disruptive active materials such as Lithium-Ion (Li-ion) and Lithium-Sulfur (Li-S) offer four-time higher energy density than Li-ion batteries. From a modeling perspective, it is crucial to describe the solid electrolyte interphase forming on the battery anode and to define its durability and long-term performance. Classic DFT, multi-physics simulations and measurements have not provided satisfactory answers particularly in terms of accuracy. Quantum computing can offer a closer characterization of the key chemical properties of battery cells such

² Problems efficiently solvable by QC belong to complexity class BQP (bounded-error quantum polynomial), the quantum analog of BPP. The relation between BQP and classical classes like NP poses many open questions. The **dynamics** of a quantum system (compute output of a quantum circuit given an initial state) is BQP-hard (Fortnow 2000), which makes it likely intractable for classical computers, but doable for quantum machines for a class of natural Hamiltonians in BQP. Inferring *global* properties of a quantum system (given a quantum circuit, is there a state that produces a desired output? What is the minimum energy eigenstate for a given Hamiltonian?) belong into QMA, a probabilistic quantum analog of NP (Aaronson 2013), and is intractable even for quantum computers. It is even possible to give physical problems that are undecidable, at least within the limit of infinite size (Cubitt, Perez-Garcia, Wolf 2015). Quantum SE needs to be aware of such peculiarities to properly ascertain the feasibility of architectures and designs by avoiding illusory, inflated expectations of potential gains.

as equilibrium cell voltages, ionic mobility and thermal stability. The quantum simulation often boils down to obtaining the ground state energies of various molecules of increasing complexity (Delgado et al. 2022); likewise, physical characteristics like dipole moments have also been calculated (Rice et al. 2021).

2.2.2 Approaches and Challenges

Programmable universal quantum computers can simulate quantum mechanical processes (Buluta, Nori 2009; Georgescu, Ashhab, Nori 2014; Altman et al. 2021; Lloyd 1996; Blatt, Roos 2012). Such simulations are specified using software (*e.g.*, using **domain-specific languages**), which takes this topic into the focus of SE. However, different approaches to quantum simulation (analog simulation, digital simulation, combinations thereof, and hybrid quantum-classical algorithms) differ in their implications. In each case, and in contrast to other forms of quantum computation, quantum simulation requires **awareness of the Hamiltonian** underlying the task (the Hamilton operator (or *Hamiltonian*) of a system, is, roughly speaking, a mathematical object³ that provides information about a physical system. It is closely related to the energy spectrum,⁴ and governs time-evolution of a quantum system. The Schrödinger equation combines Hamiltonian and quantum states, which are mathematically described by the wave function, into a differential equation).

Analog quantum simulators (Buluta, Nori 2009; Cirac, Zoller 2012) are physical systems that mimic other quantum systems (or a class of models) by closely reproducing the system's characteristics. Hence, their Hamiltonian should be as similar as possible to the simulated system. *Digital quantum simulation* is based on decomposing the Hamiltonian into operations implementable in the simulator by single- and two-qubit gate operations. This is more flexible than analog quantum simulators and enables us to overcome the limitations of the simulator system itself. Furthermore, it

³ We have been deliberately careful to avoid confusing the physical concept of a dynamical observable that can be measured with the mathematical operator/object to which it corresponds in the formal description.

⁴ Many textbooks on quantum mechanics simply state that the Hamiltonian represents the total energy of a system, sometimes requiring this as a fundamental postulate. There are reasons to avoid such strong statements, both from a fundamental perspective (in the canonical approach of replacing physical quantities in the Hamilton function H of classical mechanics with operator-valued quantities, H is always conserved, but does, as Legendre transform of the Lagrangian, not automatically equate to the sum of potential and kinetic energy; the approach to deriving a quantum Hamiltonian from energy-momentum relations delivers different results for the non-relativistic and relativistic case; and approaches based on space-time symmetries need to introduce empirical factors that relate the quantum Hamiltonian to classical energies), but also from a practical point of view that concerns the software engineering aspects of quantum simulations: It is fairly common in this field to work with *effective Hamiltonians* that describe only degrees of freedom relevant for a particular task (for instance, Spin Hamiltonians in spectroscopy, the Ligand Field Hamiltonian of coordination chemistry, or the Hückel Hamiltonian for aromatic systems, which all carry a certain relevance for quantum chemistry), and therefore do not deliver a complete energy spectrum. Correctness checks, invariants and the interpretation of results must adapt to such circumstances, and require awareness from the software side.

allows for quantum error correction and universality in a “*fully universal*” *quantum computer*. If the simulator offers a universal set of perfect quantum gates, then the model can simulate a wide class of Hamiltonians (Lloyd 1996), albeit the computing effort may vary depending on the types of gates. Some implementation technologies for QCs in use today are particularly well suited for quantum simulators. An example are Rydberg atom arrays (Morgado, Whitlock 2021; Weimer et al. 2011) that provide identical and long-lived Qubits with strong coherent interactions. To represent the physical properties of the simulated system, the properties of the simulator correspond well to these, especially when analog simulation steps are involved. At least for this aspect, this challenges the idea that abstraction layers (Bass, Clements, Kazman 2003), despite proven useful classically, can satisfactorily eliminate differences between implementation platforms.

Industrial experience with quantum simulation problems gained by some of the authors shows that the exact boundary between digital simulation and optimization is not always clear. Especially quantum-classical hybrid algorithms—most importantly, VQE (Tilly et al. 2022)—rely on optimization methods to determine observables like the ground state energy of molecules based on a physical model. It is hypothesized that VQE, which at its core is independent of the simulated problem, will provide improved modeling accuracy over classical approaches like DFT. However, engineering challenges remain such as hardware dependent **noise compensation**, an understanding of the differences between the many available variants of VQE (Fedorov et al. 2022) (requiring **problem-dependent benchmarking** (Resch, Karpuzcu 2021)), and **determining optimal quantum-classical splits**. Especially the latter topics fall within the responsibilities of SE, but it might also be possible to improve noise handling based on software-centric methods. Also, the **depth reduction of circuits** generated from Hamiltonian descriptions is an important goal, in which compilers may play a crucial role (see, e.g., (Ding et al. 2020; Kreppel et al. 2022; Schmale et al. 2022)). As with other use cases, resource usage and scalability in general need to be addressed by QSE.

Despite initial steps taken on problems of industrial scale, explorations are still in an early phase with already important collaborations emerging between large chemical and computing technology corporations (Business Value 2021). Currently, the effort of finding appropriate Hamiltonian models by far exceeds the software implementation effort; knowledge of physical principles and details by far outranks the challenges of transcribing these into quantum framework. While the modelling task in the classical domain is routinely reduced to a well-informed parametrization of canned DFT software, quantum tools—even given existing frameworks support (Qiskit contributors 2023)—require high manual programming effort.

2.2.3 Conclusions

SE tasks in quantum simulation include algorithm selection, determining the influence of mathematical/physical details on non-functional and functional properties, and comparing quantum and hybrid architectures to classical approaches and

heuristics. Many revolve, in a broader sense, around the topic of **testing**. As the goal of quantum simulation is to exceed the computational capabilities of classical approaches, this opens up new research challenges. Testing quantum simulations comprises ensuring (a) model correctness, and (b) correctness of circuits generated from the model. After establishing a Hamiltonian description of the system, empirical measurements on the actual physical system can be performed and compared to the simulation results. The resulting circuit generator is then trusted, and quantum simulation based on the generators can be used to explore the properties of novel, previously unexplored materials.

From an SE point of view, recent attempts lift established techniques for end-to-end testing of classical computations to components with probabilistic behavior (Miransky, Zhang 2019; Gerhold, Stoelinga 2018; Feng, Duan, Ying 2012; García de la Barrera et al. 2023). This includes novel notions of testing oracles based on distance measures for execution trace distributions and statistical criteria for approximating error probability by the number of repetitions of test runs. More involved quantum phenomena like superposition and entanglement of computational states are not yet properly addressed by these approaches. This, first of all, requires new abstractions concerning the notion of *observations* in testing reflecting the destructive nature of quantum measurements which obstructs established testing practices like interactive debugging (Miransky, Zhang 2019).

Further properties of quantum states and circuits are also not suited to established testing methods: As there are usually no classical control branches in quantum circuits, structural code coverage criteria are not applicable, which renders well-established, elementary software testing concepts (Sommerville 2010) useless. Likewise, localization of faults is unlike harder for quantum circuits than for classical programs, given that entangled states can intertwine arbitrary parts of a circuit and mutually influence each other. Not just the stochastic nature of quantum measurements, but also the impact of imperfection and noise in quantum circuits obstruct the definition of proper test oracles. Here, we need to distinguish unavoidable variations caused by quantum measurements from variations due to (classically) probabilistic algorithmic elements from variations induced by noise and imperfection. Distinguishing between such different probability distributions is no new challenge, but there are quantum specifics: For instance, the amount of information to be recorded for a meaningful statements (*e.g.*, by estimating the required number of samples for a desired precision and bounded error probability via Hoeffding's inequality (Pashayan, Wallman, Bartlett 2015), or randomised measurement procedures (Elben et al. 2023) that estimate quantum properties from classical observations) requires future research in QSE. Well-principled guidelines can eliminate the need for individual software engineers to be aware of such statistical peculiarities.

Other verification approaches for quantum simulation include upfront correctness validation of models (*e.g.*, finding physical invariants that can be probed with accessible measurements), equipping a model's software representation (or the representation of the simulation approach) with a formal semantics honours quantum aspects (*e.g.*, (Mauerer 2005; Bichsel et al. 2020; Cross et al. 2022; Evans et al. 2022)) that allows us to verify specific properties and correctness of generated mod-

eling circuits by decomposition techniques (see, *e.g.*, (Peham, Burgholzer, Wille 2022; Ufrecht et al. 2023)).

Summary. Quantum simulation can benefit from established means of SE to formulate and describe models of physical systems whose properties can be simulated on quantum computers. Efforts evolve more around a physical understanding of the employed models rather than programming. Validation and verification techniques, as well as architectural decomposition into quantum and classical aspects, will rely on established, yet to be adapted SE approaches.

2.3 Application Scenario 3: Embedded Quantum Computing

2.3.1 Use Cases and Examples

Embedded software systems are purpose-built for specific tasks. In contrast to general-purpose and high-performance computing systems (Application Scenario 1), embedded systems operate under restricted resources, on specific hardware platforms, and have to meet distinct quality requirements like **real-time constraints** or **safety guarantees**. Safety measures prevent from material damage and harm to individuals and deeply influence hardware and software co-design of classical embedded software (Marwedel 2011).

2.3.2 Approaches and Challenges

We recently observed a convergence between embedded systems and high-performance computing (Girbal et al. 2013), for instance, in autonomous driving, avionics, and control systems. We expect embedded systems to require even more computational resources in future applications. Hence, quantum computing may also play an important role in **hybrid embedded scenarios** by utilizing **quantum accelerators** for solving particular computational tasks (Wintersperger, Safi, Mauerer 2022). To the best of our knowledge, no approaches have been investigated so far to facilitate quality assurance techniques and tools for embedded quantum computing. Meeting these requirements in QPU accelerated hybrid systems is complicated by the dominance of iterative, probabilistic algorithms; yet, since almost all known quantum algorithms that operate on perfect error-corrected quantum systems are also inherently *probabilistic* (Lubinski et al. 2022), the problem will also extend after the NISQ area. Open research questions include how to improve understanding of termination properties and convergence toward sufficiently accurate results in iterative algorithms (Farhi, Goldstone, Gutmann 2014; Akshay et al. 2020), as well as the role of classical optimization components (Zhou et al. 2020; Streif, Leib 2020) and result degradation (but, perhaps counter-intuitively, also possible improvements (Liu et al. 2022) by imperfections and noise (Alam, Ash-Saki, Ghosh 2020; Wang et al. 2021).

In many application domains, embedded co-design development processes must achieve (safety) certifications. It is an open question how established approaches can be adapted to quantum computing including entirely novel qualification approaches aligned with QPU peculiarities. Therefore, we may expect that system engineering will play a larger role in hybrid embedded quantum computing than for classical appliances.

Prior work in safety-critical embedded systems deals with probabilistic algorithms and machine learning (*e.g.*, neural networks) in the context of unreliable hardware. Measures include redundant computation, error correction (Steiner et al. 2022), as well as more high-level concepts like safety cages (Heckemann et al. 2011) and static partitioning (Ramsauer et al. 2022), *Digital Dependability Identities* (DDI) (Reich, Schneider, et al. 2020), and *Dynamic Risk Management* (DRM) (Reich, Wellstein, et al. 2021). It is not obvious if and how these approaches can be adopted for quantum computing. It is also crucial to consider how to integrate QPUs into existing embedded development processes and infrastructures. This includes interface design (at the physical and protocol level) to ensure proper timing and co-scheduling of computational tasks offloaded to a quantum component. The integration of QPUs further impacts the software operating systems level and middleware layers. Given the strong influence of imperfection of QPUs in the near- and mid-term (Bharti et al. 2022), QPU integration will also impact co-design of hardware and algorithms to ensure computational advantages for a given set of problems. The established approaches to hardware-software co-design are currently adapted to interactions between QPUs and classical system components (G Li et al. 2021; Algaba et al. 2022), with efforts ranging from traditional embedded systems design to integration with high-performance computing (Wintersperger, Safi, Mauerer 2022; Schulz et al. 2022), (all of which also pose software engineering challenges). The feasibility of co-design decisions strongly depends on the underlying physical implementation technology, which influences the quality properties of any software executed on top.

Since embedded systems are employed in industrial and cost-sensitive domains, economic considerations are also important in a quantum setting, especially given that even in the upcoming era of fully error corrected quantum computers (but even more so in the NISQ era), different physical implementations of the computational concepts will offer different characteristics depending on their physical implementation (Wintersperger, Dommert, et al. 2023). A quantum approach with marginal improvement over existing solutions at the expense of inflating the bill of materials (or other development costs) is neither intellectually satisfying nor economically desirable. Embedded quantum SE must consider these issues.

2.3.3 Conclusions

The main challenges to enable hybrid embedded quantum computing include novel co-design principles and practices to adopt quality assurance techniques (*e.g.*, embedded systems testing) and corresponding certification processes to a quantum setting. This is particularly crucial in safety-critical application domains. In the near

term, we may expect quantum computing to find its way into in large-scale embedded systems only (*e.g.*, in CT-scanners). In contrast, the physical size of recent quantum computers is the main limiting factor for small-scaled, mobile use cases such as automotive *Electronic Compute Units* (ECU). These limitations of first-generation QPU are not quantum inherent and future quantum technology may provide quantum accelerators fitting into small, well-integrated embedded systems.

Summary. We expect that QPU, given increasing miniaturization, will be deployed as accelerators in embedded use cases. This requires applications (and extensions) of established co-design methods from embedded SE that also lean substantially towards systems engineering. Quality assurance, certification requirements, and economic and physical constraints will play pronounced roles.

3 Promises and Perils of Quantum Software Engineering

3.1 Promises and Opportunities

Application scenario 1 is aligned with classical SE for developing complete software solutions by making use of quantum cloud services, whereas application scenario 2 crosscuts classical SE and instead seeks for support of craftsmanship by individual experts. Application scenario 3 demands for principles and practices similar to systems engineering for quality-aware integration of heterogeneous software/hardware components on a computational platform. From these observations, we conclude that the work with quantum computing is, and will be, similar to the development process using embedded accelerators, such as GPUs or special-purpose hardware (see Figure 2). Similar to HW/SW co-design approaches, we expect that HW/SW/QC co-design processes will be required to split classical from quantum software parts (Dey et al. 2020; Pérez-Delgado, Perez-Gonzalez 2020). Likewise, a number of proposals have been made regarding more general questions of software architecture for quantum-classical hybrid systems, for instance (Sitdikov et al. 2023; Saurabh, Jha, Luckow 2023; Furutanpey et al. 2023).

After the diverse software parts are completed and tested as separate units (taking into account that quantum aspects bring additional challenges to reproducibility aspects (Mauerer, Scherzinger 2022)), an integration test step is required. Ideally, those steps will be embedded into continuous engineering processes (Antonino et al. 2018), *e.g.*, by making use of virtual hardware platforms or simulators for faster feedback cycles. We next discuss challenges of QSE by considering the respective SE phases. Many of these challenges have been already mentioned in recent surveys on QSE (J Zhao 2020), whereas our attempt is to relate these aspects to the insights gained from all three application scenarios described above.

Requirements Engineering. The requirements engineering phase will not fundamentally change as requirements, by definition, deal with the *What?* and not the *How?* in software projects. Hence, system-level requirements for QC do not sub-

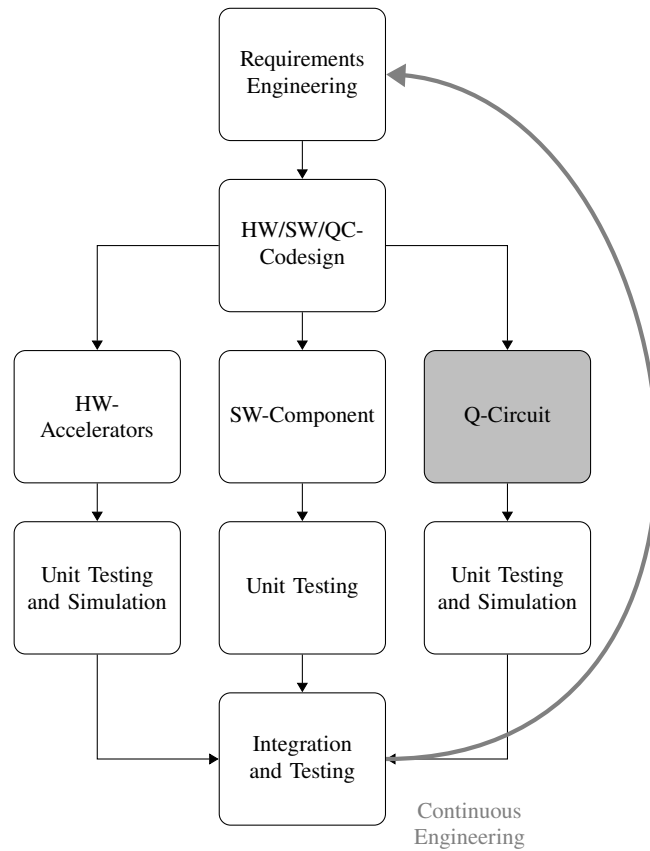


Fig. 2 Quantum Software Development Process

stantially differ from classical requirements. However, new types of non-functional requirements specific to quantum software in combination with quantum hardware might become relevant.

Systems Design/Architecture (HW/SW/QC-Co-Design). In this phase, the problem splitting between classical and quantum tasks takes place: the engineer decides which parts of the overall problem are solved by classical computations and which ones by quantum solutions. This requires architectural guidelines and patterns, as well as interface descriptions for interactions between classical and quantum data.

Programming Languages and Implementation (Q-Circuit). In this phase, algorithms need to be realized for the classical as well as for the quantum parts of a given problem. For the classical part, programming languages and compilation is well known. However, for implementing quantum algorithms, we currently rely on gate-level languages (even in case of seemingly higher-level quantum program-

ming languages like Q#). While gate-level languages are essentially the quantum-equivalent to classical assembly languages, for more efficient implementation, we need appropriate high-level quantum programming abstractions. Furthermore, we need programming guidelines and idioms, as well as design patterns for quantum programming languages. Yamaguchi, Yoshioka 2023 Design by contract for quantum software

Compilation and Deployment (Q-Circuit) Today, each quantum hardware comes with its own hardware specifics, *e.g.*, gates that can be implemented easily or at all and to which qubits these gates can apply. This requires machine-specific compilation and transpilation techniques. OpenQASM is only becoming a de-facto standard for hardware-level quantum programming. In order to allow for more efficient development and execution of quantum programs, we need a common intermediate language, *e.g.*, OpenQASM, and generic compilation techniques. This includes instruction set selection and back-end optimization that can be easily adapted and configured for specific hardware. Here, ideas for classical compiler-compilers may become useful again to automatically generate hardware-specific compilers.

This aspect naturally includes devising new methods to (statically) checking desirable properties and guarantees of quantum programs at compile time; first steps into this direction have already been taken (P Zhao, Wu, Z Li, et al. 2023; Klamroth et al. 2023; Xia, J Zhao 2023; Quetschlich, Burgholzer, Wille 2023).

Testing and Verification. In the spirit of the V-model and similar development models, the approaches in this phase complement the approaches of the respective development phases. Recent techniques for testing and verification of (partly) probabilistic hybrid software systems, may provide a conceptual foundation for assuring that the observed output behavior of quantum components are conforming to a given specification (Gerhold, Stoelinga 2018; Miranskyy, Zhang 2019; Ying 2019; Feng, Duan, Ying 2012; Abreu et al. 2022). Corresponding black-box techniques are applicable at the functional unit level as well as the system integration level of the hybrid system, by abstracting from any internal details of quantum components (application scenario 1 and 3). In contrast, in the case of a white-box setting (application scenario 2), it is not obvious how to adopt established techniques for software testing (*e.g.*, interactive debugging (Miranskyy, Zhang 2019)) and verification of quantum computations. The first step in this direction may be to find sound abstractions that properly reflect quantum-specific phenomena like superposition and entanglement of computational states and the destructive nature of quantum measurements. As always, testing and verification aim at improving software quality and minimizing the number of bugs; first steps into the direction of understanding the quantum specific aspects of these goals have been taken (P Zhao, Wu, Luo, et al. 2023).

3.2 Perils

Quantum computing will benefit from established software engineering techniques. The synthesis of both fields will likely put a few new topics on the joint research

agenda. However, there is also good reason to predict that quantum software engineering will (a) likely not radically change most established means of software engineering, and (b) will not benefit from inapt, straightforward adaptations of existing insights. In particular, we argue that this concerns the use of modeling languages and adaptations of development processes.

Albeit special-purpose quantum languages are available, most development activities in the NISQ either comprise using quantum functionalities on the API level, or constructing gate sequences that are applied on qubits. Dispatching and orchestration aspects are embedded into a classical host scripting language, typically Python (Qiskit contributors 2023; Sivarajah et al. 2020). The translation between different APIs is currently near-trivial (Schönberger, Franz, et al. 2022). Special-purpose quantum programming languages (or extensions to classical languages) promise to lift the specification or verification of quantum *algorithms* to more appropriate levels of abstraction that require less manual handling of details. We are not aware of an argument why abstraction levels that transcend algorithmic implementation details, and thus avoiding quantum specifics, would necessarily need to be crafted differently than in the classical case. Of course, it is possible to use mechanisms like UML that were intended to model software designs for describing low-level details of qubits, quantum registers, and gates. Yet it would be also possible to model classical bits, registers and electronic gates using UML in the same way; since we are not aware of any beneficial application of such a technique to the best of our knowledge, this underlines the importance of not mixing modeling techniques targeted at high levels of abstraction with low-level details. While the *design* of algorithms for quantum systems is entirely different from classical algorithms (and systematic methods range among the most challenging unsolved problems in the field), implementation details in general do almost never concern modeling at a higher level (Bass, Clements, Kazman 2003), and should therefore continue not do so in quantum software development.

Again to the best of our knowledge, using entirely non-standard development processes in specific domains is not commonly reported in the literature. Likewise, we are unaware of specially crafted software development processes—unlike architectures—that are beneficial when components like GPU accelerators or target domains like cloud deployments are considered. As we have argued in the use-case discussion above, GPUs can be seen as computational accelerators (in local appliances) or cloud resources (in distributed systems), which by analogy suggests that any such specially crafted processes will not lead to pronounced advantages. Additionally, software engineering research often finds little to no difference (Mauerer, Joblin, et al. 2021) when the implications of various forms of (social) process interactions between developers are studied for software in different domains. This insight further strengthens the hypothesis that quantum software development can be based on existing processes, and inherit the advantages and disadvantages of each approach.

Consequently, we find it unlikely that quantum software engineering in any of the scenarios described in this chapter will require entirely new development processes, or non-trivial modifications of existing approaches. Since no substantial body of

quantum software exists yet, mining quantitative empirical evidence towards one side or another will likely not be conclusive at this stage. While it cannot be ruled out that, for instance, UML will be an appropriate tool to design algorithms at gate and qubit level, or that entirely new development processes will need to be devised to implement quantum software, we call for caution before making overly ambitious statements without conclusive evidence, which could either be derived from sound ab-initio considerations, or empirically observed from mounting industrial and academic experience with creating concrete quantum software.

4 Summary and Outlook

Quantum computing is still in a very early stage with major challenges ahead. Many of these challenges have to be addressed by advancing quantum computing at the hardware level. Nevertheless, quantum computing will not only be pushed by innovations in physics, leading to advancements in quantum hardware but progress can also be expected by a pull effect caused by innovative future applications. Or, according to the aforementioned quote by Deshpande (Deshpande 2022), realizing practical applications is indeed in the domain of SE.

Nevertheless, quantum software development will not cause a revolution in SE, neither today nor in the foreseeable future. The overall aim of many SE principles (*e.g.*, separation of concerns, encapsulation, and information hiding, just to name a few) is exactly to be agnostic to diverse (existing and future) computational platforms. Hence, we should be more interested in those characteristics of quantum computations which have been exotic corner cases in SE until now but will soon become omnipresent in quantum software development.

Moreover, quantum software development today mostly happens at source code level and reaching downwards to assembly level. Quantum programming today mostly means to custom-tailor a quantum solution to a very specific instruction set of a specifically developed special-purpose quantum computer. The tendency in main-stream SE today is, however, to abstract exactly from those low-level details and instead focus on requirements and design issues. Hence, recently outdated, former core disciplines of mainstream SE research like compiler construction and instruction set architecture design will become highly relevant again.

Acknowledgements. This work is supported by the German Federal Ministry of Education and Research within the funding program quantum technologies—from basic research to market—contract numbers 13N16303 (IS) and 13N15647/13N16092 (WM). WM also acknowledges support from the High-Tech Agenda of the Free State of Bavaria. The work is further partly supported by the Austrian ministries BMK, BMAW, and the State of Upper Austria in the frame of the FFG COMET competence center INTEGRATE at SCCH as well as by BMBWF via the project QuantumReady (RR). We also acknowledge the use of IBM Quantum services and the Fraunhofer quantum computer in Ehningen for this work (MJ).

References

- Aaronson S (2013) *Quantum Computing since Democritus*. Cambridge University Press, USA ISBN: 0521199565
- Abreu R et al. (2022) “Metamorphic Testing of Oracle Quantum Programs” *2022 IEEE/ACM 3rd International Workshop on Quantum Software Engineering (QSE)*: 16–23 DOI: [10.1145/3528230.3529189](https://doi.org/10.1145/3528230.3529189)
- Akshay V et al. (Mar. 2020) Reachability Deficits in Quantum Approximate Optimization. *Phys. Rev. Lett.* 124 (9): 090504 DOI: [10.1103/PhysRevLett.124.090504](https://doi.org/10.1103/PhysRevLett.124.090504) URL: <https://link.aps.org/doi/10.1103/PhysRevLett.124.090504>
- Alam M, Ash-Saki A, Ghosh S (2020) “Design-Space Exploration of Quantum Approximate Optimization Algorithm under Noise” *2020 IEEE Custom Integrated Circuits Conference (CICC)*: 1–4 DOI: [10.1109/CICC48029.2020.9075903](https://doi.org/10.1109/CICC48029.2020.9075903)
- Algaba MG et al. (Nov. 2022) Co-Design quantum simulation of nanoscale NMR. *Phys. Rev. Res.* 4 (4): 043089 DOI: [10.1103/PhysRevResearch.4.043089](https://doi.org/10.1103/PhysRevResearch.4.043089) URL: <https://link.aps.org/doi/10.1103/PhysRevResearch.4.043089>
- Ali S, Yue T, Abreu R (Mar. 2022) When Software Engineering Meets Quantum Computing. *Commun. ACM* 65(4): 84–88 ISSN: 0001-0782 DOI: [10.1145/3512340](https://doi.org/10.1145/3512340) URL: <https://doi.org/10.1145/3512340>
- Altman E et al. (2021) Quantum simulators: Architectures and opportunities. *PRX Quantum* 2(1): 017003
- Antonino PO et al. (2018) “Enabling Continuous Software Engineering for Embedded Systems Architectures with Virtual Prototypes” *Software Architecture* ed. by CE Cuesta, D Garlan, J Pérez Springer International Publishing, Cham: 115–130 ISBN: 978-3-030-00761-4
- Awasthi A et al. (2023) *Quantum Computing Techniques for Multi-Knapsack Problems*. DOI: [10.48550/ARXIV.2301.05750](https://arxiv.org/abs/2301.05750) URL: <https://arxiv.org/abs/2301.05750>
- Bass L, Clements P, Kazman R (2003) *Software Architecture in Practice*. SEI series in software engineering Addison-Wesley ISBN: 9780321154958
- Bayerstadler A et al. (Nov. 2021) Industry Quantum Computing Applications. *EPJ Quantum Technology* 8(1): DOI: [10.1140/epjqt/s40507-021-00114-x](https://doi.org/10.1140/epjqt/s40507-021-00114-x) URL: <https://epjquantumtechnology.springeropen.com/track/pdf/10.1140/epjqt/s40507-021-00114-x.pdf>
- Becker CK-U, Gheorghe-Pop I-D, Tcholtchev N (2023) “A Testing Pipeline for Quantum Computing Applications” *Proceedings of the IEEE international conference on quantum software IEEE*
- Bharti K et al. (Feb. 2022) Noisy intermediate-scale quantum algorithms. *Rev. Mod. Phys.* 94 (1): 015004 DOI: [10.1103/RevModPhys.94.015004](https://doi.org/10.1103/RevModPhys.94.015004) URL: <https://link.aps.org/doi/10.1103/RevModPhys.94.015004>
- Bichsel B et al. (2020) “Silq: A High-Level Quantum Language with Safe Uncomputation and Intuitive Semantics” *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation PLDI 2020* Association for Computing Machinery, London, UK: 286–300 ISBN: 9781450376136

- DOI: [10.1145/3385412.3386007](https://doi.org/10.1145/3385412.3386007) URL: <https://doi.org/10.1145/3385412.3386007>
- Bittner T, Grope S (2020) “Avoiding Blocking by Scheduling Transactions Using Quantum Annealing” *Proceedings of the 24th Symposium on International Database Engineering & Applications IDEAS '20* Association for Computing Machinery, Seoul, Republic of Korea ISBN: 9781450375030 DOI: [10.1145/3410566.3410593](https://doi.org/10.1145/3410566.3410593) URL: <https://doi.org/10.1145/3410566.3410593>
- Blatt R, Roos CF (2012) Quantum simulations with trapped ions. *Nature Physics* 8(4): 277–284
- Blekos K et al. (2023) *A Review on Quantum Approximate Optimization Algorithm and its Variants*. arXiv: [2306.09198](https://arxiv.org/abs/2306.09198) [quant-ph]
- Buluta I, Nori F (2009) Quantum Simulators. *Science* 326(5949): 108–111 DOI: [10.1126/science.1177838](https://doi.org/10.1126/science.1177838) eprint: <https://www.science.org/doi/pdf/10.1126/science.1177838> URL: <https://www.science.org/doi/abs/10.1126/science.1177838>
- Business Value II for (2021) *The Quantum Decade: A Playbook for Achieving Awareness, Readiness, and Advantage*. IBM Institute for Business Value ISBN: 9781737401100 URL: https://books.google.de/books?id=MeN%5C_zgEACAAJ
- Cirac JI, Zoller P (2012) Goals and opportunities in quantum simulation. *Nature Physics* 8(4): 264–266 DOI: [10.1038/nphys2275](https://doi.org/10.1038/nphys2275) URL: <https://doi.org/10.1038/nphys2275>
- Cross A et al. (Sept. 2022) OpenQASM 3: A Broader and Deeper Quantum Assembly Language. *ACM Transactions on Quantum Computing* 3(3): ISSN: 2643-6809 DOI: [10.1145/3505636](https://doi.org/10.1145/3505636) URL: <https://doi.org/10.1145/3505636>
- Cubitt TS, Perez-Garcia D, Wolf MM (2015) Undecidability of the spectral gap. *Nature* 528(7581): 207–211 DOI: [10.1038/nature16059](https://doi.org/10.1038/nature16059) URL: <https://doi.org/10.1038/nature16059>
- D-Wave Systems Inc. (n.d.) *D-Wave Systems Leap Cloud Service*. <https://cloud.dwavesys.com/leap/> visited 2023-03-03
- De Andoin MG et al. (2022) “Comparative Benchmark of a Quantum Algorithm for the Bin Packing Problem” *2022 IEEE Symposium Series on Computational Intelligence (SSCI)* IEEE: 930–937
- Deb A, Dueck GW, Wille R (2020) Exploring the potential benefits of alternative quantum computing architectures. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 40(9): 1825–1835
- Delgado A et al. (Sept. 2022) Simulating key properties of lithium-ion batteries with a fault-tolerant quantum computer. *Phys. Rev. A* 106 (3): 032428 DOI: [10.1103/PhysRevA.106.032428](https://doi.org/10.1103/PhysRevA.106.032428) URL: <https://link.aps.org/doi/10.1103/PhysRevA.106.032428>
- Deng Y-H et al. (May 2023) Solving Graph Problems Using Gaussian Boson Sampling. *Phys. Rev. Lett.* 130 (19): 190601 DOI: [10.1103/PhysRevLett.130.190601](https://doi.org/10.1103/PhysRevLett.130.190601) URL: <https://link.aps.org/doi/10.1103/PhysRevLett.130.190601>

- Deshpande A (2022) Assessing the quantum-computing landscape. *Communications of the ACM* 65(10): 57–65
- Dey N et al. (2020) *QDLC – The Quantum Development Life Cycle*. arXiv: 2010.08053 [cs.ET]
- Ding Y et al. (2020) “Systematic Crosstalk Mitigation for Superconducting Qubits via Frequency-Aware Compilation” *2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*: 201–214 DOI: 10.1109/MICRO50266.2020.00028
- Elben A et al. (2023) The randomized measurement toolbox. *Nature Reviews Physics* 5(1): 9–24 DOI: 10.1038/s42254-022-00535-2 URL: <https://doi.org/10.1038/s42254-022-00535-2>
- Evans A et al. (2022) *MCBeth: A Measurement Based Quantum Programming Language*. arXiv: 2204.10784 [cs.PL]
- Farhi E, Goldstone J, Gutmann S (2014) *A Quantum Approximate Optimization Algorithm*. DOI: 10.48550/ARXIV.1411.4028 URL: <https://arxiv.org/abs/1411.4028>
- Fedorov DA et al. (2022) VQE method: a short survey and recent developments. *Materials Theory* 6(1): 2 DOI: 10.1186/s41313-021-00032-6 URL: <https://doi.org/10.1186/s41313-021-00032-6>
- Feng Y, Duan R, Ying M (2012) Bisimulation for Quantum Processes. *ACM Trans. Program. Lang. Syst.* 34(4): DOI: 10.1145/2400676.2400680 URL: <https://doi.org/10.1145/2400676.2400680>
- Fortnow L (2000) One Complexity Theorist’s View of Quantum Computing. *Electronic Notes in Theoretical Computer Science* 31 CATS 2000 Computing: the Australasian Theory Symposium: 58–72 ISSN: 1571-0661 DOI: [https://doi.org/10.1016/S1571-0661\(05\)80330-5](https://doi.org/10.1016/S1571-0661(05)80330-5) URL: <https://www.sciencedirect.com/science/article/pii/S1571066105803305>
- Furutanpey A et al. (2023) “Architectural Vision for Quantum Computing in the Edge-Cloud Continuum” *Proceedings of the IEEE international conference on quantum software* IEEE
- García de la Barrera A et al. (2023) Quantum software testing: State of the art. *Journal of Software: Evolution and Process* 35(4): e2419 DOI: <https://doi.org/10.1002/smr.2419> URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/smr.2419>
- Gmeinhardt FG, Wille R, Wimmer M (2021) Quantum k-community detection: algorithm proposals and cross-architectural evaluation. *Quantum Information Processing* 20(9): 302
- Georgescu IM, Ashhab S, Nori F (2014) Quantum simulation. *Reviews of Modern Physics* 86(1): 153
- Gerhold M, Stoelinga M (Jan. 2018) Model-based testing of probabilistic systems. *English Formal aspects of computing* 30(1): 77–106 ISSN: 0934-5043 URL: <https://doi.org/10.1007/s00165-017-0440-4>
- Girbal S et al. (May 2013) “On the convergence of mainstream and mission-critical markets” *2013 50th ACM/EDAC/IEEE Design Automation Conference (DAC)*: 1–10 DOI: 10.1145/2463209.2488962

- Greiwe F, Krüger T, Mauerer W (2023) “Effects of Imperfections on Quantum Algorithms: A Software Engineering Perspective” *Proceedings of the IEEE international conference on quantum software* IEEE
- Hangleiter D, Eisert J (July 2023) Computational advantage of quantum random sampling. *Reviews of Modern Physics* 95(3): DOI: [10.1103/revmodphys.95.035001](https://doi.org/10.1103/revmodphys.95.035001) URL: <https://doi.org/10.1103%2Frevmodphys.95.035001>
- Heckemann K et al. (2011) “Safe Automotive Software” *Knowledge-Based and Intelligent Information and Engineering Systems* ed. by A König et al. Springer Berlin Heidelberg, Berlin, Heidelberg: 167–176 ISBN: 978-3-642-23866-6
- Herrmann N et al. (2023) “Quantum utility—definition and assessment of a practical quantum advantage” *Proceedings of the IEEE international conference on quantum software* IEEE
- Klamroth J et al. (2023) “QIn: Enabling Formal Methods to Deal with Quantum Circuits” *Proceedings of the IEEE international conference on quantum software* IEEE
- Kole A et al. (2019) Improved mapping of quantum circuits to IBM QX architectures. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 39(10): 2375–2383
- Kreppel F et al. (2022) *Quantum Circuit Compiler for a Shuttling-Based Trapped-Ion Quantum Computer*. DOI: [10.48550/ARXIV.2207.01964](https://doi.org/10.48550/ARXIV.2207.01964) URL: <https://arxiv.org/abs/2207.01964>
- Krüger T, Mauerer W (2020) Quantum Annealing-Based Software Components: An Experimental Case Study with SAT Solving. *Q-SE@ICSE*: URL: <https://arxiv.org/abs/2005.05465>
- Lamata L et al. (2018) Digital-analog quantum simulations with superconducting circuits. *Advances in Physics: X* 3(1): 1457981 DOI: [10.1080/23746149.2018.1457981](https://doi.org/10.1080/23746149.2018.1457981) eprint: <https://doi.org/10.1080/23746149.2018.1457981> URL: <https://doi.org/10.1080/23746149.2018.1457981>
- Li G et al. (2021) “On the Co-Design of Quantum Software and Hardware” *Proceedings of the Eight Annual ACM International Conference on Nanoscale Computing and Communication* NANOCOM ’21 Association for Computing Machinery, Virtual Event, Italy ISBN: 9781450387101 DOI: [10.1145/3477206.3477464](https://doi.org/10.1145/3477206.3477464) URL: <https://doi.org/10.1145/3477206.3477464>
- Liu J et al. (Oct. 2022) *Noise can be helpful for variational quantum algorithms*. arXiv: [2210.06723](https://arxiv.org/abs/2210.06723) [quant-ph]
- Lloyd S (1996) Universal Quantum Simulators. *Science* 273(5278): 1073–1078 DOI: [10.1126/science.273.5278.1073](https://doi.org/10.1126/science.273.5278.1073) eprint: <https://www.science.org/doi/pdf/10.1126/science.273.5278.1073> URL: <https://www.science.org/doi/abs/10.1126/science.273.5278.1073>
- Lobe E, Stollenwerk T (Feb. 2022) *QUARK*. URL: <https://quantum-computing-software.gitlab.io/quark/>
- Lubinski T et al. (2022) Advancing hybrid quantum-classical computation with real-time execution. *Frontiers in Physics* 10: ISSN: 2296-424X DOI: [10.3389/fphy.2022.940293](https://doi.org/10.3389/fphy.2022.940293) URL: <https://www.frontiersin.org/articles/10.3389/fphy.2022.940293>

- Marwedel P (2011) *Embedded System Design - Embedded Systems Foundations of Cyber-Physical Systems, Second Edition*. Embedded Systems Springer ISBN: 978-94-007-0256-1 DOI: [10.1007/978-94-007-0257-8](https://doi.org/10.1007/978-94-007-0257-8) URL: <https://doi.org/10.1007/978-94-007-0257-8>
- Mauerer W (2005) *Semantics and simulation of communication in quantum programming*. DOI: [10.48550/ARXIV.QUANT-PH/0511145](https://arxiv.org/abs/quant-ph/0511145) URL: <https://arxiv.org/abs/quant-ph/0511145>
- Mauerer W, Joblin M, et al. (May 2021) In Search of Socio-Technical Congruence: A Large-Scale Longitudinal Study. *IEEE Transactions on Software Engineering* (01): 1–1 ISSN: 1939-3520 DOI: [10.1109/TSE.2021.3082074](https://www.computer.org/csdl/journal/ts/5555/01/09436025/1tJsglfkGru) URL: <https://www.computer.org/csdl/journal/ts/5555/01/09436025/1tJsglfkGru>
- Mauerer W, Scherzinger S (2022) 1-2-3 Reproducibility for Quantum Software Experiments. *Q-SANER@IEEE International Conference on Software Analysis, Evolution and Reengineering*:
- Medvidović M, Carleo G (2021) Classical variational simulation of the quantum approximate optimization algorithm. *npj Quantum Information* 7(1): 101
- Miranskyy A, Zhang L (2019) “On Testing Quantum Programs” *2019 IEEE/ACM 41st International Conference on Software Engineering: New Ideas and Emerging Results (ICSE-NIER)*: 57–60 URL: <https://doi.org/10.1109/ICSE-NIER.2019.00023>
- Morgado M, Whitlock S (June 2021) Quantum simulation and computing with Rydberg-interacting qubits. *AVS Quantum Science* 3(2): 023501 DOI: [10.1116/5.0036562](https://doi.org/10.1116/5.0036562) URL: <https://doi.org/10.1116/5.0036562>
- Nielsen MA, Chuang IL (2000) *Quantum Computation and Quantum Information*. Cambridge University Press
- Pashayan H, Wallman JJ, Bartlett SD (Aug. 2015) Estimating Outcome Probabilities of Quantum Circuits Using Quasiprobabilities. *Phys. Rev. Lett.* 115 (7): 070501 DOI: [10.1103/PhysRevLett.115.070501](https://link.aps.org/doi/10.1103/PhysRevLett.115.070501) URL: <https://link.aps.org/doi/10.1103/PhysRevLett.115.070501>
- Paudel HP et al. (2022) Quantum Computing and Simulations for Energy Applications: Review and Perspective. *ACS Engineering Au* 2(3): 151–196 DOI: [10.1021/acsengineeringau.1c00033](https://doi.org/10.1021/acsengineeringau.1c00033) eprint: <https://doi.org/10.1021/acsengineeringau.1c00033> URL: <https://doi.org/10.1021/acsengineeringau.1c00033>
- Peham T, Burgholzer L, Wille R (2022) “Equivalence checking paradigms in quantum circuit design: a case study” *DAC '22: 59th ACM/IEEE Design Automation Conference, San Francisco, California, USA, July 10 - 14, 2022* ed. by R Oshana ACM: 517–522 DOI: [10.1145/3489517.3530480](https://doi.org/10.1145/3489517.3530480) URL: <https://doi.org/10.1145/3489517.3530480>
- Perdomo-Ortiz A et al. (2015) A quantum annealing approach for fault detection and diagnosis of graph-based systems. *The European Physical Journal Special Topics* 224: 131–148
- Pérez-Delgado CA, Perez-Gonzalez HG (2020) “Towards a Quantum Software Modeling Language” *Proceedings of the IEEE/ACM 42nd International Conference on Software Engineering Workshops ICSEW'20* Association for Computing Machin-

- ery, Seoul, Republic of Korea: 442–444 ISBN: 9781450379632 DOI: [10.1145/3387940.3392183](https://doi.org/10.1145/3387940.3392183) URL: <https://doi.org/10.1145/3387940.3392183>
- Poggel B et al. (2023) “Recommending Solution Paths for Solving Optimization Problems with Quantum Computing” *Proceedings of the IEEE international conference on quantum software* IEEE
- Qiskit contributors (2023) *Qiskit: An Open-source Framework for Quantum Computing*. DOI: [10.5281/zenodo.2573505](https://doi.org/10.5281/zenodo.2573505)
- Quetschlich N, Burgholzer L, Wille R (2023) “Predicting good quantum circuit compilation options” *2023 IEEE International Conference on Quantum Software (QSW)* IEEE: 43–53
- Ramsauer R et al. (July 2022) “Static Hardware Partitioning on RISC-V - Shortcomings, Limitations, and Prospects” en (8th IEEE World Forum on Internet of Things (IEEE WFIoT2022)) DOI: <https://doi.org/10.48550/arXiv.2208.02703> URL: <https://arxiv.org/abs/2208.02703>
- Reich J, Schneider D, et al. (2020) “Engineering of Runtime Safety Monitors for Cyber-Physical Systems with Digital Dependability Identities” *Computer Safety, Reliability, and Security* ed. by A Casimiro et al. Springer International Publishing, Cham: 3–17 ISBN: 978-3-030-54549-9
- Reich J, Wellstein M, et al. (2021) “Towards a Software Component to Perform Situation-Aware Dynamic Risk Assessment for Autonomous Vehicles” *Dependable Computing - EDCC 2021 Workshops* ed. by R Adler et al. Springer International Publishing, Cham: 3–11 ISBN: 978-3-030-86507-8
- Resch S, Karpuzcu UR (July 2021) Benchmarking Quantum Computers and the Impact of Quantum Noise. *ACM Comput. Surv.* 54(7): ISSN: 0360-0300 DOI: [10.1145/3464420](https://doi.org/10.1145/3464420) URL: <https://doi.org/10.1145/3464420>
- Rice JE et al. (Apr. 2021) Quantum computation of dominant products in lithium–sulfur batteries. *The Journal of Chemical Physics* 154(13): 134115 ISSN: 0021-9606 DOI: [10.1063/5.0044068](https://doi.org/10.1063/5.0044068) eprint: https://pubs.aip.org/aip/jcp/article-pdf/doi/10.1063/5.0044068/15588046/134115_1_online.pdf URL: <https://doi.org/10.1063/5.0044068>
- Roffe J (2019) Quantum error correction: an introductory guide. *Contemporary Physics* 60(3): 226–245 DOI: [10.1080/00107514.2019.1667078](https://doi.org/10.1080/00107514.2019.1667078) eprint: <https://doi.org/10.1080/00107514.2019.1667078> URL: <https://doi.org/10.1080/00107514.2019.1667078>
- Safi H, Wintersperger K, Mauerer W (2023) “Influence of HW-SW-Co-Design on Quantum Computing Scalability” *Proceedings of the IEEE international conference on quantum software* IEEE
- Saurabh N, Jha S, Luckow A (2023) “A Conceptual Architecture for a Middleware for Hybrid Quantum-HPC Application Workflows” *Proceedings of the IEEE international conference on quantum software* IEEE
- Sax I et al. (2020) “Approximate Approximation on a Quantum Annealer” *Proceedings of the 17th ACM International Conference on Computing Frontiers*: 108–117 URL: <https://arxiv.org/pdf/2004.09267>
- Schmale T et al. (July 2022) “Backend compiler phases for trapped-ion quantum computers” *2022 IEEE International Conference on Quantum Software (QSW)* IEEE

- Computer Society, Los Alamitos, CA, USA: 32–37 doi: [10.1109/QSW55613.2022.00020](https://doi.org/10.1109/QSW55613.2022.00020) URL: <https://doi.ieeeecomputersociety.org/10.1109/QSW55613.2022.00020>
- Schönberger M, Franz M, et al. (2022) “Peel — Pile? Cross-Framework Portability of Quantum Software” *2022 IEEE 19th International Conference on Software Architecture Companion (ICSA-C)*: 164–169 doi: [10.1109/ICSA-C54293.2022.00039](https://doi.org/10.1109/ICSA-C54293.2022.00039)
- Schönberger M, Scherzinger S, Mauerer W (2023) “Ready to Leap (by Co-Design)? Join Order Optimisation on Quantum Hardware” *Proceedings of ACM SIGMOD/PODS International Conference on Management of Data*
- Schönberger M, Trummer I, Mauerer W (Aug. 2023) “Quantum Optimisation of General Join Trees” *Proceedings of the International Workshop on Quantum Data Science and Management QDSM '23*
- Schreiber FJ, Eisert J, Meyer JJ (2022) *Classical surrogates for quantum learning models*. arXiv: [2206.11740](https://arxiv.org/abs/2206.11740) [quant-ph]
- Schuld M, Sweke R, Meyer JJ (Mar. 2021) Effect of data encoding on the expressive power of variational quantum-machine-learning models. *Phys. Rev. A* 103 (3): 032430 doi: [10.1103/PhysRevA.103.032430](https://link.aps.org/doi/10.1103/PhysRevA.103.032430) URL: <https://link.aps.org/doi/10.1103/PhysRevA.103.032430>
- Schulz M et al. (July 2022) Accelerating HPC With Quantum Computing: It Is a Software Challenge Too. *Computing in Science & Engineering* 24(04): 60–64 ISSN: 1558-366X doi: [10.1109/MCSE.2022.3221845](https://doi.org/10.1109/MCSE.2022.3221845)
- Serrano MA, Perez-Castillo R, Piattini M, eds. (2022) *Quantum Software Engineering*. Springer ISBN: 978-3-031-05323-8 URL: <https://doi.org/10.1007/978-3-031-05324-5>
- Sitdikov I et al. (2023) “Middleware for Quantum: An orchestration of hybrid quantum-classical systems” *Proceedings of the IEEE international conference on quantum software* IEEE
- Sivarajah S et al. (Nov. 2020) t|ket): a retargetable compiler for NISQ devices. *Quantum Science and Technology* 6(1): 014003 doi: [10.1088/2058-9565/ab8e92](https://dx.doi.org/10.1088/2058-9565/ab8e92) URL: <https://dx.doi.org/10.1088/2058-9565/ab8e92>
- Sommerville I (2010) *Software Engineering*. 9th ed. Addison-Wesley, Harlow, England ISBN: 978-0-13-703515-1
- Steiner L et al. (2022) “An LPDDR4 Safety Model for Automotive Applications” *The International Symposium on Memory Systems MEMSYS 2021* Association for Computing Machinery, Washington DC, DC, USA ISBN: 9781450385701 doi: [10.1145/3488423.3519333](https://doi.org/10.1145/3488423.3519333) URL: <https://doi.org/10.1145/3488423.3519333>
- Stollenwerk T, Lobe E, Jung M (2019) “Flight gate assignment with a quantum annealer” *International Workshop on Quantum Technology and Optimization Problems* <https://elib.dlr.de/123777/> Springer Berlin: 99–110 doi: [10.1007/978-3-030-14082-3_9](https://doi.org/10.1007/978-3-030-14082-3_9)
- Stollenwerk T, O’Gorman B, et al. (2019) Quantum annealing applied to de-conflicting optimal trajectories for air traffic management. *IEEE transactions on intelligent transportation systems* 21(1): 285–297

- Streif M, Leib M (May 2020) Training the quantum approximate optimization algorithm without access to a quantum processing unit. *Quantum Science and Technology* 5(3): 034008 DOI: [10.1088/2058-9565/ab8c2b](https://doi.org/10.1088/2058-9565/ab8c2b) URL: <https://dx.doi.org/10.1088/2058-9565/ab8c2b>
- Tilly J et al. (2022) The Variational Quantum Eigensolver: A review of methods and best practices. *Physics Reports* 986 The Variational Quantum Eigensolver: a review of methods and best practices: 1–128 ISSN: 0370-1573 DOI: <https://doi.org/10.1016/j.physrep.2022.08.003> URL: <https://www.sciencedirect.com/science/article/pii/S0370157322003118>
- Trummer I, Koch C (May 2016) Multiple Query Optimization on the D-Wave 2X Adiabatic Quantum Computer. *Proc. VLDB Endow.* 9(9): 648–659 ISSN: 2150-8097 DOI: [10.14778/2947618.2947621](https://doi.org/10.14778/2947618.2947621) URL: <https://doi.org/10.14778/2947618.2947621>
- Ufrecht C et al. (2023) *Cutting multi-control quantum gates with ZX calculus*. DOI: [10.48550/ARXIV.2302.00387](https://arxiv.org/abs/2302.00387) URL: <https://arxiv.org/abs/2302.00387>
- Wang S et al. (2021) Noise-induced barren plateaus in variational quantum algorithms. *Nature Communications* 12(1): 6961 DOI: [10.1038/s41467-021-27045-6](https://doi.org/10.1038/s41467-021-27045-6) URL: <https://doi.org/10.1038/s41467-021-27045-6>
- Weder B et al. (2022) “Quantum software development lifecycle” *Quantum Software Engineering* Springer: 61–83
- Weimer H et al. (2011) Digital quantum simulation with Rydberg atoms. *Quantum Information Processing* 10(6): 885 DOI: [10.1007/s11128-011-0303-5](https://doi.org/10.1007/s11128-011-0303-5) URL: <https://doi.org/10.1007/s11128-011-0303-5>
- Wille R, Hillmich S, Burgholzer L (2022) Tools for quantum computing based on decision diagrams. *ACM Transactions on Quantum Computing* 3(3): 1–17
- Wintersperger K, Dommert F, et al. (2023) Neutral Atom Quantum Computing Hardware: Performance and End-User Perspective. en:
- Wintersperger K, Safi H, Mauerer W (Aug. 2022) “QPU-System Co-Design for Quantum HPC Accelerators” *Proceedings of the 35th GI/ITG International Conference on the Architecture of Computing Systems* Gesellschaft für Informatik
- Xia S, Zhao J (2023) *Static Entanglement Analysis of Quantum Programs*. arXiv: [2304.05049](https://arxiv.org/abs/2304.05049) [cs.SE]
- Yamaguchi M, Yoshioka N (2023) *Design by Contract Framework for Quantum Software*. arXiv: [2303.17750](https://arxiv.org/abs/2303.17750) [cs.CL]
- Yarkoni S et al. (2021) “Multi-car paint shop optimization with quantum annealing” *2021 IEEE International Conference on Quantum Computing and Engineering (QCE)* IEEE: 35–41
- Ying M (2019) Toward automatic verification of quantum programs. *Formal Aspects of Computing* 31(1): 3–25 DOI: [10.1007/s00165-018-0465-3](https://doi.org/10.1007/s00165-018-0465-3) URL: <https://doi.org/10.1007/s00165-018-0465-3>
- Yue T et al. (2023) “Challenges and Opportunities in Quantum Software Architecture” *Software Architecture Research Roadmaps from the Community* ed. by I Weber, R Kazman, P Pellicione Springer

- Zhao J (2020) Quantum Software Engineering: Landscapes and Horizons. *CoRR* abs/2007.07047: arXiv: [2007.07047](https://arxiv.org/abs/2007.07047) URL: <https://arxiv.org/abs/2007.07047>
- Zhao P, Wu X, Li Z, et al. (2023) *QChecker: Detecting Bugs in Quantum Programs via Static Analysis*. arXiv: [2304.04387](https://arxiv.org/abs/2304.04387) [cs.SE]
- Zhao P, Wu X, Luo J, et al. (2023) “An Empirical Study of Bugs in Quantum Machine Learning Frameworks” *Proceedings of the IEEE international conference on quantum software* IEEE
- Zhou L et al. (June 2020) Quantum Approximate Optimization Algorithm: Performance, Mechanism, and Implementation on Near-Term Devices. *Phys. Rev. X* 10 (2): 021067 DOI: [10.1103/PhysRevX.10.021067](https://doi.org/10.1103/PhysRevX.10.021067) URL: <https://link.aps.org/doi/10.1103/PhysRevX.10.021067>