

Quantum Internet Research Group
Internet-Draft
Intended status: Informational
Expires: April 2, 2021

W. Kozlowski
S. Wehner
QuTech
R. Van Meter
Keio University
B. Rijsman
Individual
A. S. Cacciapuoti
M. Caleffi
University of Naples Federico II
S. Nagayama
Mercari, Inc.
September 29, 2020

Architectural Principles for a Quantum Internet
draft-irtf-qirg-principles-05

Abstract

The vision of a quantum internet is to fundamentally enhance Internet technology by enabling quantum communication between any two points on Earth. To achieve this goal, a quantum network stack should be built from the ground up to account for the fundamentally new properties of quantum entanglement. The first realisations of quantum networks are imminent, but there is no practical proposal for how to organise, utilise, and manage such networks. In this memo, we attempt to lay down the framework and introduce some basic architectural principles for a quantum internet. This is intended for general guidance and general interest, but also to provide a foundation for discussion between physicists and network specialists.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on April 2, 2021.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Quantum information	4
2.1. Qubit	4
2.2. Multiple qubits	5
3. Entanglement as the fundamental resource	6
4. Achieving quantum connectivity	7
4.1. Challenges	8
4.1.1. The measurement problem	8
4.1.2. No-cloning theorem	8
4.1.3. Fidelity	8
4.1.4. Inadequacy of direct transmission	9
4.2. Bell pairs	9
4.3. Teleportation	10
4.4. The life cycle of entanglement	11
4.4.1. Elementary link generation	11
4.4.2. Entanglement swapping	12
4.4.3. Error Management	13
4.4.4. Delivery	16
5. Architecture of a quantum internet	16
5.1. Challenges	16
5.2. Classical communication	18
5.3. Abstract model of the network	19
5.3.1. One control plane and two data planes	19
5.3.2. Elements of a quantum network	19
5.3.3. Putting it all together	21
5.4. Network boundaries	22
5.4.1. Boundaries between different physical architectures	22
5.4.2. Boundaries between different administrative regions	22
5.4.3. Boundaries between different error management schemes	22
5.5. Physical constraints	23
5.5.1. Memory lifetimes	23

5.5.2. Rates	23
5.5.3. Communication qubits	23
5.5.4. Homogeneity	24
6. Architectural principles	24
6.1. Goals of a quantum internet	24
6.2. The principles of a quantum internet	27
7. Comparison with classical networks	29
8. Security Considerations	31
9. IANA Considerations	31
10. Acknowledgements	31
11. Informative References	31
Authors' Addresses	36

1. Introduction

Quantum networks are distributed systems of quantum devices that utilise fundamental quantum mechanical phenomena such as superposition, entanglement, and quantum measurement to achieve capabilities beyond what is possible with non-quantum (classical) networks [35]. Depending on the stage of a quantum network [8] such devices may be simple photonic devices capable of preparing and measuring only one quantum bit (qubit) at a time, all the way to large-scale quantum computers of the future. A quantum network is not meant to replace classical networks, but rather form an overall hybrid classical quantum network supporting new capabilities which are otherwise impossible to realise [23].

This new networking paradigm offers promise for a range of new applications such as secure communications [3] [4], distributed quantum computation [5], secure quantum computing in the cloud [27], quantum-enhanced measurement networks [6], or longer-baseline telescopes [32]. The field of quantum communication has been a subject of active research for many years and the most well-known application of quantum communication, quantum key distribution (QKD) for secure communications, has already been deployed at short (roughly 100km) distances [25] [24].

Fully quantum networks capable of transmitting and managing entangled quantum states in order to send, receive, and manipulate distributed quantum information are now imminent [7] [8]. Whilst a lot of effort has gone into physically realising and connecting such devices [26], and making improvements to their speed and error tolerance, there are no worked out proposals for how to run these networks. To draw an analogy with a classical network, we are at a stage where we can start to physically connect our devices and send data, but all sending, receiving, buffer management, connection synchronisation, and so on, must be managed by the application itself at a level below conventional assembly language, where no common interfaces yet exist.

Furthermore, whilst physical mechanisms for transmitting quantum states exist, there are no robust protocols for managing such transmissions.

2. Quantum information

In order to understand the framework for quantum networking, a basic understanding of quantum information is necessary. The following sections aim to introduce the bare minimum necessary to understand the principles of operation of a quantum network. This exposition was written with a classical networking audience in mind. It is assumed that the reader has never before been exposed to any quantum physics. We refer to e.g. [15] [16] for an in-depth introduction to quantum information.

2.1. Qubit

The differences between quantum computation and classical computation begin at the bit-level. A classical computer operates on the binary alphabet $\{0, 1\}$. A quantum bit, a qubit, exists over the same binary space, but unlike the classical bit, it can exist in a superposition of the two possibilities:

$$a |0\rangle + b |1\rangle,$$

where $|X\rangle$ is Dirac's ket notation for a quantum state, here the binary 0 and 1, and the coefficients a and b are complex numbers called probability amplitudes. Physically, such a state can be realised using a variety of different technologies such as electron spin, photon polarisation, atomic energy levels, and so on.

Upon measurement, the qubit loses its superposition and irreversibly collapses into one of the two basis states, either $|0\rangle$ or $|1\rangle$. Which of the two states it ends up in is not deterministic, but it can be determined from the readout of the measurement, a classical bit, 0 or 1 respectively. The probability of measuring the state in the $|0\rangle$ state is $|a|^2$ and similarly the probability of measuring the state in the $|1\rangle$ state is $|b|^2$, where $|a|^2 + |b|^2 = 1$. This randomness is not due to our ignorance of the underlying mechanisms, but rather it is a fundamental feature of a quantum mechanical system [9].

The superposition property plays an important role in fundamental gate operations on qubits. Since a qubit can exist in a superposition of its basis states, the elementary quantum gates are able to act on all states of the superposition at the same time. For example, consider the NOT gate:

$$\text{NOT} (a |0\rangle + b |1\rangle) \rightarrow a |1\rangle + b |0\rangle.$$

2.2. Multiple qubits

When multiple qubits are combined in a single quantum state the space of possible states grows exponentially and all these states can coexist in a superposition. For example, the general form of a two-qubit register is

$$a |00\rangle + b |01\rangle + c |10\rangle + d |11\rangle$$

where the coefficients have the same probability amplitude interpretation as for the single qubit state. Each state represents a possible outcome of a measurement of the two-qubit register. For example, $|01\rangle$ denotes a state in which the first qubit is in the state $|0\rangle$ and the second is in the state $|1\rangle$.

Performing single qubit gates affects the relevant qubit in each of the superposition states. Similarly, two-qubit gates also act on all the relevant superposition states, but their outcome is far more interesting.

Consider a two-qubit register where the first qubit is in the superposed state $(|0\rangle + |1\rangle)/\sqrt{2}$ and the other is in the state $|0\rangle$. This combined state can be written as:

$$(|0\rangle + |1\rangle)/\sqrt{2} \times |0\rangle = (|00\rangle + |10\rangle)/\sqrt{2},$$

where \times denotes a tensor product (the mathematical mechanism for combining quantum states together). Let us now consider the two-qubit controlled-NOT, or CNOT, gate. The CNOT gate takes as input two qubits, a control and target, and applies the NOT gate to the target if the control qubit is set. The truth table looks like

IN	OUT
00	00
01	01
10	11
11	10

Now, consider performing a CNOT gate on the state with the first qubit being the control. We apply a two-qubit gate on all the superposition states:

$$\text{CNOT } (|00\rangle + |10\rangle)/\sqrt{2} \rightarrow (|00\rangle + |11\rangle)/\sqrt{2}.$$

What is so interesting about this two-qubit gate operation? The final state is **entangled**. There is no possible way of representing that quantum state as a product of two individual qubits; they are no longer independent and the behaviour of either qubit cannot be fully described without accounting for the other qubit. The states of the two individual qubits are now correlated beyond what is possible to achieve classically. Neither qubit is in a definite $|0\rangle$ or $|1\rangle$ state, but if we perform a measurement on either one, the outcome of the partner qubit will **always** yield the exact same outcome. The final state, whether it's $|00\rangle$ or $|11\rangle$, is fundamentally random as before, but the states of the two qubits following a measurement will always be identical.

Once a measurement is performed, the two qubits are once again independent. The final state is either $|00\rangle$ or $|11\rangle$ and both of these states can be trivially decomposed into a product of two individual qubits. The entanglement has been consumed and the entangled state must be prepared again.

3. Entanglement as the fundamental resource

Entanglement is the fundamental building block of quantum networks. Consider the state from the previous section:

$$(|00\rangle + |11\rangle)/\sqrt{2}.$$

Neither of the two qubits is in a definite $|0\rangle$ or $|1\rangle$ state and we need to know the state of the entire register to be able to fully describe the behaviour of the two qubits.

Entangled qubits have interesting non-local properties. Consider sending one of the qubits to another device. This device could in principle be anywhere: on the other side of the room, in a different country, or even on a different planet. Provided negligible noise has been introduced, the two qubits will forever remain in the entangled state until a measurement is performed. The physical distance does not matter at all for entanglement.

This lies at the heart of quantum networking, because it is possible to leverage the non-classical correlations provided by entanglement in order to design completely new types of application protocols that are not possible to achieve with just classical communication. Examples of such applications are quantum cryptography [3] [4], blind quantum computation [27], or distributed quantum computation [5].

Entanglement has two very special features from which one can derive some intuition about the types of applications enabled by a quantum network.

The first stems from the fact that entanglement enables stronger than classical correlations, leading to opportunities for tasks that require coordination. As a trivial example, consider the problem of consensus between two nodes who want to agree on the value of a single bit. They can use the quantum network to prepare the state $(|00\rangle + |11\rangle)/\sqrt{2}$ with each node holding one of the two qubits. Once either of the two nodes performs a measurement, the state of the two qubits collapses to either $|00\rangle$ or $|11\rangle$, so whilst the outcome is random and does not exist before measurement, the two nodes will always measure the same value. We can also build the more general multi-qubit state $(|00\dots\rangle + |11\dots\rangle)/\sqrt{2}$ and perform the same algorithm between an arbitrary number of nodes. These stronger than classical correlations generalise to more complicated measurement schemes as well.

The second feature of entanglement is that it cannot be shared, in the sense that if two qubits are maximally entangled with each other, then it is physically impossible for any other system to have any share of this entanglement [28]. Hence, entanglement forms a sort of private and inherently untappable connection between two nodes once established.

Entanglement is created through local interactions between two qubits or as a product of the way the qubits were created (e.g. entangled photon pairs). To create a distributed entangled state, one can then physically send one of the qubits to a remote node. It is also possible to directly entangle qubits that are physically separated, but this still requires local interactions between some other qubits that the separated qubits are initially entangled with. Therefore, it is the transmission of qubits that draws the line between a genuine quantum network and a collection of quantum computers connected over a classical network.

A quantum network is defined as a collection of nodes that is able to exchange qubits and distribute entangled states amongst themselves. A quantum node that is able only to communicate classically with another quantum node is not a member of a quantum network.

More complex services and applications can be built on top of entangled states distributed by the network, see e.g. [33]

4. Achieving quantum connectivity

This section explains the meaning of quantum connectivity and the necessary physical processes at an abstract level.

4.1. Challenges

A quantum network cannot be built by simply extrapolating all the classical models to their quantum analogues. Sending qubits over a wire like we send classical bits is simply not as easy to do. There are several technological as well as fundamental challenges that make classical approaches unsuitable in a quantum context.

4.1.1. The measurement problem

In classical computers and networks we can read out the bits stored in memory at any time. This is helpful for a variety of purposes such as copying, error detection and correction, and so on. This is not possible with qubits.

A measurement of a qubit's state will destroy its superposition and with it any entanglement it may have been part of. Once a qubit is being processed, it cannot be read out until a suitable point in the computation, determined by the protocol handling the qubit, has been reached. Therefore, we cannot use the same methods known from classical computing for the purposes of error detection and correction. Nevertheless, quantum error detection and correction schemes exist that take this problem into account and how a network chooses to manage errors will have an impact on its architecture.

4.1.2. No-cloning theorem

Since directly reading the state of a qubit is not possible, one could ask the question if we can simply copy a qubit without looking at it. Unfortunately, this is fundamentally not possible in quantum mechanics [29] [30].

The no-cloning theorem states that it is impossible to create an identical copy of an arbitrary, unknown quantum state. Therefore, it is also impossible to use the same mechanisms that worked for classical networks for signal amplification, retransmission, and so on as they all rely on the ability to copy the underlying data. Since any physical channel will always be lossy, connecting nodes within a quantum network is a challenging endeavour and its architecture must at its core address this very issue.

4.1.3. Fidelity

In general, it is expected that a classical packet arrives at its destination without any errors introduced by hardware noise along the way. This is verified at various levels through a variety of error detection and correction mechanisms. Since we cannot read or copy a quantum state error detection and correction is more involved.

To describe the quality of a quantum state, a physical quantity called fidelity is used [16]. Fidelity takes a value between 0 and 1 -- higher is better, and less than 0.5 means the state is unusable. It measures how close a quantum state is to the state we have tried to create. It expresses the probability that one state will pass a test to identify as the other. Fidelity is an important property of a quantum system that allows us to quantify how much a particular state has been affected by noise from various sources (gate errors, channel losses, environment noise).

Interestingly, quantum applications do not need perfect fidelity to be able to execute -- as long as the fidelity is above some application-specific threshold, they will simply operate at lower rates. Therefore, rather than trying to ensure that we always deliver perfect states (a technologically challenging task) applications will specify a minimum threshold for the fidelity and the network will try its best to deliver it. A higher fidelity can be achieved by either having hardware produce states of better fidelity (sometimes one can sacrifice rate for higher fidelity) or by employing quantum error detection and correction mechanisms.

4.1.4. Inadequacy of direct transmission

Conceptually, the most straightforward way to distribute an entangled state is to simply transmit one of the qubits directly to the other end across a series of nodes while performing sufficient forward quantum error correction (Section 4.4.3.2) to bring losses down to an acceptable level. Despite the no-cloning theorem and the inability to directly measure a quantum state, error-correcting mechanisms for quantum communication exist [31] [36] [10]. However, quantum error correction makes very high demands on both resources (physical qubits needed) and their initial fidelity. Implementation is very challenging and quantum error correction is not expected to be used until later generations of quantum networks.

An alternative relies on the observation that we do not need to be able to distribute any arbitrary entangled quantum state. We only need to be able to distribute any one of what are known as the Bell pair states [18].

4.2. Bell pairs

Bell pair states are the entangled two-qubit states:

$$\begin{array}{l} |00\rangle + |11\rangle, \\ |00\rangle - |11\rangle, \\ |01\rangle + |10\rangle, \\ |01\rangle - |10\rangle, \end{array}$$

where the constant $1/\sqrt{2}$ normalisation factor has been ignored for clarity. Any of the four Bell pair states above will do, as it is possible to transform any Bell pair into another Bell pair with local operations performed on only one of the qubits. When each qubit in a Bell pair is held by a separate node, either can apply a series of single qubit gates to their qubit alone in order to transform the state between the different variants.

Distributing a Bell pair between two nodes is much easier than transmitting an arbitrary quantum state over a network. Since the state is known, handling errors becomes easier and small-scale error-correction (such as entanglement distillation discussed in a later section) combined with reattempts becomes a valid strategy.

The reason for using Bell pairs specifically as opposed to any other two-qubit state, is that they are the maximally entangled two-qubit set of basis states. Maximal entanglement means that these states have the strongest non-classical correlations of all possible two-qubit states. Furthermore, since single-qubit local operations can never increase entanglement, less entangled states would impose some constraints on distributed quantum algorithms. This makes Bell pairs particularly useful as a generic building block for distributed quantum applications.

4.3. Teleportation

The observation that we only need to be able to distribute Bell pairs relies on the fact that this enables the distribution of any other arbitrary entangled state. This can be achieved via quantum state teleportation. Quantum state teleportation consumes an unknown quantum state that we want to transmit and recreates it at the desired destination. This does not violate the no-cloning theorem as the original state is destroyed in the process.

To achieve this, an entangled pair needs to be distributed between the source and destination before teleportation commences. The source then entangles the transmission qubit with its end of the pair and performs a read out of the two qubits (the sum of these operations is called a Bell state measurement). This consumes the Bell pair's entanglement, turning the source and destination qubits into independent states. The measurements yields two classical bits which the source sends to the destination over a classical channel. Based on the value of the received two classical bits, the destination performs one of four possible corrections (called the Pauli corrections) on its end of the pair which turns it into the unknown quantum state that we wanted to transmit.

The unknown quantum state that was transmitted was never fed into the network itself. Therefore, the network needs to only be able to reliably produce Bell pairs between any two nodes in the network. Thus, a key difference between a classical and quantum data planes is that a classical one carries user data, but a quantum data plane provides the resources for the user to transmit user data themselves without further involvement of the network.

4.4. The life cycle of entanglement

Reducing the problem of quantum connectivity to one of generating a Bell pair has facilitated the problem, but it has not solved it. In this section, we discuss how these entangled pairs are generated in the first place, and how their two qubits are delivered to the end-points.

4.4.1. Elementary link generation

In a quantum network, entanglement is always first generated locally (at a node or an auxiliary element) followed by a movement of one or both of the entangled qubits across the link through quantum channels. In this context, photons (particles of light) are the natural candidate for entanglement carriers, called flying qubits. The rationale for this choice is related to the advantages provided by photons such as moderate interaction with the environment leading to moderate decoherence, convenient control with standard optical components, and high-speed, low-loss transmissions. However, since photons cannot be stored, a transducer must transfer the flying qubit's state to a qubit suitable for information processing and/or storage (often referred to as a matter qubit).

Since this process may fail, in order to generate and store entanglement efficiently, we must be able to distinguish successful attempts from failures. Entanglement generation schemes that are able to announce successful generation are called heralded entanglement generation schemes.

There exist three basic schemes for heralded entanglement generation on a link through coordinated action of the two nodes at the two ends of the link [19]:

- o "At mid-point": in this scheme an entangled photon pair source sitting midway between the two nodes with matter qubits sends an entangled photon through a quantum channel to each of the nodes. There, transducers are invoked to transfer the entanglement from the flying qubits to the matter qubits. In this scheme, the transducers know if the transfers succeeded and are able to herald

successful entanglement generation via a message exchange over the classical channel.

- o "At source": in this scheme one of the two nodes sends a flying qubit that is entangled with one of its matter qubits. A transducer at the other end of the link will transfer the entanglement from the flying qubit to one of its matter qubits. Just like in the previous scheme, the transducer knows if its transfer succeeded and is able to herald successful entanglement generation with a classical message sent to the other node.
- o "At both end-points": in this scheme both nodes send a flying qubit that is entangled with one of their matter qubits. A detector somewhere in between the nodes performs a joint measurement on the two qubits, which stochastically projects the remote matter qubits into an entangled quantum state. The detector knows if the entanglement succeeded and is able to herald successful entanglement generation by sending a message to each node over the classical channel.


The "mid-point source" scheme is more robust to photon loss, but in the other schemes the nodes retain greater control over the entangled pair generation.

Note that whilst photons travel in a particular direction through the quantum channel the resulting entangled pair of qubits does not have a direction associated with it. Physically, there is no upstream or downstream end of the pair.

4.4.2. Entanglement swapping


The problem with generating entangled pairs directly across a link is that efficiency decreases with channel length. Beyond a few 10s of kms in optical fibre or 1000 kms in free space (via satellite) the rate is effectively zero and due to the no-cloning theorem we cannot simply amplify the signal. The solution is entanglement swapping [18].

A Bell pair between any two nodes in the network can be constructed by combining the pairs generated along each individual link on a path between the two end-points. Each node along the path can consume the two pairs on the two links that it is connected to in order to produce a new entangled pair between the two remote ends. This process is known as entanglement swapping. Pictorially it can be represented as follows:



where X1 and X2 are the qubits of the entangled pair X and Y1 and Y2 are the qubits of entangled pair Y. The entanglement is denoted with \sim . In the diagram above, nodes A and B share the pair X and nodes B and C share the pair Y, but we want entanglement between A and C.

To achieve this goal, we simply teleport the qubit X2 using the pair Y. This requires node B to perform a Bell state measurement on the qubits X2 and Y1 which result in the destruction of the entanglement between Y1 and Y2. However, X2 is recreated in Y2's place, carrying with it its entanglement with X1. The end-result is shown below:



Depending on the needs of the network and/or application, a final Pauli correction at the recipient node may not be necessary since the result of this operation is also a Bell pair. However, the two classical bits that form the read out from the measurement at node B must still be communicated, because they carry information about which of the four Bell pairs was actually produced. If a correction is not performed, the recipient must be informed which Bell pair was received.

This process of teleporting Bell pairs using other entangled pairs is called entanglement swapping. Quantum nodes that create long-distance entangled pairs via entanglement swapping are called quantum repeaters in academic literature [18] and we will use the same terminology in this memo.

4.4.3. Error Management

4.4.3.1. Distillation

Neither the generation of Bell pairs nor the swapping operations are noiseless operations. Therefore, with each link and each swap the fidelity of the state degrades. However, it is possible to create higher fidelity Bell pair states from two or more lower fidelity pairs through a process called distillation (sometimes also referred to as purification) [34].

To distill a quantum state, a second (and sometimes third) quantum state is used as a "test tool" to test a proposition about the first state, e.g., "the parity of the two qubits in the first state is even." When the test succeeds, confidence in the state is improved, and thus the fidelity is improved. The test tool states are destroyed in the process, so resource demands increase substantially when distillation is used. When the test fails, the tested state must also be discarded. Distillation makes low demands on fidelity and resources compared to quantum error correction, but distributed protocols incur round-trip delays due to classical communication ([17]).

4.4.3.2. Quantum Error Correction

Just like classical error correction, quantum error correction (QEC) encodes logical qubits using several physical (raw) qubits to protect them from errors described in Section 4.1.3 [31] [36] [10]. Furthermore, similarly to its classical counterpart, QEC can not only correct state errors but also account for lost qubits. Additionally, if all physical qubits which encode a logical qubit are located at the same node, the correction procedure can be executed locally, even if the logical qubit is entangled with remote qubits.

Although QEC was originally a scheme proposed to protect a qubit from noise, QEC can also be applied to entanglement distillation. Such QEC-applied distillation is cost-effective but requires a higher base fidelity.

One big difference from classical error correction is the code-rate. QEC encodes a single logical qubit using many physical qubits.

4.4.3.3. Error management schemes

Quantum networks have been categorized into three "generations" based on the error management scheme they employ [10]. Note that these "generations" are more like categories; they do not necessarily imply a time progression and do not obsolete each other, though the later generations do require more advanced technologies. Which generation is used depends on the hardware platform and network design choices.

Table 1 summarises the generations.

	First generation	Second generation	Third generation
Loss tolerance	Heralded entanglement generation (bi-directional classical signaling)	Heralded entanglement generation (bi-directional classical signaling)	Quantum Error Correction (no classical signaling)
Error tolerance	Entanglement distillation (bi-directional classical signaling)	Entanglement distillation (uni-directional classical signaling) or Quantum Error Correction (no classical signaling)	Quantum Error Correction (no classical signaling)

Table 1: Classical signaling and generations

Generations are defined by the directions of classical signalling required in their distributed protocols for loss tolerance and error tolerance. Classical signalling carries the classical bits and incurs round-trip delays described in Section 4.4.3.1, hence they affect the performance of quantum networks, especially as the distance between the communicating nodes increases.

Loss tolerance is about tolerating qubit transmission losses between nodes. Heralded entanglement generation, as described in Section 4.4.1, confirms the receipt of an entangled qubit using a heralding signal. A pair of directly connected quantum nodes repeatedly attempt to generate an entangled pair until the a heralding signal is received. As described in Section 4.4.3.2, QEC can be applied to complement lost qubits eliminating the need for re-attempts. Furthermore, since the correction procedure is composed of local operations, it does not require a heralding signal. However, it is feasible only when the photon loss rate is less than 0.5.

Error tolerance is about tolerating quantum state errors. Entanglement distillation is the easiest mechanism for improved error tolerance to implement, but it incurs round-trip delays due the requirement for bi-directional classical signalling. The alternative, QEC, is able to correct state errors locally so that it does not need any classical signalling between the quantum nodes. In between these two extremes, there is also QEC-applied distillation, which requires uni-directional classical signalling.

The three "generations" summarised:

1. First generation quantum networks use heralding for loss tolerance and entanglement distillation for error tolerance. These networks can be implemented using only small, shallow quantum circuits at each node.
2. Second generation quantum networks are empowered by QEC codes for error tolerance. At first, QEC will be applied to entanglement distillation only which requires uni-directional classical signalling. Later, QEC codes will be used to create logical Bell pairs which no longer require any classical signalling for the purposes of error tolerance. Heralding is still used to compensate for transmission losses.
3. Third generation quantum networks directly transmit QEC encoded qubits to adjacent nodes, as discussed in Section 4.1.4. Elementary link Bell pairs can now be created without heralding or any other classical signalling. Furthermore, this also enables direct transmission architectures in which qubits are forwarded end-to-end like classical packets rather than relying on Bell pairs and entanglement swapping.

4.4.4. Delivery

Eventually, the Bell pairs must be delivered to an application (or higher layer protocol) at the two end-nodes. A detailed list of such requirements is beyond the scope of this memo. At minimum, the end-nodes require information to map a particular Bell pair to the qubit in their local memory that is part of this entangled pair.

5. Architecture of a quantum internet

It is evident from the previous sections that the fundamental service provided by a quantum network significantly differs from that of a classical network. Therefore, it is not surprising that the architecture of a quantum internet will itself be very different from that of the classical Internet.

5.1. Challenges

This subsection covers the major fundamental challenges building quantum networks. Here, we only describe the fundamental differences. Technological limitations are described later.

1. Bell pairs are not equivalent to payload carrying packets.

In most classical networks, including Ethernet, Internet Protocol (IP), and Multi-Protocol Label Switching (MPLS) networks, user data is grouped into packets. In addition to the user data, each packet also contains a series of headers which contain the control information that lets routers and switches forward it towards its destination. Packets are the fundamental unit in a classical network.

In a quantum network, the entangled pairs of qubits are the basic unit of networking. These qubits themselves do not carry any headers. Therefore, quantum networks will have to send all control information via separate classical channels which the repeaters will have to correlate with the qubits stored in their memory.

2. "Store and forward" vs "store and swap" quantum networks.

As described in Section 4.4.1, quantum links provide Bell pairs that are undirected network resources, in contrast to directed frames of classical networks. This phenomenological distinction leads to architectural differences between quantum networks and classical networks. Quantum networks combine multiple elementary link Bell pairs together to create one an end-to-end Bell pair, whereas classical networks deliver messages from one end to the other end hop by hop.

Classical networks receive data on one interface, store it in local buffers, then forward the data to another appropriate interface. Quantum networks store Bell pairs and then execute entanglement swapping instead of forwarding in the data plane. Such quantum networks are "store and swap" networks. In "store and swap" networks, we do not need to care about the order in which the Bell pairs were generated since they are undirected. This distinction makes control algorithms and optimisation of quantum networks different from classical ones. Note that third generation quantum networks, as described in Section 4.4.1, will be able to support a "store and forward" architecture in addition to "store and swap".

3. An entangled pair is only useful if the locations of both qubits are known.

A classical network packet logically exists only at one location at any point in time. If a packet is modified in some way, whether headers or payload, this information does not need to be conveyed to anybody else in the network. The packet can be simply forwarded as before.

In contrast, entanglement is a phenomenon in which two or more qubits exist in a physically distributed state. Operations on one of the qubits change the mutual state of the pair. Since the owner of a particular qubit cannot just read out its state, it must coordinate all its actions with the owner of the pair's other qubit. Therefore, the owner of any qubit that is part of an entangled pair must know the location of its counterpart. Location, in this context, need not be the explicit spatial location. A relevant pair identifier, a means of communication between the pair owners, and an association between the pair ID and the individual qubits is sufficient.

4. Generating entanglement requires temporary state.

Packet forwarding in a classical network is largely a stateless operation. When a packet is received, the router looks up its forwarding table and sends the packet out of the appropriate output. There is no need to keep any memory of the packet any more.

A quantum node must be able to make decisions about qubits that it receives and is holding in its memory. Since qubits do not carry headers, the receipt of an entangled pair conveys no control information based on which the repeater can make a decision. The relevant control information will arrive separately over a classical channel. This implies that a repeater must store temporary state as the control information and the qubit it pertains to will, in general, not arrive at the same time.

5.2. Classical communication

In this memo we have already covered two different roles that classical communication must perform:

- o communicate classical bits of information as part of distributed protocols such as entanglement swapping and teleportation,
- o communicate control information within a network, including both background protocols such as routing as well as signalling protocols to set up end-to-end entanglement generation.

Classical communication is a crucial building block of any quantum network. All nodes in a quantum network are assumed to have classical connectivity with each other (within typical administrative domain limits). Therefore, quantum routers will need to manage two data planes in parallel, a classical one and a quantum one. Additionally, a node must be able to correlate information between

the two planes so that the control information received on a classical channel can be applied to the qubits managed by the quantum data plane.

5.3. Abstract model of the network

5.3.1. One control plane and two data planes

Control plane protocols for quantum networks will have responsibilities similar to their classical counterparts, namely drawing the network topology, resource management, populating data plane tables, etc. They will not manipulate quantum data themselves and they operate by exchanging classical messages only. Therefore there is no separate quantum and classical control plane. There is only one network control plane.

However, as we have already mentioned earlier in this memo, there will be two data planes: a classical data plane and a quantum data plane. The classical data plane processes and forwards classical packets. The quantum data plane processes and swaps entangled pairs. Third generation quantum networks may also forward qubits in addition to swapping Bell pairs.

In addition to control plane messages, there will also be control information messages that operate at the granularity of individual entangled pairs, such as heralding messages used for elementary link generation (Section 4.4.1). In terms of functionality, these messages are closer to classical packet headers than control plane messages and thus we consider them to be part of the quantum data plane. Therefore, a quantum data plane also includes the exchange of classical control information at the granularity of individual qubits and entangled pairs.

5.3.2. Elements of a quantum network

We have identified quantum repeaters as the core building block of a quantum network. However, a quantum repeater will have to do more than just entanglement swapping in a functional quantum network. Its key responsibilities will include:

1. Creating link-local entanglement between neighbouring nodes.
2. Extending entanglement from link-local pairs to long-range pairs through entanglement swapping.
3. Performing distillation to manage the fidelity of the produced pairs.

4. Participating in the management of the network (routing, etc.).

Not all quantum repeaters in the network will be the same; here we break them down further:

- o Quantum routers (controllable quantum nodes) - A quantum router is a quantum repeater with a control plane that participates in the management of the network and will make decisions about which qubits to swap to generate the requested end-to-end pairs.
- o Automated quantum nodes - An automated quantum node is a data plane only quantum repeater that does not participate in network management. Since the no-cloning theorem precludes the use of amplification, long-range links will be established by chaining multiple such automated nodes together.
- o End-nodes - End-nodes in a quantum network must be able to receive and handle an entangled pair, but they do not need to be able to perform an entanglement swap (and thus are not necessarily quantum repeaters). End-nodes are also not required to have any quantum memory as certain quantum applications can be realised by having the end-node measure its qubit as soon as it is received.
- o Non-quantum nodes - Not all nodes in a quantum network need to have a quantum data plane. A non-quantum node is any device that can handle classical network traffic.


Additionally, we need to identify two kinds of links that will be used in a quantum network:

- o Quantum links - A quantum link is a link which can be used to generate an entangled pair between two directly connected quantum repeaters. This may include additional mid-point elements described in Section 4.4.1. It may also include a dedicated classical channel that is to be used solely for the purpose of coordinating the entanglement generation on this quantum link.
- o Classical links - A classical link is a link between any node in the network that is capable of carrying classical network traffic.

Note that passive elements, such as optical switches, do not destroy the quantum state. Therefore, it is possible to connect multiple quantum nodes with each other over an optical network and perform optical switching rather than routing via entanglement swapping at quantum routers. This does require coordination with the elementary link entanglement generation process and it still requires repeaters to overcome the short-distance limitations. However, this is a potentially feasible architecture for local area networks.

5.3.3. Putting it all together

A two-hop path in a generic quantum network can be represented as:



App - user-level application

QR - quantum repeater

EN - end-node

QL - quantum link

CC - classical channel (can consist of many classical links)

An application running on two end-nodes attached to a network will at some point need the network to generate entangled pairs for its use. This may require negotiation between the end-nodes (possibly ahead of time), because they must both open a communication end-point which the network can use to identify the two ends of the connection. The two end-nodes use the classical connectivity available in the network to achieve this goal.

When the network receives a request to generate end-to-end entangled pairs it uses the classical communication channels to coordinate and claim the resources necessary to fulfill this request. This may be some combination of prior control information (e.g. routing tables) and signalling protocols, but the details of how this is achieved are an active research question and thus beyond the scope of this memo.

During or after the distribution of control information, the network performs the necessary quantum operations such as generating entanglement over individual links, performing entanglement swaps, and further signalling to transmit the swap outcomes and other control information. Since Bell pairs do not carry any user data, some of these operations can be performed before the request is received in anticipation of the demand.

The entangled pair is delivered to the application once it is ready, together with the relevant pair identifier. However, being ready does not necessarily mean that all link pairs and entanglement swaps are complete, as some applications can start executing on an incomplete pair. In this case the remaining entanglement swaps will propagate the actions across the network to the other end, sometimes necessitating fixup operations at the end node.

5.4. Network boundaries

Just like classical networks, various boundaries will exist in quantum networks.

5.4.1. Boundaries between different physical architectures

There are many different physical architectures for implementing quantum repeater technology. The different technologies differ in how they store and manipulate qubits in memory and how they generate entanglement across a link with their neighbours. Different architectures come with different trade-offs and thus a functional network will likely consist of a mixture of different types of quantum repeaters.

For example, architectures based on optical elements and atomic ensembles [37] are very efficient at generating entanglement, but provide little control over the qubits once the pair is generated. On the other hand, nitrogen-vacancy architectures [26] offer a much greater degree of control over qubits, but have a harder time generating the entanglement across a link.

It is an open research question where exactly the boundary will lie. It could be that a single quantum repeater node provides some backplane connection between the architectures, but it also could be that special quantum links delineate the boundary.

5.4.2. Boundaries between different administrative regions

Just like in classical networks, multiple quantum networks will connect into a global quantum internet. This necessarily implies the existence of borders between different administrative regions. How these boundaries will be handled is also an open question and thus beyond the scope of this memo.

5.4.3. Boundaries between different error management schemes

Not only are there physical differences and administrative boundaries, but there are important distinctions in how errors will be managed, as described in Section 4.4.3.3, which affects the content and semantics of messages that must cross those boundaries -- both for connection setup and real-time operation. How to interconnect those schemes is also an open research question.

5.5. Physical constraints

The model above has effectively abstracted away the particulars of the hardware implementation. However, certain physical constraints need to be considered in order to build a practical network. Some of these are fundamental constraints and no matter how much the technology improves, they will always need to be addressed. Others are artefacts of the early stages of a new technology. Here, we consider a highly abstract scenario and refer to [8] for pointers to the physics literature.

5.5.1. Memory lifetimes

In addition to discrete operations being imperfect, storing a qubit in memory is also highly non-trivial. The main difficulty in achieving persistent storage is that it is extremely challenging to isolate a quantum system from the environment. The environment introduces an uncontrollable source of noise into the system which affects the fidelity of the state. This process is known as decoherence. Eventually, the state has to be discarded once its fidelity degrades too much.

The memory lifetime depends on the particular physical setup, but the highest achievable values in quantum network hardware currently are on the order of seconds [38] although a lifetime of a minute has also been demonstrated, but these qubits were not yet connected to a quantum network [39]. These values have increased tremendously over the lifetime of the different technologies and are bound to keep increasing. However, if quantum networks are to be realised in the near future, they need to be able to handle short memory lifetimes, for example by reducing latency on critical paths.

5.5.2. Rates

Entanglement generation on a link between two connected nodes is not a very efficient process and it requires many attempts to succeed [26] [14]. A fast repetition rate for Bell pair generation is achievable, but only a small fraction will succeed. Currently, the highest achievable rates of success between nodes capable of storing the resulting qubits are on the order of 10 Hz. Combined with short memory lifetimes this leads to very tight timing windows to build up network-wide connectivity.

5.5.3. Communication qubits

Most physical architectures capable of storing qubits are only able to generate entanglement using only a subset of its available qubits called communication qubits [14]. Once a Bell pair has been

generated using a communication qubit, its state can be transferred into memory. This may impose additional limitations on the network. In particular if a given node has only one communication qubit it cannot simultaneously generate Bell Pairs over two links. It must generate entanglement over the links one at a time.

5.5.4. Homogeneity

Currently all hardware implementations are homogeneous and they do not interface with each other. In general, it is very challenging to combine different quantum information processing technologies at present. Coupling different technologies with each other is of great interest as it may help overcome the weaknesses of the different implementations, but this may take a long time to be realised with high reliability and thus is not a near-term goal.

6. Architectural principles

Given that the most practical way of realising quantum network connectivity is using Bell pair and entanglement swapping repeater technology, what sort of principles should guide us in assembling such networks such that they are functional, robust, efficient, and most importantly, they work? Furthermore, how do we design networks so that they work under the constraints imposed by the hardware available today, but do not impose unnecessary burdens on future technology?

As this is a completely new technology that is likely to see many iterations over its lifetime, this memo must not serve as a definitive set of rules, but merely as a general set of recommended guidelines for the first generations of quantum networks based on principles and observations made by the community. The benefit of having a community built document at this early stage is that expertise in both quantum information and network architecture is needed in order to successfully build a quantum internet.

6.1. Goals of a quantum internet

When outlining any set of principles we must ask ourselves what goals do we want to achieve as inevitably trade-offs must be made. So what sort of goals should drive a quantum network architecture? The following list has been inspired by the history of computer networking and thus it is inevitably very similar to one that could be produced for the classical Internet [22]. However, whilst the goals may be similar the challenges involved are often fundamentally different. The list will also most likely evolve with time and the needs of its users.

1. Support distributed quantum applications

This goal seems trivially obvious, but makes a subtle, but important point which highlights a key difference between quantum and classical networks. Ultimately, quantum data transmission is not the goal of a quantum network - it is only one possible component of more advanced quantum application protocols [8]. Whilst transmission certainly could be used as a building block for all quantum applications, it is not the most basic one possible. For example, QKD, the most well known quantum application protocol, only relies on the stronger-than-classical correlations and inherent secrecy of entangled Bell pairs and does not have to transmit arbitrary quantum states [4].

The primary purpose of a quantum internet is to support distributed quantum application protocols and it is of utmost importance that they can run well and efficiently. Thus, it is important to develop performance metrics meaningful to application to drive the development of quantum network protocols. For example, the Bell pair generation rate is meaningless if one does not also consider their fidelity. It is generally much easier to generate pairs of lower fidelity, but quantum applications may have to make multiple re-attempts or even abort if the fidelity is too low. A review of the requirements for different known quantum applications can be found in [8] and an overview of use-cases can be found in [2].

2. Support tomorrow's distributed quantum applications

The only principle of the Internet that should survive indefinitely is the principle of constant change [1]. Technical change is continuous and the size and capabilities of the quantum internet will change by orders of magnitude. Therefore, it is an explicit goal that a quantum internet architecture be able to embrace this change. We have the benefit of having been witness to the evolution of the classical Internet over several decades and seen what worked and what did not. It is vital for a quantum internet to avoid the need for flag days (e.g. NCP to TCP/IP) or upgrades that take decades to roll out (e.g. IPv4 to IPv6). Therefore, it is important that any proposed architecture for general purpose quantum repeater networks can integrate new devices and solutions as they become available. It should not be constrained due to considerations for early-stage hardware and applications. For example, it is already possible to run QKD efficiently on metropolitan scales and such networks are already commercially available. However, they are not based on quantum repeaters and thus will not be able to easily transition to more sophisticated applications.

3. Support heterogeneity

There are multiple proposals for realising practical quantum repeater hardware and they all have their advantages and disadvantages. Some may offer higher Bell pair generation rates on individual links at the cost of more difficult entanglement swap operations. Other platforms may be good all around, but are more difficult to build.

In addition to physical boundaries, there may be distinctions in how errors are managed (Section 4.4.3.3). These difference will affect the content and semantics of messages that cross these boundaries -- both for connection setup and real-time operation.

The optimal network configuration will likely leverage the advantages of multiple platforms to optimise the provided service. Therefore, it is an explicit goal to incorporate varied hardware and technology support from the beginning.

4. Ensure security at the network level

The question of security in quantum networks is just as critical as it is in the classical Internet, especially since enhanced security offered by quantum entanglement is one of the key driving factors.

It turns out that as long as the underlying implementation corresponds to (or sufficiently approximates) theoretical models of quantum cryptography, quantum cryptographic protocols do not need the network to provide any guarantees about the confidentiality or integrity of the transmitted qubits or the generated entanglement. Instead, applications, such as QKD, establish such guarantees in an end-to-end fashion using the classical network in conjunction with the quantum one.

Nevertheless, whilst applications can ensure their own secure operation, network protocols themselves should be security aware in order to protect the network itself and limit disruption. Whilst the applications remain secure they are not necessarily operational or as efficient in the presence of an attacker. Security concerns in quantum networks are described in more detail in [13] [12].

5. Make them easy to monitor

In order to manage, evaluate the performance of, or debug a network it is necessary to have the ability to monitor the network while ensuring there will be mechanisms in place to

protect the confidentiality and integrity of the devices connected to it. Quantum networks bring new challenges in this area so it should be a goal of a quantum network architecture to make this task easy.

The fundamental unit of quantum information, the qubit, cannot be actively monitored as any readout irreversibly destroys its contents. One of the implications of this fact is that measuring an individual pair's fidelity is impossible. Fidelity is meaningful only as a statistical quantity which requires the constant monitoring and the sacrifice of generated Bell pairs for tomography or other methods.

Furthermore, given one end of an entangled pair, it is impossible to tell where the other qubit is without any additional classical metadata. It is impossible to extract this information from the qubits themselves. This implies that tracking entangled pairs necessitates some exchange of classical information. This information might include (i) a reference to the entangled pair that allows distributed applications to coordinate actions on qubits of the same pair, (ii) the two bits from each entanglement swap necessary to identify the final state of the Bell pair (Section 4.4.2).

6. Ensure availability and resilience

Any practical and usable network, classical or quantum, must be able to continue to operate despite losses and failures, and be robust to malicious actors trying to disable connectivity. What differs in quantum networks as compared to classical networks in this regard is that we now have two data planes and two types of channels to worry about: a quantum and a classical one. Therefore, availability and resilience will most likely require a more advanced treatment than they do in classical networks.

6.2. The principles of a quantum internet

The principles support the goals, but are not goals themselves. The goals define what we want to build and the principles provide a guideline in how we might achieve this. The goals will also be the foundation for defining any metric of success for a network architecture, whereas the principles in themselves do not distinguish between success and failure. For more information about design considerations for quantum networks see [11] [14] .

1. Entanglement is the fundamental service

The key service that a quantum network provides is the distribution of entanglement between the nodes in a network. All distributed quantum applications are built on top of this key resource. Bell pairs are the minimal entanglement building block that is sufficient to develop these applications. However, a quantum network may also distribute multipartite entangled states (entangled states of three or more qubits) [20] as this may be more efficient under certain circumstances.

2. Bell Pairs are indistinguishable

Any two Bell Pairs between the same two nodes are indistinguishable for the purposes of an application provided they both satisfy its required fidelity threshold. This observation is likely to be key in enabling a more optimal allocation of resources in a network, e.g. for the purposes of provisioning resources to meet application demand. However, the qubits that make up the pair themselves are not indistinguishable and the two nodes operating on a pair must coordinate to make sure they are operating on qubits that belong to the same Bell Pair.

3. Fidelity is part of the service

In addition to being able to deliver Bell Pairs to the communication end-points, the Bell Pairs must be of sufficient fidelity. Unlike in classical networks where errors are effectively eliminated before reaching the application, many quantum applications only need imperfect entanglement to function. However, quantum applications will generally have a threshold for Bell pair fidelity below which they are no longer able to operate. Different applications will have different requirements for what fidelity they can work with. It is the network's responsibility to balance the resource usage with respect to the applications' requirements. It may be that it is cheaper for the network to provide lower fidelity pairs that are just above the threshold required by the application than it is to guarantee high fidelity pairs to all applications regardless of their requirements.

4. Time is part of the service

With the current technology, time is the most expensive resource. It is not the only resource that is in short supply (memory, and communication qubits are as well), but ultimately it is the lifetime of quantum memories that imposes the most difficult conditions for operating an extended network of quantum nodes. Current hardware has low rates of Bell Pair generation, short

memory lifetimes, and access to a limited number of communication qubits. All these factors combined mean that even a short waiting queue at some node could be enough for the Bell Pairs to decohere. It is vital that quantum networks deliver entanglement in a timely manner. The meaning of timeliness will depend on the needs of the application (how long does it need to store the Bell pair in its own memory and/or what operations it wants to apply to it).

5. Be flexible with regards to capabilities and limitations

This goal encompasses two important points. First, the architecture should be able to function under the physical constraints imposed by the current generation hardware. Near-future hardware will have low entanglement generation rates, quantum memories able to hold a handful of qubits at best, and decoherence rates that will render many generated pairs unusable.

Second, it should not make it difficult to run the network over any hardware that may come along in the future. The physical capabilities of repeaters will improve and redeploying a technology is extremely challenging.

7. Comparison with classical networks

Creating end-to-end Bell pairs between remote end-points is a stateful distributed task that requires a lot of a-priori coordination. Therefore, a connection-oriented approach seems the most natural for quantum networks. In this section, we discuss a plausible quantum network architecture inspired by MPLS. This is not an architecture proposal, but a thought experiment to give the reader an idea of what components are necessary for a functional quantum network. We use classical MPLS as a basis as it is well known and understood in the networking community.

In connection-oriented quantum networks, when two quantum application end-points wish to start creating end-to-end Bell pairs, they must first create a quantum virtual circuit (QVC). As an analogy, in MPLS networks end-points must establish a label switched path (LSP) before exchanging traffic. Connection-oriented quantum networks may also support virtual circuits with multiple end-points for creating multipartite entanglement. As an analogy, MPLS networks have the concept of multi-point LSPs for multicast.

When a quantum application creates a quantum virtual circuit, it can indicate quality of service (QoS) parameters such as the required capacity in end-to-end Bell pairs per second (BPPS) and the required fidelity of the Bell pairs. As an analogy, in MPLS networks

applications specify the required bandwidth in bits per second (BPS) and other constraints when they create a new LSP.

Quantum networks need a routing function to compute the optimal path (i.e. the best sequence of routers and links) for each new quantum virtual circuit. The routing function may be centralized or distributed. In the latter case, the quantum network needs a distributed routing protocol. As an analogy, classical networks use routing protocols such as open shortest path first (OSPF) and intermediate-system to intermediate system (IS-IS). However, note that the definition of "shortest-path"/"least-cost" may be different in a quantum network to account for its non-classical features, such as fidelity [21].

Given the very scarce availability of resources in early quantum networks, a traffic engineering function is likely to be beneficial. Without traffic engineering, quantum virtual circuits always use the shortest path. In this case, the quantum network cannot guarantee that each quantum end-point will get its Bell pairs at the required rate or fidelity. This is analogous to "best effort" service in classical networks.

With traffic engineering, quantum virtual circuits choose a path that is guaranteed to have the requested resources (e.g. bandwidth in BPPS) available, taking into account the capacity of the routers and links and taking into account the resources already consumed by other virtual circuits. As an analogy, both OSPF and IS-IS have traffic engineering (TE) extensions to keep track of used and available resources, and can use constrained shortest path first (CSPF) to take resource availability and other constraints into account when computing the optimal path.

The use of traffic engineering implies the use of call admission control (CAC): the network denies any virtual circuits for which it cannot guarantee the requested quality of service a-priori. Or alternatively, the network pre-empt lower priority circuits to make room for the new one.

Quantum networks need a signaling function: once the path for a quantum virtual circuit has been computed, signaling is used to install the "forwarding rules" into the data plane of each quantum router on the path. The signaling may be distributed, analogous to the resource reservation protocol (RSVP) in MPLS. Or the signaling may be centralized, similar to OpenFlow.

Quantum networks need an abstraction of the hardware for specifying the forwarding rules. This allows us to de-couple the control plane (routing and signaling) from the data plane (actual creation of Bell

pairs). The forwarding rules are specified using abstract building blocks such as "creating local Bell pairs", "swapping Bell pairs", "distillation of Bell pairs". As an analogy, classical networks use abstractions that are based on match conditions (e.g. looking up header fields in tables) and actions (e.g. modifying fields or forwarding a packet to a specific interface). The data-plane abstractions in quantum networks will be very different from those in classical networks due to the fundamental differences in technology and the stateful nature of quantum networks. In fact, choosing the right abstractions will be one of the biggest challenges when designing interoperable quantum network protocols.

In quantum networks, control plane traffic (routing and signaling messages) is exchanged over a classical channel, whereas data plane traffic (the actual Bell pair qubits) is exchanged over a separate quantum channel. This is in contrast to most classical networks, where control plane traffic and data plane traffic share the same channel and where a single packet contains both user fields and header fields. There is, however, a classical analogy to the way quantum networks work. Generalized MPLS (GMPLS) networks use separate channels for control plane traffic and data plane traffic. Furthermore, GMPLS networks support data planes where there is no such thing as data plane headers (e.g. DWDM or TDM networks).

8. Security Considerations

Security is listed as an explicit goal for the architecture and this issue is addressed in the section on goals. However, as this is an informational memo it does not propose any concrete mechanisms to achieve these goals.

9. IANA Considerations

This memo includes no request to IANA.

10. Acknowledgements

The authors want to thank Carlo Delle Donne, Matthew Skrzypczyk, Axel Dahlberg, Mathias van den Bossche, Patrick Gelard, Chonggang Wang, Scott Fluhrer, Joey Salazar, Joseph Touch, and the rest of the QIRG community as a whole for their very useful reviews and comments to the document.

11. Informative References

- [1] Carpenter, B., Ed., "Architectural Principles of the Internet", RFC 1958, DOI 10.17487/RFC1958, June 1996, <<https://www.rfc-editor.org/info/rfc1958>>.

- [2] Wang, C., Rahman, A., Li, R., and M. Aelmans, "Applications and Use Cases for the Quantum Internet", draft-irtf-qirg-quantum-internet-use-cases-02 (work in progress), September 2020.
- [3] Bennett, C. and G. Brassard, "Quantum cryptography: Public key distribution and coin tossing", *Theoretical Computer Science* 560, 7-11, 2014, <<http://www.sciepub.com/reference/53249>>.
- [4] Ekert, A., "Quantum cryptography based on Bell's theorem", *Phys. Rev. Lett.* Vol. 67, Iss. 6, 1991, <<https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.67.661>>.
- [5] Crepeau, C., Gottesman, D., and A. Smith, "Secure multi-party quantum computation", *Proceedings of Symposium on Theory of Computing*, 2002, <<https://arxiv.org/abs/quant-ph/0206138>>.
- [6] Giovanetti, V., Lloyd, S., and L. Maccone, "Quantum-enhanced measurements: beating the standard quantum limit", *Science* 306(5700), 1330-1336, 2004, <<https://arxiv.org/abs/quant-ph/0412078>>.
- [7] Castelvechi, D., "The Quantum Internet has arrived (and it hasn't)", *Nature* 554, 289-292, 2018, <<https://www.nature.com/articles/d41586-018-01835-3>>.
- [8] Wehner, S., Elkouss, D., and R. Hanson, "Quantum internet: A vision for the road ahead", *Science* 362, 6412, 2018, <<http://science.sciencemag.org/content/362/6412/eaam9288.full>>.
- [9] Aspect, A., Grangier, P., and G. Roger, "Experimental tests of realistic local theories via Bell's theorem", *Phys. Rev. Lett.* 47 (7): 460-463, 1981, <<https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.47.460>>.
- [10] Muralidharan, S., Li, L., Kim, J., Lutkenhaus, N., Lukin, M., and L. Jiang, "Optimal architectures for long distance quantum communication", *Nat. Sci. Rep.* 6, 20463, 2016, <<https://www.nature.com/articles/srep20463>>.
- [11] Van Meter, R. and J. Touch, "Designing quantum repeater networks", *IEEE Communications Magazine* 51, 64-71, 2013, <<https://ieeexplore.ieee.org/document/6576340>>.

- [12] Satoh, T., Nagayama, S., Suzuki, S., Matsuo, T., and R. Van Meter, "Attacking the quantum internet", arXiv 2005.04617, 2020, <<https://arxiv.org/abs/2005.04617>>.
- [13] Satoh, T., Nagayama, S., and R. Van Meter, "The network impact of hijacking a quantum repeater", Quantum Science and Technology Vol. 3, Iss. 3, 2017, <<https://arxiv.org/abs/1701.04587>>.
- [14] Dahlberg, A., Skrzypczyk, M., Coopmans, T., Wubben, L., Rozpedek, F., Pompili, M., Stolk, A., Pawelczak, P., Kneijens, R., de Oliveira Filho, J., Hanson, R., and S. Wehner, "A link layer protocol for quantum networks", SIGCOMM '19 Proceedings of the ACM Special Interest Group on Data Communication 159-173, 2019, <<https://arxiv.org/abs/1903.09778>>.
- [15] Sutor, R., "Dancing with Qubits", Packt Publishing , 2019.
- [16] Nielsen, M. and I. Chuang, "Quantum Computation and Quantum Information", Cambridge University Press , 2011.
- [17] Bennett, C., DiVincenzo, D., Smolin, J., and W. Wootters, "Mixed state entanglement and quantum error correction", Phys. Rev. A Vol. 54, Iss. 5, 1996, <<https://arxiv.org/abs/quant-ph/9604024>>.
- [18] Briegel, H., Dur, W., Cirac, J., and P. Zoller, "Quantum repeaters: The role of imperfect local operations in quantum communication", Phys. Rev. Lett. Vol. 81, Num. 26, 1998, <<https://arxiv.org/abs/quant-ph/9803056>>.
- [19] Cacciapuoti, A., Caleffi, M., Van Meter, R., and L. Hanzo, "When Entanglement meets Classical Communications: Quantum Teleportation for the Quantum Internet", , 2019, <<https://arxiv.org/abs/1907.06197>>.
- [20] Mignant, C., Markham, D., and F. Grosshans, "Distributing graph states over arbitrary quantum networks", Phys. Rev. A Vol. 100, Iss. 5, 2019, <<https://arxiv.org/abs/1811.05445>>.
- [21] Van Meter, R., Satoh, T., Ladd, T., Munro, W., and K. Nemoto, "Path selection for quantum repeater networks", Networking Science Vol. 3, Iss. 1-4, pp 82-95, 2013, <<https://arxiv.org/abs/1206.5655>>.

- [22] Clark, D., "The design philosophy of the DARPA internet protocols", SIGCOMM '88, 1988, <<https://dl.acm.org/doi/abs/10.1145/52324.52336>>.
- [23] Van Meter, R., "Quantum Networking", ISTE Ltd/John Wiley and Sons Inc 978-1-84821-537-5, 2014.
- [24] Aguado, A., Lopez, V., Diego, D., Peev, M., Poppe, A., Pastor, A., Folgueira, J., and M. Vicente, "The engineering of software-defined quantum key distribution networks", IEEE Communications Magazine Vol. 57, Iss. 7, 2019, <<http://arxiv.org/abs/1907.00174>>.
- [25] Peev, M., Pacher, C., Alleaume, R., Barreiro, C., Bouda, J., Boxleitner, W., Debuisschert, T., Diamanti, E., Dianati, M., Dynes, J., Fasel, S., Fossier, S., Fuerst, M., Gautier, J., Gay, O., Gisin, N., Grangier, P., Happe, A., Hasani, Y., Hentschel, M., Huebel, H., Humer, G., Laenger, T., Legre, M., Lieger, R., Lodewyck, J., Loruenser, T., Luetkenhaus, N., Marhold, A., Matyus, T., Maurhart, O., Monat, L., Nauerth, S., Page, J., Poppe, A., Querasser, E., Ribordy, G., Robyr, S., Salvail, L., Sharpe, A., Shields, A., Stucki, D., Suda, M., Tamas, C., Themel, T., Thew, R., Thoma, Y., Treiber, A., Trinkler, P., Tualle-Brouiri, R., Vannel, F., Walenta, N., Weier, H., Weinfurter, H., Wimberger, I., Yuan, Z., Zbinden, H., and A. Zeilinger, "The SECOQC quantum key distribution network in Vienna", New J. Phys. Vol. 11, 2009, <<http://stacks.iop.org/1367-2630/11/i=7/a=075001>>.
- [26] Hensen, B., Bernien, H., Dreau, A., Reiserer, A., Kalb, N., Blok, M., Ruitenberg, J., Vermeulen, R., Schouten, R., Abellan, C., Amaya, W., Pruneri, V., Mitchell, M., Markham, M., Twitchen, D., Elkouss, D., Wehner, S., Taminiau, T., and R. Hanson, "Loophole-free {Bell} inequality violation using electron spins separated by 1.3 kilometres", Nature 526, 682-686, 2015, <<https://arxiv.org/abs/1508.05949>>.
- [27] Fitzsimons, J. and E. Kashefi, "Unconditionally verifiable blind quantum computation", Phys. Rev. A Vol. 96, Iss. 1, 2017, <<https://arxiv.org/abs/1203.5217>>.
- [28] Terhal, B., "Is entanglement monogamous?", IBM Journal of Research and Development Vol. 48, Iss. 1, 2004, <<https://ieeexplore.ieee.org/document/5388928>>.

- [29] Park, J., "The concept of transition in quantum mechanics", *Foundations of Physics* Vol. 1, Iss. 1, 1970, <<https://link.springer.com/content/pdf/10.1007/BF00708652.pdf>>.
- [30] Wootters, W. and W. Zurek, "A single quantum cannot be cloned", *Nature* 299, 802–803, 1982, <<https://www.nature.com/articles/299802a0>>.
- [31] Fowler, A., Wang, D., Hill, C., Ladd, T., Van Meter, R., and L. Hollenberg, "Surface code quantum communication", *Phys. Rev. Lett.* Vol. 104, Iss. 18, 2010, <<https://arxiv.org/abs/0910.4074>>.
- [32] Gottesman, D., Jennewein, T., and S. Croke, "Longer-baseline telescopes using quantum repeaters", *Phys. Rev. Lett.* Vol. 109, Iss. 7, 2012, <<https://arxiv.org/abs/1107.2939>>.
- [33] "The Quantum Protocol Zoo", <<https://wiki.veriqloud.fr/>>.
- [34] Duer, W. and H. Briegel, "Entanglement purification and quantum error correction", *Rep. Prog. Phys.* Vol. 70, Iss. 8, 2007, <<https://arxiv.org/abs/0705.4165>>.
- [35] Kimble, H., "The Quantum Internet", *Nature* 453, 1023–1030, 2008, <<http://arxiv.org/abs/0806.4195>>.
- [36] Devitt, S., Nemoto, K., and W. Munro, "Quantum error correction for beginners", *Rep. Prog. Phys.* Vol. 76, Iss. 7, 2013, <<https://arxiv.org/abs/0905.2794>>.
- [37] Sangouard, N., Simon, C., de Riedmatten, H., and N. Gisin, "Quantum repeaters based on atomic ensembles and linear optics", *Rev. Mod. Phys.* Vol. 83, Iss. 1, 2011, <<https://arxiv.org/abs/0906.2699>>.
- [38] Abobeih, M., Cramer, J., Bakker, M., Kalb, N., Markham, M., Twitchen, D., and T. Taminiau, "One-second coherence for a single electron spin coupled to a multi-qubit nuclear-spin environment", *Nat. Comm.* 9, 2552, 2018, <<https://arxiv.org/abs/1801.01196>>.
- [39] Bradley, C., Randall, J., Abobeih, M., Berrevoets, R., Degen, M., Bakker, M., Markham, M., Twitchen, D., and T. Taminiau, "A 10-qubit solid-state spin register with quantum memory up to one minute", *Phys. Rev. X* Vol. 9, Iss. 3, 2019, <<https://arxiv.org/abs/1905.02094>>.

Authors' Addresses

Wojciech Kozlowski
QuTech
Building 22
Lorentzweg 1
Delft 2628 CJ
Netherlands

Email: w.kozlowski@tudelft.nl

Stephanie Wehner
QuTech
Building 22
Lorentzweg 1
Delft 2628 CJ
Netherlands

Email: s.d.c.wehner@tudelft.nl

Rodney Van Meter
Keio University
5322 Endo
Fujisawa, Kanagawa 252-0882
Japan

Email: rdv@sfc.wide.ad.jp

Bruno Rijsman
Individual

Email: brunorijsman@gmail.com

Angela Sara Cacciapuoti
University of Naples Federico II
Department of Electrical Engineering and Information Technologies
Claudio 21
Naples 80125
Italy

Email: angelasara.cacciapuoti@unina.it

Marcello Caleffi
University of Naples Federico II
Department of Electrical Engineering and Information Technologies
Claudio 21
Naples 80125
Italy

Email: marcello.caleffi@unina.it

Shota Nagayama
Mercari, Inc.
Roppongi Hills Mori Tower 18F
6-10-1 Roppongi, Minato-ku
Tokyo 106-6118
Japan

Email: shota.nagayama@mercari.com

QIRG
Internet-Draft
Intended status: Informational
Expires: May 2, 2021

C. Wang
A. Rahman
InterDigital Communications, LLC
R. Li
NICT
M. Aelmans
Juniper Networks
October 29, 2020

Applications and Use Cases for the Quantum Internet
draft-irtf-qirg-quantum-internet-use-cases-03

Abstract

The Quantum Internet has the potential to improve application functionality by incorporating quantum information technology into the infrastructure of the overall Internet. In this document, we provide an overview of some applications expected to be used on the Quantum Internet, and then categorize them using various classification schemes. Some general requirements for the Quantum Internet are also discussed. The intent of this document is to provide a common understanding and framework of applications and use cases for the Quantum Internet.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on May 2, 2021.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. Conventions used in this document	3
3. Terms and Acronyms List	3
4. Quantum Internet Applications	5
4.1. Overview	5
4.2. Classification by Application Usage	5
4.2.1. Quantum Cryptography Applications	6
4.2.2. Quantum Sensor Applications	6
4.2.3. Quantum Computing Applications	6
4.3. Control vs Data Plane Classification	7
5. Selected Quantum Internet Use Cases	8
5.1. Secure Communication Setup	9
5.2. Secure Quantum Computing with Privacy Preservation	11
5.3. Distributed Quantum Computing	13
6. General Requirements	15
6.1. Background	15
6.2. Requirements	17
7. Conclusion	18
8. IANA Considerations	18
9. Security Considerations	19
10. Acknowledgments	19
11. Informative References	20
Authors' Addresses	23

1. Introduction

The Classical Internet has been constantly growing since it first became commercially popular in the early 1990's. It essentially consists of a large number of end-nodes (e.g., laptops, smart phones, network servers) connected by routers. The end-nodes may run applications that provide service for the end-users such as processing and transmission of voice, video or data. The connections between the various nodes in the Internet include Digital Subscriber Lines (DSLs), fiber optics, coax cable and wireless that include Bluetooth, WiFi, cellular (e.g., 3G, 4G, 5G), and satellite, etc. Bits are transmitted across the Classical Internet in packets.

Research and experimentation have picked up over the last few years for developing a Quantum Internet [Wehner]. It is anticipated that the Quantum Internet will provide intrinsic benefits such as better end-to-end and network security. The Quantum Internet will also have end-nodes, termed quantum end-nodes. Quantum end-nodes may be connected by quantum repeaters/routers. These quantum end-nodes will also run value-added applications which will be discussed later.

The connections between the various nodes in the Quantum Internet are expected to be primarily fiber optics and free-space optics. Photonic connections are particularly useful because light (photons) is very suitable for physically encoding qubits. Unlike the Classical Internet, qubits (and not classical bits or packets) are expected to be transmitted across the Quantum Internet due to the underlying physics. The Quantum Internet will operate according to unique physical principles such as quantum superposition, entanglement and teleportation [I-D.irtf-qirg-principles].

The Quantum Internet is not anticipated to replace the Classical Internet. For instance, Local Operations and Classical Communication (LOCC) tasks [Chitambar] rely on classical communications. Instead the Quantum Internet will run in conjunction with the Classical Internet to form a new Hybrid Internet. The process of integrating the Quantum Internet with the classical Internet is similar to, but with more profound implications, as the process of introducing any new communication and networking paradigm into the existing Internet. The intent of this document is to provide a common understanding and framework of applications and use cases for the Quantum Internet.

2. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

3. Terms and Acronyms List

This document assumes that the reader is familiar with the quantum information technology related terms and concepts that are described in [I-D.irtf-qirg-principles]. In addition, the following terms and acronyms are defined here for clarity:

- o Bit - Binary Digit (i.e., fundamental unit of information in a classical computer).
- o Classical Internet - The existing, deployed Internet (circa 2020) where bits are transmitted in packets between nodes to convey information. The Classical Internet supports applications which

may be enhanced by the Quantum Internet. For example, the end-to-end security of a Classical Internet application may be improved by secure communication setup using a quantum application.

- o Hybrid Internet - The "new" or evolved Internet to be formed due to a merger of the Classical Internet and the Quantum Internet.
- o Local Operations and Classical Communication (LOCC) - A method where: 1) local quantum operations (e.g., quantum measurement) are performed at one quantum node A; 2) the quantum operation result is sent to another quantum node B via classical communications; 3) the quantum node B may also perform some local quantum operations dependent on the received operation result from the quantum node A. For example, LOCC can be used to transform entangled states into other entangled states.
- o Noisy Intermediate-Scale Quantum (NISQ) - NISQ was defined in [Preskill] to represent a near-term era in quantum technology. According to this definition, NISQ computers have two salient features: (1) The size of NISQ computers range from 50 to a few hundred qubits (i.e., intermediate-scale); and (2) Qubits in NISQ computers have inherent errors and the control over them is imperfect (i.e., noisy).
- o Packet - Formatted unit of multiple related bits. The bits contained in a packet may be classical bits, or the measured state of qubits.
- o Quantum End-node - An end-node hosts user applications and interfaces with the rest of the Internet. Typically, an end-node may serve in a client, server, or peer-to-peer role as part of the application. If the end-node is part of a Quantum Network (i.e, is a quantum end-node), it must be able to generate/transmit and/or receive/process qubits. A quantum end-node must also be able to interface to the Classical Internet for control purposes and thus also be able to receive, process, and transmit classical bits/packets.
- o Quantum Computer (QC) - A quantum end-node that also has quantum memory and quantum computing capabilities is regarded as a full-fledged quantum computer.
- o Quantum Network - A new type of network enabled by quantum information technology where qubits are transmitted between nodes to convey information. (Note: qubits must be sent individually and not in packets). The Quantum Network will use both quantum channels, and classical channels provided by the Classical Internet.

- o Quantum Internet - A network of Quantum Networks. The Quantum Internet will be merged into the Classical Internet to form a new Hybrid Internet. The Quantum Internet may either improve classical applications or may enable new quantum applications.
- o Qubit - Quantum Bit (i.e., fundamental unit of information in a quantum computer). It is similar to a classic bit in that the state of a qubit is either "0" or "1" after it is measured and is denoted as its basis state $|0\rangle$ or $|1\rangle$. However, the qubit is different than a classic bit in that the qubit is in a linear combination of both states before it is measured and termed to be in superposition. The Degrees of Freedom (DOF) of a photon (e.g., polarization) or an electron (e.g., spin) can be used to encode a qubit.

4. Quantum Internet Applications

4.1. Overview

The Quantum Internet is expected to be extremely beneficial for a subset of existing and new applications. The expected applications using Quantum Internet are still being developed as we are in the formative stages of the Quantum Internet [Castelvecchi] [Wehner]. However, an initial (and non-exhaustive) list of the applications to be supported on the Quantum Internet can be identified and classified using two different schemes. Note, we do not include quantum computing applications that are purely local to a given node (e.g., quantum random number generator).

4.2. Classification by Application Usage

Applications may also be grouped by the usage that they serve into a tripartite classification. Specifically, applications may be classified according to the following usages:

- o Quantum cryptography applications - Refers to the use of quantum information technology to ensure secure communications (e.g., QKD).
- o Quantum sensors applications - Refers to the use of quantum information technology for supporting distributed sensors or Internet of Things (IoT) devices (e.g., clock synchronization).
- o Quantum computing applications - Refers to the use of quantum information technology for supporting remote quantum computing facilities (e.g., distributed quantum computing).

This is a useful classification scheme as it can be easily understood by both a technical and non-technical audience. Following are some more details.

4.2.1. Quantum Cryptography Applications

Examples of quantum cryptography applications include quantum-based secure communication setup and fast Byzantine negotiation.

1. Secure communication setup - Refers to secure cryptographic key distribution between two or more end-nodes. The most well-known method is referred to as Quantum Key Distribution (QKD) [Renner].
2. Fast Byzantine negotiation - Refers to a Quantum Network based method for fast agreement in Byzantine negotiations [Fitzi]. This can be used for the popular financial blockchain feature as well as other distributed computing features which use Byzantine negotiations.

4.2.2. Quantum Sensor Applications

Example of quantum sensor applications include network clock synchronization, radio frequency measurement, etc. These applications mainly leverage a network of entangled quantum sensors (i.e. quantum sensor networks) for high-precision multi-parameter estimation [Proctor] [Zhuang].

1. Network clock synchronization - Refers to a world wide set of atomic clocks connected by the Quantum Internet to achieve an ultra precise clock signal [Komar].
2. Radio frequency sensing - Refers to leverage connected quantum sensors to measure a broad range of radio frequencies with arbitrary frequency resolution. [Zhuang] [Fan]

4.2.3. Quantum Computing Applications

Examples of quantum computing include distributed quantum computing and secure quantum computing with privacy preservation.

1. Distributed quantum computing - Refers to a collection of remote small capacity quantum computers (i.e., each supporting a few qubits) that are connected and working together in a coordinated fashion so as to simulate a virtual large capacity quantum computer [Wehner].
2. Secure quantum computing with privacy preservation - Refers to private, or blind, quantum computation, which provides a way for

a client to delegate a computation task to one or more remote quantum computers without disclosing the source data to be computed over [Fitzsimons].

4.3. Control vs Data Plane Classification

The majority of routers currently used in the Classical Internet separate control plane functionality and data plane functionality for, amongst other reasons, stability, capacity and security. In order to classify applications for the Quantum Internet, a somewhat similar distinction can be made. Specifically some applications can be classified as being responsible for initiating sessions and performing other control plane functionality. Other applications carry application or user data and can be classified as data plane functionality.

Some examples of what may be called control plane applications in the Classical Internet are Domain Name Server (DNS), Session Information Protocol (SIP), and Internet Control Message Protocol (ICMP). Furthermore, examples of data plane applications are E-mail, web browsing, and video streaming. Note that some applications may require both control plane and data plane functionality. For example, a Voice over IP (VoIP) application may use SIP to set up the call and then transmit the VoIP user packets over the data plane to the other party.

Similarly, nodes in the Quantum Internet applications may also use the classification paradigm of control plane functionality versus data plane functionality where:

- o Control Plane - Network functions and processes that operate on (1) control bits/packets or qubits (e.g., to setup up end-user encryption); or (2) management bits/packets or qubits (e.g., to configure nodes). For example, a quantum ping could be implemented as a control plane application to test and verify if there is a quantum connection between two quantum nodes. Another example is quantum superdense encoding (which is used to transmit two classical bits by sending only one qubit). This approach does not need classical channels. Quantum superdense coding can be leveraged to implement a secret sharing application to share secrets between two parties. This secret sharing application based on quantum superdense encoding can be classified as control plane functionality.
- o Data Plane - Network functions and processes that operate on end-user application bits/packets or qubits (e.g., voice, video, data). Sometimes also referred to as the user plane. For example, a data plane application can be video conferencing, which

uses QKD-based secure communication setup (which is a control plane function) to share a secret key for encrypting and decrypting video frames.

As shown in the table in Figure 1, control and data plane applications vary for different types of networks. For a standalone Quantum Network (i.e., that is not integrated into the Internet), entangled qubits are its "data" and thus entanglement distribution can be regarded as its data plane application, while the signalling for controlling entanglement distribution be considered as control plane. But looking at Quantum Internet, QKD-based secure communication setup, which may be based on and leverage entanglement distribution, is in fact a control plane application, while video conference using QKD-based secure communication setup is a data plane application.

	Classical Internet Examples	Standalone Quantum Network Examples	Quantum Internet Examples
Control Plane	ICMP, DNS	Signalling for controlling entanglement distribution	QKD-based secure communication setup; Quantum ping
Data Plane	Web Browsing	Entanglement distribution	Video conference using QKD-based secure communication setup

Figure 1: Examples of Control vs Data Plane Classification

5. Selected Quantum Internet Use Cases

The Quantum Internet will support a variety of applications and deployment configurations. This section details a few key use cases which illustrates the benefits of the Quantum Internet. In system engineering, a use case is typically made up of a set of possible sequences of interactions between nodes and users in a particular environment and related to a particular goal. This will be the definition that we use in this section.

5.1. Secure Communication Setup

In this scenario, two banks (i.e., Bank #1 and Bank #2) need to have secure communications for transmitting important financial transaction records (see Figure 2). For this purpose, they first need to securely exchange a classic secret cryptographic key (i.e., a sequence of classical bits), which is triggered by an end-user banker at Bank #1. This results in a source quantum node A at Bank #1 to securely send a classic secret key to a destination quantum node B at Bank #2. This is referred to as a secure communication setup. Note that the quantum node A and B may be either a bare-bone quantum end-node or a full-fledged quantum computer. This use case shows that the Quantum Internet can be leveraged to improve the security of Classical Internet applications of which the financial application shown in Figure 2 is an example.

One requirement for this secure communication setup process is that it should not be vulnerable to any classical or quantum computing attack. This can be realized using QKD [ETSI-QKD-Interfaces]. QKD can securely establish a secret key between two quantum nodes, without physically transmitting it through the network and thus achieving the required security. QKD is the most mature feature of the quantum information technology, and has been commercially deployed in small-scale and short-distance deployments. More QKD use cases are described in ETSI documents [ETSI-QKD-UseCases].

In general, QKD (e.g., [BB84]) without using entanglement works as follows:

1. The source quantum node A (e.g. Alice) transforms the secret key to qubits. Basically, for each classical bit in the secret key, the source quantum node A randomly selects one quantum computational basis and uses it to prepare/generate a qubit for the classical bit.
2. The source quantum node A sends qubits to the destination quantum node B (e.g. Bob) via quantum channel.
3. The destination quantum node receives qubits and measures them based on its random quantum basis.
4. The destination node informs the source node of its random quantum basis.
5. The source node informs the destination node which random quantum basis is correct.

6. Both nodes discard any measurement bit under different quantum basis and store all remaining bits as the secret key.

It is worth noting that:

1. There are some entanglement-based QKD protocols such as [Treiber], which work differently than above steps. The entanglement-based schemes, where entangled states are prepared externally to Alice and Bob, are not normally considered "prepare-and-measure" as defined in [Wehner]; other entanglement-based schemes, where entanglement is generated within Alice can still be considered "prepare-and-measure"; send-and-return schemes can still be "prepare-and-measure", if the information content, from which keys will be derived, is prepared within Alice before being sent to Bob for measurement.
2. There are many enhanced QKD protocols based on [BB84]. For example, a series of loopholes have been identified due to the imperfections of measurement devices; there are several solutions to take into account these attacks such as measurement-device-independent QKD [ZhangPeiyu]. These enhanced QKD protocol can work differently than the steps of BB84 protocol [BB84].
3. For large-scale QKD, QKD Networks (QKDN) are required, which can be regarded as a subset of a Quantum Internet. A QKDN may consist of a QKD application layer, a QKD network layer, and a QKD link layer [QinHao]. One or multiple trusted QKD relays [ZhangQiang] may exist between the source quantum node A and the destination quantum node B, which are connected by a QKDN. Alternatively, a QKDN may rely on entanglement distribution and entanglement-based QKD protocols; as a result, quantum-repeaters/routers instead of trusted QKD relays are needed for large-scale QKD.
4. In general, there are three types of QKD solutions: 1) Basic QKD: In this case, QKD only works for two directly connected quantum nodes within a short distance or a network segment. The end-to-end security relies on some trusted nodes, which however could be attacked; 2) E2E QKD: In this case, QKD works for two faraway quantum nodes to provide the end-to-end security without relying on trusted nodes; and 3) Advanced E2E QKD: In this case, QKD leverages entanglement distribution to achieve the end-to-end security.

As a result, the Quantum Internet in Figure 2 contains quantum channels. And in order to support secure communication setup especially in large-scale deployment, it also requires entanglement generation and entanglement distribution

[I-D.van-meter-qirg-quantum-connection-setup], quantum repeaters/routers, and/or trusted QKD relays.

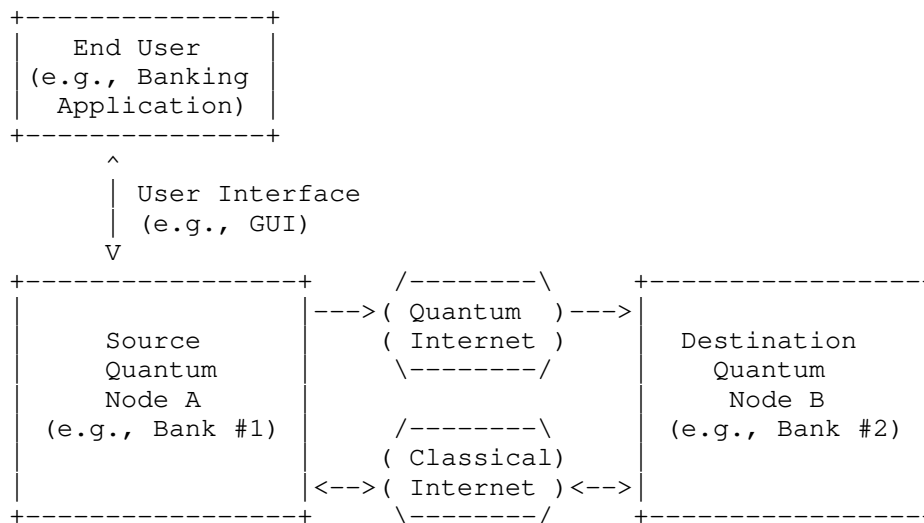


Figure 2: Secure Communication Setup

5.2. Secure Quantum Computing with Privacy Preservation

Secure computation with privacy preservation refers to the following scenario:

1. A client node with source data delegates the computation of the source data to a remote computation node.
2. Furthermore, the client node does not want to disclose any source data to the remote computation node and thus preserve the source data privacy.
3. Note that there is no assumption or guarantee that the remote computation node is a trusted entity from the source data privacy perspective.

As an example illustrated in Figure 3, the client node could be a virtual voice-controlled home assistant device like Amazon’s Alexa product. The remote computation node could be a quantum computer in the cloud. A resident as an end-user uses voice to control the home device. The home device captures voice-based commands from the end-

user. Then, the home device interfaces to a home quantum terminal node (e.g., a home gateway), which interacts with the remote computation node to perform computation over the captured voice-based commands. The home quantum terminal could be either a bare-bone quantum end-node or a full-fledged quantum computer.

In this particular case, there is no privacy concern since the source data (i.e., captured voice-based commands) will not be sent to the remote computation node which could be compromised. Protocols [Fitzsimons] for delegated quantum computing or blind quantum computation can be leveraged to realize secure delegated computation and guarantee privacy preservation simultaneously. Using delegated quantum computing protocols, the client node does not need send the source data but qubits with some measurement instructions to the remote computation node (e.g., a quantum computer).

After receiving qubits and measurement instructions, the remote computation node performs the following actions:

1. It first performs certain quantum operations on received qubits and measure them according to received measurement instructions to generate computation results (in classic bits).
2. Then it sends the computation results back to the client node via classical channel.
3. In this process, the source data is not disclosed to the remote computation node and the privacy is preserved.

In Figure 3, the Quantum Internet contains quantum channels and quantum repeaters/routers for long-distance qubits transmission [I-D.irtf-qirg-principles].

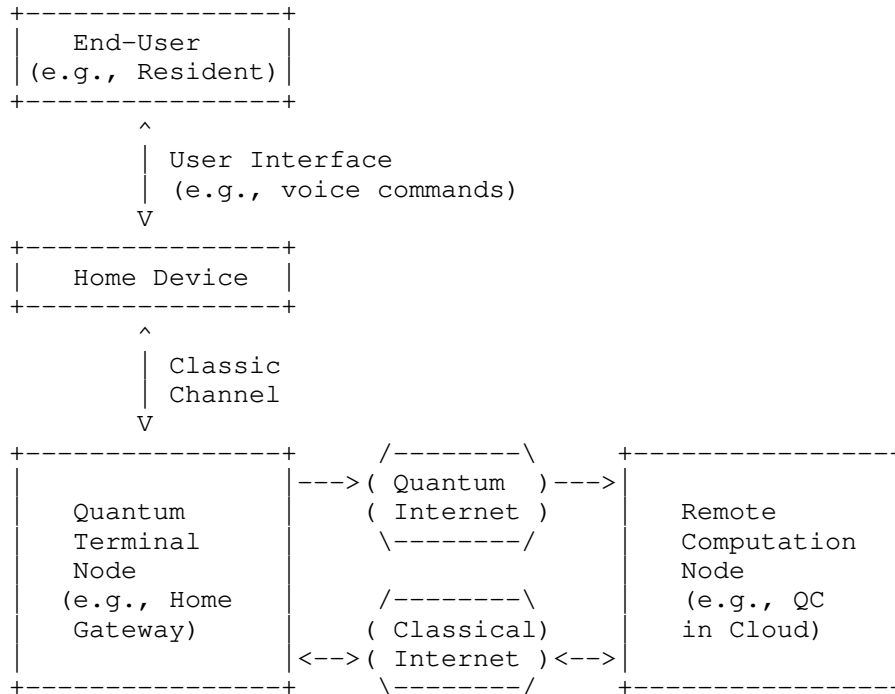


Figure 3: Secure Quantum Computing with Privacy Preservation

5.3. Distributed Quantum Computing

In this scenario, Noisy Intermediate-Scale Quantum (NISQ) computers distributed in different locations are available for sharing. According to the definition in [Preskill], a NISQ computer can only realize a small number of qubits and has limited quantum error correction. In order to gain higher computation power before fully-fledged quantum computers become available, NISQ computers can be connected via classic and quantum channels. This scenario is referred to as distributed quantum computing [Caleffi] [Cacciapuoti01] [Cacciapuoti02]. This use case reflects the vastly increased computing power which quantum computers as a part of the Quantum Internet can bring, in contrast to classical computers in the Classical Internet.

As an example, scientists can leverage these connected NISQ computer to solve highly complex scientific computation problems such as analysis of chemical interactions for medical drug development (see Figure 4). In this case, qubits will be transmitted among connected quantum computers via quantum channels, while classic control

messages will be transmitted among them via classical channels for coordination and control purpose. Qubits from one NISQ computer to another NISQ computer are very sensitive and cannot be lost. For this purpose, quantum teleportation can be leveraged to teleport sensitive data qubits from one quantum computer A to another quantum computer B. Note that Figure 4 does not cover measurement-based distributed quantum computing, where quantum teleportation may not be required.

Specifically, the following steps happen between A and B. In fact, LOCC [Chitambar] operations are conducted at the quantum computer A and B in order to achieve quantum teleportation as illustrated in Figure 4.

1. The quantum computer A locally generates some sensitive data qubits to be teleported to the quantum computer B.
2. A shared entanglement is established between the quantum computer A and the quantum computer B (i.e., there are two entangled qubits: $|q1\rangle$ at A and $|q2\rangle$ at B).
3. Then, the quantum computer A performs a Bell measurement of the entangled qubit $|q1\rangle$ and the sensitive data qubit.
4. The result from this Bell measurement will be encoded in two classical bits, which will be physically transmitted via a classical channel to the quantum computer B.
5. Based on the received two classical bits, the quantum computer B modifies the state of the entangled qubit $|q2\rangle$ in the way to generate a new qubit identical to the sensitive data qubit at the quantum computer A.

In Figure 4, the Quantum Internet contains quantum channels and quantum repeaters/routers [I-D.irtf-qirg-principles]. This use case needs to support entanglement generation in order to enable quantum teleportation, entanglement distribution or quantum connection setup [I-D.van-meter-qirg-quantum-connection-setup] in order to support long-distance quantum teleportation.

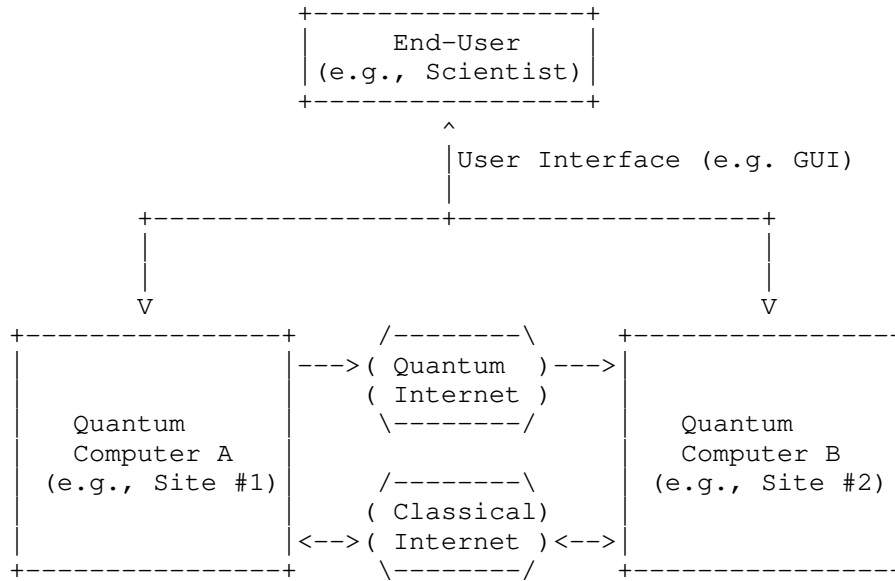


Figure 4: Distributed Quantum Computing

6. General Requirements

6.1. Background

Quantum technologies are steadily evolving and improving. Therefore, it is hard to predict the timeline and future milestones of quantum technologies as pointed out in [Grumbling] for quantum computing. Currently, a NISQ computer can achieve fifty to hundreds of qubits with some given error rate. In fact, the error rates of two-qubit quantum gates have decreased nearly in half every 1.5 years (for trapped ion gates) to 2 years (for superconducting gates). The error rate also increases as the number of qubits increases. For example, a current 20-qubit machine has a total error rate which is close to the total error rate of a 7 year old two-qubit machine [Grumbling].

On the network level, six stages of Quantum Internet development are described in [Wehner] as follows:

1. Trusted repeater networks (Stage-1)
2. Prepare and measure networks (Stage-2)
3. Entanglement distribution networks (Stage-3)

4. Quantum memory networks (Stage-4)
5. Fault-tolerant few qubit networks (Stage-5)
6. Quantum computing networks (Stage-6)

The first stage are simple trusted repeater networks, while the final stage are quantum computing (sub)networks where the full-blown Quantum Internet will be achieved. Each intermediate stage brings with it new functionality, new applications, and new characteristics. Figure 5 illustrates Quantum Internet use cases as described in this document mapped to the Quantum Internet stages described in [Wehner]. For example, secure communication setup can be supported in Stage-1, Stage-2, or Stage-3, but with different QKD solutions. More specifically:

In Stage-1, basic QKD is possible and can be leveraged to support secure communication setup but trusted nodes are required to provide end-to-end security.

In Stage-2, E2E QKD without relying on trusted nodes is possible to support secure communication setup too and the primary requirement is long-distance qubit transmission.

In Stage-3, Advanced E2E QKD can be enabled based on quantum repeater and entanglement distribution to support the same secure communication setup application.

In Stage-4, Secure quantum computing with privacy-preservation can be enabled since it needs quantum memory for multiple rounds of quantum computation.

Finally, in Stage-6, distributed quantum computing relaying more qubits can be supported.

Quantum Internet Stage	Example Quantum Internet Use Cases	Characteristic
Stage-1	Secure Comm Setup with Basic QKD	Trusted Nodes
Stage-2	Secure Comm Setup with E2E QKD	Long-distance qubit transmission
Stage-3	Secure Comm Setup with Advanced E2E QKD	Entanglement distribution
Stage-4	Secure/Blind Quantum Computing	Quantum memory
Stage-5	Higher-accuracy clock synchronization	Fault tolerance
Stage-6	Distributed quantum computing	More qubits

Figure 5: Example Use Cases in Different Quantum Internet Stages

6.2. Requirements

Some general and functional requirements on the Quantum Internet from the networking perspective, based on the above applications and use cases, are identified as follows:

1. Methods for facilitating quantum applications to interact efficiently with entanglement qubits are necessary in order for them to trigger distribution of designated entangled qubits to potentially any other quantum node residing in the Quantum Internet. To accomplish this specific operations must be performed on entangled qubits (e.g., entanglement swapping, entanglement distillation). Quantum nodes may be quantum end-nodes, quantum repeaters/routers, and/or quantum computers.
2. Quantum repeaters/routers should support robust and efficient entanglement distribution in order to extend and establish entanglement connection between two quantum nodes. For achieving this, it is required to first generate an entangled pair on each hop of the path between these two nodes.

3. Quantum end-nodes must send additional information on classical channels to aid in transmission of qubits across quantum repeaters/receivers. This is because qubits are transmitted individually and do not have any associated packet overhead which can help in transmission of the qubit. Any extra information to aid in routing, identification, etc., of the qubit(s) must be sent via classical channels.
4. Methods for managing and controlling the Quantum Internet including quantum nodes and their quantum resources are necessary. The resources of a quantum node may include quantum memory, quantum channels, qubits, established quantum connections, etc. Such management methods can be used to monitor network status of the Quantum Internet, diagnose and identify potential issues (e.g. quantum connections), and configure quantum nodes with new actions and/or policies (e.g. to perform a new entanglement swapping operation). New management information model for the Quantum Internet may need to be developed.

7. Conclusion

This document provides an overview of some expected applications for the Quantum Internet, and then details selected use cases. The applications are first grouped by their usage which is a natural and easy to understand classification scheme. The applications are then classified as either control plane or data plane functionality as typical for the classical Internet. This set of applications may, of course, naturally expand over time as the Quantum Internet matures. Finally, some general requirements for the Quantum Internet are also provided.

This document can also serve as an introductory text to persons interested in learning about the practical uses of the Quantum Internet. Finally, it is hoped that this document will help guide further research and development of the specific Quantum Internet functionality required to implement the applications and use cases described herein. To this end, a few key requirements for the Quantum Internet are specified.

8. IANA Considerations

This document requests no IANA actions.

9. Security Considerations

This document does not define an architecture nor a specific protocol for the Quantum Internet. It focuses instead on detailing use cases, requirements, and describing typical Quantum Internet applications. However, some useful observations can be made regarding security as follows.

It has been clearly identified that once large-scale quantum computing becomes reality it will be able to theoretically break many of the public-key (i.e., asymmetric) cryptosystems currently in use because of the exponential increase of computing power with quantum computing. This would negatively affect many of the security mechanisms currently in use on the classic Internet. This has given strong impetus for starting development of new cryptographic systems that are secure against quantum computing attacks [NISTIR8240].

Paradoxically, development of a Quantum Internet will also mitigate the threats posed by quantum computing attacks against public-key cryptosystems. Specifically, the secure communication setup feature of the Quantum Internet as described in Section 5.1 will be strongly resistant to both classical and quantum computing attacks against public-key cryptosystems.

A key additional threat consideration for the Quantum Internet is pointed to by [RFC7258], which warns of the dangers of pervasive monitoring as a widespread attack on privacy. Pervasive monitoring is defined as a widespread, and usually covert, surveillance through intrusive gathering of application content or protocol metadata such as headers. This can be accomplished through active or passive wiretaps, traffic analysis, or subverting the cryptographic keys used to secure communications.

Once again, the secure communication setup feature of the Quantum Internet as described in Section 5.1 will be strongly resistant to pervasive monitoring. In addition, Section 5.2 provides a method to perform remote quantum computing while preserving the privacy of the source data.

10. Acknowledgments

The authors want to thank Mathias Van Den Bossche, Xavier de Foy, Patrick Gelard, Wojciech Kozlowski, Rodney Van Meter, and Joseph Touch for their very useful reviews and comments to the document.

11. Informative References

- [BB84] Bennett, C. and G. Brassard, "Quantum Cryptography: Public Key Distribution and Coin Tossing", 1984, <<http://researcher.watson.ibm.com/researcher/files/us-bennetc/BB84highest.pdf>>.
- [Cacciapuoti01] Cacciapuoti, A. and et. al., "Quantum Internet: Networking Challenges in Distributed Quantum Computing", IEEE Network, (Early Access), 2019, <<https://ieeexplore.ieee.org/document/8910635>>.
- [Cacciapuoti02] Cacciapuoti, A. and et. al., "When Entanglement meets Classical Communications: Quantum Teleportation for the Quantum Internet", 2019, <<https://arxiv.org/abs/1907.06197>>.
- [Caleffi] Caleffi, M. and et. al., "Quantum internet: From Communication to Distributed Computing!", NANOCOM, ACM, 2018, <<https://arxiv.org/abs/1907.06197>>.
- [Castelvecchi] Castelvecchi, D., "The Quantum Internet has arrived (and it hasn't)", Nature 554, 289-292, 2018, <<https://www.nature.com/articles/d41586-018-01835-3>>.
- [Chitambar] Chitambar, E. and et. al., "Everything You Always Wanted to Know About LOCC (But Were Afraid to Ask)", Communications in Mathematical Physics, Springer, 2014, <<https://link.springer.com/article/10.1007/s00220-014-1953-9>>.
- [ETSI-QKD-Interfaces] ETSI GR QKD 003 V2.1.1, "Quantum Key Distribution (QKD); Components and Internal Interfaces", 2018, <https://www.etsi.org/deliver/etsi_gr/QKD/001_099/003/02.01.01_60/gr_QKD003v020101p.pdf>.
- [ETSI-QKD-UseCases] ETSI GR QKD 002 V1.1.1, "Quantum Key Distribution (QKD); Use Cases", 2010, <https://www.etsi.org/deliver/etsi_gs/qkd/001_099/002/01.01.01_60/gs_qkd002v010101p.pdf>.

- [Fan] Fan, L. and et. al., "Superconducting Cavity Electro-Optics: A Platform for Coherent Photon Conversion between Superconducting and Photonic Circuits", *Science Advances*, AAAS, 2018, <<https://advances.sciencemag.org/content/4/8/eaar4994.short>>.
- [Fitzi] Fitzi, M. and et. al., "A Quantum Solution to the Byzantine Agreement Problem", 2001, <<https://arxiv.org/pdf/quant-ph/0107127.pdf>>.
- [Fitzsimons] Fitzsimons, J., "Private Quantum Computation: An Introduction to Blind Quantum Computing and Related Protocols", 2017, <<https://www.nature.com/articles/s41534-017-0025-3.pdf>>.
- [Grumbling] Grumbling, E. and M. Horowitz, "Quantum Computing: Progress and Prospects", *National Academies of Sciences, Engineering, and Medicine*, The National Academies Press, 2019, <<https://doi.org/10.17226/25196>>.
- [I-D.dahlberg-11-quantum] Dahlberg, A., Skrzypczyk, M., and S. Wehner, "The Link Layer service in a Quantum Internet", draft-dahlberg-11-quantum-03 (work in progress), October 2019.
- [I-D.irtf-qirg-principles] Kozlowski, W., Wehner, S., Meter, R., Rijsman, B., Cacciapuoti, A., Caleffi, M., and S. Nagayama, "Architectural Principles for a Quantum Internet", draft-irtf-qirg-principles-05 (work in progress), September 2020.
- [I-D.van-meter-qirg-quantum-connection-setup] Meter, R. and T. Matsuo, "Connection Setup in a Quantum Network", draft-van-meter-qirg-quantum-connection-setup-01 (work in progress), September 2019.
- [Komar] Komar, P. and et. al., "A Quantum Network of Clocks", 2013, <<https://arxiv.org/pdf/1310.6045.pdf>>.

- [NISTIR8240] Alagic, G. and et. al., "Status Report on the First Round of the NIST Post-Quantum Cryptography Standardization Process", NISTIR 8240, 2019, <<https://nvlpubs.nist.gov/nistpubs/ir/2019/NIST.IR.8240.pdf>>.
- [Preskill] Preskill, J., "Quantum Computing in the NISQ Era and Beyond", 2018, <<https://arxiv.org/pdf/1801.00862>>.
- [Proctor] Proctor, T. and et. al., "Multiparameter Estimation in Networked Quantum Sensors", Physical Review Letters, American Physical Society, 2018, <<https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.120.080501>>.
- [QinHao] Qin, H., "Towards Large-Scale Quantum Key Distribution Network and Its Applications", 2019, <https://www.itu.int/en/ITU-T/Workshops-and-Seminars/2019060507/Documents/Hao_Qin_Presentation.pdf>.
- [Renner] Renner, R., "Security of Quantum Key Distribution", 2006, <<https://arxiv.org/pdf/quant-ph/0512258.pdf>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", BCP 188, RFC 7258, DOI 10.17487/RFC7258, May 2014, <<https://www.rfc-editor.org/info/rfc7258>>.
- [Treiber] Treiber, A. and et. al., "A Fully Automated Entanglement-based Quantum Cryptography System for Telecom Fiber Networks", New Journal of Physics, 11, 045013, 2009, <<https://doi.org/10.1364/OE.26.024260>>.
- [Wehner] Wehner, S., Elkouss, D., and R. Hanson, "Quantum internet: A vision for the road ahead", Science 362, 2018, <<http://science.sciencemag.org/content/362/6412/eaam9288.full>>.
- [ZhangPeiyu] Zhang, P. and et. al., "Integrated Relay Server for Measurement-Device-Independent Quantum Key Distribution", 2019, <<https://arxiv.org/abs/1912.09642>>.

[ZhangQiang]

Zhang, Q., Hu, F., Chen, Y., Peng, C., and J. Pan, "Large Scale Quantum Key Distribution: Challenges and Solutions", *Optical Express*, OSA, 2018, <<https://doi.org/10.1364/OE.26.024260>>.

[Zhuang]

Zhuang, Q. and et. al., "Physical-Layer Supervised Learning Assisted by an Entangled Sensor Network", *Physical Review Letters*, American Physical Society, 2019, <<https://journals.aps.org/prx/abstract/10.1103/PhysRevX.9.041023>>.

Authors' Addresses

Chonggang Wang
InterDigital Communications, LLC
1001 E Hector St
Conshohocken 19428
USA

Email: Chonggang.Wang@InterDigital.com

Akbar Rahman
InterDigital Communications, LLC
1000 Sherbrooke Street West
Montreal H3A 3G4
Canada

Email: rahmansakbar@yahoo.com

Ruidong Li
NICT
4-2-1 Nukui-Kitamachi
Koganei, Tokyo 184-8795
Japan

Email: lrd@nict.go.jp

Melchior Aelmans
Juniper Networks
Boeing Avenue 240
Schiphol-Rijk 1119 PZ
The Netherlands

Email: maelmans@juniper.net