

QIRG
Internet-Draft
Intended status: Standards Track
Expires: April 24, 2019

K. Kompella
M. Aelmans
Juniper Networks, Inc.
S. Wehner
QuTech
C. Sirbu
Redbit Networks
October 21, 2018

Advertising Entanglement Capabilities in Quantum Networks
draft-kaws-qirg-advent-00

Abstract

This document describes the use of link-state routing protocols on classical links in Quantum Networks. It contains proposals for additions to the IS-IS and OSPF protocols in order for them to transport relevant information for a Quantum Network, specifically, for the creation and manipulation of entangled pairs. The document will describe some of the necessary attributes and some suggestions of how this information may be used.

No Schrodinger's cats were harmed in the creation of this document.

Requirements Language

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 [2].

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 24, 2019.

Copyright Notice

Copyright (c) 2018 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
1.1. Definitions and Notation	3
2. Motivation	6
3. Theory of Operation	7
3.1. Multihop Entanglement	7
3.2. Distillation	9
3.3. Node Properties	9
3.4. Link Properties	10
4. The (Ab)use of Protocols	10
4.1. A Brief Primer on Link-state Protocols	10
4.2. Node Properties	12
4.3. Link Properties	12
5. Security Considerations	13
6. Acknowledgments	13
7. IANA Considerations	13
8. References	13
8.1. Normative References	13
8.2. Informative References	14
Authors' Addresses	15

1. Introduction

Quantum networking is an emerging field using the strange (even counterintuitive) properties of quantum mechanics to bring new, useful capabilities to networking. One of these is "entanglement" [8], where the state of a group of particles must be described as a unit -- it cannot be decomposed to the state of each particle independently. Entangled pairs (often called EPR pairs, abbreviated here as EP) of particles can be used for quantum teleportation [10] and for quantum key distribution (QKD) [14].

A Quantum Network consists of quantum nodes and links. Here, we will be concerned with controllable quantum nodes (CQN) that allow control decisions. We posit a classical network parallel to the quantum network, with classical nodes (CN) and links. A classical node is colocated with a quantum node; a classical link may be a fiber or wavelength parallel to the corresponding quantum link. Such a classical link is required by most quantum methods to create EPs deterministically or in a heralded fashion, where the creation of EPs is deterministic conditioned on a specific signal. To make useful decisions, it is desirable to augment this data to describe the capabilities and states of quantum nodes and links.

This document proposes to carry entanglement capability data as Type Length Values (TLVs) over IS-IS or OSPF link-state advertisements over the corresponding classical network. A subset of the CQNs may run quantum applications such as QKD; these nodes may want to initiate multihop EPs.

Once an EP is created, the state of one particle ("quantum bit" or qubit) of an EP can be transferred to another qubit within the same QN by a process known as swapping or a SWAP gate ([12]). Also, several pairs of imperfectly entangled qubits can be "distilled" ([13]) to fewer but "better entangled" qubits.

Long distance entanglement can be produced from piecewise short distance entanglement: Given an EP between CQN A and CQN B, and another EP between CQN B and CQN C, one can create an EP between CQN A and CQN C by a process known as an "entanglement swap". These operations can be used to manipulate EPs to improve their lifetimes or their quality, or to create multihop EPs. Physically, qubits can be realized in many ways. For example, they can be represented by the energy levels of Nitrogen Vacancy (NV) Centers in diamond ([16], [17]). Logically, a qubit can be classified as a "communication qubit", a "traveling qubit" or a "storage qubit".

This document primarily discusses the exchange of quantum capabilities over a classical network. Some illustrative examples of how these capabilities can be used in a quantum network may be given, but this document should not be considered authoritative on these procedures.

1.1. Definitions and Notation

The following terms are used in this document:

Quantum link: A quantum link is a connection transporting traveling qubits, typically photons. This could be a physical link. This document does not describe the usage of this link.

Classical link: A classical link is a connection transporting packets. This could be a physical link. The proposed extensions in this document use these links to exchange capabilities.

Controllable Quantum Node (CQN): A controllable quantum node is a quantum device consisting of at least one qubit, capable of performing (a subset of) the following operations described in detail below: storing qubits for some amount of time, performing quantum operations such as entanglement distillation and entanglement swapping, and producing entanglement between the nodes and traveling qubits. The latter are generally realized using photons over fibers or through free space.

The term controllable refers to the fact that external control in software is capable of selecting the desired operations and qubits to use. Such nodes can be quantum repeaters that allow choices of operations to be made, as well as quantum end nodes capable of executing complex application protocols [14]. Quantum repeaters that merely allow timing control, such as automatic entanglement swapping whenever qubits arrive in a specific timing interval, will not be referred to as CQN. Such automated repeaters can be seen as lying at the quantum physical layer and do not enter routing or other decision making, apart from being switched on or off, and hence are not relevant to advertisement protocols like the ones considered here.

Quantum end node (QEN): In this document, a quantum end node [14] is one of a pair of quantum nodes forming an entanglement via a sequence of zero or more CQNs. Quantum end nodes typically run a higher-layer quantum application such as QKD.

Communication qubit: A qubit is called a communication qubit if it is possible to produce entanglement between this qubit and a traveling photon. This can be done by emission from the quantum node, that is, entanglement is produced between the qubit and the photon which is emitted from the quantum node. This process has been demonstrated in a number of physical systems that can be used as quantum nodes such as NV in diamond ([16], [17]), Ion Traps ([18]) and Neutral Atoms ([19]). An example of a communication qubit is the electron spin of the NV in diamond system ([15]). Entanglement between a communication qubit and traveling photons can also be produced by absorption. Examples include atomic ensemble memories ([20]).

A communication qubit c at CQN A is denoted by $c@A$, or simply c (if the node A is understood).

Storage qubit: A qubit is called a storage qubit if the node has the capability to use this qubit as a (temporary) quantum memory, but the qubit cannot serve as a communication qubit. To make storage qubits useful a node is required to possess the ability to transfer the state of a communication qubit to a storage qubit. An example of a storage qubit is the nuclear spin in the NV in diamond system [16].

A storage qubit s at node B is denoted $s@B$.

Entangled Pair (EP): An entangled pair is a special state of two qubits, known as an EPR pair [8]. An entangled pair of qubits $c@A$ and $c@B$ is denoted $[[c@A, c@B]]$.

The process of entangling two particles $c@A$ and $c@B$ is denoted as follows:

$\text{ent}(c@A, c@B) \rightarrow [[c@A, c@B]]$

$\text{ent}(c@A, c@B)$ may take time T and succeed with probability P , and yield an entangled pair $[[c@A, c@B]]$ of fidelity F .

Fidelity: A measure of the quality of the entanglement of an EP (xref target='QFid'/>). Fidelity lies in the interval $[0, 1]$ where a higher value indicates better quality; usable fidelity values lie in the half-open interval $(0.5, 1]$.

Swap: Two qubits located in the same CQN can interchange states ([13]). For example, the states of a communication qubit and a storage qubit at A can be swapped as follows:

$\text{swap}(c@A, s@A)$

If $c@A$ was entangled with $c@B$, the result is that $s@A$ is now entangled with $c@B$.

Distillation: Distillation is the process of turning a large number of weakly entangled states into a smaller number of highly entangled states ([13]).

For example, EPs $[[c1@A, c1@B]]$ and $[[c2@A, c2@B]]$ of fidelities $F1$ and $F2$ respectively may be distilled as follows:

$\text{dist}([[c1@A, c1@B]], [[c2@A, c2@B]]) \rightarrow [[c3@A, c3@B]]$

If distillation is successful, the fidelity $F3$ of $[[c3@A, c3@B]]$ will be higher than $F1$ and $F2$.

Entanglement Swap: Given two EPs $[[c@A, c1@B]]$ and $[[c2@B, c@C]]$, one can perform an entanglement swap:

```
entSwap([[c@A, c1@B]], [[c2@B, c@C]]) -> [[c@A, c@C]]
```

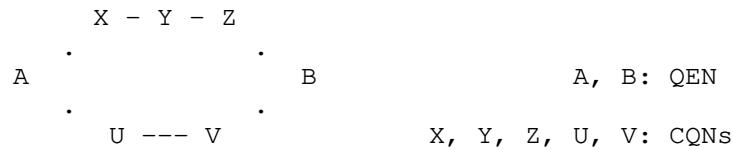
to create a new EP between $q@A$ and $q@C$. This is how "multihop" EPs are created from a sequence of "single-hop" EPs.

The swap operation can also be used within a CQN. A possible use case is when there aren't enough communication qubits to create the needed EPs. If, in the above example, B doesn't have two communication qubits $c1$ and $c2$, the following can be done:

```
ent(c@A, c@B) -> [[c@A, c@B]] # entangle
swap(c@B, s@B) -> [[c@A, s@B]] # swap EP to storage qubit
ent(c@B, c@C) -> [[c@B, c@C]] # use freed up qubit c@B
swap(s@B, c@B) -> [[c@A, c@C]] # create multihop EP
```

2. Motivation

Consider the following (very simple) quantum network consisting of QENs A and B, and CQNs X, Y, Z, U, V. The goal is to create an EP between qubits at A and at B, perhaps for the high-level task of QKD between A and B.



From A's point of view, here are a number of questions:

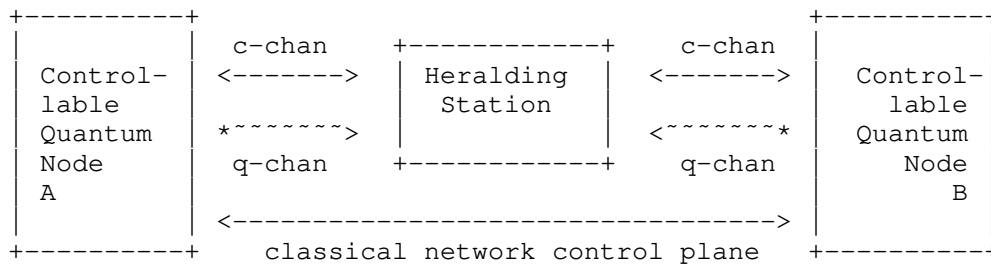
1. Is B reachable from A via quantum links that allow EP creation?
2. If so, along what sequence(s) of quantum nodes?
3. Can each pair of adjacent CQNs in this sequence form EPs? If so, how long will it take, and what fidelity can be expected?
4. If each pair of adjacent CQNs successfully forms EPs of sufficient fidelity, can these be swapped to form a multihop EP between A and B?
5. If a multihop EP between A and B were to be formed, would it be of good enough fidelity, or should a second multihop EP be formed and the two EPs distilled into one high fidelity EP? How many times should this process be repeated?

- 6. If the overall answer is Yes, should A proceed via sequence A, X, Y, Z, B, or sequence A, U, V, B?

This document aims to provide all CQNs in a quantum network with the information they need to answer such questions, and to create EPs at their desired fidelity and speed.

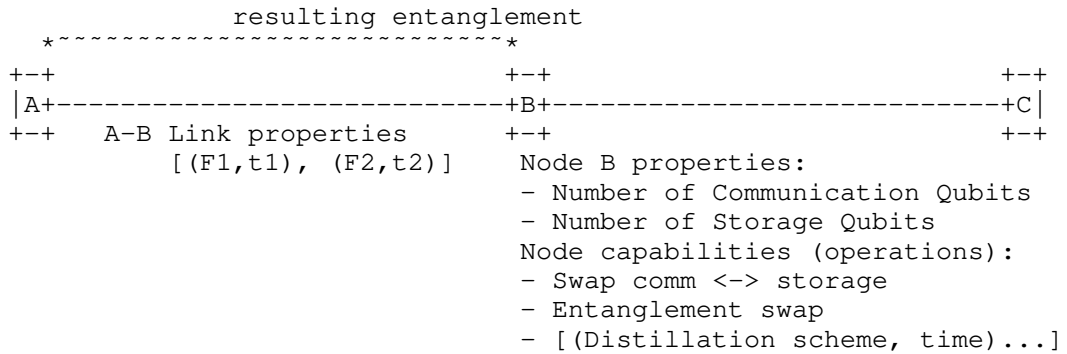
3. Theory of Operation

A CQN contains one or more communication qubits and one or more storage qubits. Many proposals exist for producing EPs between remote quantum nodes (see for example [16], [17], [18], [20]). Abstractly, these result in the generation of EPs with fidelity F after an expected time t . To give an example, we describe the generation of EPs that has been implemented in NV in diamond ([16]), and Ion Traps ([18]). The largest distance for producing long-lived entanglement is presently 1.3kms ([17]). To entangle a pair of communication qubits, the QNs send carefully timed photons towards the HS. If the process is successful, HS sends an OK message to both QNs.



The classical network control plane is of particular interest here as it would be used by the proposed protocol to advertise and exchange information about the capabilities of the CQNs to generate entanglement. This classical channel exists between all CQNs and is shared with other application specific control and data plane traffic.

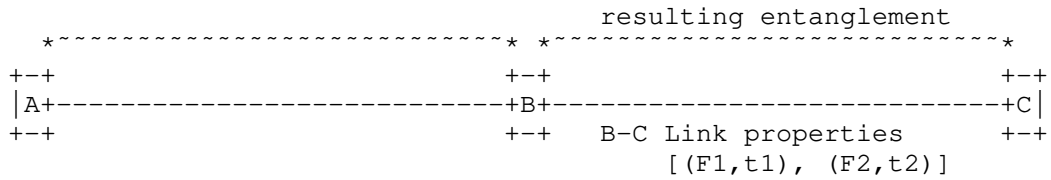
3.1. Multihop Entanglement



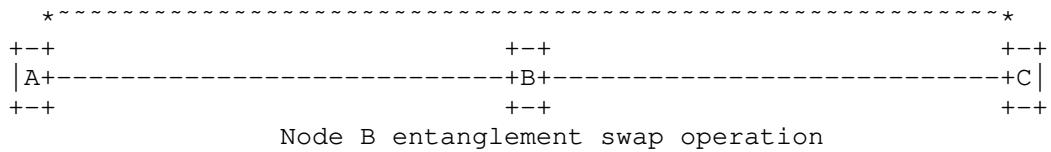
In the figure above, an example request for an entangled pair between nodes A and B will be affected by the following properties:

- o A chosen combination of F(idelity) and t(ime) duration to produce an entanglement at the respective Fidelity. These parameters roughly equate to the quality of the link, the accuracy with which the nodes can use the link, and the delay in classical networking.
- o The actual capability of nodes A and B to make use of the communication qubits.

A new EP creation between CQNs B and C will similarly be affected by the same parameters as above.



And finally, with an entanglement swap operation at node B (which is a node specific capability and has a specific duration) we end up with an A-C EP:

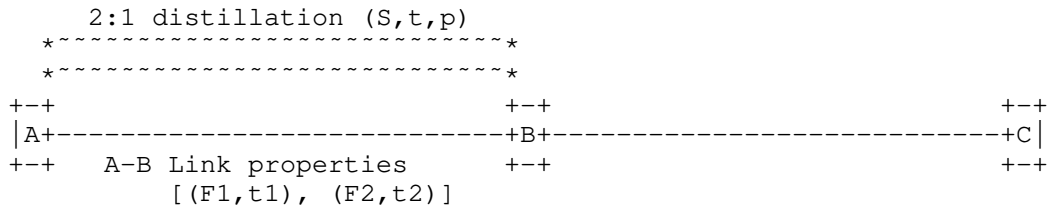


3.2. Distillation

If a pair of CQNs A and B share a number of EPs of insufficient quality, they may be combined into a single EP of higher quality by distillation. To do so, these CQNs need to agree on which distillation scheme to use before distillation can proceed. This does not necessarily need to be via communication between A and B, if one agrees upon a deterministic procedure of selecting one. This document suggests the following procedure:

1. A and B look at the distillation schemes that both advertise in common.
2. If there is none in common, stop. Distillation is not possible.
3. If there is a non-trivial subset in common, the first scheme in the node with the lower router ID is to be used by A and B.

Given a chosen distillation scheme (S,t,p) , an additional time delay will be added for the actual operation: For a 2:1 distillation scheme between nodes A and B, 2 EPs need to be produced followed by an operation on A and B that produces 1 EP. This operation will take time some expected time t , and succeed with probability p .



3.3. Node Properties

We are interested in exposing the properties of CQNs (including QENs) to allow sophisticated decision making, for example in the creation of entanglement. These properties include:

1. Number of communication qubits. The number of communication qubits determines the number of entangled pairs that the node can produce simultaneously.
2. Number of storage qubits
3. Possible operations, along with their execution time and probability of success:
 1. Swap between communication and storage qubits

2. Entanglement swap
3. List of supported distillation schemes (in order of preference).

Note that several other parameters can be advertised, such as the T1 and T2 times for a qubit's decoherence. These are omitted for now, instead just giving the decay of the fidelity of an EP. If deemed useful, T1 and T2 times can additionally be advertised.

3.4. Link Properties

A list of (F_n, t_n) pairs describing the tradeoffs of a possible entanglement produced by two nodes (the ends of said link): t_n is the time to produce an entangled pair with fidelity F_n .

4. The (Ab)use of Protocols

The routing protocols IS-IS and/or OSPF could be used in order to advertise entanglement capabilities. This section describes the additional data fields needed in order to facilitate the objective.

4.1. A Brief Primer on Link-state Protocols

This document suggests the use of a link-state protocol to distribute the capabilities of CQNs to create entanglement. This section offers a short introduction to link-state protocols for those not familiar with them.

Consider a directed graph $G=(V, E)$ with vertices (nodes) V and edges (links) E . Consider also $G'=(V', E')$; there is a 1-1 mapping from V' to V and from E' to E such that $e_{l'} = (v_{1'}, v_{2'})$ is in E' iff $e_l = (v_1, v_2)$ is in E and $v_{1'}$ maps to v_1 and $v_{2'}$ maps to v_2 . G' represents the quantum network; V' represents the set of CQNs, and E' the set of quantum links between pairs of CQNs; G represents a classical network parallel to G' ; that is, each CQN v' has a corresponding classical node v . v plays a dual role: it is the control node for v' , and proxies on behalf of v' in the link-state protocol.

The basic objective of a link-state protocol is to "flood" properties of nodes and (directed) links to all nodes in the network. This is accomplished by means of "link-state advertisements" (LSAs) that each node originates and sends to its immediate neighbors. The neighbors in turn send received LSAs to their own neighbors; this process repeats until every node receives every LSA (hence the term "flooding"). The focus of LSAs is the link properties (hence `_link-state_ advertisements`), although node properties are also advertised.

There are mechanisms to prevent looping of LSAs, and for reliable flooding. There is also a sequence number by which a more recent update of an LSA can be identified as such, and a mechanism for "aging out" LSAs belonging to nodes no longer in the network. In what follows, quantum node and link properties are added to the link-state advertisements of the corresponding classical node. Note that link properties need not be symmetric; that is, the link properties of (v, w) need not be the same as those of (w, v) .

The net result of flooding is that every node has the same picture of the network (modulo LSAs in flight); in particular, each node knows the overall topology and connectivity of the network, and can use this information to make decisions. In a classical network, such a decision could be to compute a shortest path; for the quantum network, it could be choosing a feasible path (i.e., sequence of CQNs) for a multihop entanglement. Note that a node doesn't really know when it has complete and up-to-date information about the network; LSA updates may be originated at any time. Usually, this is okay: for example, if a node v learns enough of the network to have a path to another node w , it can compute a multihop entanglement to w . Subsequent updates may provide a more optimal (or higher probability) entanglement path. There are heuristics one can apply to guess that the link-state database (LSDB) (i.e., the union of all LSAs) is complete-ish; however, as nodes (and links) can fail or disconnect, there really is no such thing as "the full LSDB".

Each node v is identified by a "router ID" (an IP address uniquely allocated to v), denoted by $\text{rid}(v)$. A link $L = (v, w)$ is identified by $(\text{rid}(v), i)$ where i is an index allocated by v for L unique for each link emanating from v . (L may also be identified by IP addresses, but we'll ignore that for now.) It is generally expected that a directed link (v, w) is matched by a link (w, v) ; if not, (v, w) is ignored from subsequent consideration; in particular, no link properties are advertised for this link by v . Note that a pair of nodes may have multiple links between them; for simplicity, the notation will not be extended to indicate this. We'll assume $\text{rid}(v') = \text{rid}(v)$ and the index allocated to a quantum link e' is the same as that of the corresponding classical link e .

Let v, w be a pair of neighboring nodes, and let $L1 = (v, w)$ and $L2 = (w, v)$ in E be directed links in opposite directions between v and w with identifiers $(\text{rid}(v), i1)$ and $(\text{rid}(w), i2)$ respectively (where $i1$ is the index allocated for $L1$ by v , and similarly for $i2$). As a first step in running a link-state protocol, v runs a "hello protocol" all its links; in particular, over $L1$. Similarly, w will run the hello protocol over $L2$. The hello protocol serves to exchange the indices $i1$ and $i2$, and thus identify $(\text{rid}(v), i1)$ as the reverse link of $(\text{rid}(w), i2)$. This allows both v and w to correlate

the link properties of L1 and L2. If the hello protocol fails between *v* and *w*, neither node includes link properties for the link in their LSAs.

Once the hello protocol has been run on all links, *v* starts the process of generating and sending its own LSA over all its links, and of receiving the current LSDB from its neighbors. Note that an LSA originated by *v* must propagate unchanged across the network; only *v* is allowed to change it (and such a change must be accompanied by updating the LSA's sequence number). Such an update is triggered by a new link coming up, an existing link going down, or a node or link property changing.

IS-IS and OSPF are in principle similar, although the details of the protocol mechanisms and encodings vary. In both protocols, a Type-Length-Value (TLV) is used to encode most node and link properties. In IS-IS, TLVs are used for all properties, and a single type of LSA is used; in OSPF, there are several types of LSAs, and many (but not all) properties are encoded as TLVs.

[1] has examples of "standard" LSAs for routing; [4] has the so-called Traffic Engineering LSAs.

4.2. Node Properties

Here, we give a protocol-independent description of quantum node properties; later documents will specify the encoding specifically for IS-IS and OSPF.

Note that the following list of node properties is a strawman; all details are subject to change, and other properties may be added as needed.

The following node properties are added to the appropriate LSA:

```
<Qubit-TLV><NCQ><NSQ>
<CS-Swap><Prob><ExecTime>
<Ent-Swap><Prob><ExecTime>
<Measure><Prob><ExecTime>
<NDistSch><DistScheme1><DistScheme2>
```

4.3. Link Properties

Only one link property is listed. It gives the time-fidelity tradeoffs of an entanglement operation as a list:

```
<N-Ent-TO><time1><fid1><time2><fid2>...
```

This is interpreted as follows: an entanglement operation may be initiated between nodes v and w over link (v, w). Depending on how fast one wants to complete (time-i), the list gives the corresponding fidelity of the resulting entanglement (fid-i). time-i is given in nanoseconds; fid-i as a number between 0 and 999999. The denominator is 1000000.

Note that this link property is symmetric, as entanglement is initiated simultaneously at v and w.

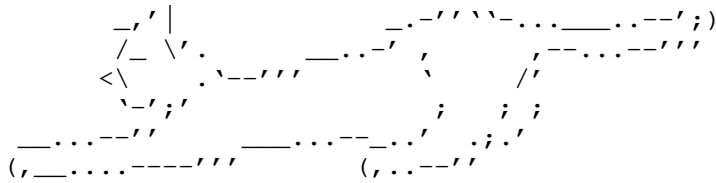
5. Security Considerations

It is not anticipated that adding these extensions to IS-IS and OSPF will present new security hazards to those protocols. Since however a common application of entangled pairs is for security purposes (such as QKD), it is worth investigating whether this application places a higher burden of security on the underlying protocols.

6. Acknowledgments

The authors would like to thank the following people for their contributions and support to this document: Vesna Manojlovic (RIPE NCC) and Axel Dahlberg (QuTech). Kompella would also like to thank Bruno Rijsman for encouraging him to learn about Quantum Computing and Networking.

Also:



7. IANA Considerations

There are no requests as yet to IANA for this document.

8. References

8.1. Normative References

[1] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", RFC 1195, DOI 10.17487/RFC1195, December 1990, <https://www.rfc-editor.org/info/rfc1195>.

- [2] 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>>.
- [3] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, DOI 10.17487/RFC3630, September 2003, <<https://www.rfc-editor.org/info/rfc3630>>.
- [4] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <<https://www.rfc-editor.org/info/rfc5305>>.
- [5] Ishiguro, K., Manral, V., Davey, A., and A. Lindem, Ed., "Traffic Engineering Extensions to OSPF Version 3", RFC 5329, DOI 10.17487/RFC5329, September 2008, <<https://www.rfc-editor.org/info/rfc5329>>.
- [6] Aggarwal, R. and K. Kompella, "Advertising a Router's Local Addresses in OSPF Traffic Engineering (TE) Extensions", RFC 5786, DOI 10.17487/RFC5786, March 2010, <<https://www.rfc-editor.org/info/rfc5786>>.

8.2. Informative References

- [7] "Qubit", <<https://en.wikipedia.org/wiki/Qubit>>.
- [8] "Quantum Entanglement", <https://en.wikipedia.org/wiki/Quantum_entanglement>.
- [9] "Quantum Network", <https://en.wikipedia.org/wiki/Quantum_network>.
- [10] "Quantum Teleportation", <https://en.wikipedia.org/wiki/Quantum_teleportation>.
- [11] "Quantum Fidelity", <https://en.wikipedia.org/wiki/Fidelity_of_quantum_states>.
- [12] Wehner, S., Elkouss, D., and R. Hanson, "Quantum internet: A vision for the road ahead", Science 362, 6412, 2018.
- [13] Rozpedek, F., Schiet, T., Thinh, L., Elkouss, D., Doherty, A., and S. Wehner, "Optimizing practical entanglement distillation", Phys. Rev. A 97, 062333, 2018, <<https://arxiv.org/abs/1803.10111>>.

- [14] "Quantum Key Distribution",
<https://en.wikipedia.org/wiki/Quantum_key_distribution>.
- [15] "Nitrogen-vacancy center",
<https://en.wikipedia.org/wiki/Nitrogen-vacancy_center>.
- [16] Humphreys, P., "Deterministic delivery of remote entanglement on a quantum network", Nature 558, 2018.
- [17] Hensen, B. and others, "Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres", Nature 526, 2015.
- [18] Hucul, D. and others, "Modular entanglement of atomic qubits using photons and phonons", Nature Physics 11(1), 2015.
- [19] Noelleke, C. and others, "Efficient Teleportation Between Remote Single-Atom Quantum Memories", Physical Review Letters 110, 140403, 2013.
- [20] Sangouard, N. and others, "Quantum repeaters based on atomic ensembles and linear optics", Reviews of Modern Physics 83, 33, 2011.

Authors' Addresses

Kireeti Kompella
Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, CA 94089
USA

Email: kireeti.kompella@gmail.com

Melchior Aelmans
Juniper Networks, Inc.
Boeing Avenue 240
Schipol-Rijk, PZ 1119
The Netherlands

Email: melchior@aelmans.eu

Stephanie Wehner
QuTech
Van der Waalsweg 122
Delft, LC 2611
The Netherlands

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

Cristian Sirbu
Redbit Networks
Dublin
Republic of Ireland

Email: cristian@redbit.network

QIRG
Internet-Draft
Intended status: Standards Track
Expires: September 28, 2019

K. Kompella
M. Aelmans
Juniper Networks, Inc.
S. Wehner
QuTech
C. Sirbu
Redbit Networks
A. Dahlberg
QuTech
March 27, 2019

Advertising Entanglement Capabilities in Quantum Networks
draft-kaws-qirg-advent-03

Abstract

This document describes the use of link-state routing protocols on classical links in Quantum Networks. It contains proposals for additions to the IS-IS and OSPF protocols in order for them to transport relevant information for a Quantum Network, specifically, for the creation and manipulation of entangled pairs. The document will describe some of the necessary attributes and some suggestions of how this information may be used.

No Schrodinger's cats were harmed in the creation of this document.

Requirements Language

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 [2].

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 September 28, 2019.

Copyright Notice

Copyright (c) 2019 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
1.1. Definitions and Notation	4
2. Motivation	6
3. Theory of Operation	7
3.1. Multihop Entanglement	8
3.2. Distillation	9
3.3. Node Properties	9
3.4. Link Properties	10
4. The (Ab)use of Protocols	10
4.1. A Brief Primer on Link-state Protocols	10
4.2. Node Properties	12
4.3. Link Properties	12
5. Security Considerations	13
6. Acknowledgments	13
7. Editorial Comments	13
8. IANA Considerations	13
9. References	14
9.1. Normative References	14
9.2. Informative References	14
Authors' Addresses	15

1. Introduction

Quantum networking is an emerging field using the strange (even counterintuitive) properties of quantum mechanics to bring new, useful capabilities to computing and networking. One of these is "entanglement" [8], where the state of a group of particles must be described as a unit -- it cannot be decomposed to the state of each particle independently. Entangled pairs (often called EPR pairs,

abbreviated here as EP) of particles can be used for quantum teleportation [10] and for quantum key distribution (QKD) [14].

A Quantum Network consists of quantum nodes and links. Here, we will be concerned with controllable quantum nodes (CQN) that allow control decisions. We posit a classical network parallel to the quantum network, with classical nodes (CN) and links. A classical node is colocated with a quantum node; a classical link may be a fiber or wavelength parallel to the corresponding quantum link. The existence of such a classical link is required by most quantum methods to create EPs deterministically or in a heralded fashion, where the creation of EPs is conditioned on a specific signal. To make useful decisions, it is desirable to augment this data to describe the capabilities and states of quantum nodes and links. At current time there is a need for classical links besides quantum links. In the future this might change into a situation where classical links will perhaps become obsolete.

This document proposes to carry entanglement capability data as Type Length Values (TLVs) over IS-IS or OSPF link-state advertisements over the corresponding classical network. A subset of the CQNs may run quantum applications such as QKD; these nodes may want to initiate multihop EPs.

Once an EP is created, the state of one particle ("quantum bit" or qubit) of an EP can be transferred to another qubit within the same QN by a process known as swapping or a SWAP gate ([12]). Also, several pairs of imperfectly entangled qubits can be "distilled" ([13]) to fewer but "better entangled" qubits.

Long distance entanglement can be produced from piecewise short distance entanglement: Given an EP between CQN A and CQN B, and another EP between CQN B and CQN C, one can create an EP between CQN A and CQN C by a process known as an "entanglement swap". These operations can be used to manipulate EPs to improve their lifetimes or their quality, or to create multihop EPs. Physically, qubits can be realized in many ways. For example, they can be represented by the energy levels of Nitrogen Vacancy (NV) Centers in diamond ([16], [17]). Logically, a qubit can be classified as a "communication qubit", a "traveling qubit" or a "storage qubit".

This document primarily discusses the exchange of quantum capabilities over a classical network. Some illustrative examples of how these capabilities can be used in a quantum network may be given, but this document should not be considered authoritative on these procedures.

1.1. Definitions and Notation

The following terms are used in this document:

Quantum link: A quantum link is a connection transporting traveling qubits, typically photons. This could be a physical link, or by means of teleportation over pre-established entanglement amongst distant network nodes. Such pre-shared entanglement effectively forms a shortcut - a virtual quantum link - which can be used exactly once. This document does not describe the usage of these links itself.

Classical link: A classical link is a connection transporting packets. This could be a physical or virtual link carried over a (MPLS) network. The proposed extensions in this document use these links to exchange capabilities.

Controllable Quantum Node (CQN): A controllable quantum node is a quantum device consisting of at least one qubit, capable of performing (a subset of) the following operations described in detail below: storing qubits for some amount of time, performing quantum operations such as entanglement distillation and entanglement swapping, and producing entanglement between the nodes and traveling qubits. The latter are generally realized using photons over fibers or through free space.

The term controllable refers to the fact that external control in software is capable of selecting the desired operations and qubits to use. Such nodes can be quantum repeaters that allow choices of operations to be made, as well as quantum end nodes capable of executing complex application protocols [14]. Quantum repeaters that merely allow timing control, such as automatic entanglement swapping whenever qubits arrive in a specific timing interval, will not be referred to as CQN. Such automated repeaters can be seen as lying at the quantum physical layer and do not enter routing or other decision making, apart from being switched on or off, and hence are not relevant to advertisement protocols like the ones considered here.

Quantum end node (QEN): In this document, a quantum end node [14] is one of a pair of quantum nodes forming an entanglement via a sequence of zero or more CQNs. Quantum end nodes typically run a higher-layer quantum application such as QKD.

Entangled Pair (EP): An entangled pair is a special state of two qubits, known as an EPR pair [8]. An entangled pair of qubits $c@A$ and $c@B$ is denoted $[[c@A, c@B]]$.

The process of entangling two particles $c@A$ and $c@B$ is denoted as follows:

$$\text{ent}(c@A, c@B) \rightarrow [[c@A, c@B]]$$

$\text{ent}(c@A, c@B)$ may take time T and succeed with probability P , and yield an entangled pair $[[c@A, c@B]]$ of fidelity F .

Fidelity: A measure of the quality of the entanglement of an EP $QFid$ [11]). Fidelity lies in the interval $[0, 1]$ where a higher value indicates better quality; usable fidelity values lie in the half-open interval $(0.5, 1]$.

Communication qubit: A qubit is called a communication qubit if it is possible to produce entanglement between this qubit and a traveling photon. This can be done by emission from the quantum node, that is, entanglement is produced between the qubit and the photon which is emitted from the quantum node. This process has been demonstrated in a number of physical systems that can be used as quantum nodes such as NV in diamond ([16], [17]), Ion Traps ([18]) and Neutral Atoms ([19]). An example of a communication qubit is the electron spin of the NV in diamond system ([15]). Entanglement between a communication qubit and traveling photons can also be produced by absorption. Examples include atomic ensemble memories ([20]).

A communication qubit c at CQN A is denoted by $c@A$, or simply c (if the node A is understood).

Storage qubit: A qubit is called a storage qubit if the node has the capability to use this qubit as a (temporary) quantum memory, but the qubit cannot serve as a communication qubit. To make storage qubits useful a node is required to possess the ability to transfer the state of a communication qubit to a storage qubit. An example of a storage qubit is the nuclear spin in the NV in diamond system [16].

A storage qubit s at node B is denoted $s@B$.

Swap: Two qubits located in the same CQN can interchange states ([13]). For example, the states of a communication qubit and a storage qubit at A can be swapped as follows:

$$\text{swap}(c@A, s@A)$$

If $c@A$ was entangled with $c@B$, the result is that $s@A$ is now entangled with $c@B$.

Distillation: Distillation is the process of turning a large number of weakly entangled states into a smaller number of highly entangled states ([13]).

For example, EPs $[[c1@A, c1@B]]$ and $[[c2@A, c2@B]]$ of fidelities $F1$ and $F2$ respectively may be distilled as follows:

```
dist([[c1@A, c1@B]], [[c2@A, c2@B]]) -> [[c3@A, c3@B]]
```

If distillation is successful, the fidelity $F3$ of $[[c3@A, c3@B]]$ will be higher than $F1$ and $F2$.

Entanglement Swap: Given two EPs $[[c@A, c1@B]]$ and $[[c2@B, c@C]]$, one can perform an entanglement swap:

```
entSwap([[c@A, c1@B]], [[c2@B, c@C]]) -> [[c@A, c@C]]
```

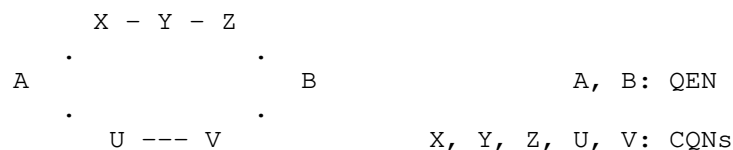
to create a new EP between $c@A$ and $c@C$. This is how "multihop" EPs are created from a sequence of "single-hop" EPs.

The swap operation can also be used within a CQN. A possible use case is when there aren't enough communication qubits to create the needed EPs. If, in the above example, B doesn't have two communication qubits $c1$ and $c2$, the following can be done:

```
ent(c@A, c@B) -> [[c@A, c@B]] # entangle
swap(c@B, s@B) -> [[c@A, s@B]] # swap EP to storage qubit
ent(c@B, c@C) -> [[c@B, c@C]] # use freed up qubit c@B
entSwap(c@A, c@B) -> [[s@B, c@C]] # create multihop EP
```

2. Motivation

Consider the following (very simple) quantum network consisting of QENs A and B, and CQNs X, Y, Z, U, V. The goal is to create an EP between qubits at A and at B, perhaps for the high-level task of QKD between A and B.



From A's point of view, here are a number of questions:

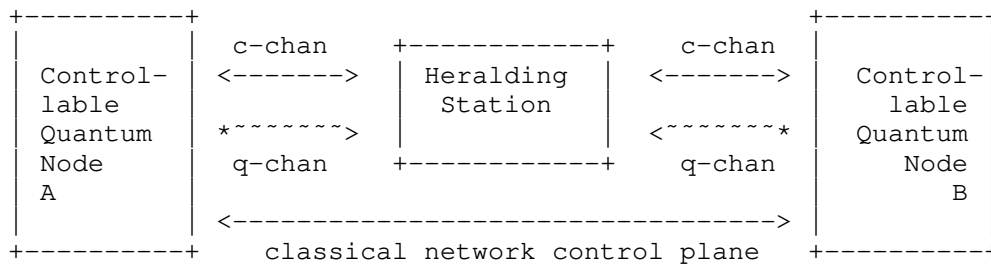
1. Is B reachable from A via quantum links that allow EP creation?
2. If so, along what sequence(s) of quantum nodes?

3. Can each pair of adjacent CQNs in this sequence form EPs? If so, how long will it take, and what fidelity can be expected?
4. If each pair of adjacent CQNs successfully forms EPs of sufficient fidelity, can these be swapped to form a multihop EP between A and B?
5. If a multihop EP between A and B were to be formed, would it be of good enough fidelity, or should a second multihop EP be formed and the two EPs distilled into one high fidelity EP? How many times should this process be repeated?
6. If the overall answer is Yes, should A proceed via sequence A, X, Y, Z, B, or sequence A, U, V, B?

This document aims to provide all CQNs in a quantum network with the information they need to answer such questions, and to create EPs at their desired fidelity and speed.

3. Theory of Operation

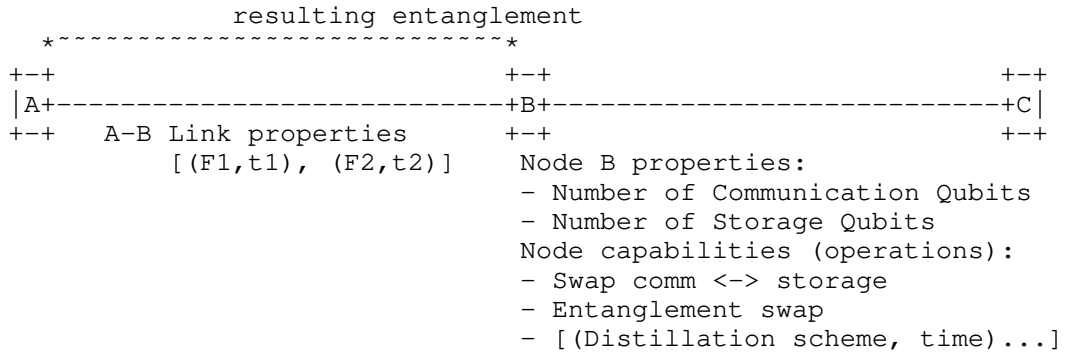
A CQN contains one or more communication qubits and one or more storage qubits. Many proposals exist for producing EPs between remote quantum nodes (see for example [16], [17], [18], [20]). Abstractly, these result in the generation of EPs with fidelity F after an expected time t . To give an example, we describe the generation of EPs that has been implemented in NV in diamond ([16]), and Ion Traps ([18]). The largest distance for producing long-lived entanglement is presently 1.3kms ([17]). To entangle a pair of communication qubits, the QNs send carefully timed photons towards the HS. If the process is successful, HS sends an OK message to both QNs.



The classical network control plane is of particular interest here as it would be used by the proposed protocol to advertise and exchange information about the capabilities of the CQNs to generate entanglement. This classical channel exists between all CQNs and is

shared with other application specific control and data plane traffic.

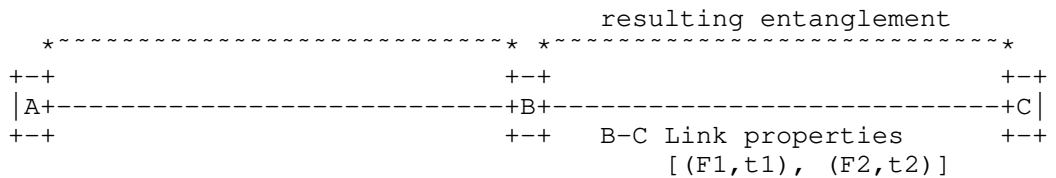
3.1. Multihop Entanglement



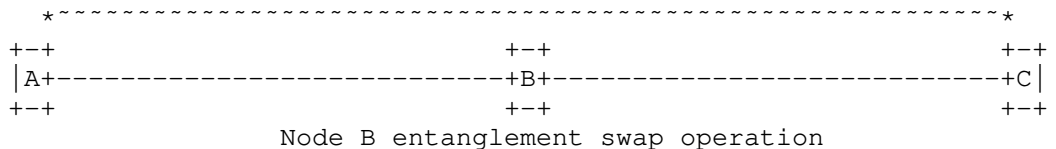
In the figure above, an example request for an entangled pair between nodes A and B will be affected by the following properties:

- o A chosen combination of F(idelity) and t(ime) duration to produce an entanglement at the respective Fidelity. These parameters roughly equate to the quality of the link, the accuracy with which the nodes can use the link, and the delay in classical networking.
- o The actual capability of nodes A and B to make use of the communication qubits.

A new EP creation between CQNs B and C will similarly be affected by the same parameters as above.



And finally, with an entanglement swap operation at node B (which is a node specific capability and has a specific duration) we end up with an A-C EP:

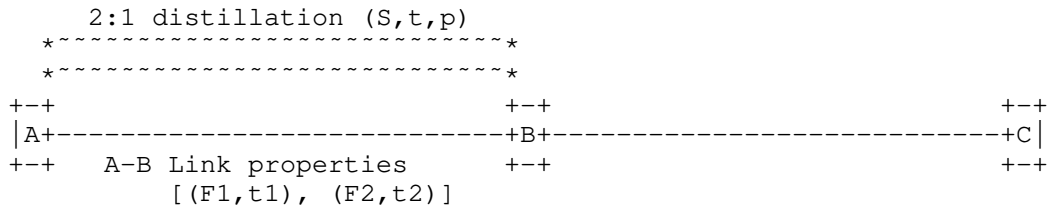


3.2. Distillation

If a pair of CQNs A and B share a number of EPs of insufficient quality, they may be combined into a single EP of higher quality by distillation. To do so, these CQNs need to agree on which distillation scheme to use before distillation can proceed. This does not necessarily need to be via communication between A and B, if one agrees upon a deterministic procedure of selecting one. This document suggests the following procedure:

1. A and B look at the distillation schemes that both advertise in common.
2. If there is none in common, stop. Distillation is not possible.
3. If there is a non-trivial subset in common, the first scheme in the node with the lower router ID is to be used by A and B.

Given a chosen distillation scheme (S,t,p) , an additional time delay will be added for the actual operation: For a 2:1 distillation scheme between nodes A and B, 2 EPs need to be produced followed by an operation on A and B that produces 1 EP. This operation will take time some expected time t , and succeed with probability p .



3.3. Node Properties

We are interested in exposing the properties of CQNs (including QENs) to allow sophisticated decision making, for example in the creation of entanglement. These properties include:

1. Number of communication qubits. The number of communication qubits determines the number of entangled pairs that the node can produce simultaneously.
2. Number of storage qubits
3. Possible operations, along with their execution time and probability of success:
 1. Swap between communication and storage qubits

2. Entanglement swap
3. List of supported distillation schemes (in order of preference).

Note that several other parameters can be advertised, such as the T1 and T2 times for a qubit's decoherence. These are omitted for now, instead just giving the decay of the fidelity of an EP. If deemed useful, T1 and T2 times can additionally be advertised.

3.4. Link Properties

A list of (F_n, t_n) pairs describing the tradeoffs of a possible entanglement produced by two nodes (the ends of said link): t_n is the time to produce an entangled pair with fidelity F_n .

4. The (Ab)use of Protocols

The routing protocols IS-IS and/or OSPF could be used in order to advertise entanglement capabilities. This section describes the additional data fields needed in order to facilitate the objective.

4.1. A Brief Primer on Link-state Protocols

This document suggests the use of a link-state protocol to distribute the capabilities of CQNs to create entanglement. This section offers a short introduction to link-state protocols for those not familiar with them.

Consider a directed graph $G=(V, E)$ with vertices (nodes) V and edges (links) E . Consider also $G'=(V', E')$; there is a 1-1 mapping from V' to V and from E' to E such that $e_{l'} = (v_{l'}, v_{l'})$ is in E' iff $e_l = (v_l, v_l)$ is in E and $v_{l'}$ maps to v_l and $v_{l'}$ maps to v_l . G' represents the quantum network; V' represents the set of CQNs, and E' the set of quantum links between pairs of CQNs; G represents a classical network parallel to G' ; that is, each CQN v' has a corresponding classical node v . v plays a dual role: it is the control node for v' , and proxies on behalf of v' in the link-state protocol.

The basic objective of a link-state protocol is to "flood" properties of nodes and (directed) links to all nodes in the network. This is accomplished by means of "link-state advertisements" (LSAs) that each node originates and sends to its immediate neighbors. The neighbors in turn send received LSAs to their own neighbors; this process repeats until every node receives every LSA (hence the term "flooding"). The focus of LSAs is the link properties (hence `_link-state_ advertisements`), although node properties are also advertised.

There are mechanisms to prevent looping of LSAs, and for reliable flooding. There is also a sequence number by which a more recent update of an LSA can be identified as such, and a mechanism for "aging out" LSAs belonging to nodes no longer in the network. In what follows, quantum node and link properties are added to the link-state advertisements of the corresponding classical node. Note that link properties need not be symmetric; that is, the link properties of (v, w) need not be the same as those of (w, v) .

The net result of flooding is that every node has the same picture of the network (modulo LSAs in flight); in particular, each node knows the overall topology and connectivity of the network, and can use this information to make decisions. In a classical network, such a decision could be to compute a shortest path; for the quantum network, it could be choosing a feasible path (i.e., sequence of CQNs) for a multihop entanglement. Note that a node doesn't really know when it has complete and up-to-date information about the network; LSA updates may be originated at any time. Usually, this is okay: for example, if a node v learns enough of the network to have a path to another node w , it can compute a multihop entanglement to w . Subsequent updates may provide a more optimal (or higher probability) entanglement path. There are heuristics one can apply to guess that the link-state database (LSDB) (i.e., the union of all LSAs) is complete-ish; however, as nodes (and links) can fail or disconnect, there really is no such thing as "the full LSDB".

Each node v is identified by a "router ID" (an IP address uniquely allocated to v), denoted by $\text{rid}(v)$. A link $L = (v, w)$ is identified by $(\text{rid}(v), i)$ where i is an index allocated by v for L unique for each link emanating from v . (L may also be identified by IP addresses, but we'll ignore that for now.) It is generally expected that a directed link (v, w) is matched by a link (w, v) ; if not, (v, w) is ignored from subsequent consideration; in particular, no link properties are advertised for this link by v . Note that a pair of nodes may have multiple links between them; for simplicity, the notation will not be extended to indicate this. We'll assume $\text{rid}(v') = \text{rid}(v)$ and the index allocated to a quantum link e' is the same as that of the corresponding classical link e .

Let v, w be a pair of neighboring nodes, and let $L1 = (v, w)$ and $L2 = (w, v)$ in E be directed links in opposite directions between v and w with identifiers $(\text{rid}(v), i1)$ and $(\text{rid}(w), i2)$ respectively (where $i1$ is the index allocated for $L1$ by v , and similarly for $i2$). As a first step in running a link-state protocol, v runs a "hello protocol" all its links; in particular, over $L1$. Similarly, w will run the hello protocol over $L2$. The hello protocol serves to exchange the indices $i1$ and $i2$, and thus identify $(\text{rid}(v), i1)$ as the reverse link of $(\text{rid}(w), i2)$. This allows both v and w to correlate

the link properties of L1 and L2. If the hello protocol fails between *v* and *w*, neither node includes link properties for the link in their LSAs.

Once the hello protocol has been run on all links, *v* starts the process of generating and sending its own LSA over all its links, and of receiving the current LSDB from its neighbors. Note that an LSA originated by *v* must propagate unchanged across the network; only *v* is allowed to change it (and such a change must be accompanied by updating the LSA's sequence number). Such an update is triggered by a new link coming up, an existing link going down, or a node or link property changing.

IS-IS and OSPF are in principle similar, although the details of the protocol mechanisms and encodings vary. In both protocols, a Type-Length-Value (TLV) is used to encode most node and link properties. In IS-IS, TLVs are used for all properties, and a single type of LSA is used; in OSPF, there are several types of LSAs, and many (but not all) properties are encoded as TLVs.

[1] has examples of "standard" LSAs for routing; [4] has the so-called Traffic Engineering LSAs.

4.2. Node Properties

Here, we give a protocol-independent description of quantum node properties; later documents will specify the encoding specifically for IS-IS and OSPF.

Note that the following list of node properties is a strawman; all details are subject to change, and other properties may be added as needed.

The following node properties are added to the appropriate LSA:

```
<Qubit-TLV><NCQ><NSQ>
<CS-Swap><Prob><ExecTime>
<Ent-Swap><Prob><ExecTime>
<Measure><Prob><ExecTime>
<NDistSch><DistScheme1><DistScheme2>
```

4.3. Link Properties

Only one link property is listed. It gives the time-fidelity tradeoffs of an entanglement operation as a list:

```
<N-Ent-TO><time1><fid1><time2><fid2>...
```

This is interpreted as follows: an entanglement operation may be initiated between nodes v and w over link (v, w). Depending on how fast one wants to complete (time-i), the list gives the corresponding fidelity of the resulting entanglement (fid-i). time-i is given in nanoseconds; fid-i as a number between 0 and 999999. The denominator is 1000000.

Note that this link property is symmetric, as entanglement is initiated simultaneously at v and w.

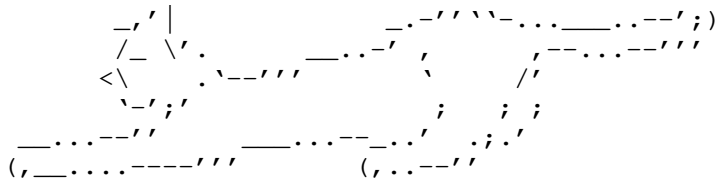
5. Security Considerations

It is not anticipated that adding these extensions to IS-IS and OSPF will present new security hazards to those protocols. Since however a common application of entangled pairs is for security purposes (such as QKD), it is worth investigating whether this application places a higher burden of security on the underlying protocols.

6. Acknowledgments

The authors would like to thank the following people for their contributions and support to this document: Vesna Manojlovic, Marcello Caleffi and Rodney Van Meter. Kompella would also like to thank Bruno Rijsman for encouraging him to learn about Quantum Computing and Networking.

Also:



7. Editorial Comments

It could be worth investigating the use of a distance-vector routing protocol to limit flooding.

8. IANA Considerations

There are no requests as yet to IANA for this document.

9. References

9.1. Normative References

- [1] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", RFC 1195, DOI 10.17487/RFC1195, December 1990, <<https://www.rfc-editor.org/info/rfc1195>>.
- [2] 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>>.
- [3] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, DOI 10.17487/RFC3630, September 2003, <<https://www.rfc-editor.org/info/rfc3630>>.
- [4] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <<https://www.rfc-editor.org/info/rfc5305>>.
- [5] Ishiguro, K., Manral, V., Davey, A., and A. Lindem, Ed., "Traffic Engineering Extensions to OSPF Version 3", RFC 5329, DOI 10.17487/RFC5329, September 2008, <<https://www.rfc-editor.org/info/rfc5329>>.
- [6] Aggarwal, R. and K. Kompella, "Advertising a Router's Local Addresses in OSPF Traffic Engineering (TE) Extensions", RFC 5786, DOI 10.17487/RFC5786, March 2010, <<https://www.rfc-editor.org/info/rfc5786>>.

9.2. Informative References

- [7] "Qubit", <<https://en.wikipedia.org/wiki/Qubit>>.
- [8] "Quantum Entanglement", <https://en.wikipedia.org/wiki/Quantum_entanglement>.
- [9] "Quantum Network", <https://en.wikipedia.org/wiki/Quantum_network>.
- [10] "Quantum Teleportation", <https://en.wikipedia.org/wiki/Quantum_teleportation>.
- [11] "Quantum Fidelity", <https://en.wikipedia.org/wiki/Fidelity_of_quantum_states>.

- [12] Wehner, S., Elkouss, D., and R. Hanson, "Quantum internet: A vision for the road ahead", *Science* 362, 6412, 2018.
- [13] Rozpedek, F., Schiet, T., Thinh, L., Elkouss, D., Doherty, A., and S. Wehner, "Optimizing practical entanglement distillation", *Phys. Rev. A* 97, 062333, 2018, <<https://arxiv.org/abs/1803.10111>>.
- [14] "Quantum Key Distribution", <https://en.wikipedia.org/wiki/Quantum_key_distribution>.
- [15] "Nitrogen-vacancy center", <https://en.wikipedia.org/wiki/Nitrogen-vacancy_center>.
- [16] Humphreys, P., "Deterministic delivery of remote entanglement on a quantum network", *Nature* 558, 2018.
- [17] Hensen, B. and , "Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres", *Nature* 526, 2015.
- [18] Hucul, D. and , "Modular entanglement of atomic qubits using photons and phonons", *Nature Physics* 11(1), 2015.
- [19] Noelleke, C. and , "Efficient Teleportation Between Remote Single-Atom Quantum Memories", *Physical Review Letters* 110, 140403, 2013.
- [20] Sangouard, N. and , "Quantum repeaters based on atomic ensembles and linear optics", *Reviews of Modern Physics* 83, 33, 2011.

Authors' Addresses

Kireeti Kompella
Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, CA 94089
USA

Email: kireeti.kompella@gmail.com

Melchior Aelmans
Juniper Networks, Inc.
Boeing Avenue 240
Schipol-Rijk, PZ 1119
The Netherlands

Email: maelmans@juniper.net

Stephanie Wehner
QuTech
Van der Waalsweg 122
Delft, LC 2611
The Netherlands

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

Cristian Sirbu
Redbit Networks
Dublin
Republic of Ireland

Email: cristian@redbit.network

Axel Dahlberg
QuTech
Van der Waalsweg 122
Delft, LC 2611
The Netherlands

Email: E.A.Dahlberg@tudelft.nl