Quantum Networks with Multiple Service Providers: Transport Layer Protocols and Research Opportunities

Maoli Liu, Jonathan Allcock, Kechao Cai, Shengyu Zhang, and John C. S. Lui

Abstract

A quantum Internet for communicating information encoded in quantum systems over large distances would enable a host of new technologies to be deployed. While exciting progress has been made in small-scale quantum networks, a global quantum Internet will require — like the classical Internet — communicating in a multiple quantum Internet service provider (multi-qISP) setting, with protocols that are oblivious to the global network topology. Here we give a brief overview of quantum networks aimed at those new to the field, present two network-oblivious transport layer protocols for enabling high fidelity communication and performing fault diagnostics, and highlight a number of research opportunities in this nascent field of research.

Introduction

Since its genesis as ARPAnet in the late 1960s, the classical internet has had a transformative impact on services and society. Now, with advances in quantum information and computing technologies gathering pace, research is underway into a new kind of communication network, namely a *quantum internet* [1, 2]. Such a network would consist of nodes that are able to share *entanglement*, a resource for transmitting quantum information by a process known as *quantum teleportation*.

The ability to send information encoded in the state of quantum systems would enable a multitude of new services such as secure communication, high precision clock synchronization, distributed and blind quantum computation and quantum telemetry. However, the creation of a quantum internet necessitates a rethink and redesign of the network protocols which underpin the classical internet, to support the features and limitations governed by quantum mechanics.

Fortunately, exciting work is already underway in this area, with progress at different layers of the quantum networking stack. At the physical layer, entanglement has been demonstrated at distances of tens of kilometers in optical fibers [3, 4]. A link layer protocol [5] has been proposed to provide robust entanglement generation service between physically connected quantum nodes. At the network layer, [6–9] discuss entanglement routing problems, and [10] proposes a quantum retransmission protocol at the transport layer.

A number of key differences between classical and quantum information impact the design of quantum networks. In contrast to classical bits, which are either in state 0 or 1, the basic unit of quantum information is the qubit which can exist as a superposition of $|0\rangle$ and $|1\rangle$ states (using *Dirac notation* to denote quantum states) at the same time. However, measurement of a qubit will collapse its state to either $|0\rangle$ or $|1\rangle$, destroying the superposition. Furthermore, by the *no-cloning theorem*, arbitrary quantum states cannot be copied. These properties have profound implications for network protocol design and preclude, for instance, the direct translation of classical protocols which rely on reading, copying and storing information at intermediate devices such as gateways or routers.

Another complicating factor is the fragility of quantum information. Quantum states are susceptible to *decoherence* (decay) via undesirable interactions with the environment, and can additionally be corrupted by errors at different stages of processing. While techniques can be used, for example, at the transport layer [10] to mitigate these errors, the general inability to copy quantum information remains a handicap. In particular, if the quantum state to be sent is particularly valuable, say, the result of a long and expensive quantum computation, a faulty transmission may necessitate the complete recomputation of the state, making the information processing highly inefficient.

As a consequence, the best-effort service model of the classical internet — predicated on the premise that copying is easy and retransmission is cheap — may not always be appropriate for quantum communication. When the information to be transmitted is valuable and hard to replicate, an almost-guaranteed delivery model based on temporarily reserved paths through the network may be more suitable.

Looking to the future, although the equivalent of a quantum internet service provider (qISP) has not yet materialized, an eventual large-scale quantum internet connecting users across the globe will also likely require communicating in a multi-qISP setting. This would necessitate protocols at various layers in the networking stack to be oblivious to the global network topology.

In this work, we consider requirements at the

Digital Object Identifier:
10.1109/MNET.001.2200151

Maoli Liu and John C.S. Lui are with the Chinese University of Hong Kong; Jonathan Allcock and Shengyu Zhang are with Tencent Quantum Laboratory, Hong Kong; Kechao Cai is with Sun Yat-Sen University, China.

0890-8044/22/$25.00 © 2022 IEEE

transport layer for such scenarios, with the aim of enabling network-oblivious high-fidelity quantum communication in a multi-qISP setting. In particular, we give:

- Protocol[1] for verifying the fidelity and rates of *virtual links* — connections corresponding to entanglement — between nodes in a quantum network.
- Protocol for qISPs to perform virtual link fidelity diagnostics to discover faults within their own networks, based on error signals received from a sender or upstream qISP. Following that, we present a number of research opportunities related to quantum communication in a multi-qISP setting.

Quantum Communication

A good introduction to quantum information from a networking perspective can be found in [5]. Here we briefly summarize the key concepts relevant to this article.

At the heart of quantum communication is the concept of *entanglement*, which is a property of the joint state of multiple qubits which can be used as a resource for transmitting states via teleportation. In teleportation, a source node *S* and a destination node *D* share an entangled pair. In addition, *S* has an information qubit in state $|\psi\rangle$. By a combination of quantum operations at *S* and *D*, coordinated via classical communication, the state $|\psi\rangle$ can be transferred to *D*, consuming the entanglement in the process.

While teleportation can be performed over arbitrary distances, and while entanglement distribution through free-space has been demonstrated over impressive distances (~ 100 km over terrestrial free-space [11] and ~ 1,200 km through satellite-to-ground links [12]), }the distribution of entanglement over large distances in fibers — likely a necessity for eventual large-scale networks — presents a major challenge. For instance, if an entangled pair is generated at *S*, and one half of the pair sent to *D* through an optical link, attenuation in the fiber leads to a probability of successful transmission that decreases exponentially with distance. Intriguingly, teleportation provides a solution to this problem by way of *quantum repeaters* placed at intermediate locations in the network.

Consider *S* and *D* too far separated to directly share entanglement reliably. A repeater node R is placed midway between S and D, and connected to them via optical fibers. If R shares one entangled pair with S, and another entangled pair with *D* (Fig. 1a), R can perform an *entanglement swapping* operation, which teleports the entangled pair half shared with *S* to *D*, leading finally to entanglement shared between *S* and *D* (Fig. 1b). In spite of *S* and *D* sharing no direct physical link, the entanglement produced by this process constitutes a *virtual link* being established between *S* and *D*. As classical communication between nodes is needed for teleportation and coordination between nodes, one typically envisages such a quantum network as lying on top of classical networks.

A large-scale quantum network can therefore be constructed from many quantum nodes and repeaters connected by physical links (e.g., optical fibers). Virtual links can then be created between any two nodes by performing entanglement swapping repeatedly along a path of physical links

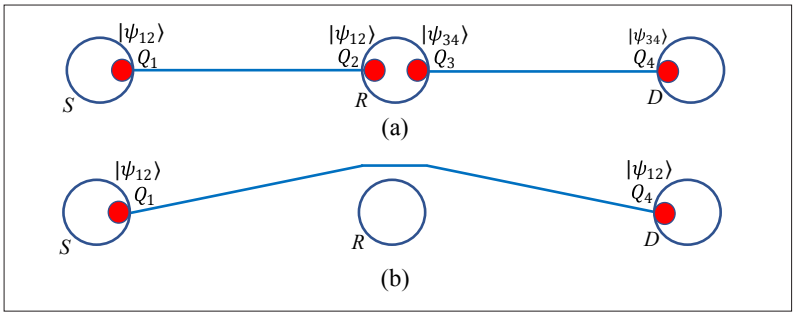

FIGURE 1. Using a quantum repeater to overcome distance restrictions and generate long-range entanglement. While the quantum network can only directly create entanglement via physical links (e.g., fiber optic connections) between nearest neighbors, after entanglement swapping, virtual links in the form of entanglement are established between distant nodes.

connecting the nodes. The selection of an optimal path for swapping is a network layer routing problem and an area of active research [6, 9, 13].

As streams of qubits may need to be sent across the network, a metric of interest is the *entanglement rate* of a virtual link associated with a path, that is, the number of entangled pairs per second that can be generated between source node and destination node via entanglement swapping along the path.

Even when teleportation between arbitrary nodes is possible, it is still a challenge to design a robust quantum internet. As with the classical internet, a layered architecture is most likely necessary to ensure reliable performance. Such a layered architecture, for example, as proposed by [5], would delegate responsibility for low quality entanglement generation to the physical layer, high quality entanglement generation between neighboring nodes to the link layer, path routing for entanglement swapping and distant entanglement generation to the network layer, and information transfer between arbitrary nodes to the transport layer.

Introduction to Network Benchmarking

There are many potential sources of error when establishing virtual links between source and destination in a quantum network. Losses in optical fibers, imperfect swapping operations, and decoherence during storage can all lead to sender and receiver sharing imperfect entangled pairs and ultimately leading to errors in transmission.

The quality of quantum state transmission through noisy channels is commonly measured by a metric known as *fidelity*. This is a value between 0 and 1 that captures how well a channel preserves states sent through it, with perfect transmission corresponding to a fidelity of 1. The *network benchmarking* protocol recently introduced by Helsen and Wehner [14] is a way of efficiently estimating fidelities of channels in a quantum network, given repeated access to those channels. This method is robust to state preparation and measurement errors, and flexible in that it can be applied to channels at various levels of the network stack.

Here we are interested in using network benchmarking as a subroutine in transport layer protocols to ensure reliable end-to-end communication. When virtual links are established, network benchmarking can be used to obtain the fidelity

[1] We use the word "protocol" to indicate that our proposed procedures consist of general quantum network system calls.

```
Input : S_ID, D_ID, f, e_rate
Output: status
paths ¬ req_links(S_ID, D_ID, qISP, f, e_rate)
good_paths, rate_set ¬ ver_links(paths, f)
if Σ_{r∈rate_set} r ≥ e_rate then
    status ¬ quantum_send(good_paths)
else
    e_rate_signal(good_paths, e_rate, qISP)
    status ¬ e_rate_failure
fidelity_Signal(paths \ good_paths, f, qISP)
release_signal(paths, qISP)
return status
```

ALGORITHM 1. Virtual Link Verification (VLV) Protocol.

of each link, so that only links which satisfy the user's fidelity requirements can be used. Note that by "fidelity of a virtual link" we mean the fidelity of the quantum channel corresponding to teleportation using the entanglement across the virtual link. Furthermore, as network benchmarking of virtual links does not require knowledge of how those virtual links were created (i.e., the entanglement swapping path), such fidelity estimation is network oblivious.

Network benchmarking estimates the average fidelity of an arbitrary channel[2] in a network based on the idea that:

- The fidelity of *depolarizing channels* — a special channel which, with probability p leaves the input state unaffected, and with probability $1 - p$ the state is replaced with the maximally mixed state, the quantum equivalent of a uniformly random bit — can be readily estimated if one has repeated access to the channel.
- Arbitrary quantum channels can be effectively transformed into depolarizing channels of the same fidelity by a process known as channel twirling, which involves random application of quantum operations from a certain operation set.

By estimating the parameter p of the depolarizing channel that results from the twirling procedure, one can deduce the fidelity of the original channel. Other commonly considered channels include the *dephasing channel*, which models decoherence due to scattering, and the *amplitude damping channel*, which models decoherence due to emission into the environment.

Network benchmarking is based on the concept of a *bounce*, which consists of:

- The source node S applying a random operation to a quantum state
- Teleporting the state to the destination node D
- D applying a random quantum operation to the received state
- Then teleporting it back to S.

A state may be bounced back and forth multiple times before S finally applies a corrective operation (which depends on the random operations applied during the bouncing) and measures the state. The average measurement outcome is then used to estimate the average fidelity of the channel. More specifically, one chooses a set of integers $\mathbf{M} = \{m_1, m_2, ..., m_M\}$ (the bounce number set) and a set of corresponding integers $\mathbf{N} = \{n_{m_1}, n_{m_2}, ..., n_{m_M}\}$, and performs the following steps (see [14] for details).

Step 1: Choose an integer $m \in \mathbf{M}$.
Step 2: Initialize the qubit at node S in a fixed state.
Step 3: For i from 1 to m:
Bounce the state from S to D and back.
Step 4: S applies a corrective operation and measures the final state.
Step 5: Repeat steps 2–4 n_m times and calculate the average survival probability b_m.
Step 6: Repeat steps 2–5 for all different values of $m \in \mathbf{M}$, and record corresponding average values b_m.
Step 7: Perform a regression on the observed data $\{b_m\}_{m\in\mathbf{M}}$ and $\mathbf{M}$ using the model $b_m = A_p{}^m + B$ to estimate the depolarizing parameter p of the twirled channel.

From p, the average channel fidelity $f = p + (1 - p)/2$ of the twirled channel is deduced, which is equal to the average fidelity f of the original, untwirled channel. The maximum number of bounces required to accurately estimate the fidelity to low variance depends on the underlying fidelity, with low fidelities requiring a smaller number of bounces.

Note that a bounce consumes two virtual links. Under the assumption of Markovianity, that is, noisy virtual links established using the same path always correspond to the same quantum channel, network benchmarking allows one to deduce the average fidelity of these virtual links. In the process, virtual links between S and D are repeatedly established along the same path and consumed.

Link Selection and Diagnostic Protocols

We now present two network-oblivious transport layer protocols. The first is a virtual link verification (VLV) protocol, that enables a sender S to verify the fidelity and rate of virtual links for sending quantum states to a destination node D. The second is a distributed fault discovery (FD) protocol for qISPs to perform entangled link fault diagnostics. Our protocols make use of the following concepts.

path_id: a unique identifier associated with each entanglement-swapping path from S to D. The *path_id* gives qISPs involved in the path access to information via the following set of API calls:

- *l_path* ¬ `local_path`(*path_id*): returns the segment of the path that passes through their part of the network. We assume that the original *path_id* is also stored as meta-data in *l_path*, so that qISPs may recover the original *path_id* by calling `parent_path`(*l_path*).
- *node_id*, *qisp_id* ¬ `ingress`(*path_id*): returns:
  - The node identifier *node_id* of the first node in the immediate upstream qISP on the path from S to D used to generate the virtual link
  - The corresponding upstream qISP identifier *qisp_id*.
- *node_id*, *qisp_id* ¬ `egress`(*path_id*): the same as `ingress` except it identifies the first node in the immediate downstream qISP.

The above API calls, as well as a number of others described below, are presented at a high level in terms of functionality without restricting to particular implementation details.

Virtual Link Verification Protocol

To send quantum information, S first requests its connected qISP to provide a number of virtual

[2] In general, network benchmarking can be applied to channels connecting multiple nodes in a network. As we are concerned with benchmarking virtual links between source node and destination node, we will restrict our attention to the 2-node case.

links to *D*. These virtual links are formed from paths selected at the network layer, and can be either pre-configured entangled paths managed by the underlying qISPs which collaboratively maintain network layer entangled paths between *S* and *D* at all times [8], or on-demand entangled paths that are dynamically generated by the underlying qISPs upon request from *S*. VLV protocol does not depend on any specific routing protocol. The only requirement is that paths used to generate virtual links cannot be changed before users complete the verification and information transmission.

The VLV protocol is depicted in Algorithm 1. In addition to the API calls described at the beginning of this section, VLV makes use of a number of other calls:

- *paths* ¬ `req_links`(*S_ID*, *D_ID*, *qISP*, *f*, *e_rate*): this API call is sent from *S* to its connected qISP, and returns a set of *path_id* values corresponding to virtual links from *S* to *D* that each individually have a link fidelity of at least *f* (or are believed by the qISPs to satisfy this), and collectively have a rate of at least *e_rate*. In order to guarantee that fidelity testing and state transmission can be performed using virtual links generated by the same paths through the network each time, we assume that the paths in the set *paths* will be reserved until the user signals to qISP to release them.
- *good_paths*, *rate_set* ¬ `ver_links`(*path_set*, *f*): this takes a set of *path_id* values and a target fidelity *f*, and returns
  – the set *good_paths* of *path_id* values with measured fidelity of at least *f*
  – the set *rate_set* of measured entanglement rates corresponding to each virtual link with *path_id* in *good_paths*.
  Note that the fidelity checking part of `ver_links` can be implemented in many ways. One option, and the one we consider below, is to use network benchmarking.
- `quantum_send`(*path_set*): initiates the transmission of the quantum message from *S* to *D*, using virtual links corresponding to paths in *path_set*; the return value is either *success* or *failure*.
- `e_rate_signal`(*path_set*, *e_rate*, *qISP*): signals to qISP that the total measured entanglement rate of links in *path_set* is below *e_rate*.
- `fidelity_signal`(*path_set*, *f*, *qISP*): signals to *qISP* that the paths in *path_set* have measured fidelities lower than *f*.
- `release_signal`(*path_set*, *qISP*) : signals to qISP that the protocol has completed, and the reserved paths can be released.

To invoke the protocol, *S* provides its source node ID (*S_ID*), the destination node ID (*D_ID*), minimum fidelity requirement (*f*) and total entanglement rate required (*e_rate*) for its quantum information transmission. S then requests a set of virtual links from its connected qISP that satisfy the fidelity and rate requirements (line 1), and calls `ver_links` to check whether the fidelity of each link corresponding to those paths is at least *f* (line 2), and also checks the measured rates of the virtual links provided.

The *path_id* values of paths whose link fidelities exceed the required threshold are stored in the variable *good_paths*, and their corresponding measured entanglement rates are stored in the variable *rate_set*. If the total entanglement rate for links in *good_paths* is above the required *e_rate*, then *S* will invoke `quantum_send` to transmit its quantum information using links corresponding to paths in *good_paths* (line 4). If not, *S* will inform its connected qISP via `e_rate_signal` (line 6), and set the status as *e_rate_failure* (line 7). *S* then calls `fidelity_signal` to inform its connected qISP about those links which are below the fidelity required (line 8), and releases the reserved paths (line 9).

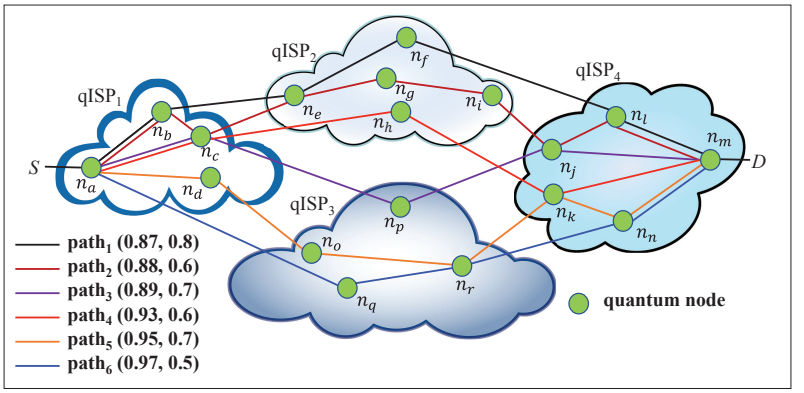

FIGURE 2. Paths corresponding to six virtual links between *S* and *D*. Each path is associated with a unique *path_id* identifier and defines a different entanglement swapping route through the multi-qISP network. The tuples labelled next to each path indicate the (fidelity,rate) of the corresponding virtual link.

If network benchmarking is used in the implementation of `ver_links`, numerous optimizations can be performed to speed up and reduce the cost of this process. Consider the example in Fig. 2 with six virtual links between *S* and *D*. One can consider at least the following approaches.

Sequential Network Benchmarking: performs benchmarking first on $link_1$, then $link_2$, and so on, and finally link6.

Parallel Network Benchmarking: performs benchmarking on all links at the same time. One can determine qualified links more quickly, with the trade-off of a higher computational cost for nodes in the network.

Progressive Parallel Network Benchmarking: here the bounce number set is partitioned into subsets. Parallel benchmarking is performed using bounce numbers from each subset in turn, with low fidelity links identified at each stage and excluded from further testing.

Fault Discovery Protocol

When the connected qISP receives the `fidelity_signal` call from *S* about the links which do not satisfy the fidelity requirement, a procedure is needed to identify and correct for any errors which caused this.

Compared with classical networks, identifying sources of error in the quantum case is significantly more complex. In addition to the multiple potential physical sources of errors, as quantum errors need not be discrete, a poor overall link fidelity could be the result of many small errors compounding at each step. Thus, the very concept of identifying *the* source of error may not be meaningful. However, under certain scenari-

```
Input : faulty_path_set, f
Output: err_code
1  (l_src, l_path, egrs_node, ds_qISP)} ¬ determine_borders (faulty_path_set)
2  foreach(l_src, l_path, egrs_node, ds_qISP) do
3      flag ¬
4          test_fidelity(l_src, egrs_node, l_path, f)
5      if(flag == TRUE then
6          /* fault assumed at downstream ISPs */
7          err_code ¬ NO_FAULT
8          path ¬ parent_path(l_path)
9          fidelity_signal(path, f, ds_qISP)
10     else
11         /* perform local node/link fidelity check */
12         err_code ¬ intl_check()
13         if err_code! = NO_FAULT then
14             /* fault located at internal network */
15             fix_intl_fault(err_code)
16         else
17             /* faults at downstream ISP */
18             err_code ¬
19             egress_link_check(l_path, egrs_node)
20             if err_code! = NO_FAULT then
21                 fix_intl_fault(err_code)
22             else
23                 fidelity_signal(path, f, ds_qISP)
```

ALGORITHM 2. Fault Discovery (FD) Protocol.

os we may make some reasonable assumptions. For instance, if a virtual link has observed fidelity significantly lower than it is expected to have then, under ordinary network operating conditions (assuming the qISPs have well-maintained networks), a single large fault may be assumed to be the most likely cause. In this case, a simple fault discovery (FD) protocol such as we give in Algorithm 2 can be used by the qISPs to locate and correct the dominant issue.

We first define several additional API calls:

- `determine_borders`(*path_set*): this call is made by the qISP and takes as argument a set *path_set* of *path_id* values. For each path_id, it returns a tuple (*l_src*, *l_path*, *egrs_node*, *ds_qISP*) where
  - *l_src* is the local node in the network controlled by the qISP that is connected to the ingress node;
  - *l_path* ¬ `local_path`(*path_id*);
  - *egrs_node*, *ds_qISP*¬ `egress`(*path_id*).
- `test_fidelity`(*l_src*, *l_path*, *egrs_node*, *f*): tests (e.g., by network benchmarking) whether the fidelity of the virtual link between *l_src* and *egrs_node* created by entanglement swapping along *l_path* is at least *f*. Again, we assume that the path is reserved for the duration of the testing.
- *err_code* ¬ `intl_check`(): instructs the qISP to perform internal checks on its network; returns an error code *err_code*, with *NO_FAULT* indicating that no internal issues were detected.
- `fix_intl_fault`(*err_code*): takes *err_code* returned by `intl_check` and instructs the qISP to perform its own proprietary fault correction.
- *err_code* ¬ `egrs_link_check`(*l_path*, *egrs_node*): instructs the qISP to perform testing of the virtual link directly connecting the *l_path* to the egress node; returns an error code *err_code* indicating the nature of any local fault detected in the creation of this virtual link, that is, that the responsibility lies with the qISP that initiated the `egrs_link_check` call.

At a high level, for each virtual link from *S* to *D* that does not have fidelity at least *f*, the user's connected qISP will make a `test_fidelity` API call on the virtual link between the source and egress node, that is, the first node in the immediate downstream qISP on the path used to generate the virtual link, returning TRUE if the fidelity of that virtual link is at least f. If this is the case, then the fault is assumed to lie in part of the network belonging to a downstream qISP, and a signal is sent to the downstream qISP to run their own FD protocol. Otherwise, the fault is assumed to lie in the original qISP's network, or in the link between the qISP and the egress node in the immediate downstream qISP, and internal fidelity checking and correction should take place. In this way, the signal to perform error testing is passed sequentially down the network until the first major error is found. Note that although the protocol calls qISPs sequentially along paths from S to D, each qISP is able to test multiple faulty paths through their own parts of the network in parallel. How best to do this to minimize the disruption to the network caused by the testing is an interesting open problem.

Evaluation

We now carry out a number of evaluations to illustrate the scope for optimization in multi-qISP protocols. We consider a network corresponding to Fig. 2, which has six paths between source na and destination n_m, with fidelities and rates given in the figure. We assume the source has set a minimum rate of *e_rate* = 2.0 and a minimum fidelity of $f = 0.9$ for transmission, which exceeds the true fidelities of paths 1, 2, 3. Simulations are performed using Netsquid [15].

We first numerically simulate network benchmarking of the six virtual links, assuming various noise models corresponding to depolarizing, dephasing, and amplitude-damping channels. Results are given in Fig. 3, and show good agreement between the true and estimated fidelities. We choose a bounce set $\mathbf{M} = \{2, 3, ..., 50\}$ and $n_m = 40$ for all $m \in \mathbf{M}$. While this is sufficient to estimate all fidelities to low variance, the total number of bounces required (50,960) is large. Performing the network benchmarking of all six links sequentially would require both a large total number of bounces, as well as a large total protocol running time, assuming that network benchmarking bounces can be performed at the rate of each link (Table 1, time is in arbitrary units). Here and in the remainder of this section, we investigate the benefits of parallelizing the testing, with results corresponding to the case of dephasing channel noise.

Parallel Testing

As mentioned earlier, parallel and progressive parallel network benchmarking can be used to reduce the total time and total number of bounces required for fidelity testing by network benchmarking. The effect of these can be seen in Table

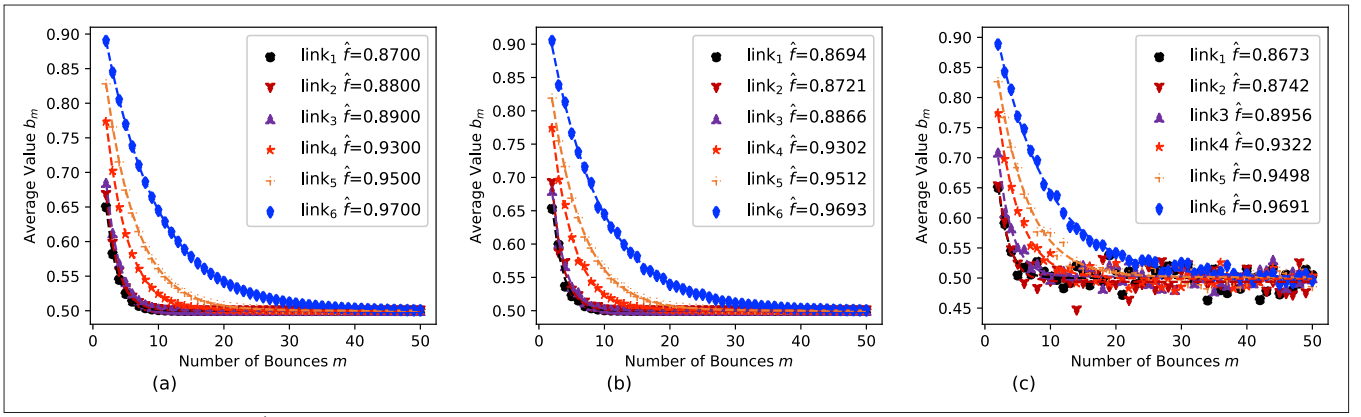

FIGURE 3. Average fidelities $\hat{f}$ of virtual links estimated by network benchmarking for noise corresponding to a: a) depolarizing channel: b) dephasing channel: c) amplitude-damping channel. True fidelities of the links 1 to 6 are {0.87, 0.88, 0.89, 0.93, 0.95, 0.97}, respectively. Bounce set $\mathbf{M}$ = {2, 3, ..., 50}, n_m = 40.

1. Parallel network benchmarking of all six links (assuming the same M, nm as before) reduces the total time required to the time needed for the link with the lowest rate. However, the total number of bounces required is the same as for sequential network benchmarking. For progressive parallel network benchmarking, we partition the bounce set into $\mathbf{M}_1$ = {2, 3, 4, ..., 20} and $\mathbf{M}_2$ = {21, 23, 24, ..., 50} , and first perform network benchmarking of all six links in parallel using $\mathbf{M}_1$ to get a first estimate of the link fidelities (recall that low fidelities can be estimated accurately using fewer bounces). These results indicate that path$_1$ and path$_2$ have fidelities below the threshold of 0.9, and can thus be excluded from further testing. Additional testing of the remaining links is performed with bounce numbers from $\mathbf{M}_2$, leading to a net reduction in both the time and total bounces compared with sequential network benchmarking.

Fault Discovery

Progressive parallel network benchmarking can also be applied to network benchmarking used in the `test_fidelity` API call in the FD protocol.

After S informs qISP$_1$ that the link fidelities associated to path$_1$, path$_2$, and path$_3$ are below f, qISP$_1$ invokes the FD protocol to check faulty paths. For simplicity, let us focus only on path$_1$. Suppose the sole error occurs at node n_l in qISP$_4$. According to FD protocol, qISP1 will check the fidelity of virtual link $n_a \leftrightarrow n_e$; qISP$_2$ will check the fidelities of virtual links $n_e \leftrightarrow n_l$ and egress links $n_e \leftrightarrow n_f$ and $n_f \leftrightarrow n_l$; and qISP$_4$ will finally check the fidelity of virtual link $n_l \leftrightarrow n_m$.

The total number of bounces for each qISP along path$_1$ is listed in Table 2, and again shows the advantage that progressive parallel network benchmarking can bring. Note that the corresponding amounts of time required depend on the rates of virtual link creation along sub-paths in the network, which we do not model here.

Open Research Questions

The prospect of quantum communication in a multi-qISP setting raises a number of interesting research opportunities:

- Our protocols require the reservation of paths for a short duration. Resource reservation is non-trivial even for the classical internet. Can we learn from protocols like RSVP and soft-state design principles to design a robust reservation protocol?
- Our protocols require a mapping between the virtual end-to-end links and the underlying paths for all involved qISPs. How can this be done in a confidential way so that it is not possible to infer the internal network structure of any qISP?
- Can statistical or online learning be leveraged to efficiently determine fidelities and entanglement rates?
- How can classical internet tools like ping and traceroute, useful for determining packet loss rate and node-to-node transmission time, best be generalized to the quantum setting?
- Conducting network benchmarking before sending quantum information consumes extra quantum resources, but helps identify links with required fidelity. How can we improve the protocols to balance the trade-off between network throughput and resources used for fidelity verification?
- How can one dynamically adjust to errors, rather than simply classifying temporarily noisy links as “faults”?

Conclusion

We envision a future quantum internet operated by multiple qISPs, which provides an added challenge to designing network protocols. In this work, we present two network oblivious transport layer protocols: a virtual link verification protocol that determines the fidelity of an end-to-end link, and a fault discovery protocol so different qISPs can distributively discovery faults within their networks. However, many open questions remain, and multi-qISP quantum communication presents an abundance of research opportunities.

Acknowledgments

The work of John C.S. Lui was supported in part by the RGC grant SRFS2122-4S02. The work of Kechao Cai was supported in part by National Key R&D Program of China under Grant 2021YFB2900200 and by the Fundamental Research Funds for the Central Universities, Sun Yat-sen University.

References

[1] H. J. Kimble, “The Quantum Internet,” *Nature*, vol. 453, no. 7198, 2008, pp. 1023–30.

[2] S. Wehner, D. Elkouss, and R. Hanson, "Quantum Internet: A Vision for the Road Ahead," *Science*, vol. 362, no. 6412, 2018, p. eaam9288.
[3] V. Krutyanskiy *et al.*, "Light-Matter Entanglement Over 50 Km of Optical Fibre," *Quantum Information*, vol. 5, no. 1, 2019, pp. 1–5.
[4] Y. Yu *et al.*, "Entanglement of Two Quantum Memories via Fibres Over Dozens of Kilometres," *Nature*, vol. 578, no. 7794, 2020, pp. 240–45.
[5] A. Dahlberg *et al.*, "A Link Layer Protocol for Quantum Networks," *Proc. ACM Special Interest Group on Data Commun., ser. SIGCOMM '19*, 2019, p. 159,173.
[6] R. Van Meter *et al.*, "Path Selection for Quantum Repeater Networks," *Networking Science*, vol. 3, no. 1, 2013, pp. 82–95.
[7] M. Pant *et al.*, "Routing Entanglement in the Quantum Internet," *Quantum Information*, vol. 5, no. 1, 2019, pp. 1–9.
[8] S. Shi and C. Qian, "Concurrent Entanglement Routing for Quantum Networks: Model and Designs," *Proc. Annual Conf. ACM Special Interest Group on Data Commun. Applications, Technologies, Architectures, and Protocols for Computer Commun., ser. SIGCOMM '20*, 2020, p. 6275.
[9] C. Li *et al.*, "Effective Routing Design for Remote Entanglement Generation on Quantum Networks," *Quantum Information*, vol. 7, no. 1, 2021, pp. 1–12.
[10] N. Yu, C.-Y. Lai, and L. Zhou, "Protocols for Packet Quantum Network Intercommunication," *IEEE Trans. Quantum Engineering*, vol. 2, 2021, pp. 1–9.
[11] J. Yin *et al.*, "Quantum Teleportation and Entanglement Distribution Over 100-Kilometre Free-Space Channels," *Nature*, vol. 488, no. 7410, 2012, pp. 185–88.
[12] J. Yin *et al.*, "Satellite-Based Entanglement Distribution Over 1200 Kilometers," *Science*, vol. 356, no. 6343, 2017, pp. 1140–44.
[13] K. Chakraborty *et al.*, "Entanglement Distribution in a Quantum Network: A Multicommodity Flow-Based Approach," *IEEE Trans. Quantum Engineering*, vol. 1, 2020, pp. 1–21.
[14] J. Helsen and S. Wehner, "A Benchmarking Procedure for Quantum Networks," arXiv preprint arXiv:2103.01165, 2021.
[15] T. Coopmans *et al.*, "Netsquid, A Network Simulator for Quantum Information Using Discrete Events," *Commun. Physics*, vol. 4, no. 1, 2021, pp. 1–15.

Biographies

Maoli Liu is a Ph.D. candidate in the Department of Computer Science & Engineering at the Chinese University of Hong Kong, under the supervision of Prof. John C.S. Lui. Her current research area includes online learning theory and quantum internet.

Jonathan Allcock is a Senior Researcher at the Tencent Quantum Laboratory in Hong Kong. Prior to that, he obtained his Ph.D. from the University of Bristol under the supervision of Prof. Noah Linden, and was a postdoctoral research fellow at the Chinese University of Hong Kong.

Kechao Cai received the Ph.D. degree at the Chinese University of Hong Kong in 2019. He is currently an assistant professor in the School of Electronics and Communication Engineering at Sun Yat-Sen University in China. His current research area includes online learning algorithms and network protocols.

Shengyu Zhang is a Distinguished Scientist and founding director of Tencent Quantum Laboratory. He obtained his Ph.D. from Princeton University in 2006, under the supervision of Prof. Andrew Yao. His research interest lies in quantum computing, algorithm designing, computational complexity, and the foundation of artificial intelligence.

John C.S. Lui is currently the Choh-Ming Li Chair Professor at the Chinese University of Hong Kong. His research interests include quantum internet, online learning theory and applications, performance modeling and evaluation of computer/communication systems. He is a fellow of IEEE and ACM.