

# Towards a distributed quantum computing ecosystem

eISSN 2632-8925  
 Received on 13th January 2020  
 Revised 7th May 2020  
 Accepted on 11th May 2020  
 E-First on 10th July 2020  
 doi: 10.1049/iet-qtc.2020.0002  
 www.ietdl.org

Daniele Cuomo<sup>1</sup> ✉, Marcello Caleffi<sup>1,2</sup>, Angela Sara Cacciapuoti<sup>1,2</sup>

<sup>1</sup>FLY: Future communications Laboratory, Department of Electrical Engineering and Information Technology, University of Naples Federico II, Naples 80125, Italy

<sup>2</sup>Laboratorio Nazionale di Comunicazioni Multimediali, National Inter-University Consortium for Telecommunications (CNIT), Naples 80126, Italy

✉ E-mail: daniele.cuomo@unina.it

**Abstract:** The Quantum Internet, by enabling quantum communications among remote quantum nodes, is a network capable of supporting functionalities with no direct counterpart in the classical world. Indeed, with the network and communications functionalities provided by the Quantum Internet, remote quantum devices can communicate and cooperate for solving challenging computational tasks by adopting a distributed computing approach. The aim of this study is to provide the reader with an overview about the main challenges and open problems arising in the design of a distributed quantum computing ecosystem. For this, the authors provide a survey, following a bottom-up approach, from a communications engineering perspective. They start by introducing the Quantum Internet as the fundamental underlying infrastructure of the distributed quantum computing ecosystem. Then they go further, by elaborating on a high-level system abstraction of the distributed quantum computing ecosystem. They describe such an abstraction through a set of logical layers. Thereby, they clarify dependencies among the aforementioned layers and, at the same time, a road-map emerges.

## 1 Introduction

Nowadays, a tremendous amount of heterogeneous players entered the quantum race, ranging from tech giants – such as IBM and Google in fierce competition to build a commercial quantum computer – to states and governments, with massive public funds to be distributed over the next years [1–4].

In 2017, the European Commission launched a €1-billion flagship program to support the quantum research for ten years starting from 2018, and it provided a first €132-million tranche during the following three years [5]. In 2018, the USA launched the National Quantum Initiative funded with \$1.2-billion over ten years and China is keeping up, investing billions to commercialise quantum technologies [6].

These huge efforts are justified by the disruptive potential of a quantum computer, beyond anything classical computers could ever achieve. Indeed, by exploiting the rules of quantum mechanics, a quantum computer can tackle classes of problems that choke conventional machines. These problems include chemical reaction simulations, optimisation in manufacturing and supply chains, financial modelling, machine learning and enhanced security [7–9]. Hence, the quantum computing has the potential to completely change markets and industries.

At the end of 2019, Google achieved the so-called *quantum supremacy* (the term was coined by Preskill in 2011 [10] to describe the moment when a programmable quantum device would solve a problem that cannot be solved by classical computers, regardless of the usefulness of the problem [4]) with a 54-qubits quantum processor, named *Sycamore*, by sampling from the output distribution of 53-qubits random quantum circuits [11]. By neglecting some performance-enhancing techniques as pointed out by IBM [12, 13], Google estimated that ‘a state-of-the-art supercomputer would require approximately 10,000 years to perform the equivalent task’ that required just 200 seconds on Sycamore.


By ignoring the noise effects and by coarsely oversimplifying, the computing power of a quantum computer scales exponentially with the number of quantum bits (qubits) that can be embedded and interconnected within [1, 2]. One of the reasons lays in a principle of quantum mechanics known as *superposition principle*.

Specifically, a classical bit encodes one of two mutually exclusive states – usually denoted as 0 and 1 – being in only one state at a certain time. Conversely, a qubit can be in an extra mode – called *superposition* – i.e. it can be in a combination of the two basic states [2, 3].

To give a flavour of the above, let us consider one of the killer applications of the quantum computing: chemical reaction simulation [14]. As highlighted in [15], the amount of information needed to fully describe the energy configurations of a relatively simple molecule such as caffeine is astoundingly large:  $10^{48}$  bits. For comparison, the estimated number of atoms on Earth is between  $10^{49}$  and  $10^{50}$  bits. Hence, describing the energy configuration of caffeine at one single instant needs roughly a number of bits comparable to 1 to 10 per cent of all the atoms on the planet. But this energy configuration description becomes suddenly feasible with a quantum processor embedding roughly 160 *noiseless* qubits, thanks to the superposition principle.

Unfortunately, qubits are very fragile and easily modified by interactions with the outside world, via a noise process known as *decoherence* [2, 3, 16]. Indeed, decoherence is not the only source of errors in quantum computing. Errors practically arise with any operation on a quantum state. However, isolating the qubits from the surrounding is not the solution, since the qubits must be manipulated to fulfil the communication and computing needs, such as reading/writing operations. Moreover, the challenges for controlling and preserving the quantum information encoded in a single qubit get harder as the number of qubits within a single device increases, due to coupling effects. In this regard, *quantum error correction* (QEC) represents a fundamental tool for protecting quantum information from noise and faulty operations [3, 17, 18]. However, QEC operates by spreading the information of one *logical* qubit into several physical qubits. Hence, solving problems of practical interest, such as integer factorisation – which constitutes one of the most widely adopted algorithms for securing communications over the classical Internet – or molecule design may require millions of physical qubits [6, 7].

Hence, on one hand researchers worldwide are leveraging on the advancement of different technologies for qubit implementation –superconducting circuits, ion traps, quantum dots and diamond vacancies among the others – and innovative QEC techniques to



**Fig. 1** Distributed quantum computing speed-up. The volume of cubes represents the ideal quantum computing power, i.e. in absence of noise and errors. As evident by comparing the power available at isolated devices versus the power achievable through clustered devices, interconnecting quantum processors via the *Quantum Internet* provides an exponential computing speed-up with respect to isolated devices

scale the number of qubits beyond two-digits. On the other hand, the *Quantum Internet*, i.e. a network enabling quantum communications among remote quantum nodes, has been recently proposed as the key strategy to significantly scale up the number of qubits [1, 2, 19–22].

In fact, the availability of such a network and the adoption of a distributed computing paradigm allows us to regard the Quantum Internet – jointly – as a virtual quantum computer with a number of qubits that scales linearly with the number of interconnected devices.

In this light, the aim of this paper is to provide the reader with an overview about the main challenges and open problems arising in the design of a *distributed quantum computing ecosystem*.

We start in Section 2 by introducing the Quantum Internet – as the fundamental underlying communication infrastructure of a distributed quantum computing ecosystem – as well as some of its unique key applications. Then we go further in Section 3, by conceptualising a high-level system abstraction of the distributed quantum computing ecosystem from a communication engineering perspective. We describe such an abstraction through a set of (logical) layers, with the higher depending on the functionalities provided by the lower ones. Thereby, we clarify dependencies among the aforementioned layers. Since, each layer of the ecosystem has some related open challenges, within Section 4 we survey such challenges and open problems. Finally, in Section 5 we conclude the paper with some perspectives.

## 2 Quantum Internet

As mentioned in Section 1, one promising approach to address the challenges arising in the realisation of large-scale quantum processors is to mimic modern high-performance computing infrastructures – where thousands of processors, memories and storage units are inter-connected via a communication network, and the computational tasks are solved by adopting a distributed approach.

To this aim, it is mandatory to design and deploy the *Quantum Internet*, which formally defines a global quantum network (we refer the reader to [23, 24] for a discussion about the differences underlying the notion of ‘*Quantum Internet*’ versus ‘*quantum network*’) able to transmit qubits and to distribute entangled quantum states (the deepest difference between classical and quantum mechanics lays in the concept of *quantum entanglement*, a sort of correlation with no counterpart in the classical world. For an in-depth introduction to quantum entanglement, we refer the reader to the classical book [17], whereas we refer the reader to [3] for a concise description) among remote quantum devices [1–3, 19–21, 23].

In fact, the availability of the corresponding underlying network infrastructure and the adoption of the distributed computing paradigm [25] allows us to regard the Quantum Internet – jointly – as a virtual quantum computer with a number of qubits that scales linearly with the number of interconnected devices. Hence, the Quantum Internet may enable an exponential speed-up [1, 25, 26] of the quantum computing power with just a linear amount of physical resources, represented by the interconnected quantum processors. Indeed, by comparing the computing power achievable with quantum devices working independently versus working as a unique quantum cluster, the gap comes out – as depicted in Fig. 1 [1].

Specifically, increasing the number of isolated devices lays to a linear speed-up, with a double growth in computational power by doubling the number of devices. Conversely, increasing the number of clustered devices provides an exponential growth, with a significant advantage clearly visible with just two interconnected devices. For instance, a single 10-qubit processor can represent  $2^{10}$  states thanks to the superposition principle, hence two isolated 10-qubit processors can represent  $2^{11}$  states. But if we *interconnect* the two processors, the resulting virtual device can represent up to  $2^{18}$  states [1], depending on the number of qubits devoted to fulfil the communication needs of the clustered processors as discussed in Section 4.2.

Before analysing with further details in Section 3 the resulting distributed quantum computing ecosystem from a communication engineering perspective, it is worthwhile to note that the Quantum Internet infrastructure enables unparalleled capabilities not restricted to the distributed computing [2, 21]. Specifically, applications such as *blind computing*, *secure communications* and *noiseless communications* have already been theorised or even experimentally verified, as recently overviewed by an Internet Engineering Task Force (IETF) Quantum Internet Draft [23].

Blind quantum computing [27–30] refers to a server-client architecture where clients can send sensitive data to server, which elaborates inputs without knowing their values. This functionality allows to achieve a twofold goal: preserving data confidentiality as well as solving tasks that are intractable for the client – that can be a classical computer – but tractable for the server, which implements the quantum paradigm.

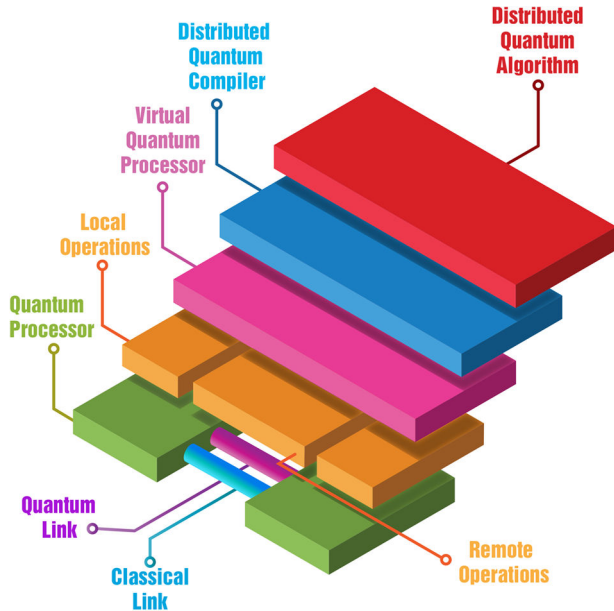
Secure communications in the quantum field refer to the class of communication protocols that exploit quantum mechanics in order to get benefits unattainable using classical Internet. For instance, in the field of quantum cryptography, researchers study strategies for sharing keys among parties in total secrecy [31–33]. Whilst quantum byzantine agreement, a protocol used by multiple entities to distributively agree on a common decision, allows to achieve the consensus in a constant number of rounds, whereas classical protocols scale polynomially with the number of processors [34].

Finally, the Quantum Internet provides the underlying infrastructure to achieve transmission rates exceeding the fundamental limits of conventional (quantum) Shannon theory [35, 36]. Specifically, by exploiting the capability of quantum particles to propagate simultaneously among multiple space-time trajectories, quantum superpositions of noisy channels can behave as perfect noiseless quantum communication channels, even if no quantum information can be successfully transmitted throughout either of the noisy component channels individually [37].


## 3 Distributed quantum computing ecosystem

The overall aim of classical distributed computing is to deal with hard computational problems by splitting out the computational tasks among several classical devices, in order to lighten the loads on single devices.

As mentioned in Section 2, with the network infrastructure provided by the Quantum Internet, this paradigm can be extended to quantum computing as well: remote quantum devices can communicate and cooperate for solving computational tasks by adopting a distributed computing approach [25, 26]. However, this extension is not trivial, since the quantum world is characterised by unconventional phenomena [1], such as *no-cloning* and



**Fig. 2** High-level system abstraction of the distributed quantum computing ecosystem. The lowest layer provides the communication/network functionalities and consists of quantum processors interconnected with both classical and quantum links. Thanks to the underlying communication infrastructure, both local and remote qubit operations can be executed. Hence, from a computing perspective, the two lowest levels concur to build a virtual quantum processor with a number of qubits that scales with the number of inter-connected physical quantum processors. The virtual processor acts as an interface for the distributed quantum compiler, which maps a quantum algorithm into a sequence of local and remote operations, hence optimising the available computing resources with respect to both the hardware and the network constraints



**Fig. 3** Coupling map of the *ibmqx3* architecture: the nodes represent the qubits while the edges represent the possibility to have interactions between two qubits, i.e. to implement the *CNOT* operation. As an instance, a *CNOT* operation can be directly executed between qubits  $q_1$  and  $q_2$  but not between qubits  $q_1$  and  $q_3$

entanglement. As a consequence, a new ecosystem must be engineered.

The infographic in Fig. 2 is a stack depicting dependencies among a possible set of layers that together provide the *distributed quantum computing ecosystem*. For the sake of clarity, we restrict our attention to a network composed by two quantum processors, directly inter-connected. Nevertheless, we can easily extend the discussion to more complex network topologies, assuming that end-to-end routing and network functionalities [38–43] are available.

Starting from bottom, in Fig. 2 we have the communication infrastructure underlying the Quantum Internet: spatially remote quantum devices able to share quantum information by means of a synergy of both classical and quantum communication resources. Indeed, as we overview in Section 4.2, the transmission of quantum information generally requires the exchange of classical information as well, hence it requires the availability of a classical network infrastructure [2, 3, 21] such as the classical Internet. This constraint has been highlighted in Fig. 2 by interconnecting the two quantum processors with both a classical and a quantum link.

By exploiting the communication functionalities provided by the lowest level, both local and remote qubit operations can be executed. Specifically, the local operations – i.e. operations between qubits stored within the same quantum processor – can be executed by exploiting the physical (or logical, whether the

quantum device should natively implement QEC functionalities [44]) controls and readouts functionalities provided by the device. Conversely, the remote operations – i.e. the operations between qubits stored at different quantum devices – pose further constraints, as discussed in Section 4.

Thanks to the abstraction provided by the two layers residing at the very bottom, we obtain a *virtual quantum processor*, where remote qubits are interconnected through virtual connections made possible by the remote operations. Clearly, remote operations likely suffer delays and error rates higher than local operations. Hence, one should prefer local operations over remote ones as much as possible, even though remote operations are unavoidable whenever the number of qubits required to perform the computational task exceeds the number of qubits available at a single device. In this light, an optimisation must be performed by the *distributed quantum compiler*, so that the different operations required by the quantum algorithm are properly allocated among the qubits of the different devices.

Finally, at the very top we have the quantum algorithm, which is completely independent and unaware of the physical/logical constraints imposed by both the hardware and network particulars, thanks to the abstraction provided by the underlying levels. We can think at this module as a minimal service for defining quantum algorithms, but also, in a wider perspective, as a platform where interesting functionalities are available – allowing, for instance, quantum machine learning [45] or quantum optimisation algorithms [46, 47].


However, we omitted several complications so far. Indeed, communication protocols intrinsically imply overhead, i.e. dedicating computing resources to deal with transmission processes and errors correction. Therefore, it is worth going further towards this discussion, expanding components of the proposed ecosystem and considering open challenges related with.

## 4 Open challenges ahead

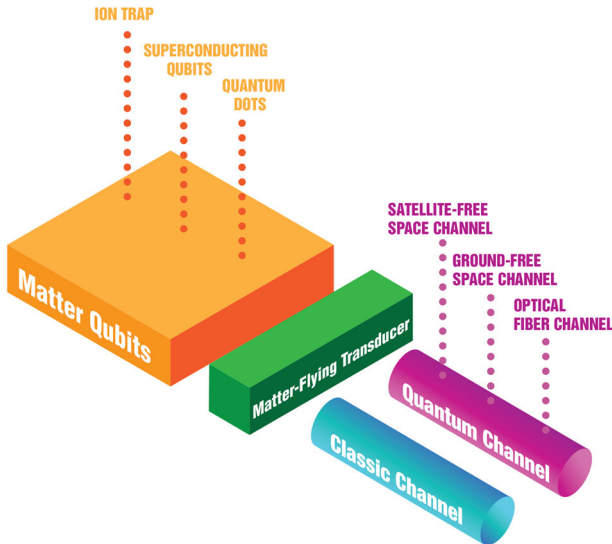
The aim of this section is to describe and to discuss some of the open problems related with the proposed distributed quantum computing ecosystem. For this, some layers are individually discussed. Specifically, in Section 4.1 we consider quantum processors, paying particular attention to drawbacks induced by local operations involving multiple qubits. In Section 4.2, we introduce the interconnection between remote quantum processors, by discussing the *quantum teleportation* as a mean to transfer quantum information between interconnected devices. After that, in Section 4.3 we discuss the gate teleportation as an alternative strategy to perform remote operations. In Section 4.4, we describe the layer responsible for abstracting and optimising the execution of quantum algorithms, based on the characteristics of the underlying system. Finally, in Section 4.4 we discuss some of the current standardisation efforts.

### 4.1 Quantum processor


The most basic element of a distributed quantum computing ecosystem can be identified with a single quantum computing device. Here the qubits are connected according to some directed and connected graph – namely, the coupling map – that accounts for the hardware limitations resulting from controlling and preserving the quantum information from decoherence and noise. As an example, Fig. 3 depicts the coupling map of the *ibmqx3* quantum processor [48], with nodes representing qubits while edges represent the possibility to have interactions between two qubits, i.e. to implement one of the fundamental quantum operations: the *CNOT* operation (the *CNOT* operation involves two qubits, referred to as control qubit and target qubit. It works as follows [49]: if the control qubit is 1, then the target qubit value is flipped. Otherwise, nothing happens). In fact, there exists a universal quantum gate set (a universal gate set is a set of gates that can implement any possible quantum operation [50]) in which the *CNOT* is the only operator involving more than one qubit. Thus, we can focus on problems related to the *CNOT* operation keeping the discourse general.



**Fig. 4** *CNOT operation between a couple of non-adjacent qubits can be implemented through a sequence of swapping operations, with each swap consisting of three CNOT (with the in-between CNOT being reverse, i.e. with target and control qubits swapped) between adjacent qubits [48]. Thus, the circuit performs a CNOT between  $q_1$  and  $q_3$ , leaving  $q_2$  unaltered. Note that  $H$  denotes the Hadamard gate mapping a basis state into an even superposition of the basis states [49], and that the quantum state stored in  $q_i$  is denoted as  $|\psi_{q_i}\rangle$ , by adopting the standard bra-ket notation*



**Fig. 5** *Pictorial representation of a matter-flying transducer, which converts matter qubits – i.e. qubits for information processing/storing – into flying qubits – i.e. qubits for information transmission – and vice versa*



**Fig. 6** *Quantum teleportation circuit. First two wires belong to source, whereas the bottom wire belongs to destination. A generic state  $|\psi\rangle$  is initially stored at the source. Once the teleportation process is completed, the original state is available at the destination, regardless of its value.  $|\Phi^+\rangle$  represents an EPR pair, that is a couple of maximally entangled qubits [55]. The result of the measurement process at the source is transmitted to the destination via a classical link. The carried classical bits are thus used for determining whether gates  $X$  and  $Z$  – corresponding to a bit- and a phase-flip, respectively [49] – must be performed to recover the original state  $|\psi\rangle$  from the EPR pair member available at the destination*

It is immediate to observe that only nodes – i.e. qubits – linked by an edge can directly interact [51]. Nevertheless, for an algorithm designer it is crucial being able to define a circuit without restrictions on interactions. Indeed, a circuit programming model can easily abstract from this restriction [17], resulting in a fully connected graph at the cost of an overhead due to indirect execution of the desired operation. For instance, let us consider to carry out a CNOT between the two non-adjacent qubits  $q_1$  and  $q_3$  of Fig. 3. By performing two state swapping operations for each edge

belonging to the shortest-path from node  $q_1$  towards node  $q_3$ , the overall result is equal to a CNOT between  $q_1$  and  $q_3$ , keeping other states unchanged. Fig. 4 clarifies the process above.

The overhead induced by the swapping operations explains how important is the topological organisation of devices and the circuit design as well.

#### 4.2 Quantum link

As mentioned in Section 3, the Quantum Internet requires both classical and quantum links. In this perspective, we make a distinction between *matter* and *flying* qubits [2], i.e. between qubits for information processing/storing and qubits for information transmission.

As regards to the matter qubits, several candidate technologies are available, each one with its pros and cons [25]. Conversely, as regards to the flying qubits, there exists a general consensus about adopting photons as qubit substrate [2]. However, heterogeneity arises by considering the different physical channels the photons propagate through, ranging from free-space optical channels (either ground or satellite free-space) to optical fibres. Thus, a transducer for matter-flying conversion is necessary as depicted in Fig. 5 and discussed with further details in [2]. Communication models need to take into account such a technological heterogeneity, with the aim of providing a black box for upper protocol layers with one common logic.


Furthermore, quantum mechanics does not allow to duplicate an unknown qubit or even simply to observe/measure it without altering the qubit. As a consequence, the communication techniques utilised to interconnect spatially remote quantum devices cannot be directly borrowed from classical communications. In this context, *quantum teleportation* is widely accepted as one of the most promising quantum communication technique between remote quantum nodes [3, 52, 53]. Quantum teleportation has been experimentally verified [54] and it requires, as depicted in Fig. 6, a pair of parallel resources (for an in-depth discussion about the quantum teleportation process, we refer the readers to [3]). One of these resources is classical: two bits must be transmitted from the source to the destination. The other resource is quantum: an entangled pair of qubits must be generated and shared between the source and the destination.

In the context of the distributed quantum computing ecosystem, quantum teleportation constitutes the foundation of a communication paradigm known as *teledata* [56], which generalises the concept of moving state among qubits to remote devices.

To provide a concrete example of the teledata concept, we further classify qubits either as *communication qubits* or as *data qubits*. Specifically, within each quantum device, a subset of matter qubits is reserved for generating entanglement. We refer to these qubits as communication qubits [23], to distinguish them from the remaining matter qubits within the device devoted to processing/storage, which we refer as data qubits.

As an example, consider two ibmqx2 architectures [57] interconnected via quantum teleportation as depicted in Fig. 7. The  $c_0, c_1$  pair is in the state  $|\Phi^+\rangle$  – that is the standard notation to denote a couple of maximally entangled qubits [17], also known as EPR pair [55]. Any kind of interaction among remote devices involves communication qubits, but not all the data qubits are connected with them. As already explained in Section 4.1, interactions between non-adjacent qubits are feasible but they imply an overhead. A solution would be adding more qubits to the communication set, but it means sacrifice further valuable resources for processing and storage. For this reason, the selection of the communication qubit set is a crucial task within the distributed quantum computing ecosystem and it implies a carefully evaluation of the trade-off between the number of data and communication qubits.

Next section shows how to exploit teleportation in order to not only send information but also perform joint operations among remote qubits.



**Fig. 7** Coupling map representing physical architecture of two quantum devices *ibmqx2* interconnected via the quantum teleportation paradigm. Nodes labelled with  $c_0$  and  $c_1$  denote communication qubits, and the dotted line indicates that they are in the entangled state  $|\Phi^+\rangle$ , i.e. they form an EPR pair. Conversely,  $q_i$  with  $i \in \{0, 1, 2, 3\}$  denote data qubits – i.e. qubits available for computing tasks

### 4.3 Teleporting gates

Distributed quantum computation requires the capability to perform quantum operations on qubits belonging to remote quantum devices.

As mentioned in the previous section, one possible solution is to resort to the *teledata* concept, by moving the quantum information from a quantum device to another via the teleportation process, through an entangled pair.

However, an entangled pair allows one to implement also a so-called teleporting gate, or *telegate* [58]. From a theoretical perspective, we have already observed that providing the  $\text{CNOT}$  operation – together with other single qubit gates – is enough to perform any kind of quantum algorithm. Therefore, returning to stack dependencies of Fig. 2, we can conceptualise a service that provides a set of remote operations based on teleported gates. Such service will directly interact with the physical system, exploiting the entanglement generation and distribution functionality [3].

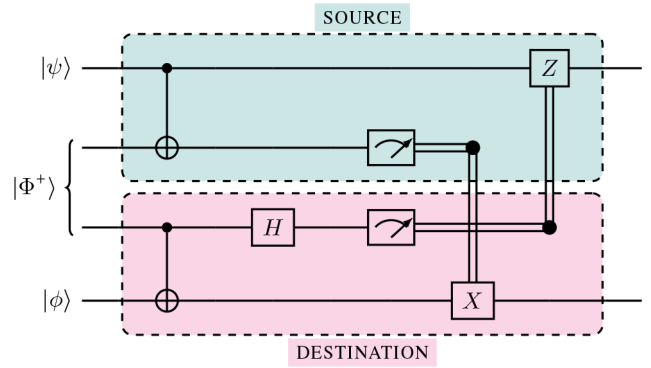
Specifically, by considering the topology depicted in Fig. 7, it is possible to implement the  $\text{CNOT}$  as the joint operation between qubits belonging to spatially remote devices. Indeed, by exploiting two communication qubits –  $c_0$  and  $c_1$  – shared between the two remote devices and storing an EPR pair, it is possible to perform a remote  $\text{CNOT}$  operation (for a more comprehensive presentation of the telegate, we refer the reader to [59]) with, for example,  $q_0$  as control qubit and  $q_4$  as target qubit. Fig. 8 shows the corresponding circuit, consisting of local  $\text{CNOT}$  operations and single-qubit operations and measurements.

### 4.4 Distributed quantum compiler

Concepts discussed so far provide some fundamental underlying communication functionalities enabling the distributed quantum computing paradigm.

Indeed, a quantum algorithm can be abstracted via a general model. Usually, such a model is the quantum circuit – where a computation is a sequence of quantum gates on a register – as those depicted in Figs. 6 and 8. Any algorithm designer may benefit from an abstraction which hides the complications due to physical features.

However, as mentioned in Section 3, remote qubits belonging to remote quantum devices are interconnected – within the virtual quantum processor – through virtual connections made possible by the remote operations as described in Sections 4.2 and 4.3. Unfortunately, remote operations suffer delays and error rates higher than local operations. The reason underlying this statement is that decoherence – i.e. a quantum-specific noise process [3] – affects each step of protocols realising remote operations – from the entanglement distribution to the gate operations. Decoherence affects local operations as well. However, since the effects of such a noise become stronger as function of the time [2, 3], remote operations, by involving further away parties, are more vulnerable. It follows that, given a particular quantum circuit describing a quantum algorithm, the distributed quantum compiler must optimise the circuit (note that we need circuit optimisation also in case of a single quantum processor [48, 60]) to reduce the overhead



**Fig. 8** Quantum telegate circuit implementing a  $\text{CNOT}$  operation between remote qubits. Specifically, the circuit performs a  $\text{CNOT}$  operation between qubits placed at different devices – say qubits  $q_0$  and  $q_4$  in Fig. 7 – through  $\text{CNOT}$ s between each qubit and the member of an EPR pair stored at the same device, followed by single-qubit gates and measurements. Note that  $|\psi\rangle$  and  $|\phi\rangle$  denotes the generic initial states stored at  $q_0$  and  $q_4$ , respectively, whereas  $|\Phi^+\rangle$  denotes the EPR pair stored by the communications qubits – say qubits  $c_0$  and  $c_1$  in Fig. 7

arising with the swap operations discussed in Section 4.1) so that the number of remote operations is minimised as much as possible to limit the decoherence effects.

Furthermore, the compiler should be able to optimise the circuit so that it can be executed, regardless of the underlying network topology. Indeed, it may be the case that two remote quantum devices are not directly connected – i.e. they do not share communication qubits – even though some remote operations between qubits of these devices are required. Hence, the compiler should optimise the corresponding quantum circuit, by minimising the entanglement swapping operations (entanglement swapping is a technique used to entangle distant nodes – without physically sending an entangled qubit through the entire distance – by *swapping* the entanglement generated at intermediate nodes [53]) among the remote devices. Finally, as discussed in Section 4.2, there exists a trade-off between data and communication qubits. The larger is the number of communication qubits in a device, the higher is the rate of remote operations achievable at the price of reducing the number of qubits devoted to computation.

### 4.5 What's next

Considering modern available technologies, it seems reasonable to envision that we will see a first attempt of interconnection among quantum computers located nearby. Likely, one of the main tech companies currently providing cloud access to isolated quantum computers – such as IBM [61], HIQ [62], Amazon Braket [63] or Azure Quantum [64] – will scale the available quantum computing power by interconnecting few quantum computers located few meters away within a *quantum farm* [2, 4]. After that, it is reasonable to envision that some of these companies will have quantum computers distributed over the world. Thus, it would not be surprising to see interconnection among quantum computers or even among quantum clusters miles away from each other.

However, interconnecting quantum devices implies the need for a communication standard. In other words, remote devices – likely to have technological differences – will need to agree on network protocols in order to exchange interpretable information. Thus, the Quantum Internet will need a logical architecture as well as the classical Internet does – i.e. with the TCP/IP Internet protocol suite being the standard *de-facto*. Research in this direction has already started, within the IETF, where researchers are trying to conceptualise the Quantum Internet as a service-oriented platform [23, 24].

## 5 Conclusions

With this paper, we have presented a layered ecosystem that outlines a possible key strategy towards large-scale quantum

processor design based on the distributed quantum computing paradigm. Within the envisioned ecosystem, the lowest layers integrate the Quantum Internet as the fundamental underlying infrastructure providing networking and communication functionalities among remote quantum devices. Conversely, the upper layers are responsible for mapping the quantum algorithm onto the underlying physical infrastructure by optimising the available computational resources as well as by accounting for the constraints induced by the hardware/network configuration.

## 6 References

- [1] Caleffi, M., Cacciapuoti, A.S., Bianchi, G.: ‘Quantum internet: from communication to distributed computing!’. Proc. of the 5th ACM Int. Conf. on Nanoscale Computing and Communication, Reykjavik, Iceland, 2018, invited paper
- [2] Cacciapuoti, A.S., Caleffi, M., Tafuri, F., et al.: ‘Quantum internet: networking challenges in distributed quantum computing’, *IEEE Netw.*, 2020, **34**, (1), pp. 137–143
- [3] Cacciapuoti, A.S., Caleffi, M., Van Meter, R., et al.: ‘When entanglement meets classical communications: quantum teleportation for the quantum internet’, *IEEE Trans. Commun.*, 2020, p. 1, invited paper
- [4] Caleffi, M., Chandra, D., Cuomo, D., et al.: ‘The rise of the quantum internet’, *Computer*, 2020, **53**, (6), pp. 67–72
- [5] Cartledge, E.: ‘Europe’s 1-billion quantum flagship announces grants’, *Science*, 2018, **362**, (6414), p. 512
- [6] Gibney, E.: ‘The quantum gold rush’, *Nature*, 2019, **574**, (7776), pp. 22–24
- [7] Bourzac, K.: ‘4 tough chemistry problems that quantum computers will solve’, *IEEE Spectr.*, 2017, **54**, (11), pp. 7–9
- [8] Schuld, M., Sinayskiy, I., Petruccione, F.: ‘An introduction to quantum machine learning’, *Contemp. Phys.*, 2015, **56**, (2), pp. 172–185
- [9] Gottesman, D., Lo, H.-K., Lutkenhaus, N., et al.: ‘Security of quantum key distribution with imperfect devices’. Int. Symp. on Information Theory, Chicago, Illinois, 2004, p. 136
- [10] Preskill, J.: ‘Quantum computing and the entanglement frontier’. 25th Solvay Conf. on Physics, Brussels, Belgium, October 2011
- [11] Arute, F., Arya, K., Babbush, R., et al.: ‘Quantum supremacy using a programmable superconducting processor’, *Nature*, 2019, **574**, (7779), pp. 505–510
- [12] Pednault, E., Gunnels, J., Maslov, D., et al.: ‘On quantum supremacy’, *IBM Res. Blog*, 2019
- [13] Pednault, E., Gunnels, J.A., Nannicini, G., et al.: ‘Leveraging secondary storage to simulate deep 54-qubit sycamore circuits’, 2019, p. arXiv:1910.09534
- [14] Cao, Y., Romero, J., Olson, J.P., et al.: ‘Quantum chemistry in the age of quantum computing’, *Chem. Rev.*, 2019, **119**, (19), pp. 10856–10915
- [15] Drouin, G.: ‘IBM Q’s Dr. Robert Sutor explains the state of the quantum computing industry’, 2018
- [16] Cacciapuoti, A.S., Caleffi, M.: ‘Toward the quantum internet: A directional-dependent noise model for quantum signal processing’. IEEE Int. Conf. on Acoustics, Speech and Signal Processing, Brighton, UK, May 2019, pp. 7978–7982
- [17] Nielsen, M.A., Chuang, I.L.: ‘*Quantum computation and quantum information*’ (Cambridge University Press, Cambridge, England, 2010)
- [18] Babar, Z., Botsinis, P., Alanis, D., et al.: ‘The road from classical to quantum codes: A hashing bound approaching design procedure’, *IEEE Access*, 2015, **3**, pp. 146–176
- [19] Kimble, H.J.: ‘The quantum internet’, *Nature*, 2008, **453**, (7198), pp. 1023–1030
- [20] Pirandola, S., Braunstein, S.L.: ‘Physics: unite to build a quantum internet’, *Nature*, 2016, **532**, (7598), pp. 169–171
- [21] Wehner, S., Elkouss, D., Hanson, R.: ‘Quantum internet: a vision for the road ahead’, *Science*, 2018, **362**, (6412), p. eaam9288
- [22] Awschalom, D., Berggren, K.K., Bernien, H., et al.: ‘Development of quantum InterConnects for next-generation information technologies’, arXiv:1912.06642, 2019
- [23] Kozlowski, W., Wehner, S., Meter, R.V., et al.: ‘Architectural principles for a quantum internet’. Internet Engineering Task Force, Internet-Draft draft-irtf-qirg-principles-03, Mar 2020, work in Progress
- [24] Wang, C., Rahman, A., Li, R.: ‘Applications and use cases for the quantum internet’. Internet Engineering Task Force, Internet-Draft draft-wang-qirg-quantum-internet-use-cases-04, Mar 2020, work in Progress
- [25] Meter, R.V., Devitt, S.J.: ‘The path to scalable distributed quantum computing’, *Computer*, 2016, **49**, (9), pp. 31–42
- [26] Yimsiriwattana, A., Lomonaco, Jr.S.J.: ‘Distributed quantum computing: A distributed shor algorithm’. Quantum Information and Computation II, Int. Society for Optics and Photonics, Orlando, Florida, 2004, vol. 5436, pp. 360–372
- [27] Broadbent, A., Fitzsimons, J., Kashefi, E.: ‘Universal blind quantum computation’. 50th Annual IEEE Symp. on Foundations of Computer Science, Atlanta, Georgia, Oct 2009, pp. 517–526
- [28] Aharonov, D., Ben-Or, M., Eban, E.: ‘Interactive proofs for quantum computations’. Proc. of Innovations in Computer Science, Beijing, China, 2010
- [29] Huang, H., Zhao, Q., Ma, X., et al.: ‘Experimental blind quantum computing for a classical client’, *Phys. Rev. Lett.*, 2017, **119**, (5), p. 050503
- [30] Huang, H., Bao, W.-S., Li, T., et al.: ‘Universal blind quantum computation for hybrid system’, *Quantum Inf. Process.*, 2017, **16**, (8), p. 199
- [31] Xu, G., Xiao, K., Li, Z.P., et al.: ‘Controlled secure direct communication protocol via the three-qubit partially entangled set of states’, *Comput. Mater. Contin.*, 2019, **58**, (3), pp. 809–827
- [32] Chen, X., Wang, Y., Xu, G., et al.: ‘Quantum network communication with a novel discrete-time quantum walk’, *IEEE Access*, 2019, **7**, pp. 13634–13642
- [33] Sutradhar, K., Om, H.: ‘Efficient quantum secret sharing without a trusted player’, *Quantum Inf. Process.*, 2020, **19**, (2), p. 73
- [34] Ben-Or, M., Hassidim, A.: ‘Fast quantum byzantine agreement’. Proc. of the thirty-seventh annual ACM symp. on Theory of computing, Baltimore, Maryland, 2005, pp. 481–485
- [35] Caleffi, M., Cacciapuoti, A.S.: ‘Quantum switch for the quantum internet: noiseless communications through noisy channels’, *IEEE J. Sel. Areas Commun.*, 2020, **38**, p. 1
- [36] Cacciapuoti, A.S., Caleffi, M.: ‘Capacity bounds for quantum communications through quantum trajectories’, 2019, p. arXiv:1912.08575
- [37] Gyongyosi, L., Imre, S., Nguyen, H.V.: ‘A survey on quantum channel capacities’, *IEEE Commun. Surv. Tutor.*, 2018, **20**, (2), pp. 1149–1205
- [38] Caleffi, M.: ‘Optimal routing for quantum networks’, *IEEE Access*, 2017, **5**, pp. 22299–22312
- [39] Gyongyosi, L., Imre, S.: ‘Multilayer optimization for the quantum internet’, *Sci. Rep.*, 2018, **8**, (1), p. 12690
- [40] Pirker, A., Dür, W.: ‘A quantum network stack and protocols for reliable entanglement-based networks’, *New J. Phys.*, 2019, **21**, (3), p. 033003
- [41] Gyongyosi, L., Imre, S.: ‘Topology adaption for the quantum internet’, *Quantum Inf. Process.*, 2018, **17**, (11), p. 295
- [42] Chakraborty, K., Rozpedek, F., Dahlberg, A., et al.: ‘Distributed routing in a quantum internet’, 2019, p. arXiv:1907.11630
- [43] Shi, S., Qian, C.: ‘Modeling and designing routing protocols in quantum networks’, 2019, p. arXiv:1909.09329
- [44] Gambetta, J.M., Chow, J.M., Steffen, M.: ‘Building logical qubits in a superconducting quantum computing system’, *NPJ Quantum Inf.*, 2017, **3**, (1), p. 2
- [45] Tacchino, F., Macchiavello, C., Gerace, D., et al.: ‘An artificial neuron implemented on an actual quantum processor’, *NPJ Quantum Inf.*, 2019, **5**, (1), pp. 1–8
- [46] Verdon, G., Arrazola, J.M., Brädler, K., et al.: ‘A quantum approximate optimization algorithm for continuous problems’, p. arXiv:1902.00409, 2019
- [47] Farhi, E., Harrow, A.W.: ‘Quantum supremacy through the quantum approximate optimization algorithm’, p. arXiv:1602.07674, 2016
- [48] Zulehner, A., Paler, A., Wille, R.: ‘An efficient methodology for mapping quantum circuits to the IBM QX architectures’, *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, 2019, **38**, (7), pp. 1226–1236
- [49] Rieffel, E.G., Polak, W.H.: ‘*Quantum computing: A gentle introduction*’ (MIT Press, Cambridge, Massachusetts, 2011)
- [50] Raychev, N.: ‘Universal quantum operators’, *Int. J. Sci. Eng. Res.*, 2015, **6**, (6), pp. 1369–1371
- [51] Linke, N.M., Maslov, D., Roetteler, M., et al.: ‘Experimental comparison of two quantum computing architectures’, *Proc. Natl. Acad. Sci.*, 2017, **114**, (13), pp. 3305–3310
- [52] Bennett, C.H., Brassard, G., Crépeau, C., et al.: ‘Teleporting an unknown quantum state via dual classical and Einstein-podolsky-rosen channels’, *Phys. Rev. Lett.*, 1993, **70**, pp. 1895–1899
- [53] Van Meter, R.: ‘Quantum networking and internetworking’, *IEEE Netw.*, 2012, **26**, (4), pp. 59–64
- [54] Ren, J.-G., Xu, P., Yong, H.-L., et al.: ‘Ground-to-satellite quantum teleportation’, *Nature*, 2017, **549**, (7670), p. 70
- [55] Einstein, A., Podolsky, B., Rosen, N.: ‘Can quantum-mechanical description of physical reality be considered complete?’, *Phys. Rev.*, 1935, **47**, pp. 777–780
- [56] Van Meter, R., Nemoto, K., Munro, W.J., et al.: ‘Distributed arithmetic on a quantum multicompiler’. 33rd Int. Symp. on Computer Architecture, Boston, Massachusetts, 2006, pp. 354–365
- [57] Qiskit: ‘Beckend information’. Available at <https://github.com/qiskit/ibmq-device-information>
- [58] Van Meter, R.: ‘Architecture of a quantum multicompiler optimized for shor’s factoring algorithm’, 2006, p. quant-ph/0607065
- [59] Chou, K.S., Blumoff, J.Z., Wang, C.S., et al.: ‘Deterministic teleportation of a quantum gate between two logical qubits’, *Nature*, 2018, **561**, (7723), pp. 368–373
- [60] Ferrari, D., Amoretti, M.: ‘Efficient and effective quantum compiling for entanglement-based machine learning on ibm q devices’, *Int. J. Quantum Inf.*, 2018, **16**, (8), p. 1840006
- [61] ‘IBM quantum experience’. Available at <https://quantum-computing.ibm.com>
- [62] ‘Quantum computing software HiQ’. Available at <https://hiq.huaweicloud.com/>
- [63] ‘Amazon braket’. Available at <https://aws.amazon.com/braket/>
- [64] ‘Azure quantum’. Available at <https://azure.microsoft.com/it-it/services/quantum/>