Available online at www.prace-ri.eu

Partnership for Advanced Computing in Europe

Quantum Computing – A European Perspective

Mikael P. Johansson [a*1], Ezhilmathi Krishnasamy [b*2], Norbert Meyer [c*3],
Christelle Piechurski [d*4]

[a]CSC – IT Center for Science, Finland, [b]University of Luxembourg, Luxembourg, [c]Poznań Supercomputing and Networking Center, Poland, [d]GENCI, France

Abstract

Quantum computers have the potential to bring forth a major breakthrough in scientific computing. The foreseen increase in computational efficiency offered by quantum computing is of such magnitude that, despite being in its infancy, it is already being coupled with traditional high-performance computing technology. Here, we give an overview of quantum computing, the present state of affairs, and future scenarios. Europe has a unique opportunity to create world-leading supercomputing infrastructures incorporating quantum technology, by capitalising on the established expertise of European HPC centres in conjunction with the emerging European quantum ecosystem. This requires dedicated and sustained funding for quantum hardware and software developments, as well as for education. In addition, coordinated efforts and support for early adoption of quantum computing in academia and industry are essential.

---

[1] mikael.johansson@csc.fi

[2] ezhilmathi.krishnasamy@uni.lu

[3] meyer@man.poznan.pl

[4] christelle.piechurski@genci.fr

Table of Contents

This technical report is part of a series of reports published in the Work Package “HPC Planning and Commissioning” (WP5) of the PRACE-6IP project. The series aims to describe the state-of-the-art and mid-term trends of the technology and market landscape in the context of HPC and AI, edge-, cloud- and interactive computing, Big Data and other related technologies. It provides information and guidance useful for decision makers at different levels: PRACE aisbl, PRACE members, EuroHPC sites and the EuroHPC advisory groups “Infrastructure Advisory Group” (INFRAG) and “Research & Innovation Advisory Group” (RIAG) and other European HPC sites. Further reports published so far on the PRACE webpage [1] cover “State-of-the-Art and Trends for Computing and Network Solutions for HPC and AI”, “Data Management Services and Storage“, “Edge Computing: An Overview of Framework and Applications”, “User Requirements influencing HPC Technologies” and “Security in an Evolving European HPC Ecosystem”.

---

[1] https://prace-ri.eu/infrastructure-support/market-and-technology-watch

1. Introduction

Quantum computing (QC) is expected to bring a new revolutionary component to the high-performance computing palette. By directly exploiting quantum mechanical phenomena to an advantage, quantum computers may solve certain computational problems more efficiently than present day supercomputers and HPC algorithms. When sufficiently mature, quantum computers could tackle problems that due to their size and complexity will forever stay beyond the reach of conventional computing alone. Similar to the advent of transistor technology in 1947, the boost in computing power provided by this new computing resource is expected to dramatically increase the impact of research and accelerate problem solution, with a very promising effect on energy consumption.

Quantum computing is expected to have an impact on practically all fields of science, research, development and innovation that utilise, or *could* utilise computational modelling. Fields include artificial intelligence and machine learning, materials science and chemistry, pharmaceutical and medical research, finance and climate modelling, *etc*. This ground-breaking technology has the potential to provide solutions to some of the most pressing challenges of our society, from accurate modelling of complex weather systems and optimisation of resource usage, to the development of novel, sustainable materials as well as more efficient and personalised drugs.

The idea of quantum computers is already forty years old. Springing from the realisation that certain types of problems, for example simulation of physical processes at atomic scale, are by their very nature extremely difficult to model on classical computers, a new computing paradigm was born. After a long period of steady but arduous growth, the advances in quantum technology are now rapid, with the pace still accelerating. Presently, quantum computing is at a stage where its power has been demonstrated by performing actual calculations that are out of reach for classical computers [Arute 2019][Zhong 2020] [Wu 2021] . These *quantum supremacy* experiments serve as proof that there are no fundamental, physical limitations that would prohibit a quantum speed-up, although the actual problems that were solved are of little practical use. There is, however, still work to be done and challenges to overcome in order for quantum computers to show *quantum advantage*, and become integral components of workflows for solving real-world problems.

To harness the full potential of the upcoming quantum revolution, constructing the hardware alone is not sufficient. In order to utilise the hardware, tailor-made algorithms and software needs to be developed. Quantum programming requires fundamental rethinking on several levels. For example, quantum physical phenomena that are absent in classical computing, like superposition and entanglement have to be exploited. Problems have to be formulated properly, and in a novel manner, in order to be amenable to computation on quantum hardware.

For boosting and catalysing quantum computing and quantum software development, mature quantum computing infrastructures are crucial. The platforms have to provide a suitable level of abstraction, so that also users without deep expertise of quantum technology can utilise the new resources. Quantum experts will develop the required low-level software libraries and tools, while experts in other domains would use these tools for solving their respective research questions. In essence, the end-user should be given the most suitable tools possible for performing the actual research he or she is an expert in.

In order to increase quantum-literacy throughout Europe, the educational aspect of quantum computing requires attention. A prerequisite for this is that platforms that provide low-barrier adoption of the technology are made available. Then, students at various levels, including professionals in fields that could either utilise or further develop quantum computing can be reached.

In this report, quantum computing is introduced by defining the key concepts in order to familiarise the reader with the technology. Next, we focus on how to provide prompt access to quantum computing to the scientific community by coupling existing supercomputers to quantum systems while highlighting the algorithms that users can rely on to address different use cases. Existing and future synergies between industrials and academics follow. Europe should consider as key to support quantum adoption, as QC can be a decisive vector of green transition and digitalisation in a period where sustainability is at the heart of many discussions. The last section will be dedicated to the quantum market and future directions to keep an eye on. The conclusion focuses on what needs to be done to pursue the European efforts, together with individual European nations.

2. Bits and Qubits

A classical, digital computer uses bits to store and process information. A bit can be either 0 or 1, and can for example be represented by the absence or presence of an electrical signal, encoding “0” or “1”, respectively. When a computer performs an operation, the values either stay the same, or change: 0 becomes 1, or 1 becomes 0.

A quantum computer exploits the laws of quantum mechanics to enhance the capability of classical bits. Quantum bits, qubits, can, like classical bits, represent the states “0” and “1”. In addition, qubits can be “0” and “1” at the same time. This is known as a quantum mechanical superposition of states. To emphasise the quantum nature of qubits, the Dirac bra-ket notation is used, with states “0” and “1” represented by $|0\rangle$ and $|1\rangle$, respectively.

In general, the quantum state $|\psi\rangle$ of a qubit is a combination of the basis states $|0\rangle$ and $|1\rangle$, defined by the coefficients α and β: $|\psi\rangle = \alpha|0\rangle+\beta|1\rangle$. The coefficients, or amplitudes, are complex numbers, not simply real numbers between 0 and 1. As the square of the amplitude of a given state corresponds to the probability of that state (the Born rule), we get a constraint on the values of α and β. The sum $|\alpha|^2 + |\beta|^2$ must equal one, that is, 100%. The state of a qubit can therefore be represented as a point on the surface of a sphere, conventionally called the Bloch sphere. Each point on the sphere corresponds to specific amplitudes of $|0\rangle$ and $|1\rangle$, that is, specific superpositions of $|0\rangle$ and $|1\rangle$, see Figure 1.

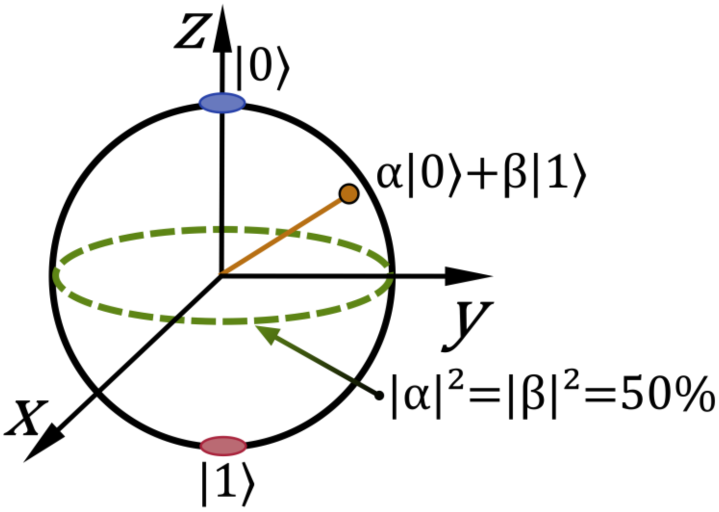

Figure 1: The Bloch Sphere. The blue dot along the z-axis represents the “north pole”, state $|0\rangle$, the red dot represents the “south pole”, state $|1\rangle$, and the dotted green line represents the “equator”, where states $|0\rangle$ and $|1\rangle$ are in equal superposition. The orange dot represents a general state of the qubit.

The Bloch sphere encapsulates the superior information content and operational flexibility of qubits with respect to classical bits. A classical bit can only take two values on the Bloch sphere, $|0\rangle$ or $|1\rangle$, and the only modification possible is to go from the north pole to the south pole, or vice versa. A qubit, on the other hand has access to the infinite set of points on the surface of the sphere; any combination of longitude and latitude, that is, a superposition of any amount of $|0\rangle$ and $|1\rangle$.

Just like ordinary bits, a qubit always returns the value “0” or “1” when read out, that is, measured, even if it would be in a superposition of both values. The result is probabilistic, with the probabilities dictated by the amplitudes. For example, a qubit at the equator of the Bloch sphere has a 50/50 chance of returning either 0 or 1 when measured; it will not return, say, 0.5. In addition, measurement destroys the superposition (collapses the wave-function in the Copenhagen interpretation): the state of the qubit is fixed to either the north or south pole of the Bloch sphere, and all information about the amplitudes is lost. This means that unlike bits, the value of qubits cannot be read mid-calculation.

In addition to encoding any value between 0 and 1 for the qubit, the complex amplitudes make it possible to describe the phase of the wavefunction. This adds another powerful feature to qubits over ordinary bits: the possibility of constructive and destructive interference. Let us consider the equator of the Bloch sphere, specifically, two points on it, the “plus” state and the “minus” state:

$$|+\rangle = \frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle;\ \ |-\rangle = \frac{1}{\sqrt{2}}|0\rangle - \frac{1}{\sqrt{2}}|1\rangle \tag{1}$$

When measuring, both qubit states will return “0” or “1” with 50% probability. They are, however, located on opposite sides of the Bloch sphere, along the x-axis, and have opposite phases. Any operation (except measurement) will affect the two states differently. For example, a rotation around the y-axis will shift the probability of measuring “0” or “1” in an opposite manner for $|+\rangle$ and $|-\rangle$.

Qubits provide a third fundamental advantage over bits, in addition to superposition and interference: entanglement. When a pair of qubits are entangled, their states are connected so that for example measuring the state of one qubit immediately fixes the state of the second qubit. Consider the two-qubit Bell state:

$$|\psi\rangle = \frac{1}{\sqrt{2}}|00\rangle + \frac{1}{\sqrt{2}}|11\rangle \tag{2}$$

Here, we have an equal superposition of two two-qubit states, one state where both qubits are “0”: $|00\rangle$ and another where both qubits are “1”: $|11\rangle$. While we do not know the value of either qubit before a measurement, we know that they must be equal. Thus, reading out one is sufficient for knowing also the value of the second qubit. If, for example, the first qubit is in state $|0\rangle$ after measurement, also the second qubit has to be $|0\rangle$. Further, performing an operation on just one of the qubits immediately affects the second qubit as well.

Physically, a qubit can be any system that can be in a superposition of two states. The ground and an excited state of an atom can, for example, represent $|0\rangle$ and $|1\rangle$. For switching the state from $|0\rangle$ to $|1\rangle$, one could use a laser light of specific frequency and duration, to provide the energy required for the atom to get excited from its ground state to an excited state (one could in principle implement ordinary bits with this scheme, too). By instead shortening the duration of the laser pulse to half of what is needed for the full flip of the qubit, we bring out the quantum: the result is an equal quantum mechanical superposition of both ground state and excited state, of both $|0\rangle$ and $|1\rangle$. It is not that the qubit *either* flipped or not, as it would be with classical objects: the qubit did both at the same time.

Several different physical implementations of qubits exist. In addition to the neutral atom scheme outlined above, for example superconducting loops, trapped ions, diamond vacancies, photonic and topological qubits are all actively developed. Presently, the different types of qubits have complementary strengths and drawbacks, and none of them are superior overall.

The main challenge for all qubit technologies is the effect of environmental noise [Cho 2020] . Noise sources, like temperature, vibrations, and cosmic radiation, interact with the qubits in an unwanted manner. This leads to decoherence, where the qubit loses its superposition, which in turn introduces errors into the calculation. In order to perform a useful calculation, the qubits need to stay in superposition for a sufficiently long time, so that enough computational operations can be performed on the qubits before measuring the result. The required coherence time depends on, among other things, the processing speed, “clock frequency” of the quantum processor. Gate operations are not perfect either, and gate errors will also affect the quality of the calculation.

Error correction schemes for mitigating the decoherence problem are actively being developed. The ultimate goal is to create perfectly functioning, so-called logical qubits. Logical qubits can be realised by combining several, on the order of a thousand, physical qubits [Google 2021]. Another approach is to employ cat qubits, named after Schrödinger’s famous feline [Mirrahimi 2014]. At least in the near-term, the reality will be that qubits are noisy and prone to errors. Even if the longest coherence times for qubits already exceed an hour [Wang 2021], this still pales in comparison to the stability of bits in classical systems, where we have become accustomed to errors affecting *any* bit of a calculation to be rarer than the lifetime of the circuits. While the errors introduced by, for example, cosmic rays need to be considered in both classical and quantum computing, the fragility of qubits is in a class of its own [Wilen 2021] .

Today, most of the work on quantum computers uses two-state systems, as in classical computing. It is perfectly possible to use more states, however. Just as classical ternary computers are based on trits instead of bits, qutrits would use three-level quantum systems to represent information, $|0\rangle$, $|1\rangle$, and $|2\rangle$ (or $|-1\rangle$, $|0\rangle$, $|1\rangle$). In general, quantum information units with more than two levels are called qudits, and have some advantages over qubits, *e.g.*, due to their ability to encode information even more densely. Taken to the extreme, discrete variables can be discarded completely, in favour of continuous-variable quantum computing [Hillmann 2020] . In what follows, we will focus on implementations using qubits, as the main ideas of quantum computing remain the same regardless of the number of quantum levels used.

3. Architectures

A rather general definition of a quantum computer is, that it is a device, that directly exploits quantum mechanical phenomena to perform a calculation. This can be implemented in several ways, and quantum computers do come in many flavours and technical implementations. Quantum computers can be grouped into the following three main categories, in order of increasing practical generality and computational power:

1. Quantum annealers,
2. Quantum simulators,
3. Universal, or general-purpose quantum computers.

Quantum annealing exploits quantum tunnelling and entanglement in order to solve a limited set of minimisation or optimisation problems. First, the qubits of the annealer are initialised to their lowest energy state, which is an equal superposition of $|0\rangle$ and $|1\rangle$. Then, the annealer applies biases to each qubit to shift its probability towards either $|0\rangle$ or $|1\rangle$. In addition, couplings between qubits are introduced, which increases or decreases the probability of two qubits to have the same value. In the end, the quantum annealer returns configurations that are close to the “energy minimum” defined by the different biases and coupling strengths. The biases and couplings are problem specific, and defined by formulating an optimisation problem as an Ising problem or through a Quadratic Unconstrained Binary Optimisation (QUBO) model. The Canadian company D-Wave has been offering quantum annealers commercially since 2011 [D-Wave] [Merali 2011].

A quantum simulator is a device that exploits superposition and entanglement to simulate model systems of real systems. This is achieved by mimicking the Hamiltonian evolution of some specific quantum system of interest on the quantum processor. This requires that the problem under study is cast into a form of a model Hamiltonian , *i.e.*, an operator corresponding to the total energy of the system, and determining its time evolution. While the problems amenable to simulation are often physical in nature, also more general optimisation problems can be implemented. As the quantum interactions between quantum particles is a built-in feature of quantum simulators, near-term quantum advantage is expected for the specific class of problems that they can describe. The quantum simulators of the French company Pasqal [Pasqal] provide both digital and analog quantum simulation capability [Henriet 2020] .

Universal quantum computers are the most diverse and potentially the most powerful class of quantum computers. They directly exploit superposition, entanglement, and wave-function interference in order to perform a calculation. A universal quantum computer can, in principle, solve any computable problem, with the additional advantage of up to exponential speed-up over classical computers and algorithms [Deutsch 1985] . Complete universality would require a sufficient number of high-quality qubits for any given problem. The term “universal” is therefore commonly used for quantum computers that operate on the same principle of generality, even if their capacity would fall short of simulating everything imaginable. The term general-purpose quantum computers is also often used for this class. The first demonstration of quantum supremacy, that is, proof that a quantum computer can perform *some* calculation faster than a classical supercomputer, was performed on Google’s general-purpose Sycamore processor [Arute 2019]. In Europe, for example the Austrian company AQT [AQT] and the Finnish company IQM [IQM] build general-purpose quantum computers based on ion traps and superconducting circuits, respectively.

Another division of QC technology can be based on the mode of operation: analog or digital. Quantum annealers are analog. Quantum simulators started out as fully analog, but, as mentioned above, can now combine digital computing elements as well. General-purpose quantum computers are digital, and use quantum gates, that is, basic logical operations for manipulating the qubits, and for achieving universality. Digital quantum computers can also benefit from performing parts of an algorithm in an analog manner, combining digital and analog blocks in quantum algorithms [Parra-Rodriguez 2020] .

Twenty years ago, DiVincenzo listed his now famous five criteria that a *general-purpose* quantum computer should fulfil [DiVincenzo 2000].

1. A scalable physical system with well characterized qubits,
2. The ability to initialize the state of the qubits to a simple fiducial state, such as $|000\ldots\rangle$,
3. Long relevant decoherence times, much longer than the gate operation time,
4. A “universal” set of quantum gates,
5. A qubit-specific measurement capability.

Note that criterion 4 cannot be fulfilled by analog quantum simulators that operate without gates. Quantum simulation without gates can in principle be universal, however [Aharonov 2007][Babbush 2014]. Also, continuous-variable quantum computing, which can be considered to be analog, comes with a universal set of quantum gates [Hillmann 2020] .

Constructing the part that performs the quantum computations, the quantum processing unit (QPU), is only the start of a full quantum hardware and software stack. A functioning architecture includes several layers above the QPU: interfaces between the classical and quantum parts; control logic and compilers that translate higher level operations or gates to specific quantum hardware; the actual quantum algorithms and quantum software; and finally, quantum computing theory [Van Meter 2013] [Fu 2016] [Bertels 2021] . Quantum error correction, when in use, is also part of the stack, both at hardware and software level. The quantum software and programming stack is just as crucial an ingredient as the actual QPU for the full stack. All components are needed in order for quantum computing to become a useful tool for doing science with. All of them are also highly non-trivial to implement.

4. Quantum Computer Emulators

It is important to have the full quantum software stack ready to take advantage of the physical QC's when they become generally available. Quantum computer emulators form an integral part of the initial stage of deploying quantum computing to a wide audience. Emulators provide access to quantum computing environments immediately, while access to real, physical quantum computers is still intermittent as physical quantum computing resources are scarce. Thus, emulators enable algorithm development ahead of access to the actual hardware.

A note on nomenclature: various definitions of a simulator and emulator in the context of quantum computing are in use in different communities. Here, we define a quantum *simulator* as a physical device used for simulating quantum mechanical systems and phenomena or problems otherwise beyond the capabilities of classical computers [Georgescu 2014]. While in principle a quantum simulator can be implemented using either universal quantum computers or with analogue devices, we follow the practice of considering quantum simulators to be analog, and in general, not universal or Turing complete. A classical device or software simulating a quantum computer, will in the text be referred to as an *emulator*. This definition will also apply in cases where the software does not necessarily model the actual physics taking place in a quantum computer.

Full emulation of a quantum computer on classical hardware is limited to a maximum of around fifty qubits, due to the exponentially increasing memory requirements of keeping track of the states of qubits. We are therefore already at the limit where the largest existing quantum computers cannot be fully simulated by classical computers anymore. Despite this limitation of using classical computers for emulating quantum computing, emulators have several advantages and features that will keep them relevant for the foreseeable future.

With sufficient hardware resources, emulators can give precision control of modelling the noise in a quantum computer. Thus, the effect of different types of noise can be studied, and bottlenecks in hardware specifications identified. Other hardware constraints, like qubit connectivity and readout errors can also be modelled, and individually assessed. By simulating the inner workings of a real quantum processing unit (QPU), the hardware itself can be improved. At the same time, algorithms can be made more resilient to noise, by, for example, optimising quantum gate operations. Debugging quantum algorithms is made simpler by the ability of reading out the full state of a qubit at any time during the execution of an algorithm. In a real quantum computer, reading the value of a qubit will return only either zero or one, and at the same time, destroy the superposition of the measured qubit. This makes debugging challenging on actual hardware. In general, new practices for code testing, augmenting existing procedures developed for classical software are needed.

Several advanced quantum computer emulators are already available and under continuous development. The Atos Quantum Learning Machine (QLM) [QLM] provides advanced simulation capabilities. From an HPC point-of-view, emulators developed for running on massively parallel architectures are of special interest. These include the Quantum Exact Simulation Toolkit (QuEST), the Jülich Universal Quantum Computer Simulator (JUQCS), and the Intel Quantum Simulator. In addition, toolkits for directly designing quantum hardware, like the open source Qiskit Metal and KQCircuits, fall in the broad category of quantum computer emulators.

Providing access and tuning the performance of quantum emulators on pre-exascale and upcoming exascale supercomputing infrastructures is crucial for extracting maximum synergy from combining HPC and QC. Having emulators play the part of actual quantum hardware in hybrid HPC+QC implementations (see next section) will speed up the development of the required interfaces and practices for connecting classical and quantum hardware into unified computing platforms. From an interconnectivity software point-of-view, whether the quantum processor is a physical device or emulated by software running on classical microchips is of little consequence.

5. Hybrid High-Performance Computing and Quantum Computing

Fault-tolerant, Large-Scale Quantum computers (LSQ) are still a technology of the future. In the present Noisy Intermediate-Scale Quantum (NISQ) era, the scientific research community can, however, already engage in quantum computing research, thanks to the recent availability of publicly accessible, physical quantum computers, in addition to the aforementioned emulators. This has enabled researchers to start developing future quantum algorithms on real hardware. There are several academic and industrial challenges that developers and application owners would face when adopting quantum computing. Such challenges are related (but not limited) to education, programmability and the availability of quantum computing systems. For decades, scientists have continued focussing on accelerating their applications through the adoption of new technological solutions, using CPU, GPU, and AI accelerators, *etc*. Accelerators in general refer to specialised processing units that handle certain computational tasks efficiently. In this sense, Quantum Processing Units (QPUs) follow an established route. The level of acceleration capabilities that quantum computing can bring in a 5 to 10-year time-frame cannot be ignored. A hybrid classical/quantum approach will allow application owners to benefit from the "best of both worlds".

There are also several challenges to overcome for quantum computers to run as separate appliances, namely user access (authentication, accessibility, environment, etc.), data access (input/output), workflow management, orchestration/allocation (batch scheduler), quantum resource management, to name a few. While these challenges need to be addressed also when coupling and properly integrating quantum systems and supercomputers, expertise and experience acquired in HPC during the last 40+ years will ease this integration. Coupling quantum simulators and computers with high-performance supercomputers through a unified cloud-mode access will allow a large part of the scientific community to become familiar with quantum computing, accelerating its adoption.

There are three main modes of integrating HPC and QC (hybrid HPC+QC), schematically represented in Figure 2: (1) stand-alone, (2) co-located, and (3) distributed. An actual implementation can use a combination of all three.

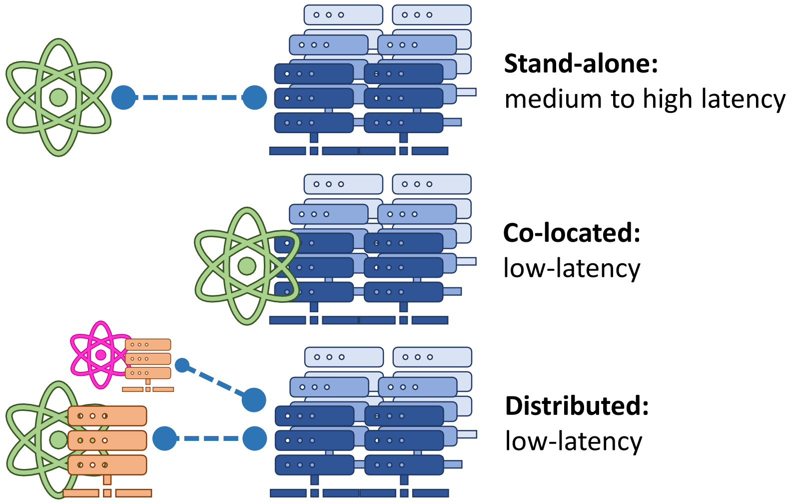

Figure 2: The three main modes of implementing hybrid HPC+QC.

The simplest manner of coupling HPC and quantum machines would be stand-alone quantum systems linked together through a communication network (Ethernet, Infiniband, etc.) to a supercomputer at an HPC centre. The HPC infrastructure then provides the programming environment and the orchestrator/resource scheduler.

For hybrid HPC+QC workflows where part of the calculation runs on classical HPC and part on the QC, latency between the classical and quantum processing units can become an issue. Algorithms that work by feedback loops, where the results of the quantum calculation are used as input for the classical part, and vice versa, can be latency-sensitive. If the execution time of the quantum and classical steps are on the scale of the network latency, then a significant part of the total execution time is spent idling. When QCs grow more powerful and allow for longer execution times for the quantum part of the algorithm, the issue of latency will diminish.

To avoid network-related latencies, an option would be to tightly couple a QPU with a CPU/GPU within a single server chassis. The host system's CPU would support QPU acceleration in, say, a PCI-e form factor, analogous to GPU accelerators, or embed a QPU within a single server. QPU and CPU could even be located on the same chip. While this could drastically reduce network latency, it would introduce additional physical design challenges related to cooling, noise suppression, electromagnetic shielding, ultra-high vacuum enclosures, etc. [Britt 2017] . This requires additional research, since this approach would require an entirely new stack of QC-specific technology that does not yet exist, rather than leveraging current, proven HPC technology.

The purpose of the classical server connected to the quantum computer or QPU, in addition to actually providing the interface between classical and quantum, is to keep the QC busy. For this, it needs the capacity for post-processing the output from the QC and preparing new input for the QC. For top-tier calculations, it will also need to connect to HPC infrastructure. It is worth noting, that the connection between the QC-connected server and the main HPC infrastructure is not expected to be latency-critical.

The server that connects the classical and quantum hardware can either be co-located with the HPC centre, or at a distance, providing a distributed infrastructure (see Figure 2). Both options enable execution of latency-sensitive algorithms, but cannot compete with the envisaged on-chip integration schemes, which would deliver even lower latencies. Co-location is somewhat simpler to implement, as all the components of the HPC+QC solutions are in close proximity. Synergies from research and system administration staff that maintain the two different architectures can also arise.

The distributed approach, where the classical server connected to the QC serves as a middleman between QC and HPC, has advantages. It for example allows QC and HPC servers to be located in spaces best suited for each architecture. The sensitivity to environmental noise requires that quantum computers are hosted in a more shielded space than classical computers. With the recent findings that cosmic rays have a larger than expected effect on qubits, it might become necessary to locate QC installations underground [Wilen 2021] .

In general, the distributed approach will enable higher modularity, increasing inclusiveness of different quantum technological solutions developed in Europe. Compared to the co-located approach, there are some additional complexities that need to be considered. For example, additional data security measures might need to be implemented. Further, scheduling has to consider that the classical server connected to the QC is separated from the HPC. When the distributed software stack is in place, connecting additional QCs into the workflow is rather straightforward. This allows QPUs in different locations to work in parallel on solving the same problem, or for cross-verification of the quantum results [Greganti 2021]

Today, there are on-going efforts in Europe to support quantum start-ups and the adoption of quantum technology by academic and industrial researchers. One such endeavour, the 4-year HPCQS pilot program, is the creation of the first (cloud-based) pan-European hybrid HPC+QC infrastructure that will integrate classical HPC infrastructure with several quantum nodes. HPCQS is an open and evolutionary infrastructure that would have the capability to grow and support diverse quantum computing solutions based on different technologies. The first building blocks of the HPCQS will rely on two 100-qubit quantum simulators from the French start-up Pasqal. By the end of 2022, these will be coupled to the Joliot-Curie (TGCC-CEA) and Juwels (JSC-FZJ) supercomputers through the Atos QLM, which will provide a hardware-agnostic programming environment for end-users.

Europe has thus started to leverage the accumulated knowledge and technological solutions found in the academic and industrial space, including HPC centres. This needs to continue, and different implementations of hybrid HPC+QC solutions need to be actively supported – as with qubit technology, the “winning” concept of combining classical and quantum computing is completely unknown.

6. Algorithms and Use Cases

As discussed in Section 2, qubits can encode much more information content than bits. The difference between bits and qubits grows more pronounced with increasing (qu)bit count. 2 bits can describe 4 different states (00, 01, 10, 11); 2 qubits can describe all 4 states *at the same time*. 3 bits can describe $2^3 = 8$ different states; 3 qubits can describe all 8 states at the same time. 20 qubits can already describe a million states simultaneously. The different states could, *e.g.*, represent inputs on which to perform a computation. With the qubits in superposition, all inputs can be processed in one run; in a classical computer, they need to be computed one by one.

There is a caveat to the massive parallelism that the qubits provide, however. In the end, a quantum computer will only provide one answer. Like ordinary bits, a qubit always returns the value 0 or 1 when read. Similarly, of the million states represented by 20 qubits in superposition, a measurement returns only one, say $|00000000011110111000\rangle$. This makes quantum computers well suited for tasks where one is interested in sifting out the “best” answer out of several possibilities. The travelling salesman problem exemplifies this: we have no interest in all the distances of all possible routes between a set of points, we only want to know the shortest route. If we *were* interested in all the routes, a quantum computer would not speed up the calculation.

Together with quantum entanglement and interference, the superposition of qubits potentially enables a very efficient solution of *certain* problems. Before getting results from a quantum computer, it has to be programmed; quantum algorithms have to be written. Quantum programming differs significantly from classical programming. The key differences are:

1. In classical computing, many basic operations are irreversible, where information is lost after the computation – *e.g.*, AND operands cannot be retrieved from the result of the operation. In quantum computing, all logical operations on the basic information units, qubits, must be reversible, where the input state can be inferred from the output, *e.g.*, NOT.
2. In classical computing, the basic bit operations are simple, either the state of the bit is unchanged or flipped. The basic operations of a quantum algorithm have to exploit the flexibility of qubits, by manipulating the qubits to move around the entire Bloch sphere.

3. In classical computing, the values of the bits can be read at any time. In quantum computing, measuring the value of a qubit collapses the superposition, thereby losing the information of its location on the Bloch sphere. For example, conditional branching inside the code (`if q0=1`) needs to account for this.
4. A classical computer is deterministic, the same input always gives the same output. A quantum computer is probabilistic by design: the results will in general be different, even with identical input values. Often, you need to run a quantum calculation several times to get statistically reliable results.
5. In classical programming, there is usually no need to consider computing errors caused by hardware. In quantum computing, the effect of noise has to be actively accounted for.

Essentially, quantum algorithms need to exploit quantum physical properties that are absent in classical computing, while considering the accompanying limitations. In order to be more efficient at a given task, a quantum algorithm has to explicitly use one or more of the quantum phenomena, superposition, entanglement, and wave interference. Then, polynomial to exponential speedups can be obtained [Montanaro 2016] Figure 3 outlines a general workflow for a quantum algorithm [Johnston 2019] .

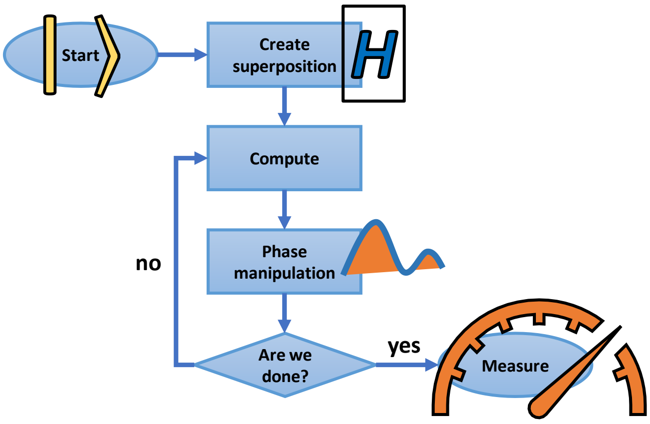

Figure 3: A general workflow for a quantum algorithm. First, an input state is prepared, and a superposition is created. After this, some computation is performed on the qubits, followed by phase manipulation. Generally, phase manipulation is used to increase the probability of the desired result. Finally, the states of the qubits are measured.

The operations for controlling the qubits depend on the quantum computer architecture. Presently, the quantum-gate programming model is the most common one for general-purpose quantum computers. Here, operations on the qubits are represented by operations, *i.e.*, gates on a virtual circuit diagram, see Figure 4. Various basic manipulations on qubits are performed in sequence, represented by symbols on the diagram. The gates can operate on several qubits at the same time, and quantum algorithms are often composed of one, two, and three-qubit operations. At the hardware level, usually only one and two qubit operations are used, while three-qubit gates can be represented as a string of one and two-qubit gates.

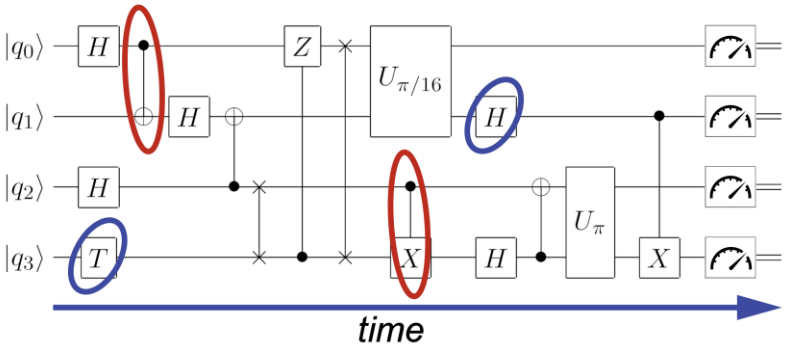

Figure 4: A fictitious quantum algorithm represented as a quantum circuit with quantum gates. The blue and red ovals encircle example one and two-qubit gates, respectively. Measurement of the qubits is at the end.

One or more of the following basic algorithmic building blocks are often present in quantum algorithms:

1. The Quantum Fourier Transform (QFT) [Camps 2021]

2. Quantum walks [Venegas-Andraca 2012]
3. Quantum search algorithms [Giri 2017],
4. Quantum algorithms for linear systems of equations [Harrow 2009] [Dervovic 2018] ,
5. Hybrid Variational Quantum Algorithms [Cerezo 2021] .

Other classes of algorithms include direct Hamiltonian simulation and optimization problems, often on quantum annealers or on quantum simulators. Derived algorithms then combine basic building blocks. The most famous example is probably Shor's algorithm for factorization of integers, a staple of public-key cryptography [Shor 1994]. The list of useful quantum algorithms with varying degrees of speedup grows continuously. The *Quantum Algorithm Zoo* maintains a comprehensive catalogue of existing quantum algorithms [Jordan 2021].

Quantum computers are excellently suited for some computational tasks, while performing poorly for other problem classes, relative to classical computers. Classical computers will always excel at certain tasks, even as universal quantum computers mature and become increasingly powerful. Tasks suitable for quantum computers can be divided into three main classes, which may overlap:

1. *simulating quantum systems*, the original driving force behind the idea of quantum computers — this class includes electronic structure theory, quantum chemistry, and materials science;
2. *optimisation problems* in general — *e.g.* logistics, portfolio optimisation, Monte-Carlo simulations;
3. *machine learning and artificial intelligence* — quantum acceleration is possible for several ML workflows, extending the applicability of quantum computing to a very diverse field of applications.

In addition to basic research efforts, there are various industrial sectors where quantum computing can solve real-world use cases: defence, aeronautics, automobile, finance, energy, travel and transportation, health, pharmacology, materials design, *etc*. Below, we highlight some typical use cases.

Solving the electronic structure of molecules and materials provides a prime example of a problem suited for quantum computers [McArdle 2020]. Today, electronic structure problems are tackled by approximating the underlying quantum-mechanical laws, and then solving these approximations on classical supercomputers. Even the approximate models become prohibitively time-consuming, when a large number of atoms or high accuracy is required. By treating the quantum mechanics of the electronic structure problem on a quantum computer, fully transforming an inherently quantum-mechanical problem to a classical computer can be avoided.

For quantum-chemistry problems, quantum algorithm development is an established and flourishing field [Cao 2019]. Electronic structure problems are central to several other fields of science as well. In the pharmaceutical sciences, the drug discovery phase is hampered by the inaccuracy of current models. With the aid of quantum computers, completely novel drug classes can be discovered. Energy storage and green energy catalysis are also based on quantum-mechanical interactions. Chemical catalysis involves the rearrangement of atoms and electrons to form new products from various reactants. Discovering new ways of synthesizing molecules and materials is of high industrial impact. Of equal importance is understanding natural processes. For example, understanding atmospheric chemistry in sufficient detail is pivotal for understanding and tackling climate change.

Even if simulating quantum systems was the original motivation for quantum computers, near-term quantum advantage is expected especially in modelling problems not directly related to quantum phenomena. Optimisation problems are abundant in, *e.g.*, finance modelling and risk assessment and analysis. In the energy sector, another use case is related to smart charging of electric vehicles to optimize the charging of thousands of electric vehicles combined with the charging-point network in an efficient way [Dalyac 2021]. The problem is a combinatorial optimization of several criteria, with increasing complexity with increased number of variables. The interest is to find the best order for the vehicles to be recharged according to users' priorities: the charging time required, charging-point availability, proximity of the terminals to vehicle location, and so forth. With increasing scale, quantum computers can excel in finding the optimal solution.

Modelling climate change at a larger scale, beyond molecules, can also benefit greatly from the increased optimisation capability of quantum algorithms. As mentioned, linear systems of equations, used for example in weather forecast models, can be solved much faster using quantum algorithms [Harrow 2009] . Accurate weather forecasts can have a direct, positive impact on food production and transport efficiency, for example.

Machine learning (ML) and artificial intelligence (AI) are used intensively for solving a large body of problems, from speech recognition, *via* targeted advertising, to optimising cancer treatments. An increasing portion of available computing capacity is spent on running AI algorithms. For this broad field, quantum computers are expected to have a large impact, and quantum machine learning (QML) is attracting much deserved academic and industrial attention [Dunjko 2020] [Huang 2021]. Quantum computers can be used both to enhance and speed up traditional machine learning procedures, as well as for studying truly quantum mechanical datasets [Uvarov 2020] .

In essence, *all* fields where computational modelling is used, will be able to benefit from quantum computing, from biology [Emani 2021] to digital humanities [Barzen 2020] At least some components of practically all complex modelling problems are amenable to quantum speed-up, if formulated in a suitable manner.

7. Academia and Industry in Europe

The European quantum ecosystem is globally strong, much due to academic researchers and start-ups. The key to set European future competitiveness resides in Europe’s capability to enable research entities and industries to exploit these new technological and complex quantum solutions in real-world use cases. Industries know the issues their sectors are facing and own business data, while academic consortia and start-ups have the expertise and know-how to support industries to identify if a use case has a “quantum” solution and if so, translate the use cases into quantum algorithms. By uniting the efforts of academia, including research and technology organisations (RTOs), and industry, the impact of quantum computing can be maximised.

Start-ups focusing on quantum computing are at the intersection of both worlds, often having an academic background. It is important to create and support start-ups that would bridge together European academia, start-ups and industries. This would help established industries to understand resource needs such as R&D, education, hiring, and financial investment, that are required for investing in quantum technologies, and thereby assess whether the investments are worth it. Start-ups, on the other hand, would gain visibility to the European community, improve their technology readiness levels (TRLs) and ease knowledge transfer to and between industrial and academic scientists.

Several European countries are already supporting quantum technology start-ups. Next, we highlight a few of them; the list is in no way exhaustive, but serves to show that such initiatives are well-received and fruitful. A recent report by the Canadian Institute for Advanced Research provides a more exhaustive overview of policy measures taken by different countries to support quantum R&D [Kung 2021].

In the Paris Region of France (Région Ile-de-France), GENCI in conjunction with Le Lab Quantique, a non-profit organisation promoting quantum computing, has created the PAck Quantique (PAQ) initiative. PAQ funds projects tackling major industrial challenges by partnering start-ups and academic consortia. In 2020, this three-year program has funded three real-world use-case projects in the energy (EDF, Total) and pharmacology (Qubit Pharmaceuticals) sectors. The success is expected to lead to an expansion of the program at the French national level. The second initiative in France is the Teratec Quantum Computing Initiative (TQCI) which aims to foster synergies between academia and industry in order to rapidly build up skills and develop know-how in the field of quantum computing, bringing together future users, technology providers and research centres.

In Finland, a collaborative project between the VTT Technical Research Centre and the globally recognised Finnish start-up IQM are constructing a quantum computer, planning to reach 50 qubits by 2024. The Finnish Quantum ecosystem in general is emerging strongly. The recently formed InstituteQ: The Finnish Quantum Institute combines expertise in quantum computation, communication, sensing and metrology, and simulation. Another recent initiative in Finland is Business-Q, which connects industry, universities, research and technology organisations, and investors. Business Finland, the government organisation for innovation funding and Helsinki Business Hub have been supporting quantum industry and start-ups for several years.

In Germany, a recent initiative, Quantum Technology and Application Consortium (QUTAC) consists of ten of the most influential German companies, who will use their resources and know-how to accelerate German innovation in the field of quantum computing [QUTAC] .

The Netherlands have also begun a quantum computing initiative, Quantum Delta NL, which has already allocated 615 M€ in funding. The Quantum Delta NL is a private-public partnership tasked with organising the implementation of the QC projects from the Netherland’s National Agenda for Quantum Technology (NAQT). The funding will be dedicated to train 2000 researchers, host three commercial R&D infrastructure by 2027 and support up to 100 start-up companies.

On a European scale, the Quantum Flagship initiative funds six projects on quantum computing and simulation [QT.EU]: AQTION – Advanced quantum computing with trapped ions; NEASQC – Next ApplicationS of Quantum Computing; OpenSuperQ – An Open Superconducting Quantum Computer; PASQuanS – Programmable Atomic Large-Scale Quantum Simulation; Qombs – Quantum simulation and entanglement engineering in quantum cascade laser frequency combs; and QLSI – Quantum Large Scale Integration in Silicon.

The European Quantum Industry Consortium (QuIC) has the objective of nurturing a fair and sustainable quantum technology business environment in Europe and ensure its global competitiveness [QuIC] .

It is of interest to identify the strengths of Europe in the field of quantum technologies and more precisely in quantum computing. Europe covers the entire domain of the quantum value-chain, from cryogenic cooling systems (*e.g.* Air Liquide, France; Bluefors, Finland), quantum communication (*e.g.* KETS, UK; SSH.COM, Finland; Cryptonext Security, France), quantum sensors (*e.g.* Qnami, Switzerland) up to quantum networks (*e.g.* QPHOX, The Netherlands; VeriQloud, France). Atos is heavily invested in quantum computing, with the Quantum Learning Machine (QLM) programming environment and QPU front-end. Academic institutions furthest along the path towards hundred-qubit quantum computers include the Wallenberg Centre for Quantum Technology (WACQT), directed from Chalmers University of Technology (Sweden) [WACQT] and the joint efforts of ETH Zürich and the Paul Scherrer Institute (Switzerland) [ETHZ 2021] . European quantum computing and quantum software start-ups and SMEs are numerous and visible on the international market. Table 1 lists a selection of these, showing the diversity of the European quantum computing domain.

European start-ups and SMEs	**Country**	**Quantum computing**	**Quantum software**
Algorithmiq	Finland		Life sciences
Alice & Bob	France	Cat qubits	
AQT	Austria	Trapped ion qubits	
C12 Quantum Electronics	France	Carbon nanotube qubits	
Cambridge Quantum Computing	UK	Quantum compiler	Chemistry, ML, finance
CEA-Leti	France	Silicon qubits	
HQS Quantum Simulations	Germany		Materials science
IQM	Finland, Germany	Superconducting qubits	HW/SW co-design, finance, KQCircuits
Multiverse Computing	Spain		Finance
ORCA Quantum Computing	UK	Photonic qubits	
Oxford Ionics	UK	Trapped ion qubits	
Oxford Quantum Computing	UK	Superconducting qubits	
Pasqal	France	Rydberg atomic qubits	Pulser
Phasecraft	UK		Materials
Qilimanjaro	Spain	Quantum annealer	
Qu & Co	The Netherlands		Chemistry & materials
Quandela	France	Photonic qubits	
Quantastica	Finland, Estonia, Serbia	Quantum emulators	Programming tools
Quantum Motion	UK	Silicon spin qubits	
Qubit Pharmaceuticals	France		Pharmacology
Rahko	UK		Drug discovery
Riverlane	UK		Deltaflow.OS quantum computer operating system
Terra Quantum	Switzerland		Cybersecurity

Table 1: Some European start-ups and SMEs actively developing quantum computing technologies.

Internationally, technology giants such as Alibaba, Amazon, Fujitsu, Google, Honeywell, IBM, Intel, Microsoft, and Huawei are investing significant resources into quantum computing. Compared to the US and Asia, the European quantum computing industrial landscape differs. At the end-user side, several large companies are involved, as noted above. Almost all of the development activity is due to start-ups and SMEs, however, with only a few incumbent firms involved. This makes the ecosystem vibrant and agile, but at the same time, vulnerable to challenges related to scaling capacity and financial stability. Therefore, pan-European efforts are needed to ensure that European start-ups have the necessary resources, incentives, and support available for continuing building their respective business cases *in Europe*. This would mitigate the risk of mergers with large non-European corporations. Incubation of a critical mass of successful start-ups is needed [Räsänen 2021]. The success of national support programmes has been demonstrated, now is the time to scale up.

8. Quantum Communication

In quantum communication, digital information is transferred using qubits, instead of bits; quantum states are transferred from one place to another [Gisin 2007]. This can be exploited in several ways. In this section, we briefly cover its application in securing digital communications, and in the context of a quantum internet.

Encryption of digital information is crucial for modern society. Online banking and e-commerce, state and company secrets, medical records and personal communication of individuals, all depend on reliable digital security. Digital encryption is based on decryption being a computationally hard task: by ensuring that breaking the encryption takes long enough, the information is practically secure. Shor's algorithm for factoring integers into prime factors [Shor 1994] showed that sufficiently advanced quantum computers can speed up the mathematical task of breaking encryption protocols. Thus, revision of commonly used encryption methods has become urgent. Quantum communication can mitigate some of the risks that quantum computing poses to digital security.

Digital encryption is based on so-called encryption keys, which are used to encrypt and decrypt information. Like all digital information, the keys can be represented as bit strings of zeroes and ones. From an encryption point-of-view, the advantage of using qubits during the transfer is, that is becomes possible to detect if someone is eavesdropping on the data transfer. This is because measurement of a quantum system, like a photonic qubit, necessarily disturbs the system. By using proper communication protocols, it thus becomes possible for two parties to exchange the secret encryption keys required for encrypting communications with the certainty that no-one has intercepted the key exchange. This is the basis for Quantum Key Distribution (QKD).

Since the Digital Assembly event in 2019, All 27 EU member states have signed the European Quantum Communication Infrastructure (EuroQCI) Declaration, committing to work together to build a pan-European, secure quantum communication infrastructure [EuroQCI] . This will be used for data transit and storage in a highly secure manner. It will ultimately link sensitive public and private communication assets all over the EU, such as banks and administrations. EuroQCI assumes both the construction of a network of terrestrial connections and the launch of a satellite segment - combining the technological potentials and capabilities of individual Member States. Although the communication mediums vary, the key motivation remains the creation of a common network that enables the use of quantum communication to improve the exchange of sensitive and classified data.

The first use of the EuroQCI infrastructure will be to provide QKD services. Work in this direction is underway in the OpenQKD EU pilot project [OpenQKD] , which connects European academia, industry and start-ups in the deployment of open QKD testbed sites around Europe. QKD is already being commercialised, with several QKD products appearing on the market. As a testament to the maturity, even satellite-based QKD is being deployed worldwide. For example, the European Space Agency (ESA) and the UK-based company ArQit are working on the QKDSat project, which aims for a satellite launch in 2023 [ESA 2021] .

While QKD offers enhanced security, it should not be considered *fully* secure. In any real-world implementation, security weaknesses will arise. For example, the French National Cybersecurity Agency [ANSSI 2020] and the National Cyber Security Centre of UK [NCSC 2020] discourage relying on QKD for ultra-secure applications. Instead of QKD, Post-Quantum Cryptography (PQC), that is, new encryption algorithms that are considered safe also against advanced quantum computers can be employed [ENISA 2021] A practical advantage of PQC over QKD is that PQC requires no new infrastructure to operate. Europe has the potential to be world-leaders in this field as well [Loesekrug-Pietri 2021]. For this, national programmes like PQC Finland [PQC.FI] need to be augmented by pan-European efforts.

Generation of truly random numbers is central to cryptography, including QKD and PQC, and is important for various computational simulation workflows and gaming. The inherent randomness of quantum systems lies at the heart of Quantum Random Number Generators (QRNGs). Several commercial QRNG systems are already available. Here too, has collaboration between academia and industry been fruitful, as exemplified by the secure, high-speed quantum random number generator developed in collaboration between the Universities of Sheffield and York, the Technical University of Denmark, and the Danish company Cryptomathic [Gehring 2021].

In quantum teleportation, information can be exchanged between two qubits over a long distance [Pirandola 2015] . First, two qubits (A and B) are entangled, by creating a Bell state (see Section 2). After this, the qubits can be transferred to separate locations. A third information qubit can then be made to interact with qubit A. This instantly affects qubit B, with information transferred through the quantum communication channel created by the entanglement. In order to make sense of the information, classical information needs to be sent from A to B as well (for example over the classical internet), which prevents superluminal information transfer.

Teleportation is the basis for a quantum internet [Castelvecchi 2018]. In addition to information transfer, a quantum internet could be used to combine the qubits of separate quantum computers. This type of parallel quantum computing would be revolutionary. By properly combining, say, two 50-qubit quantum computers, the result would in essence be a 100-qubit machine. Note the exponential increase in theoretical computing power for quantum computers: this would not *double* the computing capacity, the capacity would increase by a factor of $2^{50} = 10^{15}$, a quadrillion. The long-term ambition of the European Quantum Internet Alliance (QIA), coordinated by QuTech in the Netherlands, is to build a quantum internet that enables quantum communication applications between any two points on Earth. This is pursued by developing, integrating and demonstrating all the functional hardware and software subsystems required [QIA] .

9. Green Transition and Digitalisation

From the global energy consumption point-of-view, quantum computing offers a promising augmentation to digitalization. The fundamental logical qubit operations use very little energy [Ikonen 2017, Chiribella 2021]. The major part of the energy consumption arises from cooling of the quantum computer. The most mature quantum computer technologies need to be cooled down close to absolute zero (–273 °C). The space that needs cooling is, however, small; there is no need to cool down the entire space, only the processor itself. This is in stark contrast to a traditional data centre. A quantum computer itself emits very little heat, which further decreases the need for cooling. For example, Google’s Sycamore quantum computer setup, which was used for demonstrating quantum supremacy, used less than 26 kW [Arute 2019].

As quantum computers operate on very different computing principles compared to classical computers, comparing the energy requirements of, say, qubit and bit operations is not meaningful. Instead, one has to compare the total power consumption for running specific calculations. For problems that the D-Wave quantum computer is suited for, the energy need has been estimated to be just 1% of what a digital supercomputer consumes [D-Wave 2017] . For general-purpose quantum computers, an advantage of many orders of magnitude in energy consumption over classical supercomputers is estimated [Villalonga 2020] As the speedup offered by quantum computers is superlinear, it is expected that energy consumption grows very modestly with increasing utility of a quantum computer. Thus, the ratio of computational power to energy use will become increasingly favourable in the future.

Direct benefits of quantum computing for the green transition include enabling highly accurate and reliable solutions to problems that are intractable using traditional high-performance computing. These areas include solving the electronic structure problem of atmospheric reactions, green catalysts for clean energy production and CO_2 sequestration, environmentally friendly alternatives to chemicals, and materials for solar cells and batteries, to name a few. Optimisation problems form another class relevant from an environmental perspective, including the logistics of travel and transport, flood prediction and crop optimisation, and in general, industrial processes. A sufficiently advanced quantum computer could model and solve these problems with unprecedented accuracy and speed, far surpassing the present and future capabilities of traditional supercomputers. Quantum computers can thus directly help decrease greenhouse gas emissions in the future.

The use of quantum computers for solving other problems, not directly related to sustainability, would have an indirect but notable effect on the energy consumption of high-performance computing. Parts of most modelling problems can be solved more efficiently using quantum computing. Some of the computationally most intensive problems running on current HPC infrastructure are well suited for quantum computing. A prime example are electronic structure calculations, which consume a major part of HPC resources: at CSC, almost half of the supercomputer capacity is spent on solving electronic structure problems; on GENCI’s supercomputers, combining all problems related to chemistry, including electronic structure, represents 30% of the available computing hours. Machine learning and artificial intelligence form another, still growing application field that uses increasing amounts of energy, but at the same time is amenable to significant acceleration by quantum computers. Thus, offloading suitable tasks to quantum computers decreases the load on traditional supercomputers. This opens up the possibility to curb the increasing energy-demand of data centres. The shift towards quantum computing in HPC will lead to the production of significantly higher-quality scientific returns per Watt consumed.

As discussed in detail in the section on hybrid HPC+QC, quantum computers are not separate appliances; to extract full advantage from quantum computing, powerful, traditional HPC infrastructures are required. In this scheme, the classical supercomputer capacity stands for the vast majority of the power consumption. A distributed HPC+QC approach, where the classical HPC and the QC resources can be geographically distant is ideal, also from a sustainability point-of-view. The traditional, power-hungry HPC centres can be set up in areas where, for example, cooling needs are lower, waste heat can be readily reused, and low-carbon power sources are available (as in, e.g., the LUMI pre-exascale EuroHPC data centre in Finland), while the low-power quantum computers can be located practically anywhere, provided a high-speed internet connection is in place.

Sufficient user support actions, discussed in the next section, have a direct impact on the green transition as catalysed by QC. The earlier users are able to adopt and utilise quantum computing for their specific problems, the sooner a transition towards more impactful science, research, and development can take place. Considering the high impact that utilising quantum computers can have on solving problems crucial for society, each additional year that the inauguration can be brought forward counts.

10. Competence Development and Education

Several highly interesting technologies that could have revolutionised the world turned out to be a failure. Despite large investments in both hardware and software, including subsequent maintenance costs, some technologies have found only niche use in few areas. This is partly due to failures in engaging end-users. This has resulted in low user uptake in fields that *could* have seen a benefit, but failed to do so due to insufficient support for non-experts to port their problems and codes to new platforms and technologies.

The introduction of quantum computing along with exascale systems within a short timeframe in Europe requires both reaching production stability of the computing systems, and sufficient literacy of users to program and harness such unique computing technology. We remember from the past several use cases of effective use of new technologies and the efforts in the preparation of algorithms for:

- Vector computers in the late1980's,
- Massively Parallel Processing supercomputers (MPP) in the 1990's,
- Accelerators (GPUs) at the beginning of 2000’s,
- And the less popular use of FPGA accelerators, starting in the 1980’s.

Open, standardised, and portable programming models are of high importance for end-users in order to, as far as possible, avoid re-writing their applications on different quantum hardware implementations and architectures. It is especially important to provide a flexible way to move applications from one to another architecture automatically or at least semi-automatically with minimal effort.

The introduction of QC in Europe in a rather short time-frame is challenging. At the same time, urgency is crucial for both the advances in computational modelling that QC provides for problems of high societal impact, and for ensuring European digital sovereignty in the field of quantum computing. Therefore, the EuroHPC programme has to support the development of new algorithms and novel applications. It is imperative to:

- Develop appropriate tools to facilitate the preparation of algorithms in specific application areas, from low-level libraries and tools to high-level portable programming languages;
- Establish competence centres for users and to reinforce existing HPC competence with HPC+QC support;
- Develop programming courses that focus on the effective use of quantum computers;
- Raise awareness among all user groups, starting from young scientists, even students at various levels, to professionals in the general field of computational modelling, to decision and policy makers.

As QC will not replace classical processing, but instead, augment it, hybridisation of existing HPC centres with new QC technology is a logical next step. This requires training and competence development. The extensive experience of the Partnership for Advanced Computing in Europe (PRACE) community in creating a training network of competence centres should be leveraged. This should be done in collaboration and concord with the training activities organised by the EuroCC programme, which aims to bring together the necessary expertise to set up a cross-European network of National Competence Centres (NCCs) in HPC-related topics with 33 participating members and associated states [EuroCC] .

The PATC (PRACE Advance Training Centres) and the PRACE portal of trainings [PRACE] could be used for training courses of quantum technology prior to the establishment of pan-European quantum computing systems. PRACE organises a programme of 90+ courses annually, via its 14 PRACE Training Centres (in Austria, Belgium, Czech Rep., Finland, France, Germany, Greece, Ireland, Italy, the Netherlands, UK, Slovenia, Spain, and Sweden). This training is complemented by Seasonal Schools and special on-demand events that are run in collaboration with other projects and European Centres of Excellence (CoEs).

Only a broad process of raising awareness and competences in the future academic and industrial communities will result in an increased level of use of QC methods and the creation of new algorithms, and thereby full leverage of the technology. At the same time, it is important to mitigate the risks involved with an emerging technology. The educational programmes and degrees need to be specific enough to enable the future graduates to successfully work on quantum technologies, and at the same time, diverse enough to ensure that the acquired skill set is not lost in case the job market growth falls short of predictions and expectations.

In Europe, career-affecting educational choices are often made at the end of elementary school, roughly around the age of 15. Thus, right now, many young minds are making decisions affecting the talent pool of the matured quantum computing era. The field can attract young talent by active engagement with students at high-school, even primary school level. HPC centres and higher-education institutions can help by actively promoting awareness of QC, and showing that the field provides a viable and exciting career path.

11. Future Directions and Challenges

Quantum computing is developing with accelerating speed, with no slow-down in sight. The technology has stepped out of the academic lab, and is growing commercially [MacQuarrie 2020] The Boston Consulting Group notes that equity investments in quantum computing almost tripled in 2020 [Bobier 2021] . Market data analyst Pitchbook reports that by early September, global venture capital investments in QC for the year 2021 surpassed USD 1 billion, more than during the previous three years combined [Temkin 2021] Estimates of the future trends on the market growth of quantum computing vary, but essentially all predict continued increases in investments, revenue, and job market size. For example, in their latest forecast on quantum computing, Inside Quantum Technology IQT predicts a total revenue of roughly EUR 2 billion by 2026, a ten-fold increase compared to 2020 [IQT 2020] . McKinsey predicts a whopping trillion-euro value potential by mid-2030s [Hazan 2020] . Expectations are thus exceptionally high.

To live up to the expectations, the technology needs to demonstrate constant progress. Keeping the momentum going can be considered to be the first challenge of quantum computing. As discussed, qubits are very sensitive to noise, and are still far from functioning in an ideal manner. In addition, the present qubit count is modest. The reportedly most powerful general-purpose quantum computer, the Chinese Zuchongzhi, has 66 superconducting qubits [Wu 2021] . We are therefore presently in what has been coined the Noisy Intermediate-Scale Quantum (NISQ) era [Preskill, 2018] .

Quantum advantage is possible also with NISQ devices, and incremental speed-up for HPC is to be expected, but NISQ cannot deliver on the promise of a quantum *revolution* in computational modelling. For this, fault-tolerant quantum computers with long coherence times and higher qubit count are required. Quantum computing needs to scale up, and achieving this is anything but trivial. Despite, or perhaps rather, *due* to the challenges, large-scale quantum (LSQ) has to be kept as the ultimate goal, and work on paving the way towards it needs to be resolute. It is worth noting that the transition from NISQ to LSQ will not be a binary event. Fittingly, we will see a period with a superposition of NISQ and LSQ, with larger and less noisy quantum computers.

Several qubit technologies have been implemented and shown to work, at small scale. The main challenge is not manufacturing the individual qubits anymore, even if advances in qubit quality are still required. Now, control and calibration of ever more densely packed qubits is becoming problematic. Connecting the qubits to create, say, a million-qubit entity is supremely challenging. The cabling work required for qubit control lines alone is massive. In the end, it might turn out that connecting several smaller quantum computers over a quantum network, for example the quantum internet, will be the method of choice for scaling up QC. Already, qubits 60 metres apart have been connected to perform a concerted calculation[Daiss 2021][Hunger 2021] .

The second major challenge on the road towards large-scale quantum computing is quantum error correction (QEC) [Devitt 2013][Terhal 2015]. QEC is needed due to qubit decoherence, gate inaccuracies, readout errors, *etc.* In classical systems, errors can be mitigated by redundancy, by having multiple copies of the same data. This is not applicable in quantum computing, due to the no-cloning theorem, which states that it is impossible to create an independent and identical copy of a qubit. Instead, the information content can be spread over several entangled qubits, thus creating a logical qubit out of several physical ones. To reach fault-tolerance, the ratio of physical to logical qubits is on the order of a thousand to one. Thus, present-day quantum computers have yet to reach the milestone of implementing *one* fully fault-tolerant qubit.

To reach LSQ, classical computing plays a large role. The importance of emulators was discussed previously, as was the prospect of quantum computers for accelerating machine learning. Conversely, machine learning can also speed up quantum computers. ML schemes have already been developed for calibration of qubits[Genois 2021] , error correction schemes [Nautrup 2019][Sweke 2021] , and algorithm optimisation [Cincio 2021][Fösel 2021] .

With the increasing complexity that comes with larger quantum computers, the role of ML/AI is expected to grow. It might, for instance become necessary to use supercomputing resources for optimising and compiling quantum algorithms for specific QPUs, resources that themselves could already be of the HPC+QC variety. The interplay between classical and quantum will continue for the foreseeable future.

The software development stack for quantum computers needs to evolve, in order to facilitate quantum programming. Without the software for solving specific classes and types of problems, quantum computers are rather useless. Ready-made algorithm and subroutine libraries are sorely needed, and significant resources should be dedicated for developing real-world scientific modelling packages that incorporate quantum algorithms. Quantum software development kits and ready-made application software needs to be made available to the end-users, alongside access to the hardware. This should be coupled to the introduction of HPC+QC solutions to R&D communities. HPC+QC deployment has already started at several European HPC centres: CINECA (Italy), CSC (Finland), ICHEC (Ireland), JSC (Germany), LRZ (Germany), SURF (The Netherlands), and TGCC (France).

User engagement and uptake of quantum computing is challenging. Especially in the near-term, when gains in computing power will be incremental rather than transformative, the barrier to adoption will be high. Even for users well versed in the art of writing software for traditional HPC systems, the leap to start developing quantum applications is a long one. From the perspective of an end-user comfortable with *using* HPC resources rather than developing them, QC can seem even more esoteric. Encouraging adoption at a stage when the technology readiness level of QC is low, can even be detrimental to the long-term prospects of quantum computing. It is, however, important to continuously build up the quantum software stack. Engaging end-users early on, even when the technology is far from being a polished product, enables a constant dialogue between those developing the technology and those that will take it into practical use. To ensure early real-world quantum advantage, engineering efforts need to align with user needs.

Large-scale quantum computers are often seen as special hardware that, like the mainframes of old, would exclusively be large-scale also by size. Already, this notion is being challenged. For example, Australian-German Quantum Brilliance is developing portable quantum accelerators, envisaged to be the size of the graphics card by 2026 [Doherty 2021] . Austrian AQT is already offering quantum computers in standard 19-inch rack format [Pogorelov 2021] . The most powerful QC solutions of the future will with all probability require tailor-made hosting spaces, however, as the classical supercomputers of today. The mere feasibility of consumer-grade QPUs does illustrate how far from the experimental lab-space QC has already ventured.

We are at the verge of witnessing whether QC will begin delivering on its promises in earnest. The public roadmaps of both IBM and Google foresee reaching a million qubits by 2030. More interestingly, IBM has disclosed its shorter-term goals, including reaching the 1.000-qubit milestone by the end of 2023, with steady increase in qubit count in-between now and then [Gambetta 2020] . Reaching the milestones will certainly be challenging, and at the same time, serve as a motivation and driving force for QC development in general. In Europe, the global IT player Atos is visibly investing in QC, and many talented teams are developing QC solutions, industrialised via various start-ups. As discussed in Section 7, the European challenge is to ensure the competitiveness of the largely start-up driven quantum ecosystem on the global scene. The EU Quantum Flagship project has a plan for global stimulation of the ecosystem, which augments and complements national efforts. Sustained, sufficiently long-term EU-level support for the QC industry will be needed to ensure that European competence and capacity in the field reaches the critical mass of self-sustainability.

12. Conclusions and Summary

Quantum computing has begun to deliver on its promise of revolutionising computational modelling and supercomputing. We are still at the very early stages of this new computing era, however. In order to become a usable tool for end-users, several unsolved challenges must be overcome.

From the hardware point-of-view, the most significant issues to be addressed are scaling up the qubit count and increasing the error-resilience of quantum computers. All bets are off when it comes to which specific technologies will emerge as champions in the race towards real quantum advantage. Several different physical implementations of qubits and the controlling machinery are being explored. Most likely, some technologies will turn out to be ideal for some types of calculations, while others are better at solving other types of problems.

Although the leading quantum computing architectures and systems of the future are unknown, what we *do* understand is that the future of quantum computing is hybrid. Not only in the sense of combining classical and quantum computing in an HPC+QC platform, but also when it comes to the implementation of the quantum part. We foresee a future where several quantum technologies that presently are developed rather independently, will ultimately merge: digital and analog quantum computing; superconducting, trapped ion, neutral atom and photonic qubits; qubits, qutrits, and continuous variable data representations; quantum annealers and universal quantum computers; task-specific and general-purpose hardware; the list goes on. A comparison to classical platforms is appropriate. In the 1970s and 1980s, many supercomputers were based on vector instructions. Around 1990, microprocessors based on scalar instructions took over, and pure vector machines were mostly abandoned. By the end of the 20th century, vector instructions were reintroduced, and the microprocessors of today are hybrids of both scalar and vector architectures.

As no single technological solution is expected to dominate even in the long-term, it is imperative to explore various options and to bet on as many European Schrödinger horses as possible. This can ensure a successful build-up of a homegrown and independent European quantum technology ecosystem. The risks associated with locking down on only a few approaches at this stage are enormous: Europe could create a future situation where neither the hardware nor the know-how of the most powerful supercomputer solutions can be found within the continent.

In addition to the hardware, also the software and user interfaces for the upcoming quantum systems need to be developed. Due to a very different programming model, quantum software development requires significant effort and support. In contrast to what has been the norm for some time now, it is not only a question of porting existing codes to quantum computers, but rather fully rewriting applications and changing our mindset to "Think Quantum". A multi-tiered, don't-place-all-your-eggs-in-one-basket approach needs to be taken that allows for technological progress without having to wait for which quantum technologies will be available or become dominant in the future. Europe must forge forward in the interim in a way that allows for the development of new algorithms exploiting computational features that have never before been used in programming. This requires high-level portable programming environments. In addition, the current computational problems themselves have to be recast to a form suitable for quantum algorithms. This requires an immediate and massive educational program and support for industrial adoption. Exploiting the synergies between academia and industry is far more important than in traditional HPC. The development of widely accepted standards for programming and tools is equally important. This will be a cornerstone of the future, and Europe can be pro-active in this area – Europe must not lag behind nor depend on closed solutions, and actively support open-source efforts.

For catalysing the uptake of quantum computing, incorporating QC into existing supercomputing infrastructure is essential. For real-world problems, quantum computers will never *replace* classical computers, but rather become an integral part of HPC. Europe has the opportunity to set up unique hybrid supercomputing infrastructures by seamlessly incorporating quantum technology, and by utilising the expertise of HPC centres built up over decades. Europe is well on track in setting up individual HPC+QC hybrid infrastructures. A federated user access approach, where a login to any of the EuroHPC supercomputers would give access to a large selection of quantum computing resources, resources will allow end users to find the most suitable technologies for their specific applications. This would provide European scientists an advantage over international competition.

Efforts should be pan-European, following the already proven success of the EuroHPC Joint Undertaking for setting up petascale and pre-exascale computing facilities, and the plans for going to exascale and beyond.

As with the traditional supercomputing infrastructure, diversity is key also when setting up a distributed and federated European QC backbone. Endeavours should be concerted, but not overly concentrated. Initially, setting up at least two, and ideally three European HPC+QC hubs would represent a good balance between joint effort and risk-mitigating diversity.

The turning point for science, when research *about* quantum computers shifts towards research *using* quantum computers draws closer. Still, we need to actively remind ourselves that quantum computing is in its infancy. It needs to be nurtured by a sustained and significant support for fundamental and basic research into quantum technologies, both on the hardware and software side. Only then can European competence and competitiveness in the fields of quantum technology in general, and quantum computing in particular be ensured. The basic research on quantum technologies performed for decades in partnership between all European countries, whether members of the European Union or not, has laid the seeds for the fruits we reap today. In order to push these boundaries further, this collaboration needs to continue, with minimum restriction.

IBM and Google aim for a million qubits by 2030. Europe too needs to set its goals for quantum computing sufficiently high to keep up with global developments. Maintaining an open and inclusive spirit of international cooperation is what keeps Europe an appealing target for international investments and talent also in the future. To reach the number of qubits needed for truly disruptive quantum computing, the roughly one hundred noisy qubits that we have today have to be scaled up by four to five orders of magnitude. Achieving this feat requires dedication, sustained funding, and a worldwide pooling of talent and technology.

13. References

[Aharonov 2007] D. Aharonov *et al.*, "Adiabatic Quantum Computation is Equivalent to Standard Quantum Computation", *SIAM J. Comput.* **37** (2007) 166. https://doi.org/10.1137/S0097539705447323

[ANSSI 2020] "Should Quantum Key Distribution be Used for Secure Communications?", Agence nationale de la sécurité des systèmes d'information (2020). https://www.ssi.gouv.fr/en/publication/should-quantum-key-distribution-be-used-for-secure-communications/

[AQT] Alpine Quantum Technologies. https://www.aqt.eu/

[Arute 2019] F. Arute *et al*., "Quantum supremacy using a programmable superconducting processor", *Nature* **574** (2019) 505. https://doi.org/10.1038/s41586-019-1666-5

[Babbush 2014] R. Babbush, P. Love, A. Aspuru-Guzik, "Adiabatic Quantum Simulation of Quantum Chemistry", *Sci. Rep.* **4** (2014) 660. https://doi.org/10.1038/srep06603

[Barzen 2020] J. Barzen and F. Leymann, "Quantum humanities: a vision for quantum computing in digital humanities", *ICS Softw.-Inensiv. Cyber-Phys. Syst.* **35** (2020) 153. https://doi.org/10.1007/s00450-019-00419-4

[Bertels 2021] K. Bertels *et al.*, "Quantum Accelerator Stack: A Research Roadmap", *arXiv* 2102:02035 (2021). https://arxiv.org/abs/2102.02035

[Bobier 2021] J.-F. Bobier *et al.*, "What Happens When 'If' Turns to 'When' in Quantum Computing?", *Boston Consulting Group* report (2021). https://www.bcg.com/publications/2021/building-quantum-advantage

[Britt 2017] K.A. Britt and T.S. Humble, "High-Performance Computing with Quantum Processing Units", *J. Emerg. Technol. Comput. Syst.* **13** (2017) 39. https://doi.org/10.1145/3007651

[Cao 2019] Y. Cao *et al.*, "Quantum Chemistry in the Age of Quantum Computing", *Chem. Rev.* **119** (2019) 10856. https://doi.org/10.1021/acs.chemrev.8b00803

[Camps 2021] D. Camps, R. Van Beeumen, C. Yang, "Quantum Fourier Transform Revisited", *Numer. Linear Algebra Appl.* **28** (2021) e2331. https://doi.org/10.1002/nla.2331

[Castelvecchi 2018]. D. Castelvecchi, "The quantum internet has arrived (and it hasn't)", *Nature* **554** (2018) 289. https://doi.org/10.1038/d41586-018-01835-3

[Cerezo 2021] M. Cerezo *et al.*, "Variational quantum algorithms", *Nat. Rev. Phys.* **3** (2021) 625. https://doi.org/10.1038/s42254-021-00348-9

[Chiribella 2021] G. Chiribella, Y. Yang, R. Renner, "Fundamental Energy Requirement of Reversible Quantum Operations", *Phys. Rev. X* **11** (2021) 021014. https://doi.org/10.1103/PhysRevX.11.021014

[Cho 2020] A. Cho, "The biggest flipping challenge in quantum computing", *Science* news feature. https://doi.org/10.1126/science.abd7332

[Cincio 2021] L. Cincio *et al.*, "Machine Learning of Noise-Resilient Quantum Circuits", *PRX Quantum* **2** (2021) 010324. https://doi.org/10.1103/PRXQuantum.2.010324

[D-Wave] https://www.dwavesys.com/

[D-Wave 2017] "Computational Power Consumption and Speedup", *D-Wave White Paper Series* 14-1005A-D (2017) https://www.dwavesys.com/media/ivelyjij/14-1005a_d_wp_computational_power_consumption_and_speedup.pdf

[Daiss 2021] S. Daiss *et al.*, "A quantum-logic gate between distant quantum-network modules", *Science* **371** (2021) 614. https://doi.org/10.1126/science.abe3150

[Dalyac 2021] C. Dalyac *et al.*, "Qualifying quantum approaches for hard industrial optimization problems. A case study in the field of smart-charging of electric vehicles", *EPJ Quantum Technol.* **8** (2021) 12. https://doi.org/10.1140/epjqt/s40507-021-00100-3

[Dervovic 2018] D. Dervovic *et al.*, "Quantum linear systems algorithms: a primer", *arXiv* 1802.08227 (2018) https://arxiv.org/abs/1802.08227

[Deutsch 1985] D. Deutsch, "Quantum theory, the Church–Turing principle and the universal quantum computer", *Proc. R. Soc. Lond. A* **400** (1985) 97. http://doi.org/10.1098/rspa.1985.0070

[Devitt 2013] S.J, Devitt, W.J. Munro, K. Nemoto, "Quantum error correction for beginners." *Rep. Prog. Phys.***76** (2013) 076001. https://doi.org/10.1088/0034-4885/76/7/076001

[DiVincenzo 2000] D.P. DiVincenzo, "The physical implementation of quantum computation." *Fortschr. Phys.* **48** (2000) 771-783. DOI: https://doi.org/10.1002/1521-3978(200009)48:9/11%3C771::AID-PROP771%3E3.0.CO;2-E

[Doherty 2021] M. Doherty, "Quantum accelerators: a new trajectory of quantum computers", *Digitale Welt* **5** (2021) 74. https://doi.org/10.1007/s42354-021-0342-8

[Dunjko 2020] V. Dunjko, P. Wittek, "A non-review of Quantum Machine Learning: trends and explorations", *Quantum Views* **4** (2020) 32. https://doi.org/10.22331/qv-2020-03-17-32

[Emani 2021] P.S. Emani *et al.*, "Quantum computing at the frontiers of biological sciences", *Nat. Methods* **18** (2021) 701. https://doi.org/10.1038/s41592-020-01004-3

[ENISA 2021] "Post-Quantum Cryptography: Current state and quantum mitigation", European Union Agency for Cybersecurity (2021). https://www.enisa.europa.eu/publications/post-quantum-cryptography-current-state-and-quantum-mitigation

[ETHZ 2021] "ETH Zurich and PSI found Quantum Computing Hub", ETH Zürich press release (2021). https://ethz.ch/en/news-and-events/eth-news/news/2021/05/eth-zurich-and-psi-found-quantum-computing-hub.html

[EuroCC] https://www.eurocc-access.eu/

[EuroQCI] The European Quantum Communication Infrastructure (EuroQCI) Initiative. https://digital-strategy.ec.europa.eu/en/policies/european-quantum-communication-infrastructure-euroqci

[ESA 2021] "Quantum communication in space moves ahead", *European Space Agency*. https://www.esa.int/Applications/Telecommunications_Integrated_Applications/Quantum_communication_in_space_moves_ahead

[Fu 2016] X. Fu *et al.*, "A Heterogeneous Quantum Computer Architecture", *Proceedings of the ACM International Conference on Computing Frontiers* (2016). https://doi.org/10.1145/2903150.2906827

[Fösel 2021] T. Fösel *et al.*, "Quantum circuit optimization with deep reinforcement learning", *arXiv* 2103.07585 (2021). https://arxiv.org/abs/2103.07585

[Gambetta 2020] J. Gambetta, "IBM's roadmap for scaling quantum technology", *IBM Research* blog (2020). https://www.ibm.com/blogs/research/2020/09/ibm-quantum-roadmap

[Gehring 2021]. T. Gehring *et al.*, "Homodyne-based quantum random number generator at 2.9 Gbps secure against quantum side-information", *Nat. Commun.* **12** (2021) 605. https://doi.org/10.1038/s41467-020-20813-w

[Genois 2021] E. Genois *et al.*, "Quantum-tailored machine-learning characterization of a superconducting qubit", *arXiv* 2106.13126 (2021). https://arxiv.org/abs/2106.13126

[Georgescu 2014] I.M. Georgescu, S. Ashhab, F. Nori, "Quantum simulation", *Rev. Mod. Phys.* **86** (2014) 153. https://doi.org/10.1103/RevModPhys.86.153

[Gisin 2007] N. Gisin, R. Thew, "Quantum communication", *Nature Photon.* **1** (2007) 165. https://doi.org/10.1038/nphoton.2007.22

[Google 2021] Google Quantum AI, "Exponential suppression of bit or phase errors with cyclic error correction", *Nature* **595** (2021) 383. https://doi.org/10.1038/s41586-021-03588-y

[Giri 2017] P.R. Giri, V.E. Korepin, "A review on quantum search algorithms", *Quantum Inf. Process.* **16** (2017) 315. https://doi.org/10.1007/s11128-017-1768-7

[Greganti 2021] C. Greganti *et al.*, "Cross-Verification of Independent Quantum Devices", *Phys. Rev. X* **11** (2021) 031049. https://doi.org/10.1103/PhysRevX.11.031049

[Harrow 2009] A.W. Harrow, A. Hassidim, S. Lloyd, "Quantum Algorithm for Linear Systems of Equations", *Phys. Rev. Lett.* **103** (2009) 150502. https://doi.org/10.1103/PhysRevLett.103.150502

[Hazan 2020] E. Hazan *et al.*, "The next tech revolution: quantum computing", *McKinsey & Company* report (2020). https://www.mckinsey.com/fr/our-insights/the-next-tech-revolution-quantum-computing

[Henriet 2020] L. Henriet *et al.*, "Quantum computing with neutral atoms", *Quantum* **4** (2020) 327. https://doi.org/10.22331/q-2020-09-21-327

[Hillmann 2020] T. Hillmann *et al*., “Universal Gate Set for Continuous-Variable Quantum Computation with Microwave Circuits”, *Phys. Rev. Lett.* **125** (2020) 160501. https://doi.org/10.1103/PhysRevLett.125.160501

[Huang 2021] H.-Y. Huang *et al.*, “Power of data in quantum machine learning”, *Nat. Commun.* **12** (2021) 2631. https://doi.org/10.1038/s41467-021-22539-9

[Hunger 2021] D. Hunger, “Quantum logic at a distance”, *Science* **371** (2021) 576. https://doi.org/10.1126/science.abg1536

[Ikonen 2017] J. Ikonen, J. Salmilehto, M. Möttönen, “Energy-efficient quantum computing”, *npj Quantum Inf.* **3** (2017) 17. https://doi.org/10.1038/s41534-017-0015-5

[IQM] IQM Quantum Computers. https://www.meetiqm.com/

[IQT 2020] “Quantum Computing: A Seven-year Market Forecast”, *Inside Quantum Technology* report IQT-QCM-1020 (2020). https://www.insidequantumtechnology.com/product/quantum-computing-a-seven-year-market-forecast/

[Johnston 2019] E.R. Johnston, N. Harrigan, M. Gimeno-Segovia, “Programming quantum computers: essential algorithms and code samples”, O’Reilly Media (2019).

[Jordan 2021] S. Jordan, “Quantum Algorithm Zoo”, https://quantumalgorithmzoo.org

[Kung 2021] J. Kung, M. Fancy, “A Quantum Revolution: Report on Global Policies for Quantum Technology”, *CIFAR* report (2021). https://cifar.ca/cifarnews/2021/04/07/a-quantum-revolution-report-on-global-policies-for-quantum-technology/

[Loesekrug-Pietri 2021] A. Loesekrug-Pietri, “It’s not too late for Europe to lead the post-quantum cryptography race”, *Sifted* (2021). https://sifted.eu/articles/europe-post-quantum-cryptography/

[McArdle 2020] S. McArdle, S- Endo, A. Aspuru-Guzik, S.C. Benjamin, X Yuan, “Quantum computational chemistry”, *Rev. Mod. Phys.* **92** (2020) 015003. https://doi.org/10.1103/RevModPhys.92.015003

[MacQuarrie 2020] E.R. MacQuarrie *et al.*, “The emerging commercial landscape of quantum computing”, *Nat. Rev. Phys.* **2** (2020) 596. https://doi.org/10.1038/s42254-020-00247-5

[Merali 2011] Z. Merali, “First sale for quantum computing”, *Nature* **474** (2011) 18. https://doi.org/10.1038/474018a

[Mirrahimi 2014] M. Mirrahimi *et al.*, “Dynamically protected cat-qubits: a new paradigm for universal quantum computation”, *New J. Phys.* **16** (2014) 045014. https://doi.org/10.1088/1367-2630/16/4/045014

[Montanaro 2016] A. Montanaro, “Quantum algorithms: an overview”, *npj Quantum Inf.* **2** (2016) 15023. https://doi.org/10.1038/npjqi.2015.23

[Nautrup 2019] H.P. Nautrup *et al.*, “Optimizing Quantum Error Correction Codes with Reinforcement Learning”, *Quantum* **3** (2019) 215. https://doi.org/10.22331/q-2019-12-16-215

[NCSC 2020] “Quantum security technologies”, *National Cyber Security Centre* white paper (2020). https://www.ncsc.gov.uk/whitepaper/quantum-security-technologies

[OpenQKD] https://openqkd.eu/

[Parra-Rodriguez 2020] A. Parra-Rodriguez *et al.*, “Digital-Analog Quantum Computation”. *Phys. Rev. A* **101** (2020) 022305. https://doi.org/10.1103/PhysRevA.101.022305

[Pasqal] https://pasqal.io/

[Pirandola 2015] S. Pirandola *et al.*, “Advances in quantum teleportation”, *Nature Photon.* **9** (2015) 641. https://doi.org/10.1038/nphoton.2015.154

[Pogorelov 2021] I. Pogorelov *et al.*, “Compact Ion-Trap Quantum Computing Demonstrator”, *PRX Quantum* **2** (2021) 020343. https://doi.org/10.1103/PRXQuantum.2.020343

[PRACE] PRACE Training Portal. https://training.prace-ri.eu/

[Preskill, 2018] J. Preskill. "Quantum computing in the NISQ era and beyond." *Quantum* **2** (2018) 79. https://doi.org/10.22331/q-2018-08-06-79

[PQC.FI] Post-Quantum Cryptography Finland. https://www.pqc.fi/

[QIA] Quantum Internet Alliance. https://quantum-internet.team/

[QLM] Atos Quantum Learning Machine. https://atos.net/en/solutions/quantum-learning-machine

[QT.EU] European Quantum Flagship, https://qt.eu/about-quantum-flagship/projects/

[QuIC] The European Quantum Industry Consortium. https://quantumindustry.eu

[QUTAC] QUTAC - Quantum Technology & Application Consortium. https://www.qutac.de/

[Räsänen 2021] M. Räsänen *et al.*, "Path to European quantum unicorns", *EPJ Quantum Technol.* **8** (2021) 5. https://doi.org/10.1140/epjqt/s40507-021-00095-x

[Shor 1994] P.W. Shor, "Algorithms for quantum computation: discrete logarithms and factoring", *Proceedings 35th Annual Symposium on Foundations of Computer Science* (1994) pp. 124-134. https://doi.org/10.1109/SFCS.1994.365700

[Sweke 2021] R. Sweke *et al.*, "Reinforcement learning decoders for fault-tolerant quantum computation", *Mach. Learn.: Sci. Technol.* **2** (2021) 025005. https://doi.org/10.1088/2632-2153/abc609

[Temkin 2021] M. Temkin, "Investors bet on the technologically unproven field of quantum computing", *Pitchbook* news & analysis (2021). https://pitchbook.com/news/articles/quantum-computing-venture-capital-funding

[Terhal 2015] B.M. Terhal, "Quantum error correction for quantum memories", *Rev. Mod. Phys.* **87** (2015) 307. https://doi.org/10.1103/RevModPhys.87.307

[Uvarov 2020] A.V. Uvarov, A.S. Kardashin, J.D. Biamonte, "Machine learning phase transitions with a quantum processor", *Phys. Rev. A* **102** (2020) 012415. https://doi.org/10.1103/PhysRevA.102.012415

[Van Meter 2013] R. Van Meter, C. Horsman, "A Blueprint for Building a Quantum Computer", *Commun. ACM* **56** (2013) 84. https://doi.org/10.1145/2494568

[Venegas-Andraca 2012] S.E. Venegas-Andraca. "Quantum walks: a comprehensive review", *Quantum Inf. Process.* **11** (2012) 1015–1106. https://doi.org/10.1007/s11128-012-0432-5

[Villalonga 2020] B. Villalonga *et al.*, "Establishing the quantum supremacy frontier with a 281 Pflop/s simulation", *Quantum Sci. Technol.* **5** (2020) 034003. https://doi.org/10.1088/2058-9565/ab7eeb

[WACQT] Wallenberg Centre for Quantum Technology. https://wacqt.se

[Wang 2021] P. Wang *et al.*, "Single ion qubit with estimated coherence time exceeding one hour", *Nat. Commun.* **12** (2021) 233. https://doi.org/10.1038/s41467-020-20330-w

[Wilen 2021] C.D. Wilen *et al.*, "Correlated charge noise and relaxation errors in superconducting qubits", *Nature* **594** (2021) 369. DOI: https://doi.org/10.1038/s41586-021-03557-5

[Wu 2021] Y. Wu *et al.* "Strong quantum computational advantage using a superconducting quantum processor", *arXiv* 2106.14734 (2021). https://arxiv.org/abs/2106.14734

[Zhong 2020] H.-S. Zhong *et al.* "Quantum computational advantage using photons", *Science* **370** (2020) 1460. https://dx.doi.org/10.1126/science.abe8770

14. List of acronyms

AI	Artificial Intelligence
CPU	Central Processing Unit
EU	European Union
GPU	Graphics Processing Unit
HPC	High Performance Computing
HPC+QC	(Hybrid) High Performance Computing and Quantum Computing
I/O	Input/Output
IP	Intellectual Property
kW	Kilowatt
LSQ	Large-Scale Quantum
ML	Machine Learning
MPP	Massively Parallel Processing
NISQ	Noisy Intermediate-Scale Quantum
OS	Operating System
PQC	Post-Quantum Cryptography
QC	Quantum Computing
QEC	Quantum Error Correction
QFT	Quantum Fourier Transform
QKD	Quantum Key Distribution
QLM	Quantum Learning Machine
QML	Quantum Machine Learning
QUBO	Quadratic Unconstrained Binary Optimisation
QPU	Quantum Processing Unit
QRNG	Quantum Random Number Generator
SME	Small and Medium-sized Enterprise
TRL	Technology Readiness Level

Acknowledgements

This work was financially supported by the PRACE project funded in part by the EU's Horizon 2020 Research and Innovation programme (2014-2020) under grant agreement 823767.