

Synchronization Control-Plane Protocol for Quantum Link Layer

Brandon Ru, Winston K.G. Seah, Alvin C. Valera

School of Engineering and Computer Science

Victoria University of Wellington

Wellington 6140, New Zealand

Email: {rubran, winston.seah, alvin.valera}@ecs.vuw.ac.nz

Abstract—Heralded entanglement generation between nodes of a future quantum internet is a fundamental operation that unlocks the potential for quantum communication. In this paper, we propose a decentralized synchronization protocol that operates at the classical control-plane of the link layer, to navigate the coordination challenges of generating heralded entanglement across few-qubit quantum network nodes. Additionally, with quantum network simulations using NetSquid, we show that our protocol achieves lower entanglement request latencies than a naive distributed queue approach. We observe a sixfold reduction in average request latency growth as the number of quantum network links increases. The Eventual Synchronization Protocol (ESP) allows nodes to coordinate on heralded entanglement generation in a scalable manner within multi-peer quantum networks. To the best of our knowledge, this is the first decentralized synchronization protocol for managing heralded entanglement requests.

Index Terms—heralded entanglement, quantum link layer, quantum control plane, quantum internet, quantum network protocol

I. INTRODUCTION

Quantum networks harness quantum mechanical phenomena to achieve superior performance compared to classical networks and present opportunities for novel use cases. Fundamentally, they enable the distribution of quantum states over long distances, facilitating the development of quantum communication protocols like quantum key distribution (QKD) [1], which offers theoretically unbreakable encryption. Quantum networks will also enable distributed quantum computing, where complex problems are more efficiently solved than any single quantum computer could on its own [2].

Central to the operation of quantum networks is the phenomenon of entanglement, where pairs or groups of particles become interconnected in such a way that the state of one particle instantaneously affects the state of the other, regardless of the distance separating them [3]. Entanglement is a fundamental resource of quantum networks that leverages correlations between particles to form a quantum link between nodes [4]. It is a critical resource for quantum communication, enabling protocols like quantum teleportation, which allows the transfer of quantum states between distant locations without moving the physical particles themselves.

However, quantum networks operate under a set of constraints that are markedly different from those of classical networks. Quantum bits, or qubits, are highly susceptible to

decoherence and loss, requiring error correction and fault-tolerant mechanisms that are still in developmental stages. The no-cloning theorem, which states that it is impossible to create an identical copy of an arbitrary unknown quantum state, imposes fundamental limits on the way information can be duplicated and distributed across the network and complicates the design of quantum communication protocols [5].

Despite their advanced capabilities, quantum networks cannot operate entirely independently and must be integrated with classical networks. Classical infrastructure plays a crucial role in tasks such as clock synchronization, control of quantum devices, and error correction. Classical communication is also necessary for the post-processing steps in quantum key distribution and for coordinating entanglement distribution across the network. Therefore, the successful deployment of quantum networks will rely on a hybrid approach that combines both quantum and classical technologies.

In this paper, we will build off the quantum link layer definition and protocol definition from Dahlberg et al. [6], and propose a control-plane protocol for synchronization. Although there is a mixed consensus regarding the definition of the link layer across other quantum network stack proposals [7] [8] [9], it is agreed that entanglement is a fundamental resource core to quantum networking.

The proposed control-plane protocol is designed for synchronizing nodes in a multi-peer quantum network that wish to generate heralded entanglement. Fig. 1 illustrates an example of a quantum network topology the protocol could operate within. The protocol operates in a decentralized manner, without a central scheduler, and offers the opportunity for concurrent entanglement generation. Although other schemes exist for generating entanglement between separated nodes, recent research consisting of experimental trials in multi-peer quantum networks [10] [11] [12] continue to use a heralded approach, which offers greater scalability and control to the end nodes.

This paper makes the following key contributions:

- 1) **Decentralized Synchronization Protocol:** We introduce a control plane protocol designed for synchronizing heralded entanglement requests. Our quantum network simulations in multi-node topologies display a six-fold reduction in mean latency growth when compared to our baseline distributed queue protocol.

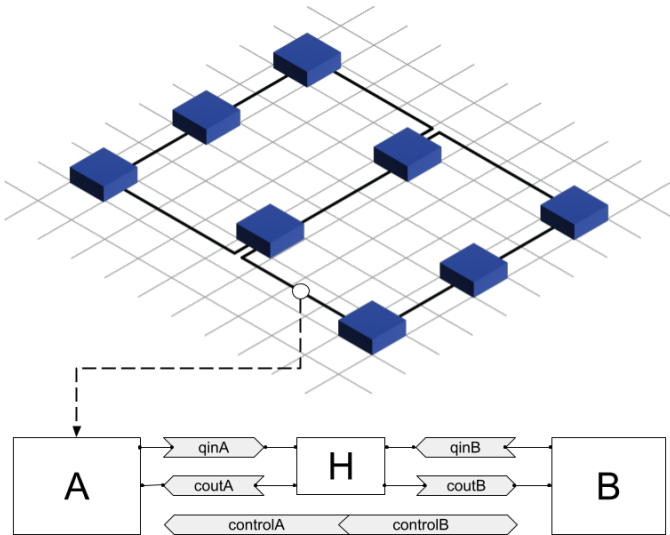


Fig. 1: Example of an eight-node quantum network topology. Each pair of nodes generates heralded entanglement using the single-photon protocol, sending entangled qubits (photons) to the heralding station H using inbound fiber-optic links. Entanglement success is probabilistic and heralded by outbound classical connections. Coordinating these entanglement requests involves an additional control plane topology.

- 2) **Simulated Benchmarks for DQP:** In our paper, we also extend the original DQP implementation in [6] beyond a two-node network, and test it in a multi-peer quantum network to measure its performance.
- 3) **Control Plane Definition:** We provide a clear definition of the control plane within the quantum link layer, emphasizing the interdependent relationship between classical and quantum networking required within the link layer.

The rest of the paper is organized as follows. The background of quantum entanglement generation, the constraints of quantum networking, and related work are outlined. The design details of the Eventual Synchronization Protocol (ESP) are presented in Section V. The performance of ESP is outlined, evaluated and analysed in Section VI. Conclusions are in Section VII.

II. THEORETICAL BACKGROUND

A. Quantum Bits (Qubits) and Entanglement

Qubits are the foundational building blocks of quantum systems. They are two-state quantum systems, and represent the linear superposition of a ground state and excited state (0 or 1). Quantum states are sensitive to the environment and suffer from **decoherence**. Decoherence is the loss of coherence in a qubit, which is the fundamental property that allows the qubit to exist in a superposition of states. The causes of decoherence are specific to the physical implementation of the qubit, but vary depending on the material and temperature [13]. In general, interactions with the environment and the passage of time will introduce decoherence into a qubit.

Quantum entanglement is another fundamental phenomenon where two or more particles are linked in a way such that the quantum state of one particle cannot be independently described without considering the state of the others. As quantum entanglement is a phenomenon that does not depend upon distance or a physical connection, maximally entangled (Bell State) pairs of qubits can be used to create a quantum link between network nodes. This correlative phenomenon between Bell State pairs is also crucial to facilitate quantum teleportation, the transfer of a qubit along this entangled link.

B. Heralded Entanglement

Remote entanglement using emitted photons carrying a quantum state (single-photon protocol) is well-suited for entanglement generation between network nodes [14]. Situated between the nodes is a *heralding station* H , which is connected to the end nodes with an optical fibre link. Each end node, A and B , prepares a *communication qubit* in the superposition state $\sqrt{\alpha}|0\rangle + \sqrt{1-\alpha}|1\rangle$, where α is known as the *bright state population* of the qubit. Excitation of the qubit in the state $|0\rangle$, the *bright state*, produces a joint entangled quantum state consisting of the communication qubit and a photon. Then, the role of the heralding station, H , is to act as a beam splitter and interfere with the incoming photons from A and B . Depending on the measurement outcome of the photons at H , the communication qubits at A and B are projected into one of the Bell State pairs, $|\Psi\rangle^+$ or $|\Psi\rangle^-$, thereby generating an entangled pair of qubits between adjacent network nodes.

However, this heralding process is probabilistic and only succeeds with a small probability. This entanglement process is known as *heralded entanglement generation* [15], and requires photons to only travel half the distance of a link and a stateless intermediate node. However, observe that **both** end nodes A and B are involved in this process, as they need to send their photons to H . This suggests the need for a synchronization protocol to ensure end nodes coordinate before they attempt entanglement generation.

C. Entanglement Fidelity

Physical systems are not perfect and the entanglement generation is subject to noise. That is, a generated entanglement under perfect conditions would be a Bell State, but in reality is slightly different. *Entanglement fidelity*, denoted by F , is a purely quantum metric that measures the quality of the generated entanglement, where $F \in [0, 1]$ and $F = 1$ is the ideal state. Entanglement fidelity is of importance as it directly affects the functionality of higher level services that consume these entanglements. For example, it is well known that there are strict bounds on the quantum bit error rate (QBER) in quantum key distribution (QKD) protocols, to ensure safety against eavesdropping. Lower entanglement fidelities used in QKD would result in a higher QBER, meaning entanglement fidelity is an important quality-of-service parameter that should be accounted for during entanglement generation. In practice, only entanglements where $F > 0.5$ contain useful amounts of entanglement for use [16].




Fig. 2: Relationship of the entanglement generation probability (p_{success}), minimum entanglement fidelity (F_{\min}), and bright state population (α). This relationship is used to model the physical layer entanglement generation process within our simulations. In this example, we obtain p_{success} from a $t = 120$ (seconds) simulation on 2×2 network topology.

The relationship between the bright state population of the single-photon heralded entanglement generation protocol, α , and entanglement fidelity, F is also of significance [14]. Empirical simulated results are displayed in Fig. 2 and they underpin the physical layer of our network simulations. Heralded entanglements generated with lower values of α for the initial communication qubit have a higher final entanglement fidelity and vice versa. However, there is also a tradeoff to consider as there is a positive relationship between α and p_{success} , the probability of a successful entanglement occurring at H between the incoming photons. If each entanglement attempt requires time to execute, then there is a clear fidelity-time tradeoff, where higher fidelity entanglements require longer times to generate and vice versa.

D. Quantum Link Layer

Entanglement Generation Protocol (EGP) is an explicit definition of a protocol for robust entanglement generation [6]. EGP manages the requests between adjacent nodes that wish to generate an entanglement, and is responsible for a wide range of tasks including scheduling, timekeeping, quantum memory management, and interfacing with a physical layer protocol.

One such physical layer protocol, the Midpoint Heralding Protocol (MHP), uses the single-photon scheme and is responsible for polling the higher EGP layer and attempting heralded entanglement between an adjacent node. Importantly, MHP contains no decision making elements and operates with discrete time steps at the nanosecond scale (i.e. the physical layer operates without state in an automated manner).

Hence, EGP relies on a synchronization protocol to ensure the remote node of the heralded entanglement attempt is aware of the request and it can succeed. This is because both end nodes of the entanglement must simultaneously send photons to the heralding station, H , and is a general problem in heralded entanglement generation.

III. RELATED WORK

Coordination is an inherent part of the quantum link layer and is achieved with various control plane protocols. The generation of heralded entanglement requires end nodes of the link to interact with the midpoint heralding station. For example, in EGP, a distributed queue protocol (DQP) is used to coordinate requests between layers [6]. The distributed queue is implemented using a master-follower architecture, and requests are placed into the queue by the end nodes of a link using a two-way handshake. In contrast, our solution uses a decentralized architecture where all end nodes are equal, and communicate with one another.

Skrzypczyk [17] and Pompili et al. [18] suggest a centralized approach is more appropriate, as a centralized scheduler can produce dynamic time division multiple access (TDMA) schedules for the network, given it has an overarching view of entanglement request details. This setup is similar to software-defined networking (SDN) schemes in classical networks, and allows the scheduler to account for quality-of-service requirements. However, even though the central scheduler can be implemented in a fault-tolerant distributed manner, it still poses as additional overhead to a local quantum network and raises scalability concerns. Our proposed protocol currently differs as it does not perform any explicit scheduling of requests, and operates with a first-in-first-out (FIFO) strategy.

Another control-plane protocol, RuleSet [16] tackles this coordination problem directly by introducing a declarative protocol that allows nodes to execute coordinated actions. Rules are distributed to nodes that need to perform shared actions, and consist of *Conditions* (which are conditional statements) and *Actions* which reserve and allocate resources to other nodes. However, it is not explicitly clear how this resource allocation and locking occurs, as there can be contention between nodes. The proposed solution hopes to offer a potential strategy for addressing this resource allocation, as it offers a means to synchronize requests.

In addition to explicit control plane protocol research, there are experimental trials to generate entanglement in a scalable manner within multi-peer quantum networks. For example, Hermans et al. [11] performed qubit teleportation between non-adjacent nodes in a quantum network by first generating entanglement with a shared central node and performing entanglement swapping. Each node is optically connected in a chain, and utilize the single-photon protocol to generate entangled pairs using a communication qubit.

Another recent endeavour has also shown promising experimental results on a metropolitan scale between three-nodes, maximally separated by 12.5 km, from Liu et al [10]. The single-photon heralded entanglement scheme is also used, and each of the quantum end nodes are connected to a central server. The central server contains a multi-input two-output optical switch that allows the nodes to share a heralding station. However, hot-swapping between the pairs wishing to generate entanglement is required at the switch, and is akin to a time-sharing multiple access protocol.

At a classical level, there also exist protocols for managing contention and coordination. These include centralised lock managers which have a global view of requests and are able to provide desirable qualities such as the use of scheduling strategies [19] [20], and starvation and deadlock prevention techniques [21]. Distributed protocols, such as the Ricart-Agrawala Algorithm [22] and Lamport’s Distributed Mutual Exclusion Protocol [23] also exist and whilst are effective for traditional distributed systems, do not fully transfer into the context of managing entanglement generation. In a quantum network, a node may receive multiple simultaneous requests to generate entanglement with different nodes, requiring a mechanism to defer and redirect attention to other nodes if a request is rejected, which is not adequately handled by these protocols.

IV. PROBLEM FORMULATION

A quantum network, represented by $G_Q = (V, E_Q)$, consists of quantum links, E_Q , between quantum nodes, V , forming the logical entanglement topology of the network. That is, E_Q is equivalent to the set of possible entanglements that can be generated. During heralded entanglement, both nodes involved in the request must be aware of the request details and synchronized. This requires a separate control-plane network $G_C = (V', E_C)$, where each node must be able to make an entanglement request, $V \subseteq V'$, but E_C and E_Q can be different.

In this model, we assume each quantum node can only participate in entanglement generation with exclusively one remote node, due to strict timing constraints from entanglement decoherence. That is, an invariant of the system is:

$$\sum_{j \in V, j \neq i} x_{ij} \leq 1 \quad \forall i \in V \quad (1)$$

where x_{ij} indicates active entanglement generation between node i and node j . The objective is to design a control plane protocol for efficient entanglement generation that operates on a suitable G_C configuration, given the constraints of the quantum network G_Q .

Requests arrive sequentially from higher layers of the quantum network stack, denoted by the set $M = \{r_1, r_2, \dots, r_n\}$, where each request r_k is defined as:

$$r_k = (i_k, j_k, f_k, p_k, t_k) \quad (2)$$

where:

- $i_k, j_k \in V$ are the initiator and remote nodes.
- f_k denotes the fidelity requirement for the entanglement.
- p_k is the number of entangled pairs requested.
- t_k represents the time at which the request is made, with sequential arrivals such that $t_1 \leq t_2 \leq \dots \leq t_n$.

In this paper, our objective is to minimize the average request latency over M within the network, through the design of an appropriate link-layer protocol that communicates using the control plane G_C . For comparison, we use the DQP protocol from [6] as our baseline. In our implementation,

TABLE I: Protocol Messages in ESP

Message	Description
<i>SYN</i>	Initial request to synchronize and acquire resources
<i>SYNACK</i>	Request to synchronize can be met at remote node
<i>SYNREJ</i>	Request to synchronize cannot be met at remote node
<i>ACK</i>	Request is valid at initiator node, complete.
<i>NACK</i>	Request is no longer valid at initiator node, reject.
<i>WAKE</i>	Unblock request at remote node.

TABLE II: Protocol States in ESP

State	Description
IDLE	Node has not initiated any synchronization.
SYN_SENT	Node has sent out the initial <i>SYN</i> message.
SYN_RECEIVED	Node has received a <i>SYN</i> message and accepts it.
WAKE_SENT	Node has sent out an unblocking signal.
READY	Resource reserved. Synchronization complete.
BUSY	Resource is in use.

DQP adopts a centralized scheduler, S , to manage the shared view of requests. Each node makes requests and obtains updates from S and it follows that $G_C = (V \cup S, E_C)$ where $E_C = \{(i, S) \mid i \in V\}$. S maintains a shared queue of all requests and updates, enabling nodes to synchronize properly before initiating entanglement generation on a given edge in E_Q . We choose this approach as our baseline because a centralized method has been proposed in [17], [18], and is a natural extension of the link-local implementation in [6] to a multi-peer quantum network. It functions similarly to a centralized lock manager in a distributed system, offering a straightforward yet efficient means of managing coordination and contention.

V. EVENTUAL SYNCHRONIZATION PROTOCOL (ESP)

This section presents the Eventual Synchronization Protocol (ESP), a decentralized control-plane protocol for heralded entanglement. Fig. 3 illustrates the general control flow of ESP using a finite-state machine, while Table I lists the ESP messages, Table II details the states and transitions of the protocol’s finite-state machine, and Table III outlines the relevant implementation details.

A. Overview

In ESP, each quantum node is treated as a resource that must be locked during the entanglement process. The control plane protocol should ensure that heralded entanglement generation only occurs between adjacent nodes which are properly synchronized. Specifically, for an edge $(i, j) \in E_Q$, node i must acquire the *Lock* on node j to initiate entanglement. Otherwise, the entanglement process cannot proceed if both nodes in the edge (i, j) are not participating. Entanglement requests are placed into a main queue Q_{main} and are processed sequentially by Algorithm 1.

Instead of using a centralized scheduler S , in ESP, requests for entanglement generation are managed through the control plane network $G_C = (V, E_C)$ where $E_C = E_Q$. That is, the control plane network is equivalent to the logical quantum network topology. Synchronization requests are initiated with the *SYN* message, and put the initiator into the *SYN_SENT* state and if the remote node is *IDLE*, into the *SYN_RECEIVED*

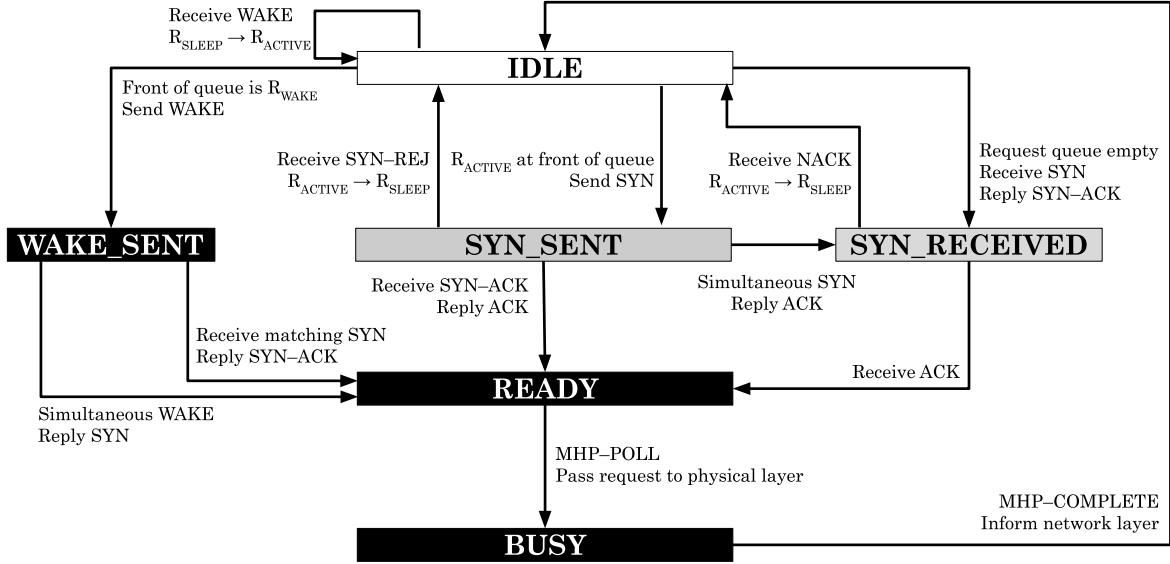


Fig. 3: ESP Finite State Machine Representation

Algorithm 1 ESP Request Processor

For each quantum node, running continuously:

```

1: if  $|Q_{ready}| > 0$  then
2:    $R \leftarrow Poll(Q_{ready})$ 
3:   Execute  $R$  at physical layer            $\triangleright$  Non-blocking
4:   On completion, Release(Lock)
5: else if  $|Q_{main}| > 0$  then
6:    $P \leftarrow Poll(Q_{main})$ 
7:   if  $P$  is of type active then
8:     Send SYN
9:      $S_{sent} \leftarrow P$ 
10:  else if  $P$  is of type wake then
11:    Send WAKE
12:     $S_{wake} \leftarrow P$ 
13:  end if
14: else
15:   Block until new request in  $Q_{ready}$  or  $Q_{main}$ 
16: end if

```

state. In Fig. 4, nodes B and D are in the state SYN_SENT as they attempt to make entanglements with node A. If a node receives multiple SYN requests such as node A, it chooses the request that arrives first and enters the SYN_RECEIVED state, and rejects the remaining options.

A key concept in ESP is handling requests that cannot immediately acquire the remote node. For example, consider the request to generate entanglement along the edge (i, j) . If node i makes the request and node j is busy, then node i puts the request to sleep and places it into S_{sleep} , essentially making the latter inactive. The protocol will exit the SYN_SENT state and return to IDLE, allowing the node to process other pending requests without blocking. A given node is considered busy if it has an acquired *Lock*, that is, its state is in SYN_RECEIVED, WAKE_SENT, READY, or BUSY.

TABLE III: Notation used in Protocol Design

Notation	Description
S_{wake}	Set of requests where WAKE has been sent.
S_{sent}	Set of requests where SYN has been sent.
S_{sleep}	Set of requests that are sleeping.
Q_{ready}	Queue of requests that are synchronized. Ready for physical layer to process.
Q_{main}	Priority queue of requests that require synchronization. Requests sorted in ascending order by arrival number.
$Lock$	Protected lock on physical layer.

Nodes that reject requests with SYN_REJ due to resource contention asynchronously signal the original request initiator when they become available again by queuing a wake-up request, R_{wake} into the main queue Q_{main} . Upon processing of an R_{wake} request, the node will enter the WAKE state. Then, upon receiving a WAKE message, a given node will remove a sleeping request from S_{sleep} and place it back into Q_{main} . In Fig. 4, nodes E and F are busy generating entanglements, and after node F finishes this process, it processes the R_{wake} request to node I, which it must have rejected prior. Furthermore, the WAKE process cannot be interrupted by external SYN requests, to ensure starvation from priority inversion cannot occur. That is, if node E decides to make another set of entanglements with node F while it is waking node I up, it will reject node E. This mechanism ensures that the original initiator is notified when previously unavailable resources are freed, thus achieving *eventual synchronization* and an efficient management of entanglement requests.

Requests are assigned a local *arrival number*, A , which is strictly sequentially increasing with each new request. This value is decided upon arrival of any type of request, including any wake requests, R_{wake} . Requests are maintained in Q_{main} , which is a priority queue such that requests with the lowest arrival numbers are processed first. When a sleeping request is reactivated and placed into Q_{main} , its arrival number remains

unchanged. This ensures that requests that exit sleep are prioritized above newer ones to avoid “starvation” or indefinite waiting.




Fig. 4: Example protocol state snapshot of a nine-node (3×3) quantum network topology G_Q . E, F, G, and H are BUSY. C and I are IDLE. B and D are in SYN_SENT. F has a queued R_{wake} request for I.

Each request also contains a minimum start time, t_{min} , as included in the original link-local protocol (DQP) [6]. t_{min} accounts for the *propagation time* of the link to prevent the entanglement attempt process from starting prematurely. This value is decided once the resource is reserved at the remote node, that is, when the state of the protocol enters a READY state at the remote node. This minimum start time is crucial because if the remote node initiates the entanglement process too early, it leads to unnecessary attempts that can ultimately impact the final fidelity of the produced entanglement.

The decentralized nature of ESP allows each node to manage its own state, eliminating the need for a centralized scheduler. This design scales effectively for larger networks with multiple quantum nodes and can be adapted to handle unstable network conditions through mechanisms such as timeouts, retry attempts, and expiration of stale requests.

VI. PERFORMANCE EVALUATION

In this section, the performance of ESP is compared against traditional network performance metrics as it operates at the control plane. Performance is compared against a baseline protocol, DQP, which consists of a centralized subnet-local scheduler.

A. Performance Metrics

- 1) Request Latency, L is defined as:

$$L = NL_{finish} - NL_{start} \quad (3)$$

where NL_{start} is the initial timestamp of requesting entanglement from the network layer, and NL_{finish} is the final timestamp where all entangled pairs have been created.

- 2) Request Jitter, J is defined as the standard deviation of request latencies over a given time period:

$$J = \sqrt{\frac{1}{n} \sum_{i=1}^n (L_i - \bar{L})^2} \quad (4)$$

where L_i is the request latency of the i -th request, \bar{L} is the mean request latency, and n is the total number of requests in the given time period.

- 3) Scaled Latency, SL is defined as:

$$SL = \frac{L(R)}{|R|} \quad (5)$$

where $|R|$ is the number of entanglements requested in R , therefore, this can be treated as a normalized latency.

- 4) Scaled Network Throughput, T , is defined as:

$$T = \frac{\sum_{i=1}^t |E_i|}{E_Q \cdot t} \quad (6)$$

where the numerator is the total number of entanglements that have been generated in the time period t , E_Q is the number of edges in the topology G_Q , and the denominator is the product of the number of edges and the time period t .

- 5) Busy Time, B is given as:

$$B = PL_{finish} - PL_{start} \quad (7)$$

where PL_{start} is when R is passed to the physical layer for entanglement generation, and PL_{finish} is when the generation finishes. It does not include any queueing or transmission time.

- 6) Queueing Time, W is given as:

$$W = PL_{start} - NL_{start} \quad (8)$$

where this time difference measures the conjunctive time R has spent waiting in Q_{main} and Q_{ready} , which begins when the request is created, NL_{start} , and ends when the request is passed to the physical layer, PL_{start} .

B. Simulation Setup

The performance of ESP is evaluated and compared using NetSquid (Python / C++), a discrete-event quantum network simulator developed by QuTech [24]. Simplified versions of the original physical and link layer protocols, MHP and EGP, [6] are used to bootstrap our control-plane protocols for testing. Simulation settings are shown in Table IV. We simulate entanglement generation in a quantum network, considering various parameters such as qubits per node, heralding station position, link length, photon detection window, and physical layer clock frequency. Requests for entanglement are generated using a bootstrap network protocol, representing an edge in the G_Q topology. The simulation models the connections between nodes, quantum memory, processors, and the heralding station, ensuring accurate representation of the quantum network’s physical and logical structures.

We employ a stochastic simulation approach. For each undirected edge in the topology, a request R is made to generate an entanglement from one of the end nodes (randomly chosen). Once R is satisfied, we sample from an exponential distribution to determine a backoff time period before making another request, t_b . This approach ensures that the entanglement requests follow a Poisson process, providing a realistic

TABLE IV: Simulation Parameters

Parameter	Value
Qubits per Node	6
Simulation Time	120 seconds
Heralding Station Position	$L/2$
Link Length	2 km
Photon Detection Window	20 ns
Physical Layer Clock Frequency	Uniformly [450, 550] ns
Request Arrival Period	$\lambda = 50$ ms
Request Size	Uniformly [1, 6]




Fig. 5: Network topologies used in the evaluation of ESP.

simulation of network behavior [25]. Finally, the size of each request, $|R|$ is uniformly sampled from $[1, 6]$, denoting the number of entanglements requested.


C. ESP Evaluation

We perform our simulated evaluations on a range of network topologies G_Q [26], including a set of $N \times M$ Manhattan grids (2×2 , 2×3 , 3×3), complete graphs (K_2 , K_3 , K_4 , $K_{2,3}$), star graphs ($K_{1,3}$, $K_{1,4}$), friendship graphs (F_2 , F_3), and wheel and cyclic graphs (W_6 , C_5 , C_6). These topologies are illustrated in Fig 5. For each network topology, we also evaluate the impact of entanglement fidelity F , and obtain results from the range $[0.50, 0.90]$, with 0.05 increments. In this evaluation, we compare the performance of ESP against the baseline DQP. Simulations are performed on a distributed computing cluster (ECS Grid) for total 368 simulated minutes (≈ 4.42 billion physical layer cycles), across 1288 wall-clock hours. Table V and Fig. 6 show an overview of these results across the different topologies for $F = 0.75$. An example time series graph for a $t = 120$ network simulation on a 2×2 grid G_Q is depicted in Fig. 7.

D. Impact of Network Size - G_Q

Fig. 8 illustrates how the average request latency \bar{L} of ESP compares to that of DQP as the number of edges $|E_Q|$ in the network increases. As the network size grows, the total number of requests entering the system also increases, since the arrival rate λ depends on each edge $(i, j) \in E_Q$. In DQP, the distributed queue processes requests only from the front of the queue and does not allow requests that could execute without contention to proceed. In contrast, ESP does not use a shared queue, which alleviates this bottleneck. Each node in ESP

Topology	Protocol	Mean Latency (ms)	Mean Jitter (ms)
$K_{1,3}$	DQP	29.48	20.28
	ESP	29.89	20.92
$K_{1,4}$	DQP	39.48	26.12
	ESP	39.50	25.79
$K_{2,3}$	DQP	66.31	34.57
	ESP	34.51	24.56
K_2	DQP	18.82	13.89
	ESP	18.75	13.47
K_3	DQP	29.31	20.32
	ESP	28.57	20.58
K_4	DQP	67.24	35.62
	ESP	44.80	30.57
2×2	DQP	39.19	25.50
	ESP	27.83	20.28
2×3	DQP	84.61	39.49
	ESP	33.03	24.32
3×3	DQP	180.11	48.86
	ESP	39.77	29.69
C_5	DQP	50.47	29.48
	ESP	27.67	19.78
C_6	DQP	65.91	33.95
	ESP	28.08	20.26
F_2	DQP	64.15	32.95
	ESP	38.11	28.04
F_3	DQP	119.40	43.39
	ESP	53.84	41.40
W_6	DQP	138.70	44.33
	ESP	54.34	38.97

TABLE V: Comparison of ESP and DQP across various G_Q topologies (Fidelity = 0.75), measured in milliseconds.Fig. 6: Scaled network throughput T (entanglements / s) for each graph (Fidelity = 0.75), comparing DQP and ESP.

can independently execute requests if they are synchronized, allowing for concurrent entanglement generation across the overall network. The slope values for DQP and ESP are $m = 15.32$ and $m = 2.55$, respectively, with R^2 Pearson coefficients of 0.96 and 0.54. This indicates that ESP scales more efficiently than DQP, maintaining significantly reduced ($\sim 6 \times$ lower) latency values as the network expands.

Fig. 9 shows the relationship between average request jitter

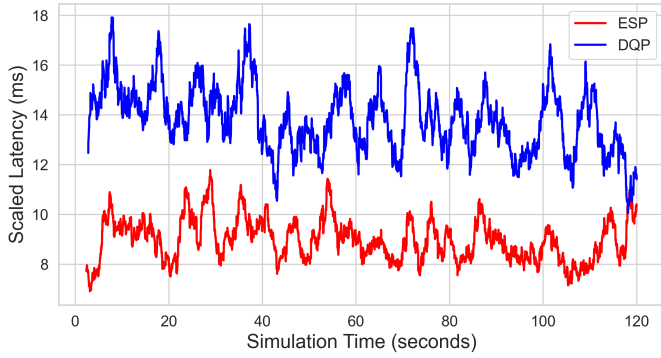


Fig. 7: Time series of scaled latency SL (ms) over a simulation period of $t = 120$ seconds for the 2×2 G_Q topology. A rolling average with $k = 256$ is applied due to the volume of requests processed.




Fig. 8: Relationship between the total number of edges $|E_Q|$ in G_Q and the mean request latency \bar{L} , with dashed regression lines for each protocol (Fidelity = 0.75).

\bar{J} for ESP and DQP as the network size increases. Given the probabilistic nature of entanglement generation, some level of constant jitter is anticipated. We can see this through the intercept values for ESP and DQP, $c = 12.46$ and $c = 13.77$ respectively, which are similar due to this inherent jitter. The results indicate that while mean jitter increases with network size for both ESP and DQP, with slope values of $m = 3.34$ and 2.02 respectively, ESP exhibits a slightly smaller rate of increase. This is likely due to ESP's inherently lower mean request latency, which contributes to more stable performance and reduced jitter as the network expands. However, as both m values are relatively similar and smaller in magnitude compared to c , we conclude the majority of the jitter still originates from the entanglement generation process and not the control plane protocol.

E. Impact of Matching Number of Network - $\nu(G_Q)$

In graph theory, a *matching* is a set of edges in a graph where no two edges share a common vertex. For a quantum network, a matching represents a set of simultaneous requests that can be executed without conflicts. The *matching number* of a graph, denoted $\nu(G_Q)$, is defined as the size of the largest possible matching in the graph G_Q . For example, in the graph shown in Fig. 4 (3×3), the set of edges performing




Fig. 9: Relationship between the total number of edges $|E_Q|$ in G_Q and the mean request jitter \bar{J} , with dashed regression lines for each protocol (Fidelity = 0.75).

entanglement generation, $\{(E, F), (G, H)\}$, represents a valid matching of G_Q , where $\nu(G_Q) = 4$. As $\nu(G_Q)$ increases, ESP can potentially exploit this by handling more concurrent requests within the network. Fig. 10 illustrates that ESP does not show substantial benefits in networks with $\nu(G_Q) = 1$ and may even suffer from performance degradation due to additional overhead ($K_2, K_3, K_{1,3}, K_{1,4}$). This behaviour is observed in smaller graph topologies and star graphs. However, for networks where $\nu(G_Q) > 1$, ESP demonstrates significant reductions in latency over DQP by efficiently allocating concurrent requests.

F. Impact of Entanglement Fidelity F

As shown in Fig. 2, higher entanglement fidelity of a request reduces $p_{success}$ at the physical layer, leading to an increase in busy time. This means each entanglement takes longer to generate, spreading a greater load through the network as queues take longer to process. Fig. 11 demonstrates this effect, illustrating the range of request latency between $F = 0.5$ and $F = 0.9$. In smaller networks, the impact on \bar{L} is less noticeable due to minimal time spent queuing. For instance, in graphs K_2, K_3 , and $K_{1,3}$, the ranges are similar between DQP and ESP. However, in larger graphs, the effect is significantly amplified, and increased entanglement fidelity results in larger increases in \bar{L} . Results show that ESP outperforms DQP in managing the effects of increased fidelity requests within the overall network by producing a tighter range of request latencies.

VII. CONCLUSIONS AND FUTURE WORK

In this paper, we developed the Eventual Synchronization Protocol (ESP), designed for managing heralded entanglement requests in a multi-node quantum network. To the best of our knowledge, this is the first decentralized control-plane protocol for the quantum link layer. Our protocol operates in a distributed manner, enabling concurrent entanglement generation across the network. This approach results in a substantial increase in system throughput compared to our baseline protocol (DQP), which relies on a centralised scheduler. The advantages of our protocol are particularly notable in larger network topologies, where it effectively reduces the overall network impact associated with high-fidelity requests.

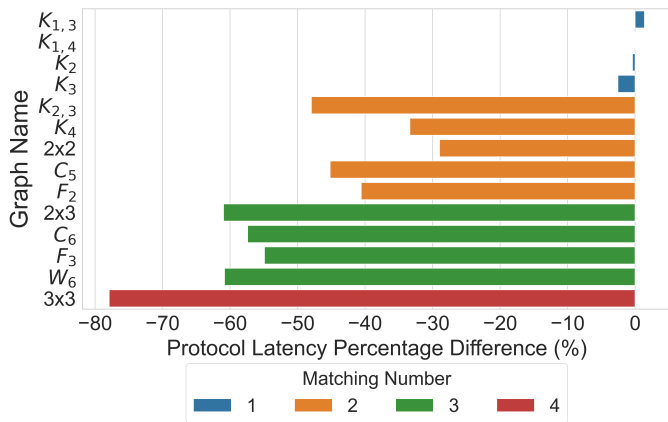


Fig. 10: Illustrative results showing the percentage differences in mean request latency \bar{L} (ms) between ESP and DQP, categorized by the quantum network's matching number $\nu(G_Q)$.

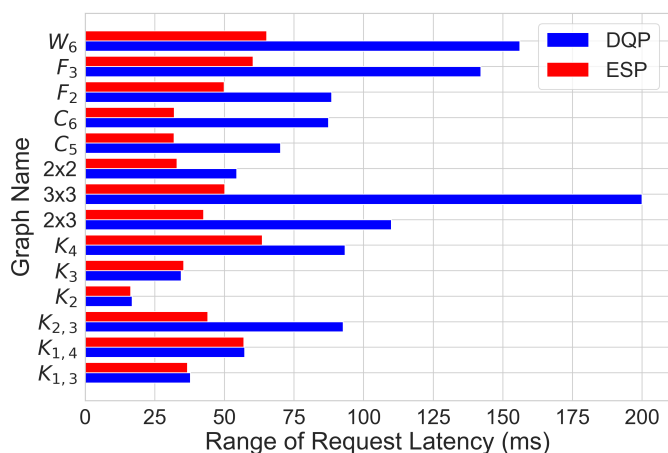


Fig. 11: Bar chart illustrating the range of request latency (ms) from $F = [0.5, 0.9]$ for each graph, comparing DQP and ESP.

We observe a sixfold reduction in average request latency growth as the number of quantum network links increases.

Future research will focus on improving the robustness of the protocol. This will involve investigating failure modes, addressing unstable network conditions, and validating certain distributed properties. Additionally, we plan to explore its potential as a resource reservation protocol, aiming to further optimize quantum network performance and the efficient management of entanglement resources.

ACKNOWLEDGMENT

This work is supported by the Royal Society of New Zealand through a James Cook Research Fellowship.

REFERENCES

[1] C. H. Bennett and G. Brassard, "Quantum cryptography: Public key distribution and coin tossing," *Theoretical Computer Science*, vol. 560, p. 7–11, Dec. 2014.
 [2] M. Caleffi *et al.*, "Distributed Quantum Computing: a Survey," 2022.

[3] A. Einstein, M. Born, and H. Born, *The Born-Einstein Letters: Correspondence Between Albert Einstein and Max and Hedwig Born from 1916-1955, with Commentaries by Max Born*. Macmillan, 1971.
 [4] W. Kozłowski *et al.*, "Architectural Principles for a Quantum Internet," RFC 9340, Mar. 2023.
 [5] W. K. Wootters and W. H. Zurek, "A single quantum cannot be cloned," *Nature*, vol. 299, no. 5886, pp. 802–803, 1982. [Online]. Available: <https://doi.org/10.1038/299802a0>
 [6] A. Dahlberg *et al.*, "A link layer protocol for quantum networks," in *Proceedings of the ACM Special Interest Group on Data Communication*, ser. SIGCOMM '19. ACM, Aug. 2019.
 [7] J. Illiano, M. Caleffi, A. Manzalini, and A. S. Cacciapuoti, "Quantum internet protocol stack: A comprehensive survey," *Computer Networks*, vol. 213, p. 109092, Aug. 2022.
 [8] A. Pirker and W. Dür, "A quantum network stack and protocols for reliable entanglement-based networks," *New Journal of Physics*, vol. 21, no. 3, p. 033003, Mar. 2019.
 [9] R. Van Meter, T. Ladd, W. Munro, and K. Nemoto, "System design for a long-line quantum repeater," *IEEE/ACM Transactions on Networking*, vol. 17, no. 3, p. 1002–1013, Jun. 2009.
 [10] J.-L. Liu *et al.*, "Creation of memory–memory entanglement in a metropolitan quantum network," *Nature*, vol. 629, no. 8012, pp. 579–585, May 2024.
 [11] S. L. N. Hermans, M. Pompili, H. K. C. Beukers, S. Baier, J. Borregaard, and R. Hanson, "Qubit teleportation between non-neighbouring nodes in a quantum network," *Nature*, vol. 605, no. 7911, p. 663–668, May 2022.
 [12] M. Pompili *et al.*, "Realization of a multinode quantum network of remote solid-state qubits," *Science*, vol. 372, no. 6539, p. 259–264, Apr. 2021.
 [13] H. Zarrabi, S. Hajihosseini, M. Fardmanesh, and S. I. Mirzaei, "Case study of decoherence times of transmon qubit," 2023.
 [14] S. L. N. Hermans, M. Pompili, L. D. Santos Martins, A. R-P Montblanch, H. K. C. Beukers, S. Baier, J. Borregaard, and R. Hanson, "Entangling remote qubits using the single-photon protocol: an in-depth theoretical and experimental study," *New Journal of Physics*, vol. 25, no. 1, p. 013011, Jan. 2023.
 [15] H. Bernien *et al.*, "Heralded entanglement between solid-state qubits separated by three metres," *Nature*, vol. 497, no. 7447, p. 86–90, Apr. 2013.
 [16] T. Matsuo, C. Durand, and R. Van Meter, "Quantum link bootstrapping using a ruleset-based communication protocol," *Physical Review A*, vol. 100, no. 5, Nov. 2019.
 [17] M. Skrzypczyk and S. Wehner, "An architecture for meeting quality-of-service requirements in multi-user quantum networks," 2021.
 [18] M. Pompili *et al.*, "Experimental demonstration of entanglement delivery using a quantum network stack," *npj Quantum Information*, vol. 8, no. 1, Oct. 2022.
 [19] C. A. Waldspurger, "Lottery and stride scheduling: Flexible proportional-share resource management," USA, Tech. Rep., 1995.
 [20] M. Thombare, R. Sukhwani, P. Shah, S. Chaudhari, and P. Raundale, "Efficient implementation of multilevel feedback queue scheduling," in *Proceedings of the International Conference on Wireless Communications, Signal Processing and Networking (WiSPNET)*, 2016, pp. 1950–1954.
 [21] H. Jula, D. Tralamazza, C. Zamfir, and G. Candea, "Deadlock immunity: enabling systems to defend against deadlocks," in *Proceedings of USENIX OSDI*, ser. OSDI'08, 2008, p. 295–308.
 [22] G. Ricart and A. K. Agrawala, "An optimal algorithm for mutual exclusion in computer networks," *Commun. ACM*, vol. 24, no. 1, p. 9–17, jan 1981. [Online]. Available: <https://doi.org/10.1145/358527.358537>
 [23] L. Lamport, "Time, clocks, and the ordering of events in a distributed system," in *Concurrency: the Works of Leslie Lamport*. Association for Computing Machinery, 2019, pp. 179–196.
 [24] T. Coopmans *et al.*, "Netsquid, a network simulator for quantum information using discrete events," *Communications Physics*, vol. 4, no. 1, Jul. 2021.
 [25] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson, "On the self-similar nature of ethernet traffic," *SIGCOMM Comput. Commun. Rev.*, vol. 23, no. 4, p. 183–193, oct 1993.
 [26] D. Gries and F. B. Schneider, *A Theory of Graphs*. New York, NY: Springer New York, 1993, pp. 423–460. [Online]. Available: https://doi.org/10.1007/978-1-4757-3837-7_20