

Received January 31, 2020; revised February 28, 2020; accepted February 28, 2020; date of publication April 22, 2020; date of current version June 3, 2020.

Digital Object Identifier 10.1109/TQE.2020.2981074

Quantum Computer Architecture Toward Full-Stack Quantum Accelerators

KOEN BERTELS¹, A. SARKAR, T. HUBREGTSEN, M. SERRAO,
A.A. MOUEDENNE, A. YADAV, A. KROL, I. ASHRAF,
AND C. GARCIA ALMUDEVER

Quantum Computer Architecture Lab, Delft University of Technology, 2628 CD Delft, The Netherlands

Corresponding author: Koen Bertels (k.l.m.bertels@tudelft.nl)

ABSTRACT This article presents the definition and implementation of a quantum computer architecture to enable creating a new computational device—a quantum computer as an accelerator. A key question addressed is what such a quantum computer is and how it relates to the classical processor that controls the entire execution process. In this article, we present explicitly the idea of a quantum accelerator that contains the full stack of the layers of an accelerator. Such a stack starts at the highest level describing the target application of the accelerator. The next layer abstracts the quantum logic outlining the algorithm that is to be executed on the quantum accelerator. In our case, the logic is expressed in the universal quantum-classical hybrid computation language developed in the group, called OpenQL, which visualized the quantum processor as a computational accelerator. The OpenQL compiler translates the program to a common assembly language, called cQASM, which can be executed on a quantum simulator. The cQASM represents the instruction set that can be executed by the microarchitecture implemented in the quantum accelerator. In a subsequent step, the compiler can convert the cQASM to generate the eQASM, which is executable on a particular experimental device incorporating the platform-specific parameters. This way, we are able to distinguish clearly the experimental research toward better qubits, and the industrial and societal applications that need to be developed and executed on a quantum device. The first case offers experimental physicists with a full-stack experimental platform using realistic qubits with decoherence and error-rates, whereas the second case offers perfect qubits to the quantum application developer, where there is neither decoherence nor error-rates. We conclude the article by explicitly presenting three examples of full-stack quantum accelerators, for an experimental superconducting processor, for quantum accelerated genome sequencing and for near-term generic optimization problems based on quantum heuristic approaches. The two later full-stack models are currently being actively researched in our group.

INDEX TERMS Quantum computing, parallel architectures, parallel programming, quantum entanglement.

I. INTRODUCTION

The history of computer architecture dates back various decades and has been very evolving. An important extension is the emergence of accelerators [1] as specialized processing units to which the host processor offloads suitable computational tasks. Recently, computer architecture research is getting more focused on quantum computing. In the next 5–10 years of quantum computer development, it does not make sense to talk about quantum computing in the sense of a universal, Turing computer that can be applied in any kind of application domain. Given the recent insights leading to, e.g., noisy intermediate-scale quantum (NISQ) technology as expressed in [2], we are much more inclined to believe that the

first industry-based and societal relevant application will be a hybrid combination of a classical computer and a quantum accelerator. It is based on the idea that any end-application contains multiple computational kernels and the properties of these parts are better executed by a particular accelerator, which can be, as shown in Fig. 1, either field-programmable gate arrays (FPGA), graphics processing units (GPU), neural processing units such as Google's tensor processing unit, etc. The formal definition of an accelerator is indeed a coprocessor linked to the central processor that is capable of accelerating the execution of specific computational intensive kernels, as to speed up the overall execution according to Amdahl's law. We now add two classes of quantum accelerator as

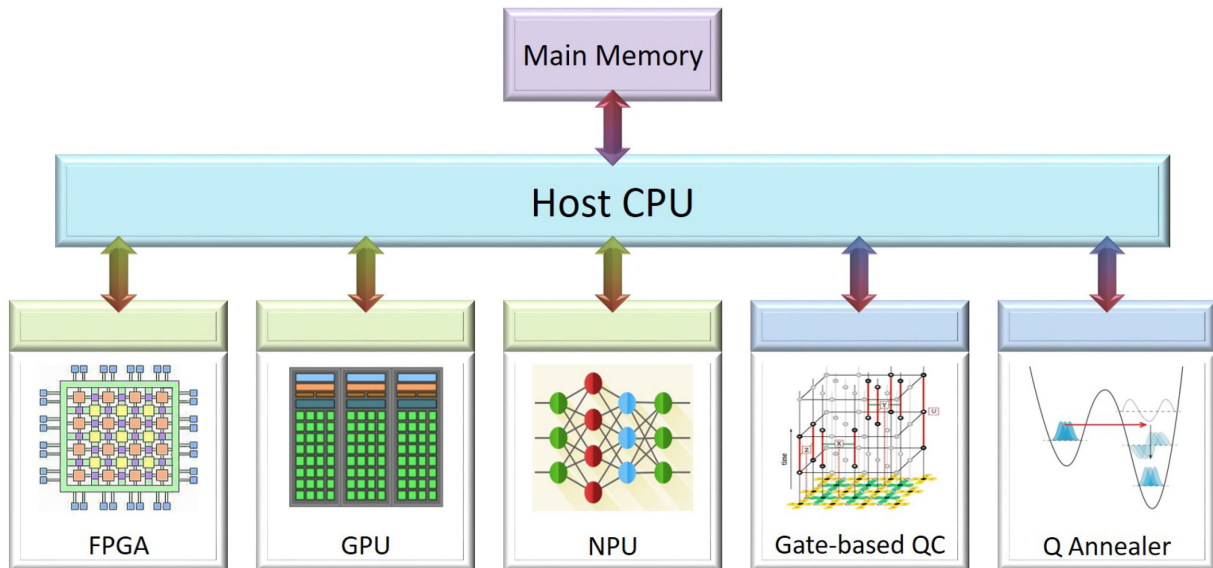


FIGURE 1. System architecture with heterogeneous accelerators.

additional coprocessors. The first one is based on quantum gates and the second is based on quantum annealing. The classical host processor keeps the control over the total system and delegates the execution of certain parts to the available accelerators.

Computer architectures have evolved quite dramatically over the last couple of decades. The first computers that were built did not have a clear separation between compute logic and memory. It was only with von Neumann's idea to separate and develop these distinctly that the famous von Neumann architecture was born. This architecture had for a long time a single processor and was driven forward by the ever increasing number of transistors on the chip, which doubled every 18 months. In the beginning of the 21st century, the single cores became too complex and did not provide any substantial processing improvement. This led to the incorporation of multiple cores. The homogeneous multicore processor dominated the processor development for a couple of years but companies such as IBM and Intel started understanding that heterogeneity is the right way forward to improve the compute power. GPUs and FPGAs are seen as natural extensions of the computer architecture, implying that the quantum accelerator would be a logical next step.

In the quantum computing world, there exist two important challenges. The *first* is to have enough numbers of good quality qubits in the experimental quantum processor. The current competing qubit technologies include ion traps, majoranas, semiconducting and superconducting qubits, nitrogen-vacancy-centers, and even graphene. Improving the overall status of the qubits is challenging as these suffer from decoherence that introduces errors when performing quantum gate operation. It is only when the quantum physical community overcomes those challenges that the quantum

accelerator will be a widespread adopted solution. This direction is shown in the left picture of Fig. 2 where different quantum technologies are depicted in the lowest layer. The *second* challenge is to formulate, at a high level, the quantum logic that companies and other organizations need to be able to use high-performance accelerators for certain computations that can only run on the quantum device. This requires a long-term investment in terms of people and technical know-how from companies that want to pursue this direction and reap the benefits. The right part of Fig. 2 shows the industrial commitment to think about the required quantum logic that can be executed using the full stack, evaluated and tested on a quantum simulator. It is important to emphasize that the qubits are called perfect qubits that do not decohere or have any other kind of errors generated by them. With the emergence of huge amounts of data, commonly called big data, it is understood that this paradigm is not scalable to superlarge datasets. The key factor is the huge amount of data that needs to be processed by multiple computing cores, which is a very tough problem to solve. The data communication between the cores is a very difficult programming problem and the data management problem is substantially slowing down the overall performance.

Based on our group's research since 2004 [1] and as shown in Fig. 2, an important concept that we have been implementing in the quantum computing world is the implementation of a full stack for a quantum accelerator as will be described later in this article. The basic philosophy of any accelerator is that a full stack needs to be defined and implemented. The past 10–15 years have shown a large number of accelerators that were developed as part of any modern computer architecture. It always consists of the same following layers: it starts at the highest level describing the logic that needs to be mapped on the accelerator. Examples are video processing,

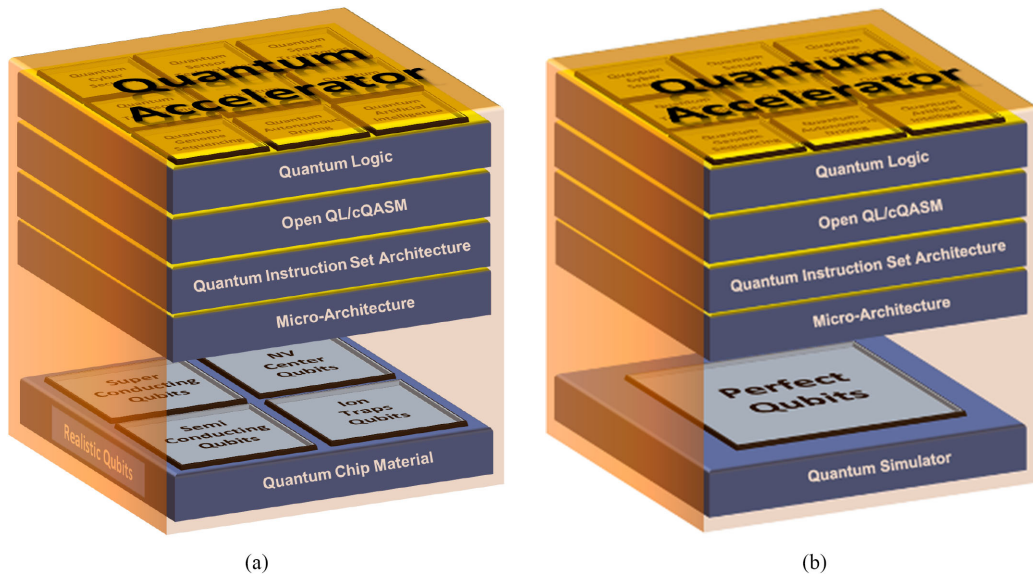


FIGURE 2. Two approaches for full-stack quantum accelerators. (a) Experimental full-stack with realistic qubits. (b) Simulated full-stack with perfect qubits.

security, matrix computation, etc. These application-specific algorithms can be defined in various languages such as C++ or Fortran. In the case of FPGAs, these algorithms are translated into VHDL or Verilog. In the case of GPUs, the language is often formulated using mathematics or other libraries and translated by the compiler to an assembly language that can be mapped on the GPU architecture. Especially in the case of FPGAs, there is no standard microarchitecture on which the VHDL or Verilog can be executed. Such an architecture needs to be developed for every application that needs to be accelerated. The final layer is a chip-based implementation of the microarchitecture combined with the hardware accelerator blocks that are needed.

A. BACKGROUND

One of the first proposals on quantum computing was written by Feynman in 1982 [3], which launched a world-wide research on quantum computing focusing on important low-level challenges leading to the development of superconducting qubits, ion trap qubits, or spin-qubits. He formulated the use of quantum computers as an important scientific instrument to allow us to understand the quantum phenomena that quantum physics tries to understand. The design of proof-of-concept quantum algorithms and their analysis with respect to their theoretical complexity improvements over classical algorithms has also received some attention. However, we still need substantial progress in either of those domains. Qubits with a sufficiently long coherence time combined with a true quantum killer application are still crucial achievements on which the community is working. These are vital to demonstrate the exponential performance increase of quantum over conventional computers *in practice* and are urgently needed to convince quantum skeptics

about the usefulness of quantum computing such that it can become a mainstream technology within the coming 10–15 years. However, as we will describe in this article, we need much more before any kind of computational device can be developed, which ultimately connects the algorithmic level with the physical chip. What is needed involves a compiler, run-time support, and more importantly a microarchitecture that executes a well-defined set of quantum instructions.

An interesting and quite high-level kind of description was published in 2013 [4]. Meter and Horsman describe their understanding of the blueprint of a quantum computer. They correctly emphasized the need to look at computer engineering to better understand what the similarities and differences are between quantum and classical computing. As mentioned before, the most important difference is the substantially higher error rate that qubits and quantum gates (10^{-3}) have compared to CMOS technology (10^{-15}). Guaranteeing fault-tolerant (FT) computation can easily consume more than 90% of the actual computational activity. The second difference focuses on the nearest-neighbor (NN) constraint, which imposes that two-qubit gates can only be applied if the qubits reside next to each other. The no-cloning theorem prohibits copying quantum states. The way that two-qubit gates are applied requires the two qubits to be sufficiently close to each other. They also describe a hierarchical layered structure but rather than defining these layers in terms of more computer engineering concepts, the schema is more expressed in terms of the different, relevant fields and research domains. Examples are quantum error correction (QEC) theory, programming languages, FT implementation, etc. There are also other mechanisms with undefined time costs that are necessary to make FT-quantum computing (hopefully) efficient and performing. Examples are state distillation for

ancilla factories and the emergence of a wide variety of defects and errors, which all impose an additional burden on the microarchitecture and the corresponding run-time management.

An older but conceptually quite similar paper was published by DiVincenzo in 2000 [5]. This article outlines the following five criteria needed to build a quantum computer:

- 1) a scalable physical system with well-characterized qubits;
- 2) the ability to initialize the state of the qubits to simple fiducial state;
- 3) long relevant coherence times;
- 4) a universal set of quantum gates;
- 5) a qubit specific measurement capability.

Two additional criteria needed for quantum communication are the ability to interconvert stationary and flying qubits and the ability to transmit flying qubits between specified locations. Considering currently available quantum processors, we could say that they already comply to DiVincenzo’s criteria, and thus, we already have a quantum computer. However, an important and missing criterion is the number of qubits that we need for any kind of reasonable application. Depending on the application domain, the estimates of the number of qubits goes from relatively low, such as a couple of hundreds, to several billions. Being less critical, we could say that the first criterion explicitly formulates the size of the system, which is still a very considerable challenge to compute in a reliable way.

The rest of the article is structured as follows. First, we describe the various layers such as application, algorithmic logic, programming language and OpenQL compiler, microarchitecture, mapping, and simulator. Each layer is positioned with respect to real, realistic, and perfect qubits. Three examples of quantum accelerators are presented next. Finally, our vision on various aspects of quantum computing is discussed.

II. QUANTUM FULL-STACK

In the context of quantum accelerator development, the same full-stack approach is adopted for either perfect or realistic qubits. The execution can be either on an experimental quantum chip or on the QX simulator. The highest level starts at the end-user application for which a part of that application is developed in a quantum language, such as OpenQL. The quantum part of any industrial or societal application can be executed on any kind of available quantum prototype. For any quantum logic that is specified, a specific and target-related microarchitecture needs to be defined and used. We present the considerations for the various layers in this section. Besides gate-based quantum computing approach, we also include the quantum annealer based system/simulator in Fig. 3, as we currently investigate the components of all types of architectures currently in the market. We first introduce here the different kinds of qubit models that we support at this state of research in the quantum computer engineering

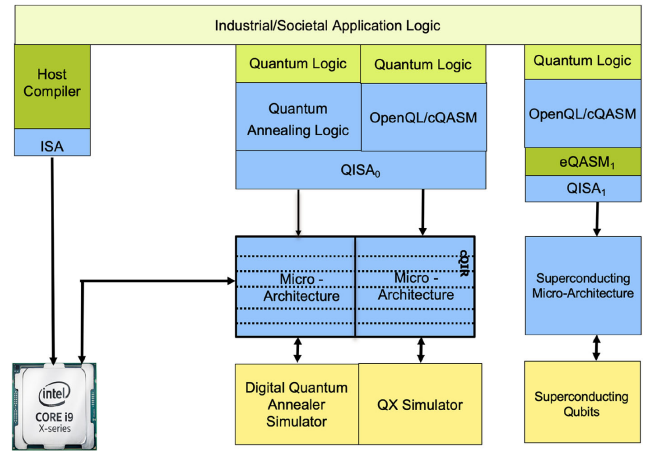


FIGURE 3. Full-stack execution.

field. The real, realistic, and perfect qubits are presented here, which can be used for either purely experimental or purely application development perspective.

A. REAL, REALISTIC, AND PERFECT QUBITS

An important concept that is introduced for our line of research is the use of three kinds of qubits, namely real, realistic, and perfect qubits. In this section, we define them in detail and how these relate to each other.

Real qubits: The first qubit type is the experimental qubit, called the real qubit, which refers to experimentally realized systems with challenges, such as decoherence and error-rates. These features need to be substantially improved for any commercially available quantum device. The real qubits are investigated by the experimental quantum physicists community. The goal is to improve the quality of the real qubits such that these become more easy to scale to large numbers and allow for a pragmatic microarchitectural control. This implies that there is a need to study how long the qubits can stay in a particular state and maintain their fidelity, called the coherence time. Most of the real qubits go to the ground state in a very short time (ranging from micro- to milliseconds) after these are created in a particular state. Adding to that, all the quantum gates that need to be applied to the qubits generate errors. In quantum gate operations, the errors and the error-rates need to be better than the current 10^{-2} rates.¹ There are currently many quantum technologies experimenting to produce good quality qubits for reasonable quantum computation. The use of real qubits is very important as the physicists need to understand the dynamic and static behavior of the qubits under different circumstances. Many large companies implement physical system for quantum computing, such as IBM, Google, Rigetti, D-Wave Systems, IonQ, etc. However, the quality as well as the number of these qubits is very limited and the decoherence and

¹We will limit ourselves now to quantum gates but will introduce later the quantum annealing approach.

error-rates as mentioned before are currently problematic for application development as these tend to influence the overall result that the quantum device is computing.

Realistic qubits: Realistic qubits represent the third dimension in Fig. 2 and any computer architecture needs functionality to continuously monitor the quantum system to detect and recover possible errors, as we describe here. For quite a long period, the focus has been mostly on planar surface codes (SC) as it was considered one of the most promising QEC codes for short-term implementations and for scalability concerns in the FT era and manufacturing. Qubits are generally manufactured in a regular 2-D lattice connectivity with only NN interactions. The array comprises two kinds of qubits, namely the data and ancilla qubits. Data qubits are used to store the quantum information for the computation, whereas ancilla qubits are helper qubits that are used to detect bit-flip and phase-flip errors by performing error syndrome measurements. This implies that after every sequence of quantum gates, the system needs to measure out its state and interpret those measurements to see if an error has been produced. Given the constraints of the coherent qubit lifetime, it implies that a very large graph needs to be processed and interpreted in real time such that any error can be identified. Measurements themselves can be erroneous and, therefore, need to be repeated multiple times before a final conclusion is reached. In 2018, Preskill [2] introduced a counterargument to this approach because SC requires too many ancilla-qubits for logical protection. This led to the reinitiation of the small-codes, which were first defined almost 20 years ago. The impact on the system architectural and compiler level is yet unclear but this is currently the focus of a lot of research.

Perfect qubits: Companies, governments, and other organizations interested in building a quantum accelerator need to evaluate the availability of quantum computing resources in terms of quantum algorithms and have a way to test the correctness of the quantum logic. To serve these needs, we use perfect qubit, such that any of the erroneous behavior arising due to qubit quality can be avoided during application development phase. These qubit modeled in the simulator do not decohere and stay in ideal state required for the algorithm. Using these perfect qubits guarantees that the end-users can verify and check the algorithm that they are working on and test if the computed results have a meaning that can be easily interpreted. We are not the only ones who use this but it is a very clear concept that separates the two directions that we are investigating in the Quantum Computer Architecture Lab. As explained earlier, we introduce in OpenQL, a datatype that represents the perfect qubit that has a more stable behavior than the realistic qubits. Whether or not the NN constraint applies is a discretion of the designer. The compiler may or may not compute a route for the qubits. These decisions are based on the requirement and maturity of the application development stage before translating to realistic experimental testing.

B. INDUSTRIAL AND SOCIETAL QUANTUM APPLICATION LOGIC

The highest layer in the full stack focuses on the application that needs to be developed for any organization. On current, modern architectures, there are a large number of initiatives developed that run on either the FPGA, the GPU, or the tensor processing unit as the accelerator platform. When envisioning the quantum accelerator idea, many similar topics are well suited, such as security, artificial intelligence, autonomous driving, genome sequencing, sensors, and trajectories for aeroplanes and rockets. For the three application examples that we are currently developing, we assume the use of perfect qubits such that the focus can be completely given to the algorithm logic and the microarchitecture design.

- 1) We research algorithms for accelerating quantum genome sequencing (QGS). These are motivated by the application of gene therapy and personalized medication for every single individual on earth. The treatment will be based on every person's DNA profile that has to be generated by extensive computational processing of the reads from sequencing devices.
- 2) The other example that we will discuss in this article is for optimization problems pervasive in operations research based on the traveling salesman problem (TSP). It is expressed as a quadratic unconstrained binary optimization (QUBO) problem and can be solved both on the gate-based model or the annealing model.
- 3) We are also working on a quantum accelerator model in collaboration with our research partners in the automotive industry focusing on autonomous and electrical cars. For confidentiality agreements, we do not go into any detail of this project.

Given the potential of quantum acceleration, this top-down approach is necessary to understand how investing in the development of quantum computing has the potential to become a world-wide technology that can be used by every country, organization, or individual. In Section III, we will present the three accelerators in more detail.

C. QUANTUM LOGIC

For this section, we always consider perfect qubits. The highest level is the application layer where a potential end-user of the quantum compute power instructs what exactly needs to be computed. Quantum computing promises to become a computational game changer, allowing the calculation of various algorithms much faster (in some cases exponentially faster) than their classical counterparts. Especially, applications requiring manipulation of a large set of data items to produce a statistical answer are very suitable to be processed by quantum computers, which we call in this article quantum accelerators. Currently, there is no generally acknowledged or accepted functional domain where quantum technology would be the game changer. Potential promising domains include physical system simulation, cryptography,

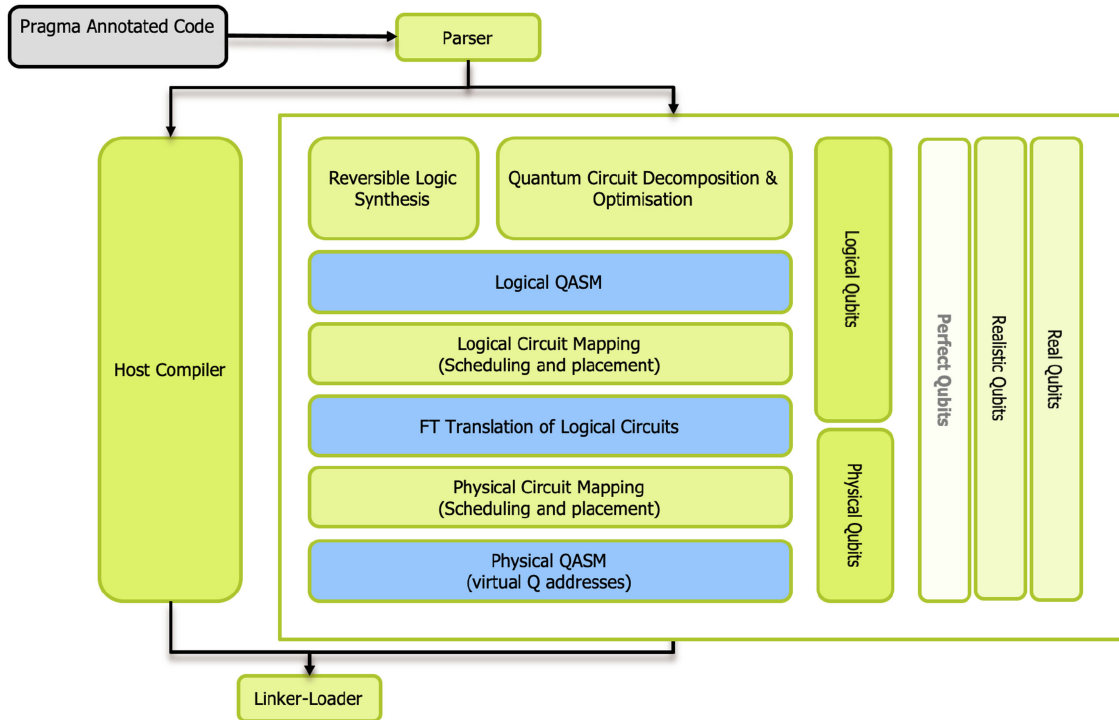


FIGURE 4. Compiler infrastructure.

and machine learning. Evidently, the cryptography domain is a clear candidate as algorithms such as Shor’s factorization showed that potentially a quantum computer can break any RSA-based encryption, as it leads to finding the prime factors of the public key [6] based on which the private key can be easily calculated. However, the cryptography domain has actively establishing a new research theme, namely the postquantum cryptography such that the attacks emerging from such a compute power can be rebuffed.

Another potential application area is the biological domain where chemistry, medication, and pharmacology belong to. We focus on one such candidate application of genome sequence reconstruction. For instance, quantum computational power would be imperative if we aim want to compute the DNA-profile of every human being in the world, which takes around one week on a large network of very powerful servers for one person’s DNA. With the availability of enough qubit capacity, the entire parallel input dataset can be evolved simultaneously as a superposition of a wave function.² This particular property makes it possible to perform the computation of the entire dataset in parallel. This kind of computational acceleration provides a promising approach to address the computational challenges of DNA analysis algorithms. The essence of accelerating sequence reconstruction is the ability to run parallel search operations on the short reads obtained from sequencing an individual DNA from a sequencing machine onto an already available reference of the organism. In recent years, GPU, FPGA, and cluster computing frameworks such as Hadoop and Spark have been used to reduce the total run-time. Potentially, quantum computation

offers a fundamentally different way to address the enormous volume of data by employing superposition of reads in the search process, thereby reducing the memory requirement maybe even exponentially. The quantum search primitive (Grover’s search) itself is provably optimal [7] over any other classical or quantum unstructured search algorithm. The rather modest quadratic speedup in cycles, however, becomes extremely relevant for industrial application due to the total CPU run-time involved in the big data manipulation (in order of 1000 s of CPU hours [8] for a single human DNA sequence reconstruction).

D. PROGRAMMING LANGUAGE, COMPILER, AND RUN-TIME SUPPORT

The quantum algorithms and applications presented in the previous section can be described using a high-level programming language, such as Q# [9], Scaffold [10], QisKit [11], Quipper [12], or OpenQL [13], and compiled into a series of instructions that belong to the (quantum) instruction set architecture. Many of other languages that are available in the world are mostly developed for supporting the quantum physical experiments, such as Forest [14], CirQ [15], Strawberry Fields [16], XACC [17], and Ocean Software [18].

Consistent with our distinction between perfect, realistic, and real qubits, the compiler is capable of adapting to the requirements of the end-user. So, there is an option that

²By our estimate, given the size of the human genome and currently available sequencers, the number of qubits required will be around 150 logical qubits.

translates the qubits in perfect, realistic, or real manner. As shown in Fig. 4, the compiler infrastructure for such a heterogeneous system consists of the classical compiler for the host processor combined with the quantum compiler. It is important to note that the architectural heterogeneity where classical processors are combined with different accelerators, such as the quantum accelerator, imposes a specific compiler structure where each compiler part can target the different instruction sets and ultimately generates one binary file that can be executed on different instruction set architectures. For the computer architecture envisioned in our research, any high-level implementation of the system application will consist of two interleaved types of logic: the classical logic that will be executed by the microarchitecture of the controlling processor and the quantum logic that will be mapped onto the quantum processor. The quantum logic can be encapsulated by classical language structures, such as decision and loop constructs. The microarchitecture extracts the quantum part and send it to the quantum processor.

As we adopt the quantum circuit model as a computational model, the quantum compiler translates the quantum logic into quantum circuits for which reversible circuit design, quantum gate decomposition, and circuit mapping are needed. The output of this compiler is a series of instructions, expressed in a quantum assembly language, such as cQASM, which belongs to the defined instruction set architecture.³ The definition of a shared quantum assembly language is a key challenge such that there is uniformity in the algorithmic descriptions of different research groups.

- 1) *Real qubits*: The OpenQL compiler can generate code that physicists can use for testing the behavior of the qubits, taking all kinds of errors and decoherence into account. An important exercise is to examine the FT of the quantum circuits. A central issue for any quantum technology is its fragility, implying that the qubit superposition state disappears quite rapidly. First, the coherence time of real qubits is extremely short. For example, superconducting qubits may lose their information in tens of microseconds [19], [20]. Second, quantum operations are unreliable with error rates around 0.1% [21]. As mentioned earlier, in January 2018, Preskill [2] emphasizes that early stage quantum computers should be based on NISQ technology with much less ancilla qubits for QEC activities. QEC is more challenging than classical error correction, due to the no-cloning theorem, which states that (unknown) quantum states cannot be copied. This makes the classical way of creating several copies of the same bit impossible. In addition, quantum errors are continuous and any measurement will destroy the information stored in qubits. The basic idea of QEC techniques is to use several physical imperfect qubits to compose more

reliable units called *logical qubits* based on a specific QEC code [22]–[28]. This is what scientists looking at physical implementations of qubits have been doing such that it is relatively simple to generate and test a super- or semiconducting qubit.

- 2) *Realistic qubits*: Similar to real qubits, it is also possible to simulate the behavior of realistic qubits such that we have a better understanding of the impact of realistic error models, better error-rates, and longer coherence times on the overall quantum circuit performance, the microarchitecture needed to control them, etc. Therefore, there is an option to compile for realistic qubits such that the duration of a quantum gate operation is shorter or less error-prone. It can also lead to better investigation of the qubit plane topological constraint and the associated routing algorithm required for multi-qubit gate operations.
- 3) *Perfect qubits*: The compiler can also target the use of perfect qubits. As defined earlier, which implies that these qubits live as long as they are needed and have principally no error-rates in the quantum gates that are executed. Depending on the state of the execution platform, connectivity constraints can be imposed for mapping and routing. When we generate everything in terms of perfect qubits, that also implies that there is no separation anymore between logical and physical qubits as there is no requirement for error coding.

E. QUANTUM MICROARCHITECTURE

Any computer has a series of instructions that can be executed on the dominant processor. To this purpose, any kind of processor has a particular architecture capable of executing any sequence of the legitimate instructions. This also holds for the quantum processor, which also has a series of instructions that it can execute, some of which are classical logic and others are the quantum instructions that will be executed on the quantum chip. So the quantum accelerator will consist of two components: the classical and digital microarchitecture part that has a classical processor to execute part of the accelerator logic and the quantum chip that contains the qubits that need to be executed in an analog way.

Essential to any kind of computational device is the presence of one or multiple computer architectures that are responsible for executing the instructions that are delegated to the coprocessor. The architecture of a machine connects the physical hardware to the applications that can run (on that hardware) and dictates how instructions are executed. This is also true for the case of a quantum accelerator. For the quantum algorithms to be understood by the quantum accelerator, a low-level representation of the quantum instructions is required that the classical control hardware of the quantum chip can understand. This is known as the quantum instruction set architecture (QISA). The content of the QISA can be modified for each accelerator logic that needs to be implemented. One example of a microarchitecture is given in Fig. 5. For any microarchitecture, there are a number of properties that we

³QASM is one candidate for such a language and was originally produced by Nielsen and Chuang to generate the LaTeX figures for the quantum circuits for their book.

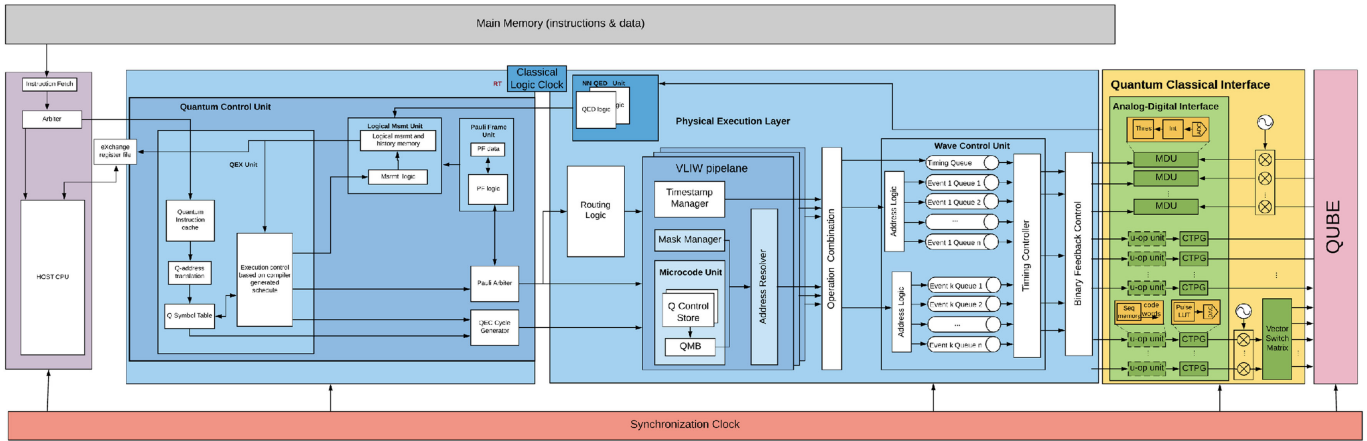


FIGURE 5. Example of a general-purpose microarchitecture.

have to estimate, such as the appropriate instruction-length, pipeline depth (for parallel quantum gates), and targeting multiple control channels per single instruction. Based on these principles, the basic blocks are constructed, such as timing control unit and the microcode instruction set of the overall microarchitecture.

- 1) *Real and realistic qubits*: To accommodate quantum processor development, we look at the experimental algorithms that the physics community are interested in, such as randomized (single and double) qubit gates. This phase would also comprise of hardware assessment and characterization to meet the timing-precision and signal synchronization requirements for a specific qubit-technology. In a later phase, the experimental implementation will need to include error-correcting codes in the pipeline. A system-on-chip running a quantum error-decoder would enable faster development and debugging capabilities for QEC on hardware. Area utilization and power consumption of such a firmware would become a necessary consideration at this point, depending on the size of decoders. The development and testing of this platform would be done on both the QX simulator and the physical quantum processing unit. Mapping of the quantum circuit also needs to be addressed as part of the compilation process.
- 2) *Perfect qubits*: We do not yet have a full implementation of the microarchitecture for logic expressed in terms of the perfect qubits. Later in this article, we present a tentative microarchitecture for QGS, which is one of the accelerators that we are working on. It is important to define the QISA needed for QGS and fine-tune the corresponding microarchitectural blocks needed to execute the quantum instructions on the QX simulator.

F. MAPPING OF QUANTUM CIRCUITS

Mapping of quantum circuits is considered in two different contexts: the first is when applied on small real quantum

processors and the second one targets a simulation engine that addresses larger number of qubits. Depending on the test objective, we can either take into account large number of qubits or stay at a small scale and closer to the experimental state of the art.

- 1) *Real qubits*: When targeting a real quantum processor, the mapping of circuits is an important topic as described in [29] and [30]. The circuit description of the algorithms does not usually consider a physical location of the qubits and assumes that any kind of interaction between qubits is possible. However, real qubits need to be placed on a specific physical qubit layout that will limit the possible interactions between these, leading to an increase of the circuit latency. It is, therefore, important to optimize the mapping process that includes the following.
 - a) *Scheduling of operations*: The parallelism of current quantum algorithms is pretty limited but applying classical scheduling methods and techniques, the inherent parallelism of the logical qubits can be exploited. Depending on the chosen QEC, different constraints apply to the scheduling problem. For instance, in defect-based SC, single-control multitarget CNOT gates are possible, whereas planar-based surface only supports single-control single-target CNOT gates. Furthermore, other limitations such as the number of available frequencies to control the qubits can also affect the scheduling process and restrict the parallelism.
 - b) *Placement and routing of qubits*: As mentioned before, most of the current quantum technologies are pursuing a 2-D array of qubits with only NN-interactions. This means that 2-qubit (physical) operations are only possible between adjacent qubits. It also impacts the placement of logical qubits. For instance, a CNOT between two planar-based SC qubits can theoretically be performed transversally, i.e., applying pairwise

CNOT gates to each pair of data qubits in the sublattices. However, it is not possible to implement such a transversal gate in a 2-D array requiring techniques such as lattice surgery [31] where planar-based SC qubits still need to be placed next to each other. Finally, not all qubits can be placed in the necessary adjacent positions. Therefore, some of them will have to be moved or routed for which the compiler will insert a MOVE-operation for the run-time routing logic.

- 2) *Realistic qubits*: This is the quantum processor where the qubits are not yet fully realized in any experimentally designed qubit processor. Realistic qubits imply that we are focusing on experimental processors but have modified some parameters in the overall design to understand the impact, for instance, a different topology, a different error distribution, the number of qubits, etc.
 - a) *Scheduling of operations*: Assuming that we also have parallelism between the qubits when executing a quantum circuit, we have to understand what the scheduling is of qubits and, for instance, how CNOT-gates need to be implemented to have a successful execution of the quantum gates. Comparison between the behavior of defect-based SC qubits single-control multitarget CNOT gates and planar-based surface that only supports single-control single-target CNOT gates needs to be investigated.
 - b) *Placement and routing of qubits*: Even with realistic qubits, we still have the challenge to take the NN-interaction constraints into account. Preskill's paper and talk against Surface Code and in favor of NISQ brings to our awareness the limitations on the experimental physicists. Small codes maybe be more relevant in this regime. Similar to real qubits, it also impacts the placement of logical qubits. It is important to understand if we have similar constraints as the real qubits when we are using the realistic qubit paradigm. We also need to understand if we need a similar MOVE-instruction to put the qubits close to each other.
- 3) *Perfect qubits*: When the algorithmic behavior and content is not yet defined, which is the case in most of the situations, it is important to be able to use perfect qubits that are more reliable and predictable than the experimental ones, as that have no decoherence and execute reliably the quantum gates of the quantum circuit.
 - a) *Scheduling of operations*: With perfect qubits, we have the freedom to impose or relax similar kind of restrictive scheduling instructions on their behavior.
 - b) *Placement and routing of qubits*: Also for this feature, it depends on how much freedom the

algorithm designers needs to experiment with the algorithm they are designing. The more restrictive we are in the placement and routing, the more difficult it becomes. In a more relaxed situation, the designer enjoys more possibilities to experiment and test the algorithm.

G. QX SIMULATOR

The QX simulator, as shown in Fig. 3, was developed in our group as a platform to simulate quantum operations on either realistic or perfect qubits. The QX engine can execute any quantum logic expressed in OpenQL and translated by the compiler to cQASM, the common quantum assembly language. The assumed microarchitectural layer encapsulating the QX simulator executes the cQASM instructions by sending the quantum instruction to QX, which then executes it, measures the qubit states, and sends back the results to the microarchitecture. The QX simulator is scalable based on the underlying host processor, and is capable of simulating with up to 35 fully entangled qubits on a laptop PC, which are either perfect or realistic. The main advantage of a platform like QX is to provide application developers, computer scientists, and computer engineers the tools to model and test designs before experimental implementation on quantum processors. A order of 50 fully entangled qubits already give a lot of possibilities to test the application in a proof-of-concept simulation. We can also use the different kinds of qubits that we presented in this article.

- 1) *Realistic qubits*: Whenever we are interested in running quantum circuits on real hardware, we need to be able to introduce error models for the qubit or gate operations, at the simulation level of realistic qubits. Current quantum error rates do not go beyond 10^{-2} , so there is a need to understand the impact of error rates in the order of $10^{-5}/10^{-6}$. The errors will have an effect on the real qubits as well as the quantum gates. Using the QX simulator on such realistic qubits, we can investigate beyond simplistic error models such as the depolarizing model (where every quantum gate is followed by some error, drawn from a uniform distribution of the different errors than can follow the Pauli gates X, Y, or Z). It can be extended to other error distributions, which are more realistic sketching the extensions that the quantum physics research community need to address.
- 2) *Perfect qubits*: For application development, there is the need to execute the quantum logic to verify the computed results of the algorithm in the functional sense. The QX simulator is capable of assuming the nonemergence of errors. The current stage of research on QGS algorithm uses the QX simulator in this mode of development. In principle, any universal quantum logic can be executed on the simulator, the result can be measured and fed back to the microarchitecture.

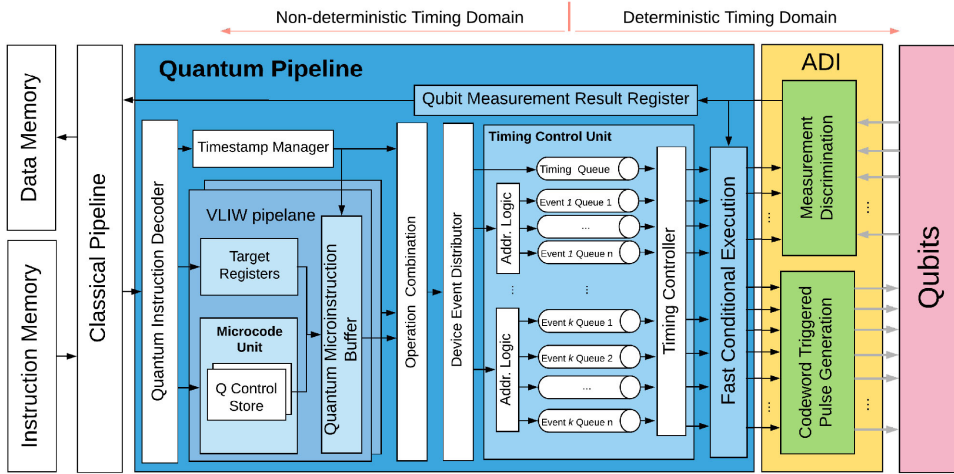


FIGURE 6. Experimental implementation of the microarchitecture for superconducting (real) qubits.

III. THREE FULL-STACK ARCHITECTURE EXAMPLES

In this section, we present and briefly describe three implementations of the full stack. The first was developed for the experimental design of superconducting qubits, the second is being implemented for accelerating genome sequencing on quantum logic, whereas the third application involves optimization problems. The full stack, as shown in Fig. 2, is used as the basic structure.

A. FULL STACK FOR REAL, SUPER/SEMICONDUCTING QUBITS

Here, we present the developed microarchitecture for the superconducting quantum chip based on an experimental implementation of all the components that were defined and needed for the quantum research collaborations in our department. The end-to-end pipeline involves writing an algorithm, up to sending the analog pulses to the qubits. It starts with a high-level quantum algorithm, which is useful for the physicists. We have been focusing on randomized benchmarking experiments for one or two qubits, which was written in OpenQL.

The code is translated by the OpenQL compiler into our version of the quantum assembly language, cQASM. As a logical extension of cQASM, the compiler then translates that version to an executable QASM, called eQASM, which supports, in principle for any quantum technology, taking low-level information into account, such as gate times, topology, etc. It basically means that there is a second back-end compiler pass that translates cQASM into the eQASM version.

Based on the cQASM code, the compiler generates the eQASM instructions that can be executed by the microarchitecture, as shown in Fig. 6 [32]. The eQASM is then executed and at run-time translated into the horizontal microcode version, which ultimately sends the micro-operations to the queues.⁴ From that level on, the timing execution

requirements are very strict and need to be precise up to the nanosecond level. The code words that are generated by the microcode unit will ultimately be translated in an analog pulse and sent to the qubit chip.

This microarchitectural demonstration was done for two quite different quantum technologies: one for the superconducting qubit chip and one for the semiconducting quantum chip. The specific combination of the microarchitecture design parameters, the c/eQASM compiler passes and the microcode unit proved very useful. Especially, the last two options allowed us to retarget the same microarchitecture to two different quantum technologies and the only changes that were needed are the configuration file for the compiler and the implementation of the microcode unit needed for the specific quantum technology to make sure the analog pulses, stored in the analog–digital interface (ADI), were available.

B. FULL STACK FOR QGS ON PERFECT QUBITS

Genome sequencing involves taking fragments of the DNA (called short reads) from the sequencing machines and stitching them together to reconstruct the original genome of the individual. Reconstruction can either be carried out by aligning these reads to an already available reference genome, or in a *de novo* assembly manner. This requires the algorithmic primitive of searching an unstructured database or graph-based combinatorial optimization. Translating such quantum kernels to an efficient implementation on a quantum accelerator requires in-depth tuning of both an architecture-aware quantum algorithm and the underlying microarchitecture.

We have obtained initial results from combining domain-specific modification on the Grover’s search [33] and quantum associative memory [34] approaches. This new alignment algorithm, described and analyzed in [35], has been tested on the QX simulator platform. The reference DNA is sliced and stored as indexed entries in a superposed quantum database giving exponential increase in capacity. The designed algorithm [36] considers inherent read errors in

⁴Described in a recently submitted paper by Fu et al. [32].

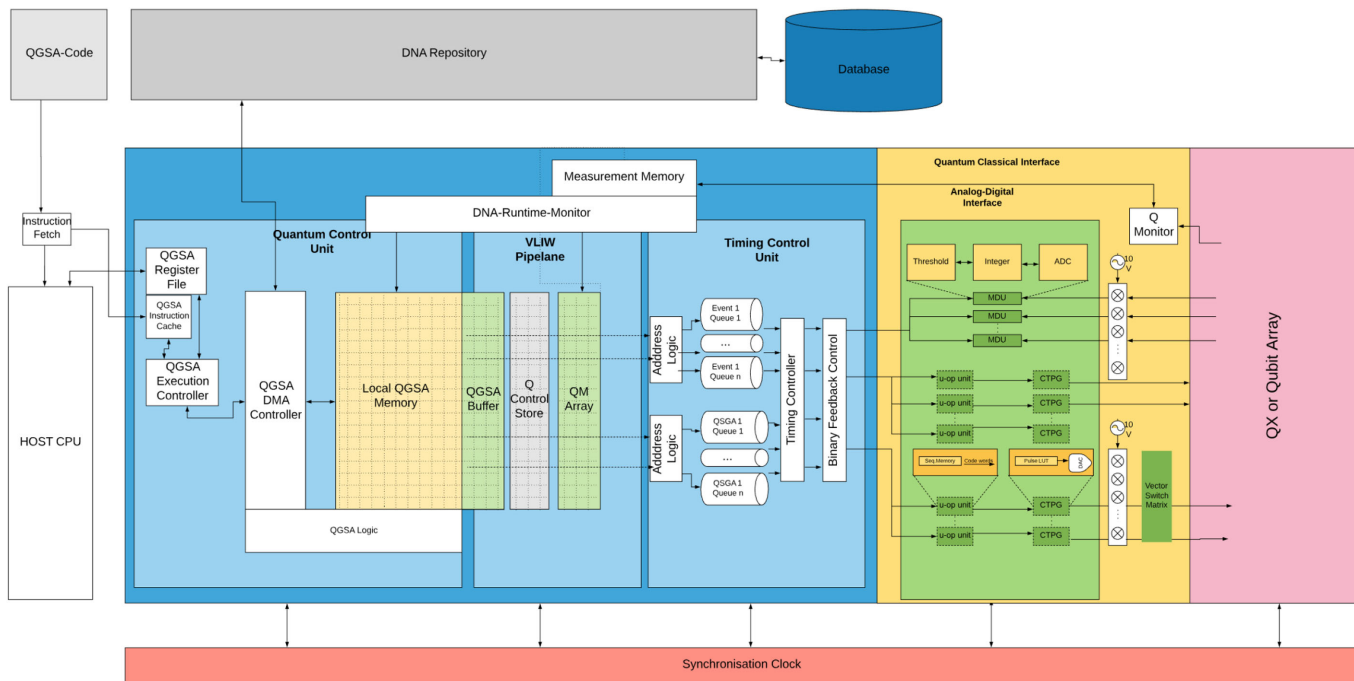


FIGURE 7. New microarchitecture for the QGS accelerator.

the sequence, incorporating the requirement for approximate optimal matching. A quantum search on the database amplifies the measurement probability of the nearest match to the query and thereby of the corresponding index. Due to the reference database and index, being entangled, the closest-match index can be estimated. Current explorations involves designing optimization algorithm for genomics applications using near-term quantum machine learning (QML) primitives like the quantum approximate optimization algorithm (QAOA).

As already mentioned, the proposed quantum accelerator will not be a standalone machine, but rather a quantum coprocessor that will be part of a heterogeneous system in which classical processors are connected to the quantum accelerator. Each processor will have its own instruction set. A first tentative view of the QGS microarchitecture is shown in Fig. 7.

There is need for run-time support to coordinate the activities of the different microarchitectural components and, as discussed, be responsible for the run-time routing of qubit states for two-qubit gates. In the quantum accelerator, the executed instructions generally flow through modules from left to right. The pink block on the right of the figure represents the QX simulation platform or an implementation of a quantum chip on which the test-runs of the QGS algorithms will be performed. The rest of the large (blue) block represents the microarchitecture. The DNA datasets is to be retrieved from an external classical database and transported to a local memory in the quantum accelerator. The size of the local memory will depend on the capabilities of the QX simulator platform and how that information is encoded. This research is based

on the large-scale microarchitecture simulation platform that we have already developed. Using the QX simulator platform makes it possible to rapidly develop hardware prototypes and verify their behavior and performance before an FPGA implementation is started. The set of queues will be relevant for feeding the DNA information to the qubit chip and for defining how the quantum gates are applied.

In a specific qubit plane topology, qubits will have to move around so that two-qubit gates can be applied on adjacent qubits. It is a prevailing idea that quantum compilers generate technology-dependent instructions [10], [37], [38]. However, not all technology-dependent information can be determined at compile time, because some information is only available at run-time due to hardware limitations, for instance qubits that need to be re-calibrated.

For testing the functionality of the algorithm, we use artificial DNA sequences that preserve the statistical and entropic complexity of the base pairs in biological genomes; yet in a reduced size so that they can be efficiently simulated in a classical architecture with qubit limitations. This implies understanding which run-time and, thus, routing support will be necessary to make sure that the quantum accelerator always has enough data to process and that they are in adjacent positions when necessary.

From an algorithmic perspective, near-term quantum optimization algorithms employ the variational principle, where a shallow parameterized quantum circuit is iterated multiple times while the parameters are optimized by a classical optimizer in the host CPU. This model of hybrid quantum-classical (HQC) algorithms requires fast feedback between the quantum accelerator and the real-time circuit/instruction

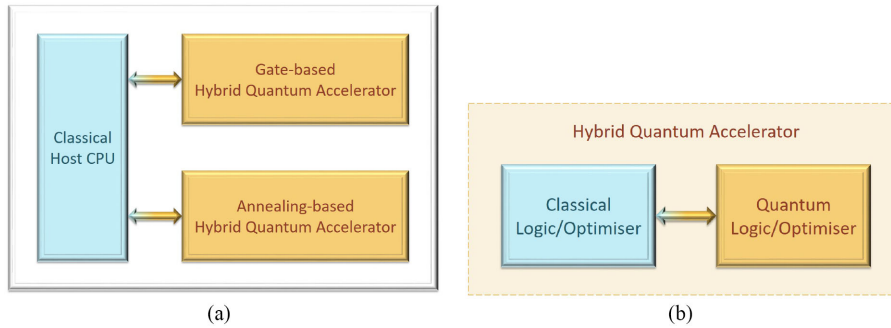


FIGURE 8. Model for near-term quantum-accelerated optimization. (a) Quantum accelerator model. (b) Hybrid quantum accelerators.

generator (i.e., the compiler and the microarchitecture). Since most quantum algorithms expect a statistical central tendency over multiple measurement, the expected probability of the solution state can be calculated inside the quantum accelerator itself, aggregating the measurements over multiple runs.

C. FULL STACK FOR QUANTUM OPTIMIZATION ON HYBRID QUANTUM ACCELERATORS

Optimization problems are ubiquitous and well-suited for near-term quantum acceleration. In this stack model, a generic execution model for optimization problems is considered, as shown in Fig. 8(a). Near-term quantum processors will be limited in size (number of qubits), quality (noisy operations), power (connectivity and controllability of every qubit), as well as the length of reliable computation (decoherence). To work with these constraints and still achieve a quantum advantage over pure-classical computation, the quantum application community is in favor of a hybrid approach, where some parts of the computation is carried out on classical logic.

The application is modeled in the classical host CPU, and translated to a quantum representation using a quantum programming language (such as OpenQL). The entire application software would generally consist of one or more quantum kernels (which are suitable for acceleration) and classical pre/postprocessing that are required to produce the final result of the problem. The quantum kernels are loaded to the hybrid quantum accelerator using a hybrid quantum representation (such as cQASM 2.0).

We consider two different types of quantum computation models for optimization: the gate-based and the annealing based methods. Both models can solve an optimization task encoded as a QUBO model, as discussed later.

The hybrid quantum accelerator typically has two processing elements, as shown in Fig. 8(b). The part that can benefit computationally from quantum effects, such as superposition, entanglement and tunneling, are offloaded to the quantum logic. Since near-term quantum processors cannot run a long computation, the entire process is generally split into small chunks of quantum circuits/anneals that can be carried out in burst, measured, and restarted based on the obtained

results. The classical logic keeps track of this progress and suggests the quantum logic the parameters for the next trial run.

The optimization problem is modeled as a QUBO expressed by: minimize $y = x^T Q x$, where x is a vector of binary decision variables ($x_i \in \{0, 1\}$) and Q is a (symmetric or upper triangular) square matrix of constants. Quantum annealers use the Ising model of spin variables (with the binary variables taking the values of $\{-1, +1\}$) as the computational model. This is isomorphic to the QUBO model and can, thus, be easily translated to an implementation on the annealer for estimating the minimum energy state of the spins. QUBO models can also be solved on gate-based quantum systems using the QAOA. QAOA is a variational algorithm where the classical optimizer specifies a low-depth quantum circuit to find the lowest energy configuration of a problem Hamiltonian. We believe that the choice of the quantum accelerator is dependent on the specific energy landscape of the application, as well as the characteristics of the quantum systems (e.g., annealers can process larger problem sizes, whereas gate models allow longer coherence times).

A specific use-case, we consider here, is the optimization problem called TSP. TSP falls under the NP-hard class (thus outside BQP), so the time to find the exact solution scales exponentially also on a quantum computer with respect to the problem size. Often a good suboptimal solution is admissible, thus heuristic algorithms of much lesser complexity can be employed. Our choice of TSP is motivated by its usefulness in many industrial applications in the domains of planning, scheduling, logistics, packing, DNA sequencing, network protocols, telescope control, VLSI testing, etc.

Given a complete graph $G = (V, E)$ with weights w_{ij} on the edge $i \leftrightarrow j$, the TSP aims to find a (directed/undirected) Hamiltonian cycle of minimum weight, i.e., a cycle that visits all nodes (cities) of the graph and such that the sum of the edge weights (travel cost) is minimum. Intuitively, given the ordered pairwise distance between cities, the TSP involves finding the shortest route that visits every city once. The order in which these cities are visited is not constrained. In our example, shown in Fig. 9, we search the shortest route between four cities in the Netherlands. The TSP graph is

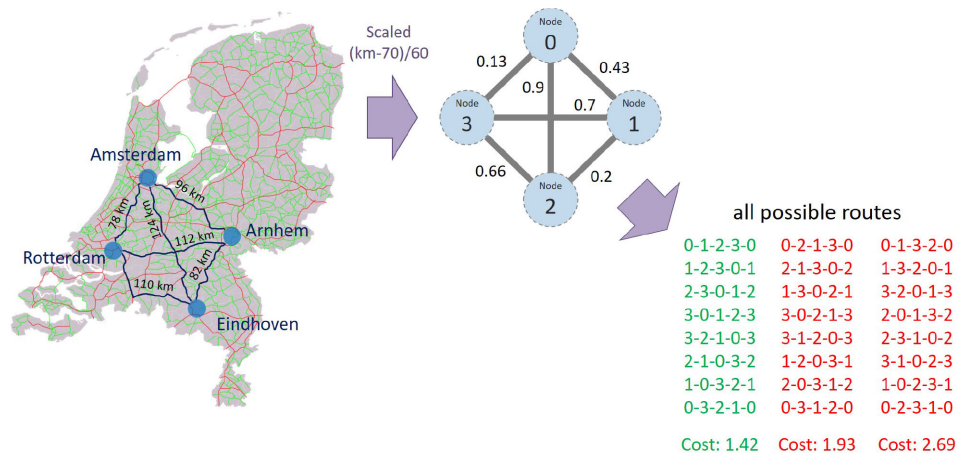


FIGURE 9. Route-planning reduction to TSP graph.

made from the scaled Euclidean distance We enumerate all possible solutions and find an optimal solution for this TSP with a cost of 1.42 (as shown in green).

Since the total number of visits (time IDs) equals the total number of nodes (city IDs); the total possible combinations of “(c, t)” is square of the the number of cities. The QUBO interactions (for the Q matrix off-diagonal entries) denote pairs of two nodes that can/cannot coexist and the associated reward/penalty. The interactions are categorized as follows.

- 1) Every node must be assigned.
- 2) Same node assigned to two different time slot is penalized.
- 3) Same time slot assigned to two different nodes is penalized.
- 4) The additional cost including an edge in the route to two consecutive time slots is the weight of the edge in the TSP tour.

We need 16 qubits to encode the example TSP into a QUBO.

When mapping the QUBO to a realistic hardware (such as D-Wave 2000Q, or IBM 20 qubit System One), the connectivity of the qubits in the physical topology is important. The embedding and mapping process considerably increases the number of required qubits and also the quality of the solution. In the traveling sales person, for example, which was given earlier, the highest number of cities that can be solved on a D-Wave 2000Q machine is nine. The amount of qubits needed to solve the problem grows as N^2 and finding embedding for the case with ten cities will fail in most (if not all) cases. On Fujitsu’s digital annealer, where it is fully connected (no embedding), we should be able to solve 90 cities. Error correction and routing for gate-based models adds further overhead in number of required qubits and operations. In classical computation, however, the current record for exact solutions to the problem, using branch and bound

algorithms is 85 900 cities. Heuristics, such as Monte Carlo methods, are used for larger inputs.

IV. HARDWARE AND SOFTWARE LONG-TERM VISION

There are different ways of building a computer and the way it is currently done is to combine multiple heterogeneous multicore processors. There are several models of quantum computation. The theoretical models, such as the quantum circuit model, adiabatic quantum computing, measurement-based (cluster state) quantum computation, and topological quantum computing, are equivalent to each other within polynomial time reduction. One of the most popular and by far the most extensively developed is the circuit model for gate-based quantum computation. This is the conceptual generalization of Boolean logic gates (e.g. AND, OR, NOT, NAND, etc.) used for classical computation. The gate set for the quantum counterpart allows a richer diversity of states on the complex vector space (Hilbert space) formed by qubit registers. The quantum gates, by their unitary property, preserves the two-norm of the amplitude of the states thereby undergoing a deterministic transformation of the probability distribution over bit strings. The power of quantum computation stems from this exponential state space evolving in superposition while interacting by interference of the amplitudes.

Most of the quantum computer that are made today are based on superconducting qubits, but in the past, there have been attempts on ion traps and semiconducting qubits are becoming very popular. We are just starting to reach the 50 qubit mark in processors but are way below the required coherence. The big system is shown in Fig. 3, where we include both the quantum annealer and the quantum gate accelerator. The same holds for the microarchitecture, for which, the components need to be developed.

Full connectivity: An important limitation that is not yet solved in any scalable way is the connectivity between the qubits, as for two-qubit gates the qubits need to be close to each other. It means that there is direct connectivity only in

the neighborhood of each qubits. This has important implications on the initial mapping of the qubits on the topology and especially the routing of the qubits to a location close to the other. Evidently, the kind of logical qubit one uses is very important. That is also an open issue currently brought to light by Preskill's paper [2] stating that SCs are too expensive. It suggests to move to small codes where much less qubits are needed to create a logical qubit. That is also why we have introduced the notion of a perfect qubit such that some of the complexities and problems can be abstracted away for the application developer.

Fig. 3 shows our long-term schema of what a quantum computer can look like in the two directions that are currently being explored, quantum gates-based and the quantum annealing approach. To give an overview of what is available on the market is very difficult as there are no commercially available computer systems that can be used in any reasonable way. The market can be split in two parts: Companies that are building a quantum-gate-based computer and ones that are focusing much more on optimization problems that can be solved with quantum annealing.

A. QUANTUM GATE-BASED COMPUTERS

Gate-based quantum algorithms are designed such that the solution states interfere constructively, whereas the nonsolutions interfere destructively, biasing the final probability distribution in favor of reading out the solution(s). However, the error rates are still around $10^{-2/-3}$ and need to be substantially improved.

- 1) IBM: They offer a quantum processor up to 53 qubits. The qubits have all the normal error behavior but they can be programmed. The specific thing is that IBM has not yet looked at any microarchitectural control of the physical level.
- 2) Intel: It is looking at both semi- and superconducting qubits but are in essence more interested in the semi-conducting qubit processor. The essence is fixing a lot on the qubit production, partly supported by a solid microarchitecture.
- 3) Microsoft: It has some preference for the majorana-based approach, but they still have to make the first qubit based on that quasi-particle. They are very active in the software development.
- 4) Alibaba: It is a strong player in this field, and they have a quantum lab that focuses on a range of activities going from the development of a quantum processor, quantum-classical algorithms up to simulation of quantum physics.
- 5) Google: It is also one of the leaders in superconducting qubits (John Martinis's team).
- 6) Rigetti: It is a start-up focusing on the superconducting quantum processor. They advance well, but there is not yet any applicable processor in the market even though there is a processor that can be used for some testing purposes.

- 7) Xanadu: The team focuses on continuous variable quantum computing based on photonics of squeezed light.

B. QUANTUM ANNEALING BASED COMPUTERS

Quantum annealing has a slightly different software stack than gate-model quantum computers and must be interpreted as a more limited edition of a quantum accelerator based on quantum gates algorithms. Instead of a quantum circuit, the level of abstraction is the classical Ising model, i.e., the problem we are interested in solving must be in this form. Just like superconducting gate-model quantum computers, superconducting quantum annealers also suffer from limited connectivity. It means that we have to find a graph minor embedding, combining several physical qubits into a logical qubit. Finding an embedding is NP-hard in itself, so probabilistic heuristics are normally used.⁵ We make a distinction between companies that offer a quantum computer such as D-Wave or a quantum-inspired computer such as Fujitsu. QNNcloud is a third offer based on neural-network and optical based quantum mechanisms.

- 1) D-Wave: The technology is up to 2000 superconducting qubits (in 2018), compared to the less than 100 qubits on gate-model quantum computers. D-Wave systems has been building superconducting quantum annealers for over a decade. D-Wave systems company offers an open source suite called Ocean, which can be used to make small examples of applications that can be executed on a D-Wave computer.
- 2) Fujitsu: It has invested in the development of a digital annealer. They offer a quantum-inspired computer and not a quantum computer. It is meant for the same kind of optimization problems that D-Wave can handle (QUBO problems). They currently offer 8192 nodes with full-connectivity and a programming interface but it is not clear and not known how the quantum-inspired accelerator works.
- 3) Hitachi: Similar to Fujitsu, Hitachi is also specializing in making a quantum accelerator based on quantum annealing using semiconducting qubits. More information can be found on the URL site mentioned here.⁶
- 4) QNNcloud: It is a company that offers a neural-network-based optical quantum computer where the neurons can be put in superposition and quantum-measurement circuits. A quantum optics implementation by QNNcloud uses a coherent Ising model, having different restrictions from superconducting architectures.
- 5) IQBit: It develops general-purpose algorithms for quantum computing hardware, primarily focused on computational finance, materials science, quantum

⁵Reference to the QAnnealing workflow: Open source software in quantum computing—arxiv.org/abs/1812.09167.

⁶https://www.hitachi.com/rd/portal/contents/story/cmos_annealing2/index.html

chemistry, and the life sciences. While there is a plethora of quantum computing languages, frameworks, and libraries for the gate-model, quantum annealing is less well-established. Their IQloud platform is focused on mapping optimization problems into QUBO format necessary to compute with quantum annealing processors and similar devices from Fujitsu, D-Wave, Hitachi, and NTT (QNNcloud) while their QEMIST platform is focused on advanced materials and quantum chemistry research with universal quantum computing processors.

C. QUANTUM PROGRAMMING LANGUAGES

A last component of the offering is related to the programming language that can be used.

- 1) Qiskit [11]: It is IBM's open-source quantum computing software development framework for leveraging available quantum processors. It consists of Terra (core compiler and libraries for quantum programming), Aer (noise modeling and noise-free simulators), Ignis (characterization of errors and QEC), and Aqua (applications).
- 2) Forest [14]: It is Rigetti's SDK that includes a simulator and cloud connection utilities.
- 3) Cirq [15]: It is Google's open-source quantum framework for experimenting with NISQ algorithms on near-term quantum processors. Also, the OpenFermion platform helps in translating problems in chemistry and materials science into quantum circuits.
- 4) Strawberry Fields [16]: It is Xanadu's full-stack Python library for designing, simulating, and optimizing quantum optical circuits. Xanadu also offers the Blackbird quantum programming language and the PennyLane QML platform.
- 5) XACC [17]: A vendor-independent solution is XACC, an extensible compilation framework for HQC computing architectures.
- 6) Ocean Software [18]: It is a quantum macroassembler for D-Wave systems from Los Alamos National Laboratories. It fills a gap in the software ecosystem for D-Wave's adiabatic quantum computers by shielding the programmer from having to know system-specific hardware details while still enabling programs to be expressed at a fairly low-level of abstraction.
- 7) OpenQL: The language that was discussed in this article earlier.

V. TOWARD IN-MEMORY COMPUTING

In-memory computing is becoming increasingly important as a new computer architecture. Rather than moving the huge amounts of data around to the logic, it is much more meaningful to move the logic around and keep the data as local as possible without moving it around, using, for instance, innovative technology, such as memristors. Memristors were theoretically defined already several decades ago by Leon Chua,

but recently the semiconductor manufacturers are seriously investigating their production. The key idea of a memristor is that it can be used to store data but also to make calculations. This is why memristors are an ideal candidate for making an in-memory architecture. The concept of in-memory computing is described in a paper where the concept is illustrated using memristor-based devices [39]. The main advantage of memristors is that they can be used both to store information and to work on it. So an intelligent merging of logic with data storage is the key of an in-memory architecture. It is a completely new way of designing algorithms and computing systems and it is far from evident what the design rules are that are needed to fully exploit the in-memory computing potential.

The link with quantum computing is very straight: the quantum logic is directly applied on the qubits and the qubits do not need to be transported to any quantum arithmetic and logical unit (ALU) before being processed. In quantum computing, the routing of qubit states is, therefore, also a very important problem. The qubits need to be put on the quantum chip in a way that the movement of qubit states is as minimal as possible. Also what routing protocols will be used for any quantum chip is a big open area of research in quantum computer engineering. Currently, in any of the semiconducting or superconducting quantum implementations the interaction between qubits has a NN constraint. That induces the need for deciding where to map and how to route the qubits used in the algorithm on the quantum chip. This qubit routing is an important and illustrative example of what in-memory quantum computing actually means. When adopting an in-memory computing architecture, a crucial challenge is to decide on the placement of the data that needs to be processed and to have a programming language and compiler, such that the appropriate logic can place close to the data. Any kind of algorithm will have data that the algorithm is changing to get a result and it is quite unlikely that there is no dependence between any of those data items. What that implies is that intermediate results will have to move around in the architecture such that it reaches the place where that result is used in the next computational step. Even though in-memory puts all the data in some kind of memory, those data items still have to move around such that a final result can be computed by the classical host-CPU. From a quantum physics point of view, the main challenges are the coherence of the qubits, the fidelity of the operations and the overall error rate of the quantum computation, involving both the qubits as well as their operation and the involved error-corrections. This is already being sufficiently studied by the quantum community but there are also clearly other challenges that need to be researched as soon as possible.

One of the main problems is the error-proneness of the qubit behavior, which consumes up to 90% of the (quantum) computer time. As explained, the routing and moving around of qubit states is a very important challenge. So any progress the physics community is making in that respect is

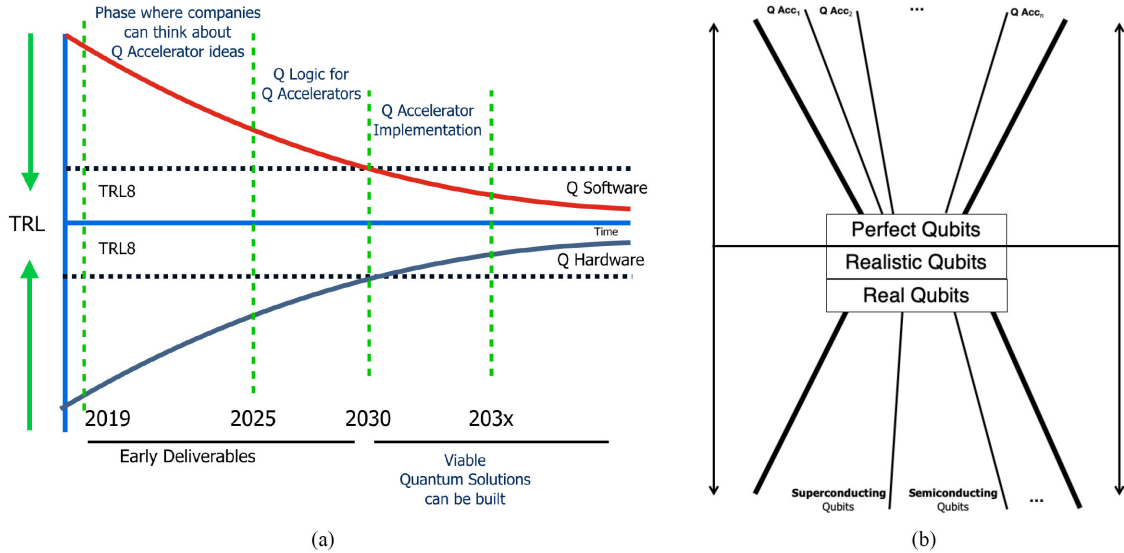


FIGURE 10. Quantum computer development future projections. (a) Development time frame. (b) Structural division between perfect and realistic qubits.

extremely important as it will reduce substantially the pressure on the microarchitecture and the overall system design. Brennen *et al.* [40] present a quantum computer architecture that addresses the important problem of qubit state routing for NN two-qubit gate execution. They use an idea from the von-Neumann architecture of classical machines such as a quantum bus, which is a refreshable entanglement resource to connect distant memory nodes. The overall approach is at the level of entanglement purification and qubit pairs with different fidelities. Given that a quantum computation on qubits complies to the same overall in-memory computing logic, that particular architecture is definitely interesting for any quantum device. The challenges involved with in-memory computing are, therefore, the same as for quantum computing. The underlying technology are not memristors or other technology but any of the quantum technologies and require also a full-stack integration of the different layers. In that sense, the quantum computing research should be based very much on the basic principles of in-memory computing.

VI. FUTURE PROSPECTS

It is very important that companies and other organizations start investing as soon as possible in quantum technology. Fig. 10 shows a projection of when different parts of software and hardware development will be required to create an efficient quantum computer. The distinction is made between the use of quantum accelerators and that of manufacturing a quantum chip. In general, any commercial or other organization is interested in new technology if the technology readiness level (TRL) is high enough. If we adopt the same levels as for classical technology, the TRL needs to have reached level 8 and that is sketched in the red and black line that are shown in Fig. 10(a). There are four vertical, green-dotted lines to illustrate three moments leading to the last phase where we assume there is enough software or hardware

maturity that can be used for any accelerator one wants to build. Phase I focuses on the reflection by the organization on the concrete need that exists and for which a quantum accelerator logic can be developed. Phase II resembles the team members brainstorming on the logic for the quantum accelerator. They will express that logic in OpenQL and develop some prototype microarchitecture and executed the logic on the QX-simulator. Phase III then focuses exclusively on the actual implementation and execution of the quantum accelerator logic, whether on an experimental quantum chip or on the QX simulator. This is the moment when the top and low curves can be combined in a real quantum prototype of the accelerator. Fig. 10(b) represents the way that the two lines of research are currently separated and which will be joined in maybe over the next decade. The division was used in this article where we made the distinction between the use of perfect and realistic qubits and how that determines the different layers in the full stack.

VII. CONCLUSION

Over the last couple of decades, quantum computing has been a 1-D research effort focusing on understanding how to make coherent qubits and how to implement the different universal quantum gate sets on any of the multiple quantum approaches. As far as computer architectural choices were made, the community has been focused very much on the von Neumann computer architecture and defined qubits in terms of memory and processing qubits. However, computer engineering as a field has understood by now that this approach never scales to the size needed for handling, for instance, the Big Data volumes that worldwide are being generated and collected. Two approaches seem to be very promising: the first comes from the accelerator community and involves the full-stack integration of the different layers that are needed to build the quantum accelerator. The use of

perfect qubits in that context makes sense as the end-users of any quantum accelerator can focus their reasoning on the quantum logic of the application and verify it through some implementation of the microarchitecture and the execution of the quantum instructions on the quantum simulator. The second option is to use the full stack for the control of, for instance, superconducting and semiconducting qubits with a microcode layer where we translate any kind of common QASM into an operational set of microinstructions, for a meaningful adoption of existing computer technology. It is very difficult to predict the performance improvement of a quantum computational device but that it will be much higher than any existing computational technology is clear. It also depends on the quantum application that is being looked at and the way the qubits are manufactured. Research is still needed for at least a decade before the full-integration effects become visible and verifiable.

REFERENCES

- [1] S. Vassiliadis, S. Wong, G. Gaydadjiev, K. Bertels, G. Kuzmanov, and E. M. Panainte, "The MOLEN polymorphic processor," *IEEE Trans. Comput.*, vol. 53, no. 11, pp. 1363–1375, Nov. 2004, doi: [10.1109/TC.2004.104](https://doi.org/10.1109/TC.2004.104).
- [2] J. Preskill, "Quantum computing in the NISQ era and beyond," *Quantum*, vol. 2, 2018, Art. no. 79, doi: [10.22331/q-2018-08-06-79](https://doi.org/10.22331/q-2018-08-06-79).
- [3] R. P. Feynman, "Simulating physics with computers," *Int. J. Theor. Phys.*, vol. 21, pp. 467–488, 1982, doi: [10.1007/BF02650179](https://doi.org/10.1007/BF02650179).
- [4] R. Van Meter and C. Horsman, "A blueprint for building a quantum computer," *Commun. ACM*, vol. 56, pp. 84–93, 2013, doi: [10.1145/2494568](https://doi.org/10.1145/2494568).
- [5] D. P. DiVincenzo, "The physical implementation of quantum computation," *Fortschritte Phys.*, vol. 48, no. 9–11, pp. 771–783, 2000, doi: [10.1002/1521-3978\(200009\)48:9/11<771::AID-PROP771>3.0.CO;2-E](https://doi.org/10.1002/1521-3978(200009)48:9/11<771::AID-PROP771>3.0.CO;2-E).
- [6] P. W. Shor, "Algorithms for quantum computation: Discrete logarithms and factoring," in *Proc. 35th Annu. Symp. Foundations Comput. Sci.*, 1994, pp. 124–134, doi: [10.1109/SFCS.1994.365700](https://doi.org/10.1109/SFCS.1994.365700).
- [7] C. Zalka, "Grover's quantum searching algorithm is optimal," *Phys. Rev. A*, vol. 60, 1999, Art. no. 2746, doi: [10.1103/PhysRevA.60.2746](https://doi.org/10.1103/PhysRevA.60.2746).
- [8] E. J. Houtgast, V.-M. Sima, K. Bertels, and Z. Al-Ars, "Hardware acceleration of BWA-MEM genomic short read mapping with longer read length," *Comput. Biol. Chem.*, vol. 75, pp. 54–64, 2018, doi: [10.1016/j.compbiolchem.2018.03.024](https://doi.org/10.1016/j.compbiolchem.2018.03.024).
- [9] K. Svore *et al.*, "Q#: Enabling scalable quantum computing and development with a high-level DSL," in *Proc. Real World Domain Specific Lang. Workshop*, ACM, 2018, pp. 7:1–7:10, doi: [10.1145/3183895.3183901](https://doi.org/10.1145/3183895.3183901).
- [10] A. J. Abhari *et al.*, "Scaffold: Quantum programming language," Tech. Rep. TR-934-12, Princeton Univ., Princeton, NJ, USA, 2012.
- [11] "QISKit—Quantum programming," 2020. [Online]. Available: <https://www.qiskit.org>
- [12] A. S. Green, P. L. Lumsdaine, N. J. Ross, P. Selinger, and B. Valiron, "An introduction to quantum programming in Quipper," in *Proc. Int. Conf. Reversible Comput.*, Springer, 2013, pp. 110–124.
- [13] N. Khammassi *et al.*, "OpenQL 1.0: A quantum programming language for quantum accelerators," QCA Tech. Rep. 8, QCA Lab, Delft University of Technology, Delft, The Netherlands, 2018.
- [14] "Rigetti forest software development kit," 2017. [Online]. Available: <http://docs.rigetti.com/en/stable/>
- [15] "Cirq—Software development kit," 2017. [Online]. Available: <https://github.com/quantumlib/Cirq>
- [16] "Strawberry fields—Photonic quantum computing," 2017. [Online]. Available: <https://strawberryfields.readthedocs.io/en/stable/>
- [17] "XACC quantum computing," 2018. [Online]. Available: <https://xacc.readthedocs.io/en/latest/>
- [18] "Ocean software for quantum computing," 2016. [Online]. Available: <https://ocean.dwavesys.com>
- [19] D. Riste *et al.*, "Detecting bit-flip errors in a logical qubit using stabilizer measurements," *Nature Commun.*, vol. 6, 2015, Art. no. 6983, doi: [10.1038/ncomms7983](https://doi.org/10.1038/ncomms7983).
- [20] A. Córcoles *et al.*, "Demonstration of a quantum error detection code using a square lattice of four superconducting qubits," *Nature Commun.*, vol. 6, 2015, Art. no. 6979, doi: [10.1038/ncomms7979](https://doi.org/10.1038/ncomms7979).
- [21] J. Kelly *et al.*, "State preservation by repetitive error detection in a superconducting quantum circuit," *Nature*, vol. 519, pp. 66–69, 2015, doi: [10.1038/nature14270](https://doi.org/10.1038/nature14270).
- [22] D. A. Lidar and T. A. Burn, *Quantum Error Correction*. Cambridge, U.K.: Cambridge Univ. Press, 2013.
- [23] P. W. Shor, "Scheme for reducing decoherence in quantum computer memory," *Phys. Rev. A*, vol. 52, 1995, Art. no. R2493, doi: [10.1103/PhysRevA.52.r2493](https://doi.org/10.1103/PhysRevA.52.r2493).
- [24] A. Steane, "Multiple-particle interference and quantum error correction," in *Proc. Roy. Soc. London A: Math., Phys., Eng. Sci.*, vol. 452, 1996, Art. no. 2551, doi: [10.1098/rspa.1996.0136](https://doi.org/10.1098/rspa.1996.0136).
- [25] A. R. Calderbank and P. W. Shor, "Good quantum error-correcting codes exist," *Phys. Rev. A*, vol. 54, 1996, Art. no. 1098, doi: [10.1103/PhysRevA.54.1098](https://doi.org/10.1103/PhysRevA.54.1098).
- [26] D. Gottesman, "Class of quantum error-correcting codes saturating the quantum hamming bound," *Phys. Rev. A*, vol. 54, 1996, Art. no. 1862, doi: [10.1103/PhysRevA.54.1862](https://doi.org/10.1103/PhysRevA.54.1862).
- [27] H. Bombin, and M. A. Martin-Delgado, "Topological quantum distillation," *Phys. Rev. Lett.*, vol. 97, 2006, Art. no. 180501, doi: [10.1103/PhysRevLett.97.180501](https://doi.org/10.1103/PhysRevLett.97.180501).
- [28] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, "Surface codes: Towards practical large-scale quantum computation," *Phys. Rev. A*, vol. 86, 2012, Art. no. 032324, doi: [10.1103/PhysRevA.86.032324](https://doi.org/10.1103/PhysRevA.86.032324).
- [29] C. Lin, S. Sur-Kolay, and N. K. Jha, "PAQCS: Physical design-aware fault-tolerant quantum circuit synthesis," *IEEE Trans. Very Large Scale Integr. Syst.*, vol. 23, no. 7, pp. 1221–1234, Jul. 2015, doi: [10.1109/TVLSI.2014.2337302](https://doi.org/10.1109/TVLSI.2014.2337302).
- [30] M. J. Dousti and M. Pedram, "Minimizing the latency of quantum circuits during mapping to the ion-trap circuit fabric," in *Proc. Conf. Des., Autom., Test Europe*, 2012, pp. 840–843, doi: [10.5555/2492708.2492917](https://doi.org/10.5555/2492708.2492917).
- [31] C. Horsman, A. G. Fowler, S. Devitt, and R. Van Meter, "Surface code quantum computing by lattice surgery," *New J. Phys.*, vol. 14, 2012, Art. no. 123011, doi: [10.1088/1367-2630/14/12/123011](https://doi.org/10.1088/1367-2630/14/12/123011).
- [32] X. Fu *et al.*, "eQASM: An executable quantum instruction set architecture," in *Proc. IEEE Int. Symp. High Perform. Comput. Architecture*, 2019, pp. 224–237, doi: [10.1109/HPCA.2019.00040](https://doi.org/10.1109/HPCA.2019.00040).
- [33] L. K. Grover, "Quantum mechanics helps in searching for a needle in a haystack," *Phys. Rev. Lett.*, vol. 79, 1997, Art. no. 325, doi: [10.1103/PhysRevLett.79.325](https://doi.org/10.1103/PhysRevLett.79.325).
- [34] D. V. P. Wang *et al.*, "Artificial associative memory using quantum processes," in *Proc. Joint Conf. Inf. Sci.*, 1998, vol. 2, pp. 218–221.
- [35] A. Sarkar, Quantum algorithms for pattern-matching in genomic sequences, M.Sc. thesis, Quantum and Comput. Eng., Delft Univ. Technol., Delft, The Netherlands, 2018.
- [36] A. Sarkar, Z. Al-Ars, C. G. Almudever, and K. Bertels, "An algorithm for DNA read alignment on quantum accelerators," 2019, *arXiv:1909.05563*.
- [37] K. M. Svore, A. V. Aho, A. W. Cross, I. Chuang, and I. L. Markov, "A layered software architecture for quantum computing design tools," *Computer*, vol. 39, pp. 74–83, 2006, doi: [10.1109/MC.2006.4](https://doi.org/10.1109/MC.2006.4).
- [38] T. Häner, D. S. Steiger, K. Svore, and M. Troyer, "A software methodology for compiling quantum programs," *Quantum Sci. Technol.*, vol. 3, 2018, Art. no. 020501, doi: [10.1088/2058-9565/aaa5cc](https://doi.org/10.1088/2058-9565/aaa5cc).
- [39] S. Hamdioui *et al.*, "Memristor based computation-in-memory architecture for data-intensive applications," in *Proc. Des., Autom., Test Europe Conf. Exhib.*, 2015, pp. 1718–1725, doi: [10.7873/DATE.2015.1136](https://doi.org/10.7873/DATE.2015.1136).
- [40] G. K. Brennen, D. Song, and C. J. Williams, "Quantum-computer architecture using nonlocal interactions," *Phys. Rev. A*, vol. 67, 2003, Art. no. 050302, doi: [10.1103/PhysRevA.67.050302](https://doi.org/10.1103/PhysRevA.67.050302).