

# Resource Allocation in Quantum Networks for Distributed Quantum Computing

Claudio Cicconetti

*Institute of Informatics and Telematics  
National Research Council  
Pisa, Italy*

Marco Conti

*Institute of Informatics and Telematics  
National Research Council  
Pisa, Italy*

Andrea Passarella

*Institute of Informatics and Telematics  
National Research Council  
Pisa, Italy*

**Abstract**—The evolution of quantum computing technologies has been advancing at a steady pace in the recent years, and the current trend suggests that it will become available at scale for commercial purposes in the near future. The acceleration can be boosted by pooling compute infrastructures to either parallelize algorithm execution or solve bigger instances that are not feasible on a single quantum computer, which requires an underlying *Quantum Internet*: the interconnection of quantum computers by quantum links and repeaters to exchange entangled quantum bits. However, Quantum Internet research so far has been focused on provisioning point-to-point flows only, which is suitable for (e.g.) quantum sensing and metrology, but not for distributed quantum computing. In this paper, after a primer on quantum computing and networking, we investigate the requirements and objectives of smart computing on distributed nodes from the perspective of quantum network provisioning. We then design a resource allocation strategy that is evaluated through a comprehensive simulation campaign, whose results highlight the key features and performance issues, and lead the way to further investigation in this direction.

**Index Terms**—Distributed Quantum Computing, Quantum Internet, Quantum Routing

## I. INTRODUCTION

Quantum Computing (QC) exploits the properties of matter at very small scale to solve some problems much faster than a classical counterpart. Even though QC has been theorized 40 years ago [24], only recently the technology evolution and a spur of investments have made it possible to obtain practical results and speculate about approaching mass deployments [28]. QC is being already used in the chemical and pharmaceutical industry, while new applications are being progressively unlocked in material science, Machine Learning (ML) and engineering optimization, production and logistics, post-quantum security [25]. Essentially, the computational advantage of QC stems from the properties of superposition and entanglement of the *qubits* (i.e., the “quantum bits”), which we review in Sec. II.

We can expect that the computational power of a single QC will remain relatively limited in the near future, due to scalability issues in maintaining a very stable and controlled environment to cope with the flimsy nature of qubits. For this reason, computation speed-up can be sought by distributing the execution over multiple QCs, which requires an end-to-end entanglement of the qubits they use across geographical distances, in other terms the realization of the **Quantum**

**Internet** [34]. The latter, in fact, is receiving attention and investments: when deployed, it will allow QCs to run algorithms in a distributed fashion, thus benefiting the users from shared compute capacity across multiple systems, much like what happened with grid computing in the not-so-distant past.

In this work we focus on a problem that is ancillary to distributed QC and has received little attention so far: the allocation of resources in the Quantum Internet among multiple quantum computers based on the characteristics of the underlying quantum network. The properties of distributed QC applications are fundamentally different from those of point-to-point quantum applications, which have been studied intensively by the scientific community in relation to quantum sensing [23] and Quantum Key Distribution (QKD) [22]. In particular, distributed QC applications have a more elastic nature along two directions: (i) a given host may have multiple communication peers, as a quantum computer may offload its computation to many nodes in a pool, not just one; and (ii) the rate of entangled qubits exchanged is not constant, but rather a quantum computer may desire to consume as many of them as possible in a greedy manner, which leads to fairness concerns. *Our main contribution is to explore these novel requirements from the point of view of network resource allocation, while also proposing a practical solution, which we evaluate through simulations with the goal of identifying performance trade-offs in terms of internal and external system properties.*

The rest of this paper is structured as follows. In Sec. II we provide an introduction to quantum computing and networking. We review the related work on routing in the Quantum Internet in Sec. III. We then describe the system model adopted and our scientific contribution in Sec. IV. Simulation results are presented in Sec. V, while Sec. VI concludes the paper.

## II. BASICS OF QUANTUM COMPUTING AND COMMUNICATION

In this section we provide a bird’s-eye view of the fundamental principles of quantum computing and communication, with the goal of making the paper readable even without prior knowledge on these topics. We suggest the following textbooks to the interested reader for a complete illustration of quantum computing [19] and networking [31].

As already introduced briefly, the unit of computation of QC is the *qubit*. While a classical bit represents a binary piece

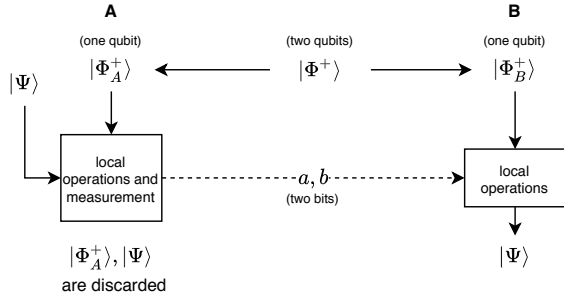


Fig. 1. Quantum teleportation scheme.  $|\Psi\rangle$  is the qubit to be teleported from A to B, while  $|\Phi_A^+\rangle$  and  $|\Phi_B^+\rangle$  are two maximally entangled qubits in a Bell state that are consumed in the process;  $(a, b)$  are classical bits that are read by A and used by B to decide which local operations to apply.

of information, e.g., head (0) or tail (1) as the result of flipping a coin, a qubit can be in a *superposition* of the two opposite states, which can be visualized as the coin flipping in the air: until it drops, it is impossible to say what its value. In conventional bra-ket notation this is expressed as  $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ , where  $|0\rangle$  and  $|1\rangle$  are the two extreme levels of the systems (head or tail in the example) and  $\alpha, \beta$  are complex numbers representing their respective probability amplitudes. Measuring the qubit causes its state to collapse to one of the two possible values, thus effectively converting the qubit into a classical bit; once a qubit has been measured its quantum information cannot be recovered anymore (i.e., once the coin has dropped, it is not possible to restore it to its status while flipping). Quantum algorithms are executed by preparing the input qubits in known initial states, based on the configuration of the real-life problem that one is going to solve, then applying quantum gates expressing quantum logic operations, and finally measuring the outcomes, which gives the output as classical information.

Another important property of qubits is that they can be *entangled*. When a set of qubits is entangled, they express a high correlation that has no classical counterpart (and, in fact, is deeply counterintuitive), which persists even though they are separated in space. For instance, let us consider two qubits in the maximally entangled *Bell state*:

$$|\Phi^+\rangle := \frac{1}{\sqrt{2}} (|00\rangle + |11\rangle). \quad (1)$$

When performing a measurement on any of the two qubits, there is equal probability to get 0 or 1, but if the measurements of the first qubit gives (e.g.) 0, then measuring the second qubit will certainly give a 0; this remains true even if the second measurement is done later or at an arbitrary distance.

Another distinguishing feature of qubits is that they *cannot* be copied: a qubit must be either transferred *as is* or “consumed” through (irreversible) measurement. However, it is possible to *teleport* a qubit from a quantum system (A) to another (B) as follows – see scheme in Fig. 1. The two quantum systems must have exchanged an entangled pair of qubits in the Bell state beforehand (call them  $|\Phi_A^+\rangle$  and  $|\Phi_B^+\rangle$ ); then, A performs local operations on  $|\Phi_A^+\rangle$  and the qubit

to teleport  $|\Psi\rangle$ , followed by a measurement, which destroys both. Such a measurement provides a result, expressed as two classical bits  $(a, b)$ , which are transferred to B. Finally, B applies local operations on  $|\Phi_B^+\rangle$ , which transform it exactly into  $|\Psi\rangle$ . It is important to note that the duration of teleporting is limited by the time required to transfer classical information from A to B, hence it does not violate the laws of physics that prohibit faster-than-light travel (even for information).

Teleportation is a crucial operation involved in the creation of quantum networks. Quantum computers store so-called *matter qubits* in quantum memories, which are interconnected by logic that makes them interact via quantum gates for the execution of algorithms. To extend the execution of a quantum circuit beyond a single computer, it is necessary to create an entanglement between one of its qubits with another, physically located in the quantum memory of a remote peer. This can be achieved through so-called *flying qubits*, typically realized by encoding the state into the polarization of photons, which can be transmitted efficiently over fiber optic cables [30] or in free-space using satellite links [16].

However, flying qubits fade with distance, which makes it impractical to deploy a global scale quantum network without intermediate devices that can extend the range of single links: the **quantum repeaters**. A quantum repeater is a device that performs a measurement, called *entanglement swapping*, between two (flying) qubits resulting in the end-to-end entanglement of two (matter) qubits, as illustrated in Fig. 2, following the same principle of teleportation discussed above. The entanglement swapping can be repeated along a chain of quantum repeaters between two end points, thus extending the range of quantum networks for the creation of end-to-end entangled pairs of qubits arbitrarily, at least in theory. In practice, channel impairments and measurement imperfections reduce the quality of the end-to-end entanglement generated. A widely used metric to measure these collective effects is the *fidelity*, which is a relative measure between the actual state of a quantum system and its desired state; the fidelity is defined in the range  $[0, 1]$ , where 1 means perfection. Practical quantum algorithms can tolerate some degree of fidelity degradation. A comprehensive review of quantum teleportation for the realization of the Quantum Internet can be found in [4].

The evolution of quantum networks is commonly seen in incremental steps, sometimes referred as “generations” (term inspired from cellular networks), where the first generation (1G) is characterized by quantum repeaters operating according to the scheme above. Future generations will embed error correction to ensure high-fidelity entangled pairs at long distance and with many intermediate nodes [18], but we do not consider them as they are very far from being available in practice —indeed, not even 1G quantum repeaters are available today as commercial off-the-shelf products.

Finally, the attentive reader might have noticed that entanglement swapping does not lead to the distribution of arbitrary quantum states  $|\psi\rangle$ , but only Bell states as in Eq. (1). Such an assumption does *not* restrict the general applicability of a network of quantum repeaters to realize distributed QC, because




Fig. 2. Entanglement swapping scheme. A pair of maximally entangled qubits  $|\Phi^+\rangle_{A1}, |\Phi^+\rangle_{B1}$  is generated thanks to the measurement operation done by an intermediate quantum repeater, followed by local operations on, e.g., B.

it can be shown that *any arbitrary state can be obtained from a Bell state through a sequence of local operations only*.

### III. RELATED WORK

The literature on quantum networking and distributed QC is not vast: even though the basic ingredients have been known since a long time ago—consider for instance the seminal paper by Bouwmeester *et al.* on quantum teleportation [2] published on Nature in 1997—only recently there have been investments in an order of magnitude sufficient for technology to take off. This revamped interest has triggered new research activities in this area, briefly reviewed below.

In general terms, the problem of **quantum routing** is formulated as follows: given a network of quantum nodes (repeaters or computers) and a set of traffic flows identified by their sources, destinations, and application requirements (e.g., the minimum fidelity), find the “best” paths that fulfill the constraints. Some works have studied the problem by reusing the findings in the area of routing in classical networks. Van Meter *et al.* proposed a quantum version of the famous Dijkstra’s shortest path algorithm, which was shown to give very good performance with an appropriate selection of the routing metric that considers the specific properties of quantum networks [33]. More recently, Caleffi *et al.* have proposed a slightly less efficient variation of Dijkstra’s algorithm that can work with non-isotonic routing metrics, which they have advocated to provide superior performance in selected use cases [5]. Dijkstra’s algorithm is also the subject of [7], where the authors lay some mathematical foundations that allow them to derive upper bounds of performance in specific network topologies, including grid and ring.

A different direction is explored by Pant *et al.*, who studied the distribution of routing information to the nodes [20]; for this they propose a time-slotted approach: in the first part of the slot every repeater tries to create a local entanglement with all its neighbors, then in the second part the paths are established as instructed by a centralized authority. One interesting aspect of the paper is that multiple paths are selected for the same (source, destination) to maximize the rate of end-to-end Bell pairs. We have also adopted this time-slotted model in [8], where we have investigated the issue of “scheduling” of traffic flows, i.e., determining the order in which to assign paths to pending requests, in case the network

resources are not sufficient to serve them all. This problem is called “distribution” in [13], where the authors formulate it as an Integer Linear Programming (ILP), for which they derive closed formula performance bounds in the case of a homogeneous chain of quantum repeaters. The issue is also addressed in [6], where the authors have proposed to split the overall quantum routing problem in two to reduce the computational complexity: first, they determine the rates achievable by the traffic flows under the given network constraints using an approach based on multi-commodity flow optimization, then they map these rates to paths. The paper adopts a network model using probabilistic entanglement swapping, which we reuse in this work (described in Sec. IV-A).

An important reference for our study is [17], where the authors study the allocation strategy of traffic flows for which the paths have been pre-determined: they do so by borrowing the fairness concept from data networks and re-using traditional algorithms from the relevant literature. In our paper, we also borrow from the same literature, though we apply the concepts to a different class of applications, as it will be clear in the next section. *As a matter of fact, all the scientific works cited above have focused on point-to-point traffic flows, while in Sec. IV-B we introduce a different type of traffic that is more suitable to model distributed QC, with distinguishing features that do not allow the reuse of state-of-the-art solutions.* Rather, we claim that any existing routing/allocation/scheduling solutions should work *in parallel* to our proposed scheme to provide an effective resource allocation to each of the two traffic classes.

In addition to mere routing aspects, system-wide studies have also been published. We mention [32], which is a compendium of several previous studies from the same authors that illustrates an overall architecture of the Quantum Internet, also including application, protocol, and deployment aspects at a high level. On the other hand, other works have focused on specific components, which are complementary to the research activity presented, e.g., [37] on congestion control in transport protocols and [11] on the link layer, with a focus on hardware and physical-layer considerations.

Furthermore, some research groups have been working to define the basic principles of **distributed QC**. Parekh *et al.* have defined an elegant framework for the parallel execution of a broad class of quantum algorithms on multiple nodes [21], both using remote entanglement and with Local Operations

and Classical Communication (LOCC) only, also studying in depth three classes of algorithms: variational quantum eigensolver, low-depth quantum amplitude estimation, and quantum k-means clustering. In [10] the authors address the problem of the efficient compilation of circuits for distributed QC by considering that some gate operations will be executed remotely, hence with much different latency and reliability than on-chip operations. The research of Dahlber *et al.* moved in the same direction and went as far as defining a set of low-level instructions (called NetQASM) for distributed QC systems seamlessly supporting local and remote gates [12]. *These works confirm that there is a growing interest in distributed QC, which is a motivation for our work.*

#### IV. TRAFFIC CLASSES IN QUANTUM NETWORKS AND RESOURCE ALLOCATION

In this section we illustrate the novel contribution of our work. We first introduce the system model used for quantum networks (Sec. IV-A). Then, in Sec. IV-B, we elaborate on two different classes of traffic, namely constant-rate flows and distributed QC applications. For the latter we propose a resource allocation solution in Sec. IV-C.

##### A. System model

In this work we adopt the model proposed in [6], which focuses on some fundamental characteristics of quantum networks, which will remain true for at least the 1G of quantum repeaters (described in Sec. II). We consider a network of quantum devices, called *nodes*, interconnected via *links*. Without loss of generality, we do not distinguish between quantum computers vs. repeaters. The links are directional and they are assumed to have a fixed capacity, in terms of the rate at which Bell pairs can be generated, indicated as  $C_{i,j} > 0$ , where  $i$  and  $j$  are the two nodes connected. For instance,  $C_{1,2} = 8$  means that the quantum physical-layer devices between node 1 and 2 are such that 8 maximally entangled Bell pairs will be generated every second; we recall that qubits cannot be stored, thus all unused Bell pairs will be necessarily discarded. The network can then be represented as a weighted graph  $\mathcal{G}$  where the vertices  $\mathcal{V}$  are the nodes, the edges  $\mathcal{E}$  are the links, and the weights are the link capacities. In practice, depending on the technology used, the initial fidelity of a Bell pair is a stochastic process with expectation  $F_{ij}^{\text{init}}$ , which could be used to decorate the edges further; for simplicity of notation, in this work we assume that  $F_{ij}^{\text{init}} = \bar{F}$  for all  $(i,j) \in \mathcal{E}$ .

As already explained, 1G quantum repeaters do not use quantum error corrections, i.e., the measurement operation (see Fig. 2), which enables entanglement swapping, may fail according to some (as of yet) unknown probabilistic process: when this happens, the qubits have to be discarded and end-to-end entanglement fails accordingly. If linear optics are used to implement this operation, the failure probability is *at least* 50% [26]. Let us consider two nodes  $i$  and  $k$  attempting repeatedly end-to-end entanglement with a rate  $R$  through node  $j$ , which has  $q_j$  average measurement success probability. The entanglement will only succeed a fraction of the times




Fig. 3. Example quantum network, with  $q = 0.5$  and  $F = 0.95$ .

( $q_j$ ), which means that the rate of usable Bell pairs will be  $r = R \cdot q_j$ . We call  $r$  net rate, and  $R$  gross rate. The net rate is available to the quantum applications, while the gross rate is that consumed on network resources, which has to be available along the path in accordance with the link capacities:  $\frac{r}{q_j} \leq \min\{C_{ij}, C_{jk}\}$ . If there were  $N$  intermediate nodes, the net rate would decrease exponentially:  $r = R \cdot \prod_{j=1}^N q_j$ , which is a well-known undesirable property of quantum networks. Again, to keep the notation simple, we assume that the average measurement success probability of all nodes in the network is equal, i.e.,  $q_j = q$  for all  $j \in \mathcal{V}$ , which gives:

$$R = \frac{r}{q^L}, \quad (2)$$

where  $L \geq 1$  is the length of the path  $p$  between the two nodes (in fact, with  $L = 0$  there is no entanglement swapping and it is  $R = r$ ). Also the fidelity is known to degrade exponentially with the path size, according to the following formula [3] (under some simplifying assumptions):

$$F = \frac{1}{4} + \frac{3}{4} \left( \frac{4\bar{F} - 1}{3} \right)^{L+1}. \quad (3)$$

An example network is shown in Fig. 3. Assume we want to create a sequence of end-to-end entanglements between nodes 1 and 6 with rate 1 Bell pair/s. The shortest path, in number of hops, is through nodes 3 and 5, but the path size would be  $L = 2$ , which entails  $R = r/q^L = 1/(0.5)^2 = 4$ , which exceeds the minimum capacity of the edges along the path, that is 2. The only viable path is through 2, 3, and 4 (blue walk in the figure, with  $L = 3$ ), requiring a gross rate  $R = 8$ , which is available on all the edges. Note that going through this longer route consumes more network resources, but it is inescapable in this case. Instead, a path between nodes 1 and 5 with  $r = 1$  is feasible passing through 3 (green walk in the figure), because with  $L = 1$  the gross rate is just 2, which is compatible with the capacities along the edges (1,3) and (3,5).

##### B. Quantum network traffic classes

So far, the problem of routing demands in quantum networks has been investigated only for a single class of applications, i.e., those essentially defined by source, destination, and Bell rate (and possibly minimum fidelity). Typical applications

TABLE I  
CHARACTERISTICS OF CONSTANT-RATE APPLICATIONS (FLOWS) VS.  
DISTRIBUTED QC APPLICATIONS (APPS).

	Flows	Apps
Nodes involved	source + destination	host + set of peers
Rate requirement	constant Bell rate	elastic Bell rate
Fidelity requirement	strict	flexible
Session duration	user-driven	depends on allocation
QoS	guaranteed upon admission	best-effort

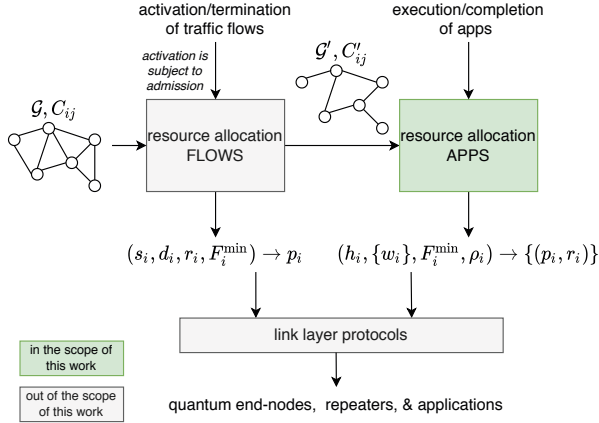


Fig. 4. Overview of the architectural elements for the coexistence of the two traffic classes considered in this work for quantum networks.

that fit this model are: QKD, i.e., the secure exchange of a shared secret between two parties, and quantum sensing, for the accurate measurement of physical quantities. For those applications, both the rate and fidelity requirements are to be intended as strict guarantees: if not matched, then the application cannot work properly, but if they are exceeded then the performance does not get any better. Therefore, it is implicit that admission of new flows of this type must be subject to an admission control scheme to verify that resources are available in the network.

Distributed QC is another class of applications, which to the best of our knowledge has been totally neglected so far, from the point of view of quantum routing. Applications in this class do not have a minimum required rate, but rather they can greedily consume as many Bell pairs as provided, under the assumption that end-to-end entanglement is the limiting factor for the execution of quantum circuits on distributed infrastructures, which appears to be reasonable in the near future based on current technology trends. Furthermore, a single host wishing to offload its computation to other QCs may want to do so in a pool of peers, rather than point-to-point: this adds a degree of freedom in the routing decision as one does not only have to select a path, but also choose which of the possible destinations to activate for a given application. In the remainder of this paper, we call for short *flows* the constant-rate traffic flows vs. *apps* the flexible distributed computing applications. A comparison of their main distinguishing features is reported in Table I.

We can speculate that in a real quantum network flows and apps will coexist, much like latency-sensitive constant bit-rate applications (e.g., digitized voice) coexist with best-effort

elastic TCP applications in the Internet today. Only time will tell whether one type of traffic will dominate or which business model will emerge, but for now we argue that it makes sense to give priority to flows: since their requirements are inflexible, the rate of flows that fail admission can be minimized by considering all the network resources available for this class of traffic. Accordingly, we propose a high-level view of the overall resource allocation, illustrated in Fig. 4. In the figure the difference between flows and apps is once more evident: a flow is identified by  $(s, d, r, F^{\min})$ , i.e., source, destination, rate, and minimum fidelity, and it is mapped to a single path  $p$ ; an app is identified by  $(h, \{w\}, F^{\min}, \rho)$ , i.e., host, set of candidate peers (depending on the commercial agreements between QC infrastructures), minimum fidelity, and priority, and it is mapped to a set of paths, each with a given net rate. We have identified the priority (as a numerical weight) as a means to provide differentiated service; for example, a higher priority may be acquired at extra service costs or it can be an internal parameter that increases while the resources are not used. The green box in the figure, i.e., the resource allocation of apps, is elaborated further in the next section, while the gray boxes are covered already by prior works: *resource allocation of flows* is investigated in [5]–[8], [13], [17], [33], whereas *link layer protocols* are explored in [11], [20], [37].

### C. Resource allocation of apps

For apps, the resource allocation problem can be expressed in general terms as follows: given a set of apps, find the set of paths, and corresponding rate, for each possible (host, peer), that maximizes a given objective function. In this form, it resembles the classical problem of resource allocation in data networks (e.g., [1]), even though there are important new features: i) the gross capacity decreases as the path length increases (more resources needed) and ii) so does the fidelity (which may eventually result in the path becoming unfeasible). Nonetheless, we believe that the notion of *fairness* widely studied in data networks can be swiftly reused in this context as the resource allocation goal. To this end, we propose an algorithm for the resource allocation of apps inspired from well-known Weighted Round Robin (WRR) [15]. *The basic idea is to pre-compute, for each app, all the possible paths towards the requested destinations (up to a maximum of  $k$  paths per destination, to keep the computational complexity low), then to allocate resources to apps in round-robin, for fairness reasons, where the maximum amount of gross rate an app is given in each round is a fraction of the round size  $\phi$  proportional to its weight.* The algorithm is explained with the help of the pseudo-code in Algorithm 1. Full implementation details can be found in the online repository of the simulator source code, publicly available as described in Sec. V.

The algorithm takes as input the residual graph  $\mathcal{G}'$  after all the possible flows have been allocated and a set  $\mathcal{A}$  of apps each characterized by  $(h_i, \mathcal{W}_i, F_i^{\min}, \rho_i)$ , with  $i \in \mathcal{A}$ ; it produces as output the set of paths  $\mathcal{F}$  that are allocated to the apps, where any app may have multiple paths each with a different net/gross rate. The algorithm has two system parameters:  $k \in$

---

**Algorithm 1:** Resource allocation of apps.

---

```

1  $\mathcal{G}'$  is the residual graph after all the possible flows have
   been allocated
2  $\mathcal{A}$  is the set of apps to be provisioned, with each app  $i$ 
   characterized by  $(h_i, \mathcal{W}_i, F_i^{\min}, \rho_i)$ 
3  $\forall i \in \mathcal{A} : \mathcal{P}_i \leftarrow \emptyset$ 
4  $\mathcal{L} \leftarrow \{i | i \in \mathcal{A}\}$ 
5  $a \leftarrow \text{last}(\mathcal{L})$ 
6  $\delta \leftarrow 0$ 
7  $\mathcal{F} \leftarrow \emptyset$ 
8 for  $i \in \mathcal{A}$  do
9   for  $j \in \mathcal{W}_i$  do
10     $\mathcal{P}_i \leftarrow \mathcal{P}_i \cup \text{find\_paths}(h_i, w_{ij}, k, \mathcal{G}')$ 
11   end
12 end
13 while  $\mathcal{L} \neq \emptyset$  do
14   if  $\delta = 0$  then
15      $a \leftarrow \text{next}(a, \mathcal{L})$ 
16      $\delta \leftarrow \phi \frac{\rho_a}{\sum_{i \in \mathcal{A}} \rho_i}$ 
17   end
18    $p \leftarrow \emptyset$ 
19   while  $p = \emptyset \wedge \mathcal{P}_a \neq \emptyset$  do
20      $p \leftarrow \text{find\_shortest\_path}(\mathcal{P}_a)$ 
21     if  $\exists e \in p | e \notin \mathcal{G}'$  then
22        $\text{remove\_path}(p, \mathcal{P}_a)$ 
23     end
24   end
25   end
26   if  $p = \emptyset$  then
27      $a \leftarrow \text{remove\_app}(a, \mathcal{L})$ 
28      $\delta \leftarrow \phi \frac{\rho_a}{\sum_{i \in \mathcal{A}} \rho_i}$ 
29   else
30      $R = \min \{\delta, \min_{e \in p} C_e\}$ 
31      $\text{add}(a, p, R, \mathcal{F})$ 
32      $\text{remove\_capacity}(R, p, C_e)$ 
33      $\text{update\_graph}(\mathcal{G}', C_e)$ 
34      $\delta \leftarrow \delta - R$ 
35   end
36 end
37 return  $\mathcal{F}$ 

```

---

$\mathbb{N}$  and  $\phi$ , in Bell pairs/s, which we call *round size*, already introduced. The following working variables are used:

- $\mathcal{L}$  is the *active list*, i.e., the set of apps that can be still provisioned: it is initialized with all the input apps (line 4) and then slowly depleted (line 27) each time an app  $a$  does not have further paths available in the residual graph;
- $a$  is the current app considered for allocation, initialized with the last app in the active list (line 5);
- $\delta$ , also called *credit* below, is the maximum amount of gross capacity that can be allocated to the current app  $a$ ;
- $\mathcal{P}_a$  is the set of paths available for the current app  $a$ , which is initialized at the beginning of the algorithm (lines 8–12) and then depleted when one path  $p$  is not feasible anymore in the residual graph (line 22).

**Pseudo-code walk through.** Lines 1–12 cover the **input definition and initialization** of the working variables, as explained above. The rest of the procedure is a loop that iterates over the active list until it is empty (condition line 13.) Within the main loop (lines 14–35),  $a$  is the identifier of the current app, which

TABLE II  
FUNCTIONS USED IN THE PSEUDO-CODE OF ALGORITHM 1.

$\text{last}(\mathcal{L})$	return the last app in the active list
$\text{next}(a, \mathcal{L})$	return the subsequent app to $a$ in the active list $\mathcal{L}$ , wrapping around when the end is reached
$\text{remove\_app}(a, \mathcal{L})$	remove the app $a$ from $\mathcal{L}$ and return the next app in the list, wrapping around when the end is reached
$\text{find\_paths}(s, d, k, \mathcal{G})$	returns the set of (up to) $k$ paths from node $s$ to node $d$ in $\mathcal{G}$
$\text{find\_shortest\_path}(\mathcal{P})$	return the shortest path in $\mathcal{P}$
$\text{remove\_path}(p, \mathcal{P})$	remove the path $p$ from the set $\mathcal{P}$
$\text{add}(a, p, R, \mathcal{F})$	assign the path $p$ to app $a$ with gross rate $R$ in the output $\mathcal{F}$ ; if the same path $p$ is already present in $\mathcal{F}$ for app $a$ , then $R$ is simply added to the gross rate already reserved
$\text{remove\_capacity}(R, p, C_e)$	remove the capacity $R$ from the capacities $C_e$ along the path $p$
$\text{update\_graph}(\mathcal{G}, C_e)$	update $\mathcal{G}$ by removing the edges with vanishing capacity in $C_e$

will be either assigned a path or removed from the active list. The **first step** (lines 14–17) is to move to the next app  $a$  if the credit  $\delta$  of the current one is exhausted, in which case the credit of the new app is initialized as  $\delta = \phi \frac{\rho_a}{\sum_{i \in \mathcal{A}} \rho_i}$ , i.e., a fraction of the round size  $\phi$  proportional to its priority. The **second step** (lines 18–25) is to identify one possible path to be assigned to the current app  $a$ ; this is done by selecting the shortest path  $p$  from  $\mathcal{P}_a$ , which contains the list of possible paths from  $h_a$  to any of its possible destinations in  $\mathcal{W}_a$ . The shortest path  $p$  is the one that requires the least amount of resources among those available for the current app, according to Eq. (2), and gives the maximum fidelity, according to Eq. (3). Since  $\mathcal{P}_a$  was determined with the initial graph  $\mathcal{G}'$ , and the latter will be updated as new paths are assigned to the apps, it is possible that  $p$  is no longer feasible under the updated graph: in this case, the path is removed from  $\mathcal{P}_a$  (line 22) and the loop is restarted. When the loop is done (line 25), the path  $p$  is either valid or there are not anymore feasible paths for the current app  $a$ . The two cases are handled in the **third and last step** (lines 26–35). If the app  $a$  does not have any more feasible paths (line 26), it is removed from the active list (line 27) and the credit for the next app is updated (line 28). Otherwise, if the path  $p$  is valid, the gross rate  $R$  is first determined as the minimum of the residual credit  $\delta$  and the minimum capacity along the path (line 30). Then, the path is added to the output (line 31), the selected gross rate  $R$  is removed from the capacities of all the edges along the path (line 32), the residual graph  $\mathcal{G}'$  is updated to reflect the new capacities (line 33), and the credit is updated (line 34). The functions called in Algorithm 1 are defined in Table II.

We now analyze the worst-case time complexity of Algorithm 1, separately for the initialization (lines 8–12) and the main loop (lines 13–35). With regard to the **initialization**, if  $A$  is the number of apps and  $W$  the average number of peers per host, we execute  $A \cdot W$  times the algorithm `find_paths`. The latter can be implemented with Yen’s algorithm to determine the  $k$  shortest paths in a graph [36], which with the help of a Fibonacci heap to keep edges sorted [14] is known to

have complexity  $\mathcal{O}(kV(E + V \log V))$ . A trade-off exists: if a large value of  $k$  is used then more paths will be found and the algorithm will be slower, but in the main loop there will be a wider choice of paths to be considered for each host-destination pair. Regarding the **main loop**, an exact upper bound of the time complexity is more difficult to derive, because it depends on the specific network topology and capacities. In particular, the number of iterations of lines 13–36 may depend on either the smallest capacity of edges along the candidate path or the credit given to each app, which in turn depends on the value of the system parameter  $\phi$ . Therefore, we can expect that another trade-off exists: smaller values of  $\phi$  lead to a slower algorithm, but also to a better fairness. This will be confirmed by the simulation results in the next section.

## V. PERFORMANCE EVALUATION

In this section we evaluate the performance of the resource allocation solution in Sec. IV-C via simulation. In the research community some simulation tools have been used and made publicly available to evaluate the performance of quantum networks. For example NetSquid<sup>1</sup> [9] and SeQuENCe<sup>2</sup> [35] are Python event-driven simulators that allow the user to customize the basic building blocks of quantum networks for the evaluation of protocols (e.g., quantum routing) and distributed applications (e.g., QKD); QuISP<sup>3</sup> [27], instead, is a full-fledged network simulator based on Omnet++<sup>4</sup> designed to study the dynamic behavior of large-scale quantum networks, e.g., congestion and stability aspects. However, to the best of our knowledge, none of the public simulators available so far address network resource provisioning, which is the subject of our work. Therefore, we have implemented a custom simulator, developed in C++ and using the Boost Graph Library [29], which is available as open source under a MIT license on GitHub<sup>5</sup>; for full reproducibility, we have also included the scripts to run the experiments, as well as the artifacts obtained and the Gnuplot files to produce the plots.

Like in [13], we use a Poisson Point Process (PPP) to generate the position of an average of  $\mu$  nodes in a flat square grid with edge size 60 km; a link is added between two nodes with probability  $p_{\text{link}}$  if their Euclidean distance is smaller than a threshold  $\tau$ . The capacity of each link is drawn from a r.v. uniformly distributed between 1 Bell pair/s and 400 Bell pairs/s, as in [6]. The initial fidelity of Bell pairs is  $\bar{F} = 0.95$ . Based on calibration experiments (not shown but available in the repository) we have selected  $p_{\text{link}} = 0.5$  and two possible values for  $\mu$  (50 and 100) and  $\tau$  (15 km and 20 km), which yield network characteristics reported in Table III.

We used a Monte Carlo approach: for any combination of the parameters under study, we simulated 10,000 drops with randomly generated networks and random workload. The latter

TABLE III  
CHARACTERISTICS OF THE SCENARIOS (AVERAGE VALUES).

	$\mu = 100,$ $\tau = 15$ km	$\mu = 100,$ $\tau = 20$ km	$\mu = 50,$ $\tau = 15$ km	$\mu = 50,$ $\tau = 20$ km
Capacity	159k	257k	47k	67k
Num edges	792	1282	232	331
Max degree	16	24	9	13
Diameter	8	5	10	6

TABLE IV  
METRICS USED FOR THE EVALUATION.

Metric	Definition
net rate	sum of the net rates assigned to peers, in Bell pairs/s
fidelity	Eq. (3) weighted on the net rates assigned to peers
visits	loop counter of Step 1-Step 6 in Sec. IV-C
fairness	Raj Jain's index on net rates $(\bar{r})^2 / (r^2)$

is made of a number of apps that depends on the experiment, with host node ( $h$ ) selected uniformly from all the nodes, and a set of peers of cardinality  $W$  sampled from all the nodes that can be reached by  $h$  in at most  $D$  hops. We set  $F^{\min} = 0$  for all apps, so that fidelity does not constrain resource allocation and can be evaluated *a posteriori*. Finally, it is always  $\rho = 1$ . The impact of constrained fidelity and heterogeneous priorities is left for future work. The metrics used are reported in Table IV. Statistical significance has been verified for all the metrics in the experiments performed, but we do not include error bars in plots for better readability.

We start with a first experiment where we increasingly add new apps until we meet a target residual capacity, defined as the sum of the  $C_{ij}$  of all the links at the end of the resource allocation. In Fig. 5 (left) we show the number of apps reached vs. the normalized residual rate from 0.1 (= only 90% of the capacity remains unused at the end of the allocation) to 1 (= fully used capacity). As can be seen, all the curves increase very steeply (the y-axis is in logarithmic scale) for low residual capacity: this shows that, despite the elastic allocation of resources to apps, it is very difficult to exhaust resources in a quantum network. This is because *reaching the capacity of a link through multiple hops requires an exponential increase of the resources used*, as in Eq. (2), which is in stark contrast with classical data networks. In the right part of Fig. 5 we report the fidelity. Interestingly, all the network combinations lead to a non-monotone pattern: when there is abundance of resources (low x-axis values) then only the best (i.e., shortest) paths are used; at intermediate loads, resource allocation tends to explore also longer paths, as  $F^{\min} = 0$  in the experiments; finally, with many applications, the only viable solution to reach all capacity is via short paths, which on average increases the fidelity again.

In a second batch of experiments we study the performance trade-offs of four key parameters by varying each one in a range while keeping the others constant to keep the analysis tractable. In particular, we simulate  $W \in \{2, \dots, 20\}$ ,  $D \in$

TABLE V  
QUALITATIVE RESULTS WITH 100 APPS VARYING  $W, D, k, \phi$ .

Metric	$W$	$D$	$k$	$\phi$
net rate	↑ linear	↓ asynt	↓ slight	≈
fidelity	↑ steep	↓ asynt	≈	≈
visits	↑ linear	↓ asynt	↑ linear	↓ exp
fairness	↑ steep	↓ asynt	≈	≈ then ↓ exp

<sup>1</sup><https://netsquid.org/>

<sup>2</sup><https://github.com/sequence-toolbox/SeQuENCe>

<sup>3</sup><https://github.com/sfc-aqua/quisp>

<sup>4</sup><https://omnetpp.org/>

<sup>5</sup><https://github.com/cciconnetti/quantum-routing/>, tag v1.0.

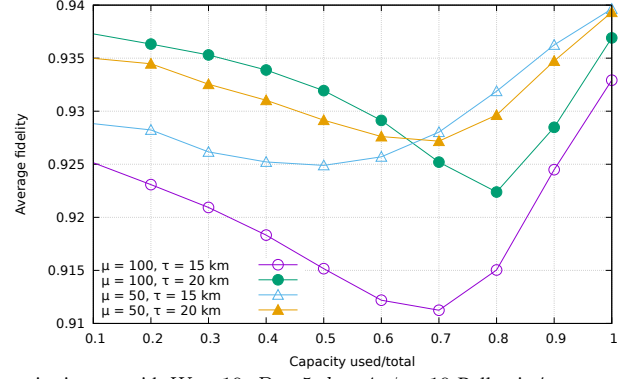



Fig. 5. Experiment with increasing apps until a target residual capacity is met, with  $W = 10$ ,  $D = 5$ ,  $k = 4$ ,  $\phi = 10$  Bell pairs/s.





Fig. 6. Net rate when increasing  $W$  (left) and  $D$  (right).


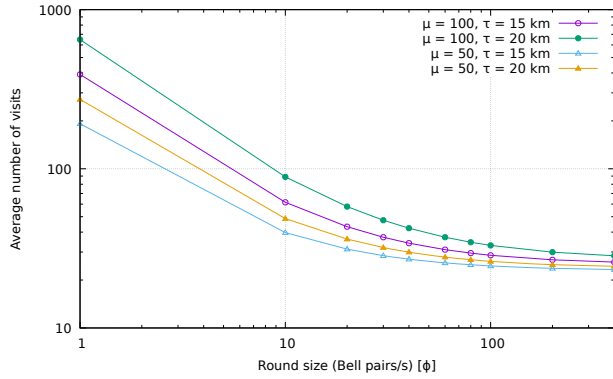


Fig. 7. Number of visits (left) and fairness Jain's index (right) when increasing  $\phi$ .

$\{1, \dots, 10\}$ ,  $k \in \{2, \dots, 20\}$ , and  $\phi/100 \in \{1, \dots, 400\}$  Bell pairs/s; the constant values are:  $W = 10$ ,  $D = 5$ ,  $k = 4$ , and  $\phi/100 = 10$  Bell pairs/s. We report in Table V how the increase of each key parameter affects the performance, from which we gather the following remarks:

- Having a larger pool of peers (i.e., large  $W$ ) is generally beneficial in terms of all the user-oriented performance metrics, but the resource allocation time complexity increases, so if execution time is a bottleneck then a trade-off arises. The net rate is shown in the left part of Fig. 6, which also indicates a clear correlation with the network density.
- The peers should be as close as possible to their respective hosts. In particular, as can be seen in Fig. 6 (right part), adding peers with distance above 4-5 does not yield no-

ticeable improvements in terms of the net rate, for all the combinations of the network parameters. We note that in practice an app's peers may depend on external factors, such as contracts between QC infrastructures and interface/technology compatibility limitation, so this parameter cannot be considered fully under the control of the quantum network operator or the users.

- $k$  can be small, as this yields both a higher net rate *and* reduces time complexity.
- As expected, a trade-off exists between the time complexity and fairness in the choice of  $\phi$ , which is a purely internal parameter of the algorithm. This is studied in quantitative terms in Fig. 7, suggesting that intermediate values achieve an excellent compromise.



## VI. CONCLUSIONS

In this work we have first provided an introduction to the emerging topics of quantum computing and networking. Then, we have raised awareness on distributed QC, which is a class of quantum applications that has not received significant attention so far compared to point-to-point applications, like QKD and sensing. We have elaborated on the distinguishing features of the two classes and proposed a resource allocation scheme for distributed QC that takes into account the fundamental aspects of end-to-end entanglement in quantum networks. Its performance has been evaluated with simulations, which have led to the identification of critical trade-offs, which will have to be studied in future research for efficient deployment and runtime optimization of quantum networks. Further open research areas are: the use of purification to increase fidelity at the expense of capacity; modeling distributed QC applications to understand their characteristic time scales and requirements; integration with link layer protocols.

## ACKNOWLEDGMENT

Work co-funded by EU, *PON Ricerca e Innovazione* 2014–2020 FESR/FSC Project ARS01\_00734 QUANCOM, and European High-Performance Computing Joint Undertaking (JU) under grant agreement No 101018180 HPCQS. The paper reflects only the authors' view and the funding agencies are not responsible for any use that may be made of its content.

## REFERENCES

- [1] D. Bertsekas and R. Gallager. *Data networks (2nd ed.)*. Prentice-Hall, Inc., USA, 1992.
- [2] D. Bouwmeester, J.-W. Pan, K. Mattle, M. Eibl, H. Weinfurter, and A. Zeilinger. Experimental quantum teleportation. *Nature*, 390(6660):575–579, Dec. 1997.
- [3] H.-J. Briegel, W. Dür, J. I. Cirac, and P. Zoller. Quantum repeaters for communication. *arXiv:quant-ph/9803056*, Mar. 1998. arXiv: quant-ph/9803056.
- [4] A. S. Cacciapuoti, M. Caleffi, R. Van Meter, and L. Hanzo. When Entanglement Meets Classical Communications: Quantum Teleportation for the Quantum Internet. *IEEE Transactions on Communications*, 68(6):3808–3833, June 2020.
- [5] M. Caleffi. Optimal Routing for Quantum Networks. *IEEE Access*, 5:22299–22312, 2017.
- [6] K. Chakraborty, D. Elkouss, B. Rijsman, and S. Wehner. Entanglement Distribution in a Quantum Network: A Multicommodity Flow-Based Approach. *IEEE Transactions on Quantum Engineering*, 1:1–21, 2020.
- [7] K. Chakraborty, F. Rozpedek, A. Dahlberg, and S. Wehner. Distributed Routing in a Quantum Internet. July 2019. arXiv: 1907.11630.
- [8] C. Cicconetti, M. Conti, and A. Passarella. Request Scheduling in Quantum Networks. *IEEE Trans. on Quantum Engineering*, 2, 2021.
- [9] T. Coopmans, R. Kneijens, A. Dahlberg, D. Maier, L. Nijsten, J. de Oliveira Filho, M. Papendrecht, J. Rabbie, F. Rozpedek, M. Skrzypczyk, L. Wubben, W. de Jong, D. Podareanu, A. Torres-Knoop, D. Elkouss, and S. Wehner. NetSquid, a NETWORK Simulator for QUantum Information using Discrete events. *Communications Physics*, 4(1):164, Dec. 2021.
- [10] D. Cuomo, M. Caleffi, K. Krsulich, F. Tramonto, G. Agliardi, E. Prati, and A. S. Cacciapuoti. Optimized compiler for Distributed Quantum Computing. *arXiv:2112.14139 [quant-ph]*, Dec. 2021.
- [11] A. Dahlberg, M. Skrzypczyk, T. Coopmans, L. Wubben, F. Rozpedek, M. Pompili, A. Stolk, P. Pawelczak, R. Kneijens, J. de Oliveira Filho, R. Hanson, and S. Wehner. A link layer protocol for quantum networks. In *Proc. ACM SIGCOMM*, pages 159–173, Beijing China, Aug. 2019.
- [12] A. Dahlberg, B. van der Vecht, C. D. Donne, M. Skrzypczyk, I. t. Raa, W. Kozłowski, and S. Wehner. NetQASM – A low-level instruction set architecture for hybrid quantum-classical programs in a quantum internet. *arXiv:2111.09823 [quant-ph]*, Nov. 2021. arXiv: 2111.09823.
- [13] W. Dai, T. Peng, and M. Z. Win. Optimal Remote Entanglement Distribution. *IEEE Journal on Selected Areas in Communications*, 38(3):540–556, Mar. 2020.
- [14] M. Fredman and R. Tarjan. Fibonacci Heaps And Their Uses In Improved Network Optimization Algorithms. In *25th Annual Symposium on Foundations of Computer Science, 1984.*, pages 338–346, Oct. 1984.
- [15] M. Katevenis, S. Sidiropoulos, and C. Courcoubetis. Weighted round-robin cell multiplexing in a general-purpose atm switch chip. *IEEE Journal on Selected Areas in Communications*, 9(8):1265–1279, 1991.
- [16] S. Khatri, A. J. Brady, R. A. Desporte, M. P. Bart, and J. P. Dowling. Spooky action at a global distance: analysis of space-based entanglement distribution for the quantum internet. *npj Quantum Information*, 7(1):4, Dec. 2021.
- [17] C. Li, T. Li, Y.-X. Liu, and P. Cappellaro. Effective routing design for remote entanglement generation on quantum networks. *npj Quantum Information*, 7(1):10, Dec. 2021.
- [18] S. Muralidharan, L. Li, J. Kim, N. Lütkenhaus, M. D. Lukin, and L. Jiang. Optimal architectures for long distance quantum communication. *Scientific Reports*, 6(1):20463, Apr. 2016.
- [19] M. A. Nielsen and I. L. Chuang. *Quantum computation and quantum information*. Cambridge University Press, Cambridge ; New York, 10th anniversary ed edition, 2010.
- [20] M. Pant, H. Krovi, D. Towsley, L. Tassiulas, L. Jiang, P. Basu, D. Englund, and S. Guha. Routing entanglement in the quantum internet. *npj Quantum Information*, 5(1):25, Dec. 2019.
- [21] R. Parekh, A. Ricciardi, A. Darwish, and S. DiAdamo. Quantum Algorithms and Simulation for Parallel and Distributed Quantum Computing. *arXiv:2106.06841 [quant-ph]*, June 2021. arXiv: 2106.06841.
- [22] I. Pedone, A. Atzeni, D. Canavese, and A. Liroy. Toward a Complete Software Stack to Integrate Quantum Key Distribution in a Cloud Environment. *IEEE Access*, 9:115270–115291, 2021.
- [23] S. Pirandola, B. R. Bardhan, T. Gehring, C. Weedbrook, and S. Lloyd. Advances in photonic quantum sensing. *Nature Photonics*, 12(12):724–733, Dec. 2018.
- [24] J. Preskill. Quantum computing 40 years later. *arXiv:2106.10522 [quant-ph]*, June 2021. arXiv: 2106.10522.
- [25] Quantum Technology and Application Consortium. Industry quantum computing applications. *EPJ Quantum Technology*, 8(1):25, Dec. 2021.
- [26] N. Sangouard, C. Simon, H. de Riedmatten, and N. Gisin. Quantum repeaters based on atomic ensembles and linear optics. *Reviews of Modern Physics*, 83(1):33–80, Mar. 2011.
- [27] R. Satoh, M. Hajdušek, N. Benchasattabuse, S. Nagayama, K. Teramoto, T. Matsuo, S. A. Metwalli, T. Satoh, S. Suzuki, and R. Van Meter. QuISP: a Quantum Internet Simulation Package. *arXiv:2112.07093 [quant-ph]*, Dec. 2021. arXiv: 2112.07093.
- [28] J. Sevilla and C. J. Riedel. Forecasting timelines of quantum computing. *arXiv:2009.05045 [quant-ph]*, Sept. 2020. arXiv: 2009.05045.
- [29] J. Siek, L.-Q. Lee, and A. Lumsdaine. *The boost graph library*. Addison-Wesley Professional, 2002 edition.
- [30] N. Tomm, A. Javadi, N. O. Antoniadis, D. Najer, M. C. Löbl, A. R. Korsch, R. Schott, S. R. Valentin, A. D. Wieck, A. Ludwig, and R. J. Warburton. A bright and fast source of coherent single photons. *Nature Nanotechnology*, Jan. 2021.
- [31] R. Van Meter. *Quantum networking*. Networks and telecommunications series. ISTE ; Wiley, 2014.
- [32] R. Van Meter, R. Satoh, N. Benchasattabuse, T. Matsuo, M. Hajdušek, T. Satoh, S. Nagayama, and S. Suzuki. A Quantum Internet Architecture. *arXiv:2112.07092 [quant-ph]*, Dec. 2021. arXiv: 2112.07092.
- [33] R. Van Meter, T. Satoh, T. D. Ladd, W. J. Munro, and K. Nemoto. Path Selection for Quantum Repeater Networks. *Networking Science*, 3(1-4):82–95, Dec. 2013. arXiv: 1206.5655.
- [34] S. Wehner, D. Elkouss, and R. Hanson. Quantum internet: A vision for the road ahead. *Science*, 362(6412):eaam9288, Oct. 2018.
- [35] X. Wu, A. Kolar, J. Chung, D. Jin, T. Zhong, R. Kettimuthu, and M. Suchara. SeQuENCE: a customizable discrete-event simulator of quantum networks. *Quantum Science and Tech.*, 6(4):045027, Oct. 2021.
- [36] J. Y. Yen. Finding the K Shortest Loopless Paths in a Network. *Management Science*, 17(11):712–716, 1971. Publisher: INFORMS.
- [37] Y. Zhao and C. Qiao. Quantum Transport Protocols for Distributed Quantum Computing. May 2021. arXiv: 2105.08109.