

Quantum Internet Research Group
Internet-Draft
Intended status: Informational
Expires: August 23, 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.
February 19, 2021

Architectural Principles for a Quantum Internet
draft-irtf-qirg-principles-06

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 August 23, 2021.

Copyright Notice

Copyright (c) 2021 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.	The control and 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	24
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.	A thought experiment inspired by classical networks	29
8.	Security Considerations	31
9.	IANA Considerations	31
10.	Acknowledgements	31
11.	Informative References	32
	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 [\[37\]](#). Depending on the stage of a quantum network [\[8\]](#) such devices may range from 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 [\[24\]](#).

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 [\[28\]](#), quantum-enhanced measurement networks [\[6\]](#), or higher-precision, long-baseline telescopes [\[34\]](#). 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 [\[26\]](#) [\[25\]](#).

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 [\[27\]](#), 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, called 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 may not be deterministic, but can be determined from the readout of the measurement. The measurement result is a classical bit, 0 or 1, corresponding to $|0\rangle$ and $|1\rangle$ 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 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:

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 [28], 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 [29]. 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. [35]

4. Achieving quantum connectivity

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

4.1.1. 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.

4.1.1.1. 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.

4.1.1.2. No-cloning theorem

Since directly reading the state of a qubit is not possible, one could ask if we can simply copy a qubit without looking at it. Unfortunately, this is fundamentally not possible in quantum mechanics [[30](#)] [[31](#)].

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.

4.1.1.3. 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 [33] [32] [38] [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 [19].

4.2. Bell pairs

Bell pair states are the entangled two-qubit states:

$$|00\rangle + |11\rangle, |00\rangle - |11\rangle, |01\rangle + |10\rangle, |01\rangle - |10\rangle,$$

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 node 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 [18]. Quantum state teleportation consumes an unknown qubit 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 qubit state that we wanted to transmit. This requirement to communicate the measurement read out over a classical channel unfortunately means that entanglement cannot be used to transmit information faster than the speed of light.

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 [20]:

- 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.

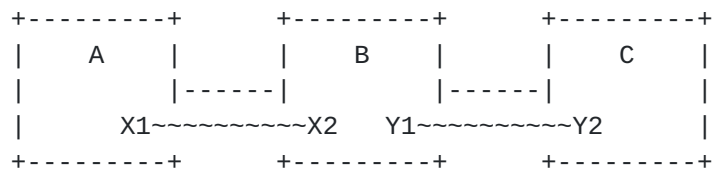
The "mid-point source" scheme is more robust to photon loss, but in the other schemes the nodes retain greater control over the entangled pair generation.

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

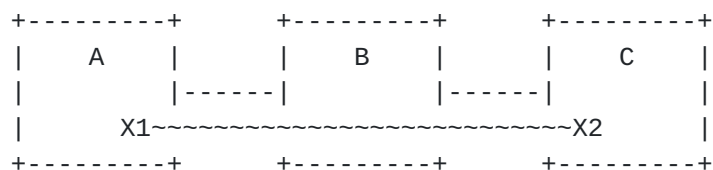
The problem with generating entangled pairs directly across a link is that efficiency decreases with channel length. Beyond a few 10s of kms in optical fibre or 1000 kms in free space (via satellite) the rate is effectively zero and due to the no-cloning theorem we cannot simply amplify the signal. The solution is entanglement swapping [\[19\]](#).

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 [19] 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) [36].

To distil 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](#) [33] [32] [38] [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.

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 possible only when the photon loss rate from transmission to measurement is less than 50%.

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 even with a limited set of available quantum gates.
2. Second generation quantum networks improve upon the first generation with QEC codes for error tolerance (but not loss 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 does a lookup in 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.](#) The control and data planes

Control plane protocols for quantum networks will have many responsibilities similar to their classical counterparts, namely drawing the network topology, resource management, populating data plane tables, etc. Most of these protocols do not require the manipulation of quantum data and can operate simply by exchanging classical messages only. There may also be some control plane functionality that does require the handling of quantum data, e.g. a quantum ping [2]. As it is not clear if there is much benefit in defining a separate quantum control plane given the significant overlap in responsibilities with its classical counterpart, the question of whether there should be a separate quantum control plane is beyond the scope of this document.

However, the data plane separation is much more distinct and 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 the network control plane. 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.

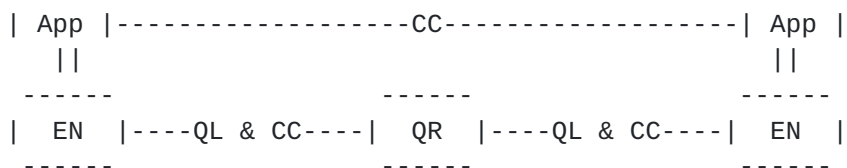
Additionally, we need to identify two kinds of links that will be used in a quantum network:

- 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 [[39](#)] 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 [[27](#)] 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 affect the content

and semantics of messages that must cross those boundaries -- both for connection setup and real-time operation [42]. 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 [40] although a lifetime of a minute has also been demonstrated for qubits not connected to a quantum network [41] (as of 2020). 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 [27] [14]. 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 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 quantum networking 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 [23]. 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, entanglement-based 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. The architecture 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) [\[21\]](#) 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 an expensive resource

Time 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 some of 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 a Bell pair to decohere or result in an end-to-end pair below an application's fidelity threshold. Therefore, managing the idle time of qubits holding live quantum states should be done carefully. Ideally by minimising the idle time, but potentially also by moving the quantum state for temporary storage to a quantum memory with a longer lifetime.

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, the architecture 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. A thought experiment inspired by classical networks

To conclude, we discuss a plausible quantum network architecture inspired by MPLS. This is not an architecture proposal, but rather 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.

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 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 [\[22\]](#).

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-04](#) (work in progress), January 2021.
- [3] Bennett, C. and G. Brassard, "Quantum cryptography: Public key distribution and coin tossing", *Theoretical Computer Science* 560, 7-11, 2014, <<http://www.sciencedirect.com/science/article/pii/S0304397514000011>>.
- [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] 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>>.
- [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., Knegjens, 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] Bennett, C., Brassard, G., Crepeau, C., Jozsa, R., Peres, A., and W. Wootters, "Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels", Phys. Rev. Lett. Vol. 70, Iss. 13, 1996, <<https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.70.1895>>.
- [19] 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>>.

- [20] 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>>.
- [21] 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>>.
- [22] 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>>.
- [23] Clark, D., "The design philosophy of the DARPA internet protocols", SIGCOMM '88, 1988, <<https://dl.acm.org/doi/abs/10.1145/52324.52336>>.
- [24] Van Meter, R., "Quantum Networking", ISTE Ltd/John Wiley and Sons Inc 978-1-84821-537-5, 2014.
- [25] 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>>.
- [26] 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>>.

- [27] 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>>.
- [28] Fitzsimons, J. and E. Kashefi, "Unconditionally verifiable blind quantum computation", Phys. Rev. A Vol. 96, Iss. 1, 2017, <<https://arxiv.org/abs/1203.5217>>.
- [29] Terhal, B., "Is entanglement monogamous?", IBM Journal of Research and Development Vol. 48, Iss. 1, 2004, <<https://ieeexplore.ieee.org/document/5388928>>.
- [30] 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>>.
- [31] Wootters, W. and W. Zurek, "A single quantum cannot be cloned", Nature 299, 802-803, 1982, <<https://www.nature.com/articles/299802a0>>.
- [32] 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>>.
- [33] Jiang, L., Taylor, J., Nemoto, K., Munro, W., Van Meter, R., and M. Lukin, "Quantum repeater with encoding", Phys. Rev. A Vol. 79, Iss. 3, 2009, <<https://arxiv.org/abs/0809.3629>>.
- [34] 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>>.
- [35] "The Quantum Protocol Zoo", <<https://wiki.veriqloud.fr/>>.
- [36] 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>>.

- [37] Kimble, H., "The Quantum Internet", Nature 453, 1023-1030, 2008, <<http://arxiv.org/abs/0806.4195>>.
- [38] 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>>.
- [39] 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>>.
- [40] 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>>.
- [41] 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>>.
- [42] Nagayama, S., Choi, B., Devitt, S., Suzuki, S., and R. Van Meter, "Interoperability in encoded quantum repeater networks", Phys. Rev. A Vol. 93, Iss. 4, 2016, <<https://arxiv.org/abs/1508.04599>>.

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