On the Specification and Analysis of Quantum Repeater Protocols

Lorenzo La Corte*, Anita Buckley*, Pavel Chuprikov† and Patrick Eugster*
* *Computer Systems Institute, Università della Svizzera italiana (USI), Lugano, Switzerland*
† *LTCI, Télécom Paris, Institut Polytechnique de Paris, Palaiseau, France*

***Abstract*—The distribution of high-quality entanglement over long distances remains a challenge due to losses in quantum channels. As a solution, quantum repeater protocols generate entanglement over short segments and progressively extend its range using entanglement swapping, while using distillation to compensate for entanglement degradation. Numerous studies analyze repeater protocols through tailored tools that are tied to specific protocol types and network assumptions. In this work, we propose a formal language to specify quantum repeater protocols in a unified manner, modeling the concurrent distribution of entanglement across multiple pairs of network nodes. We consider delays arising from failures and contention for resources, and account for the effects of noise and decoherence on the quality of entanglement. Furthermore, we provide a prototype tool based on our language to precisely track the rate and fidelity of the distributed entangled pairs, supporting heterogeneous hardware configurations and constraints on available quantum memories. Finally, we demonstrate the effectiveness of our approach by expressing common protocols from the literature and evaluating their performance.**

***Index Terms*—Quantum networks, entanglement distribution, protocol specification, performance evaluation**

I. Introduction

Quantum networks leverage entanglement to enable distributed quantum computation and secure communication [1]. They bring numerous advantages over classical networks, enhancing the capabilities of existing applications and enabling new ones, e.g., multi-party computation [2], blind quantum computing [3], and quantum key distribution [4]. The crucial service provided by quantum networks is the distribution of high-quality entangled pairs of qubits. However, this is a challenging task due to exponential losses in optical fibers, which preclude the direct transmission of quantum information over long distances [5], [6]. Furthermore, the no-cloning theorem implies the impossibility of copying unknown quantum states [7], which prevents the use of classical "store-and-forward" techniques to distribute entanglement. To tackle this problem, quantum repeater protocols distribute entanglement by leveraging *quantum repeaters* interposed between the network end-nodes. A protocol starts with entangled pairs being generated at short distances between end-nodes and repeaters; then, repeaters leverage their specialized hardware to perform entanglement swapping, which consumes two adjacent links to produce a longer link of reduced *fidelity* (quality of the entangled state). To maintain high-quality entanglement, different approaches exist. First-generation repeater networks use probabilistic entanglement distillation [8], which consumes two entangled pairs to produce a new pair of higher fidelity. Later generations employ quantum error correction (QEC) to address operational errors and losses [9]. Furthermore, several works [10]–[12] study cutoff conditions, which impose thresholds on the quality of the pairs. This work focuses on probabilistic protocols (without QEC) combining operations like entanglement generation, swapping, and distillation, with cutoff conditions on the quality of entanglement.

Due to the nature of quantum mechanics, it is challenging to design probabilistic protocols that distribute high-fidelity pairs at sufficient rates. In particular, all these operations are stochastic heralded primitives, meaning that they have a high probability of failure and rely on classical signals to acknowledge their outcome. Delays introduced by failures and heralding are compounded by the nature of qubits, which are scarce and short-lived resources, exponentially decohering in quantum memories with time.

For these reasons, many efforts have been devoted to extracting insights on the optimal design of quantum repeater protocols. To evaluate their performance, tools from the literature leverage simulations, analytical expressions, or probabilistic semantics. Although flexible and easy to implement, numerical simulation tools [13], [14] can be computationally expensive and it can be hard to estimate confidence intervals. Therefore, techniques have been proposed to optimize simulations; for instance, Iñesta *et al.* [11] formulate entanglement distribution as a Markov decision process (MDP), to find optimal strategies for designing protocols; however, their study is limited to chains of nodes with identical (*homogeneous*) hardware. Similarly, Haldar *et al.* [12] combine MDP with reinforcement learning to discover improved entanglement distribution strategies, considering both homogeneous and heterogeneous chains. Both works include cutoffs, but neither considers entanglement distillation nor non-linear topologies. As an alternative to simulations, several analytical and numerical models have been proposed in the literature. Li *et al.* [10] develop closed-form expressions and an efficient algorithm

Work partially supported by Hasler Foundation grant 2024-09-27-175 and ANR-22-EXES-0013 grant (project STeP2/F4C).

© 2026 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.
Published version: https://doi.org/10.1109/QCNC69040.2026.00042

to evaluate nested (*doubling*) protocols, focusing on homogeneous repeater chains. On top of that, La Corte *et al.* [15] consider heterogeneous chains and other swapping schemes. However, their framework is also restricted to linear topologies and does not express *reactive* strategies, such as swap-as-soon-as-possible (SWAP-ASAP) — described, for example, by the analytical expressions devised by Goodenough *et al.* [16].

Orthogonal to these works, Buckley *et al.* design the PBKAT [17] language to specify quantum repeater protocols. Although general, PBKAT does not implement a physical model for noise and decoherence, nor does it evaluate performance metrics, i.e., quality of service (QoS).

Building upon PBKAT, this work proposes a framework to faithfully express quantum repeater protocols, and to track their performance in terms of the number of entangled pairs distributed per second (*rate*) and their quality (*fidelity*).

Our contributions are the following. First, we develop an algebraic language to formally specify repeater protocols, considering arbitrary numbers of nodes, topologies, and hardware constraints. We consider protocols capable of *concurrently* distributing entangled pairs to different locations. Second, based on the semantics of the language, we provide a high-precision numerical evaluation tool that, given a protocol and a set of constraints, computes the rate and fidelity of entanglement distribution. Lastly, we showcase the practical usefulness of our framework, expressing common protocols from the literature, validating our tool against analytical results, and evaluating a concurrent protocol on a star topology.

The rest of the paper is organized as follows. Section II introduces quantum repeater protocols, the physical model considered, and how protocols are evaluated in light of probabilistic behaviors and non-determinism. Section III formalizes our language, detailing the abstractions that model the behavior of protocols. Section IV showcases the expressiveness of our language by describing common swapping schemes from the literature. Section V presents numerical evaluations, validating our tool and using it to analyze resource contention in a protocol that concurrently distributes entanglement at multiple locations. Section VI concludes the paper with outlook.

II. Preliminaries

This work considers arbitrary topologies (represented as graphs, such as the 4-node chain and star in Figure 1) of heterogeneous quantum networks — where all nodes and channels may have different hardware capabilities.

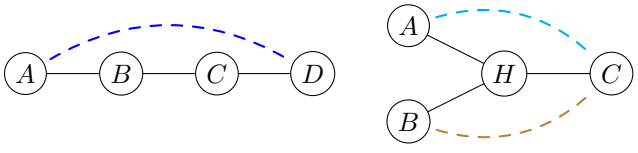


Fig. 1. Different 4-node network topologies, a chain (left) and a star (right). Nodes are connected by quantum channels, drawn as black solid lines. Entangled pairs at end-nodes are represented as colored dashed lines.

A. Repeater Protocols for Entanglement Distribution

We divide a protocol into rounds, i.e., discrete time windows in which network operations (e.g., generation, swapping, and distillation) run synchronously [18]. These operations, which we call *actions*, can be composed in parallel within the same round or in sequence through different rounds. Allowing for actions to run in parallel is crucial, as network operations are commonly executed simultaneously to tame decoherence. We account for the time it takes for an action to complete, also considering the time to classically herald its outcome.

At the start of a round, actions operate on entangled pairs present in the network, acquiring and consuming available pairs to produce new ones. If the set of pairs required to execute an action is not available, the action cannot be executed and no resources are consumed. If an action executes but fails, the acquired pairs are still consumed without producing anything. Thus, we allow for any action consuming the set of required pairs to probabilistically output new or no entangled pairs. Furthermore, we allow for protocols to repeatedly attempt actions, which is a common strategy to increase the probability of distributing entanglement [18].

We consider possible competition (*contention*) for resources, emerging when actions running in parallel try to acquire the same entangled pairs. Typically, contention is resolved by using pre-determined strategies that define the scheduling order of actions (e.g., choosing at random). Following the approach of Buckley *et al.* [17], we model the order of execution of actions non-deterministically — by leaving non-determinism unresolved, we track all the possible choices of which action has priority over the contended resources.

Non-deterministic behaviors are particularly relevant for complex topologies and protocols distributing more than one entangled pair concurrently, due to multiple user requests. Furthermore, contention is amplified when considering resource-constrained networks, e.g., where nodes have limited memory capacity; therefore, we also model upper bounds on the number of entangled pairs that network nodes can hold at the same time. These realistic aspects cannot be neglected when evaluating real-world quantum networks.

B. Physical Model

An entangled pair is represented as a Werner state [19]

$$\rho(w) = w\,|\Phi^+\rangle\langle\Phi^+| + (1-w)\frac{\mathbb{I}}{4}, \tag{1}$$

where $|\Phi^+\rangle$ is the pure state and $\mathbb{I}/4$ is the maximally mixed two-qubit state. The Werner parameter $w \in [0,1]$ indicates the fidelity of ρ to the pure state $|\Phi^+\rangle$, since $F = (1+3w)/4$.

We consider networks of end-nodes and repeaters, which are uniquely identified by labels (e.g., A or A^1). Thus, we represent an entangled pair between nodes A and B as the unordered pair $A{\sim}B$.

Nodes in the network can be connected through quantum channels $i \in [1,2,\ldots,n]$ of possibly heterogeneous (integer) lengths L_i. Accordingly, we define the unit of time as

$$t_{\text{unit}} = \frac{L_0}{c} \quad \text{with} \quad L_0 = \gcd(\{L_1, L_2, \ldots, L_n\}), \tag{2}$$

where c is the speed of light and L_0 is the greatest common divisor of the elements of the channel lengths set. Note that in our model, the notion of the time unit does not coincide with the notion of a round, e.g., two actions running in parallel within a round may take different time to complete.

When generated, an entangled pair has an initial Werner parameter w_0 representing its quality. Entangled pairs along different channels i may be generated with different probabilities $p_{\text{ge}}^{(i)}$ and initial Werner parameter $w_0^{(i)}$. Whenever an entangled pair is stored in imperfect memories of nodes A and B for time t, its Werner parameter w gets reduced:

$$w' = w\,\gamma^{AB}(t)\,, \tag{3}$$

where the *decoherence factor* γ is defined as

$$\gamma^{AB}(t) = \exp\left(-t/t_{\text{coh}}^{A}\right)\ \exp\left(-t/t_{\text{coh}}^{B}\right), \tag{4}$$

with t_{coh} being the coherence times of the involved memories.

For entanglement swapping, different repeater nodes may have different probabilities of success. The Werner parameter of the pair in the output depends on the initial qualities w_1 and w_2 of the input pairs and on the times they wait in memory:

$$w_{\text{sw}} = w_1' w_2', \tag{5}$$

where w' reflects the decoherence of pairs as in (3).

For entanglement distillation, we consider the BBPSW protocol [8], where both the success probability and the output Werner parameter depend on the quality and waiting times of the input pairs (decohering with time as above) [20]:

$$p_{\text{di}} = \frac{1 + w_1' w_2'}{2} \quad \text{and} \quad w_{\text{di}} = \frac{w_1' + w_2' + 4w_1' w_2'}{6p_{\text{di}}}. \tag{6}$$

C. Tracking the Quantitative Behavior of Protocols

Through the following example, we introduce key language components that allow us to track the execution of a protocol and compute its output — a probability distribution over the multisets of entangled pairs that the protocol may produce.

We consider the star network in Figure 1 with nodes A, B, C and a center node (hub) H, and describe a protocol that aims to concurrently generate end-to-end pairs $A{\sim}C$ and $B{\sim}C$. For simplicity of presentation, Figure 2 only captures the probabilities (without timestamps and Werner parameters) of producing $A{\sim}C$ and $B{\sim}C$ within one iteration of the protocol.

In round 1, each pair $A{\sim}H$, $B{\sim}H$, and $C{\sim}H$ is generated with equal probability $p_{\text{ge}} = \frac{1}{3}$. The output is a probability distribution over the 8 possible *multisets* of the generated pairs:

$$\tfrac{8}{27}\cdot\emptyset + \tfrac{4}{27}\cdot\{\!\{A{\sim}H\}\!\} + \cdots + \tfrac{1}{27}\cdot\{\!\{A{\sim}H, B{\sim}H, C{\sim}H\}\!\}$$

Round 2 depends on the output of round 1. For instance, when only $A \sim H$ and $C \sim H$ are generated in round 1, they are swapped in round 2 into $A \sim C$ with probability $p_{\text{sw}}^{H} = \frac{2}{5}$. On the other hand, when all $A \sim H$, $B \sim H$, and $C \sim H$ are successfully generated in round 1, both swap actions are attempted in parallel: one requires $A{\sim}H$ and $C{\sim}H$, while the other requires $B{\sim}H$ and $C{\sim}H$. Since only one copy of $C \sim H$ is available to both actions, this leads to contention,

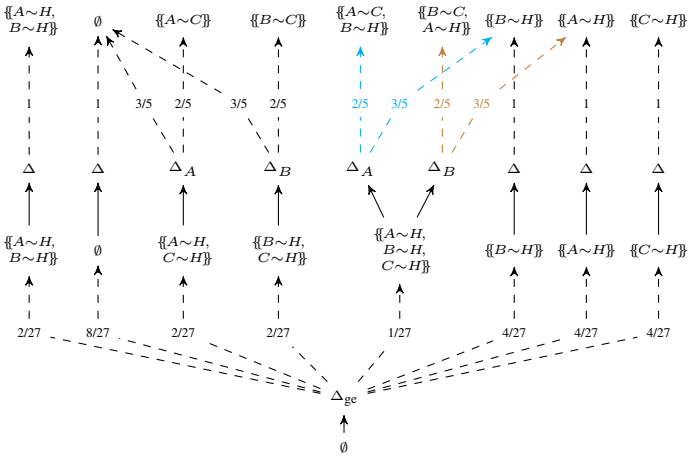


Fig. 2. Execution trace of a protocol distributing end-to-end entangled pairs $A{\sim}C$ and $B{\sim}C$ in the 4-node star network of Figure 1. Non-deterministic choices for the prioritized execution of swaps involving either $A{\sim}H$ or $B{\sim}H$ are shown respectively in cyan and brown.

which we model as non-deterministic choices (shown in cyan and brown in Figure 2) on which swap has priority. Thus, the output of round 2 is a *convex set* containing two distributions over the possible output entangled pairs [17], one where the swap producing $A{\sim}C$ was attempted, and one where the swap producing $B{\sim}C$ was attempted.

In Section III, we build upon the above protocol example by augmenting the notion of entangled pairs (and multisets) with their timestamp of production and Werner parameter, which allows us to quantitatively track the QoS of protocols. Following the guidelines of the Quantum Internet Research Group [21], our work targets two figures of merit for the evaluation: rate, i.e., number of entangled pairs distributed per second, and fidelity, i.e., their average quality — which can be computed from the augmented output of a protocol.

III. Protocol Specification

This section introduces the syntax for protocol specification (cf. Figure 3) and semantics that model protocol execution in quantitative terms. As a first step, we formalize the abstractions needed to express common network protocols.

A. Network State

An entangled pair may have additional classical data associated to it. The notion of entangled pairs ($ep \in \mathsf{EP}$) introduced so far only carries the location of qubits (e.g., $A{\sim}B$); however, this makes multiple pairs produced at the same end-nodes indistinguishable. To enable reasoning about QoS, we must keep track of an entangled pair's timestamp t and fidelity w at the time of production. A *tagged entangled pair* $A{\sim}B_{(t,w)}$ ($ep^{\dagger} \in \mathsf{EPt}$) carries these quantitative metrics.

Following the work of Buckley et al. [17], we define *network state* $a \in \mathcal{M}(\mathsf{EP})$ to be a multiset of entangled pairs present in the network, e.g., $a = \{\!\{A \sim B, B \sim C\}\!\}$. In order to properly track the evolution of network states over time, we also annotate multisets of entangled pairs with their current *clock* information, $\{\!\{-\}\!\}_T$. For example, we write a

Network objects

Nodes	$\mathsf{N} ::=$	$A, B, C, \ldots$
Entangled pairs	$\mathsf{EP} \ni ep ::=$	$\mathsf{N}{\sim}\mathsf{N}$
Time units	$\mathbb{N} \ni t, T ::=$	$1, 2, \ldots$
Werner parameter	$[0,1] \ni w ::=$	$w_1, w_2, \ldots$
Tagged pairs	$\mathsf{EPt} \ni ep^\dagger ::=$	$\mathsf{N}{\sim}\mathsf{N}_{(t,w)}$
Multisets	$\mathcal{M}(\mathsf{EP}) \ni a ::=$	$\{\!\{ep_1, \ldots, ep_k\}\!\}$
	multiset union	$a \uplus b$
Tagged multisets	$\mathcal{M}\mathsf{t}(\mathsf{EPt}) \ni a^\dagger ::=$	$\{\!\{ep_1^\dagger, \ldots, ep_k^\dagger\}\!\}_T$
	tagged multiset union	$a^\dagger \uplus b^\dagger$

Guards & expressions

$$\mathsf{BExp} \ni \alpha, \beta ::= 0 \mid 1 \mid a \mid \alpha \wedge \beta \mid \alpha \vee \beta \mid \bar{\alpha}$$

$$\mathsf{Exp} \ni e, f ::= \pi \triangleq r \triangleright (\Sigma\, p{\cdot}o^\dagger) \mid \alpha \mid e\,;\, f \mid e \parallel f \mid e \circ f \mid e +_\alpha f \mid e^{(\alpha)}$$

0	*false / abort*
1	*true / skip*
a	*atomic guard*
$\wedge, \vee$	*and, or*
$\bar{\alpha}$	*not*
$;$	*sequential composition*
$\parallel$	*parallel composition*
$\circ$	*ordered composition*
$+_\alpha$	*if α then e else f*
$e^{(\alpha)}$	*while α do e*

Fig. 3. Syntax. An action $\pi \triangleq r \triangleright (\Sigma\, p{\cdot}o^\dagger)$ has two parameters, $r \in \mathcal{M}(\mathsf{EP})$ is a multiset containing the required entangled pairs and $\Sigma\, p{\cdot}o^\dagger$ is a probability distribution over tagged output multisets of entangled pairs. When expressing protocols we use concrete guards, e.g., $\exists_{A\sim B}$ checks the presence of $A{\sim}B$.

tagged multiset as $a^\dagger = \{\!\{A{\sim}B_{(t_1,w_1)}, B{\sim}C_{(t_2,w_2)}\}\!\}_T$, with $a^\dagger \in \mathcal{M}\mathsf{t}(\mathsf{EPt})$. (We write t for pairs' timestamps and T for multisets' clocks. For details on how pairs timestamps and multisets' clocks are updated at run-time, see Section III-C.)

B. Actions

Actions are abstractions for quantum network operations. We think of actions as protocol primitives for manipulating entangled pairs. Semantically, we represent an action π as a function of the form $\pi \triangleq r \triangleright (\Sigma\, p{\cdot}o^\dagger)$, described as follows. To execute, an action π requires a set r of entangled pairs. We use r instead of $r^\dagger$ as only the location of pairs matters, e.g., distillation requires two copies of the same pair ep, regardless of their tags. When applied to a multiset $a^\dagger$, the action requires entangled pairs $r \subseteq a$, and produces a multiset $o^\dagger$ with probability p. Thus, the probability distribution $\Sigma\, p \cdot o^\dagger$ ranges over the possible outputs. Based on the required multiset r, entangled pairs $r^\dagger \subseteq a^\dagger$ are consumed by the action. Importantly, the probabilities p and the tags of $o^\dagger$ may depend on the tags of $r^\dagger$. In the case when there are insufficient pairs in the input, i.e., $r \not\subseteq a$, the action is not executed.

For example, assume that the star network in Figure 1 is at time $T = 2$ in the state $a^\dagger = \{\!\{A{\sim}H_{(1,0.9)}\}\!\}_2$, meaning that it only contains $A{\sim}H$ generated at time 1 with the initial Werner parameter $w_0 = 0.9$. If, in the next round, we attempt to probabilistically generate $B{\sim}H$, this action may either succeed producing the network state $\{\!\{A{\sim}H_{(1,0.9)}, B{\sim}H_{(3,0.9)}\}\!\}_3$ or fail yielding $\{\!\{A{\sim}H_{(1,0.9)}\}\!\}_3$. (Notice that the time taken by the failed attempt is still accounted for in the clock). Thus, each outcome occurs with a probability and advances the network state's clock according to a well-defined rule. In what follows, we formalize these aspects in order to capture the behavior of the most common actions from the literature.

An action π takes time t_{c}^π to complete, which accounts for local operations and classical communication of measurement outcomes and heralding. The completion time depends on the action and the distance between the involved nodes, e.g., $t_{\mathrm{c}}^{\mathrm{ge}\langle A\sim B\rangle}$ and $t_{\mathrm{c}}^{\mathrm{ge}\langle B\sim D\rangle}$ may differ. We omit the superscript when actions can be inferred from the context. Thus, each action π may produce entangled pairs at time t, with

$$t = T' = T + t_{\mathrm{c}}^\pi\,, \tag{7}$$

where T is the clock's time when the action is invoked (at the start of the round), and T' is the updated clock when the action completes (even in case of failure).

a) Entanglement generation: We model the generation of $A{\sim}B$ through photon transmission as an action,

$$\mathsf{ge}\langle A{\sim}B\rangle_{p_{\mathrm{ge}},w_0} \triangleq \emptyset_T \triangleright \left(p_{\mathrm{ge}}{\cdot}\{\!\{A{\sim}B_{(t,w_0)}\}\!\}_{T'} + (1-p_{\mathrm{ge}}){\cdot}\emptyset_{T'}\right)$$

where p_{ge} is the probability of successfully generating $A{\sim}B$ and w_0 is the initial pair quality (Werner parameter). The pair is produced at time t, defined in (7), with the completion time of generation depending on the distance between A and B, accounted in terms of the units t_{unit} and L_0 (see Section II-B).

b) Entanglement swapping: We model a swap as an action producing $A{\sim}C$ by consuming $A{\sim}B$ and $B{\sim}C$,

$$\mathsf{sw}\langle A{\sim}C \mathbin{@} B\rangle_{p_{\mathrm{sw}}^B,t_{\mathrm{coh}}^A,t_{\mathrm{coh}}^B,t_{\mathrm{coh}}^C} \triangleq \{\!\{A{\sim}B_{(t_1,w_1)}, B{\sim}C_{(t_2,w_2)}\}\!\}_T \triangleright p_{\mathrm{sw}}^B{\cdot}\{\!\{A{\sim}C_{(t,w)}\}\!\}_{T'} + (1-p_{\mathrm{sw}}^B){\cdot}\emptyset_{T'}$$

with the probability p_{sw}^B depending on the hardware at node B. The output Werner parameter w is computed as follows. The action starts at time T when $A{\sim}B$ and $B{\sim}C$ decohered for $(T-t_1)$ and $(T-t_2)$ time units, respectively; accordingly, the quality w_{sw} of $A{\sim}C$ is computed by (5). The Bell state measurement that forms $A{\sim}C$ is assumed to be instantaneous, and is followed by the classical heralding and transmission of measurement outcomes, taking $t_{\mathrm{c}}^{\mathrm{sw}}$ time units, which depends on the maximum between the channel lengths L_{AB} and L_{BC}. Accordingly, t and T' are obtained by (7). In addition to the decay of input pairs, the output pair $A{\sim}C$ also decoheres for $t_{\mathrm{c}}^{\mathrm{sw}}$ time units, resulting in the output quality:

$$w = w_1\, w_2\, \gamma^{AB}(T-t_1)\, \gamma^{BC}(T-t_2)\, \gamma^{AC}(t_{\mathrm{c}}^{\mathrm{sw}})\,.$$

c) Simultaneous entanglement swapping: We model an action producing the pair $A^1{\sim}A^k$ by simultaneously swapping $k-1$ pairs $A^1{\sim}A^2$, $A^2{\sim}A^3$, …, $A^{k-1}{\sim}A^k$:

$$\mathsf{sw}\langle A^1{\sim}A^k \mathbin{@} A^2, ..., A^{k-1}\rangle_{p_{\mathrm{sw}}^{A^2},\ldots,p_{\mathrm{sw}}^{A^{k-1}},t_{\mathrm{coh}}^{A^1},\ldots,t_{\mathrm{coh}}^{A^k}} \triangleq \{\!\{A^1{\sim}A^2{}_{(t_1,w_1)}, ..., A^{k-1}{\sim}A^k{}_{(t_{k-1},w_{k-1})}\}\!\}_T \triangleright p_{\mathrm{sw}}{\cdot}\{\!\{A^1{\sim}A^k{}_{(t,w)}\}\!\}_{T'} + (1-p_{\mathrm{sw}}){\cdot}\emptyset_{T'}$$

The action succeeds with probability $p_{\mathrm{sw}} = p_{\mathrm{sw}}^{A^2}{\cdot}p_{\mathrm{sw}}^{A^3}\cdots p_{\mathrm{sw}}^{A^{k-1}}$. The output Werner parameter w takes into account the decoherence of the respective entangled pairs, and is equal to:

$$w_1 \cdots w_{k-1}\gamma^{A^1A^2}(T-t_1)\cdots\gamma^{A^{k-1}A^k}(T-t_{k-1})\gamma^{A^1A^k}(t_{\mathrm{c}}^{\mathrm{sw}}),$$

where $t_{\mathrm{c}}^{\mathrm{sw}}$ is proportional to the total distance between the involved nodes, and is used to obtain t and T' in (7).

d) Entanglement distillation: We model distillation as an action consuming two entangled pairs at the same locations to probabilistically return one of higher quality or fail,

$$\mathsf{di}\langle A{\sim}B\rangle_{t^A_{\text{coh}},t^B_{\text{coh}}} \triangleq \{\!\!\{ A{\sim}B_{(t_1,w_1)},\ A{\sim}B_{(t_2,w_2)} \}\!\!\}_T \triangleright$$
$$p_{\text{di}}\cdot\{\!\!\{ A{\sim}B_{(t,w)} \}\!\!\}_{T'} + (1-p_{\text{di}})\cdot\emptyset_{T'}$$

with the success probability p_{di} depending on the input entangled pairs. In particular, we use (6) to compute p_{di} and w_{di} by plugging in the decayed Werner parameters:

$$w_1' = w_1\gamma(T-t_1) \qquad \text{and} \qquad w_2' = w_2\gamma(T-t_2)\,.$$

Furthermore, we consider ($t_{\text{c}}^{\text{di}} = L^{AB}/c$) as the time for which the distilled pair $A{\sim}B$ decoheres, used to compute t by (7) and the final output quality $w = w_{\text{di}}\gamma(t_{\text{c}}^{\text{di}})$.

C. Protocols

To additionally control the execution of actions, protocols can include *guards* testing for available entangled pairs, e.g., to guide their allocation to actions that require them. In our language, guards are predicates over multisets of entangled pairs. Guards 0 and 1 are "always false" and "always true" predicates, respectively. We write protocols as expressions consisting of actions and guards, combined by *conditionals* and *loops*, as well as *operators* for sequential, parallel, and ordered composition. For example, the conditional "generate $A{\sim}B$ if not yet present, skip otherwise" is expressed as $\mathsf{ge}\langle A{\sim}B\rangle +_{\nexists_{A\sim B}} 1$ with $\nexists_{A\sim B} = \overline{\{\!\!\{A{\sim}B\}\!\!\}}$.

Sequential composition $(e; f)$ models transitioning through rounds by first applying e to the input multiset of entangled pairs and then applying f to each multiset produced by e. Parallel and ordered composition govern the execution of actions that occur synchronously within a round. Parallel composition $(e \parallel f)$ allows running e and f in parallel, possibly triggering resource competition (as shown in Section II-C). By contrast, ordered composition $(e \circ f)$ imposes that e has priority over f in accessing the currently available entangled pairs.

The outputs of actions happening synchronously within a round are combined by multiset union and by multiplying their probabilities, as illustrated below. The clock accounts for the action taking the longest time in that round, also considering actions that executed but failed, i.e., the clock of the combined multiset $a^\dagger \uplus b^\dagger$ is $\max(T_{a^\dagger}, T_{b^\dagger})$. For example, consider the chain in Figure 1, and a protocol generating in parallel pairs $A{\sim}B$ and $B{\sim}D$ of quality w_0 with probabilities $1/4$ and $1/10$. The actions respectively take 1 and 2 time units to complete and produce outputs:

$$\tfrac{3}{4}\,\emptyset_1 + \tfrac{1}{4}\,\{\!\!\{A{\sim}B_{(1,w_0)}\}\!\!\}_1 \qquad \text{and} \qquad \tfrac{9}{10}\,\emptyset_2 + \tfrac{1}{10}\,\{\!\!\{B{\sim}D_{(2,w_0)}\}\!\!\}_2.$$

Thus, the output of $\mathsf{ge}\langle A{\sim}B\rangle \parallel \mathsf{ge}\langle B{\sim}D\rangle$ is:

$$\tfrac{27}{40}\,\emptyset_2 \;+\; \tfrac{9}{40}\,\{\!\!\{A{\sim}B_{(1,w_0)}\}\!\!\}_2 +$$
$$\tfrac{3}{40}\,\{\!\!\{B{\sim}D_{(2,w_0)}\}\!\!\}_2 \;+\; \tfrac{1}{40}\,\{\!\!\{A{\sim}B_{(1,w_0)}, B{\sim}D_{(2,w_0)}\}\!\!\}_2.$$

D. Network Constraints

The final output of each round is, in general, a probability distribution $\Sigma\, p\cdot o^\dagger$ over tagged network states. In practice, network states are further constrained by *network constraints*. In this work, we consider two constraints: (i) a *time-cutoff*, i.e., the maximum allowed age of entangled pairs; and (ii) a *network capacity*, i.e., the maximum number of entangled pairs that any two respective nodes can hold.

In what follows, we describe how constraints are enforced on a network state $o^\dagger$. We apply the time-cutoff by defining

$$\text{cut}_\tau(o^\dagger) \;=\; \{\!\!\{ A^i{\sim}A^j_{(t_{ij},w_{ij})} \in o^\dagger \;\mid\; (T-t_{ij}) \le \tau \}\!\!\}_T\,,$$

dropping all pairs with age exceeding the threshold τ. Then, we cap the resulting multiset by the network capacities m_{ij}, which bound the number of entangled pairs between each pair of nodes A^i and A^j that the network can handle. Let n_{ij} be the multiplicity of $A^i{\sim}A^j$ in o, disregarding the tags. We define the capacity capping operator

$$\lfloor o \rfloor \;=\; \{\!\!\{ \min\big(n_{ij}, m_{ij}\big) \times A^i{\sim}A^j \}\!\!\}\,,$$

dropping $n_{ij} - m_{ij}$ pairs exceeding capacity. In practice, we drop the worst pairs in terms of time and quality. Thus, we can combine the time-cutoff and capacity as $\lfloor\, \text{cut}_\tau(o^\dagger)\,\rfloor$.

Similarly, we can define a fidelity-cutoff, e.g., by changing the condition of $\text{cut}_\tau(o^\dagger)$ to be $w\,\gamma(T-t) \ge w_{\text{cut}}$.

IV. Case Studies

In this section, we demonstrate the expressivity of our language by expressing repeater protocols belonging to the most common swapping schemes from the literature, following the classification in the survey of Abane *et al.* [22, Sec. III.B].

We use while loops (see Figure 3) to represent that an action is attempted until it succeeds, e.g., $\mathsf{ge}\langle A{\sim}B\rangle^{(\nexists_{A\sim B})}$ means "while there is no $A{\sim}B$, keep generating". For readability, we omit the hardware parameters, e.g., p_{ge}, w_0, p_{sw}, t_{coh}.

a) Sequential: In this scheme, swap actions are performed in sequence, joining entangled pairs sequentially from left to right. Consider a protocol on the chain $A-B-C-D$, aiming to distribute the end-to-end entangled pair $A{\sim}D$:

$$\mathsf{p1} = (\mathsf{ge}\langle A{\sim}B\rangle \parallel \mathsf{ge}\langle B{\sim}C\rangle)^{(\nexists_{A\sim B}\vee\nexists_{B\sim C})}$$
$$\mathsf{p2} = \big(\mathsf{p1}\,;\, \mathsf{sw}\langle A{\sim}C \,@\, B\rangle\big)^{(\nexists_{A\sim C})}$$
$$\mathsf{p3} = \big((\mathsf{p2} \parallel \mathsf{ge}\langle C{\sim}D\rangle^{(\nexists_{C\sim D})})\,;\, \mathsf{sw}\langle A{\sim}D \,@\, C\rangle\big)^{(\nexists_{A\sim D})}$$

The negation of the predicate $(\nexists_{A\sim B} \vee \nexists_{B\sim C})$ is the exit condition for the while loop in p1, exiting once both pairs $A{\sim}B$ and $B{\sim}C$ are generated. Then, p2 attempts swapping until it produces $A{\sim}C$. The generation of $C{\sim}D$ happens in parallel with p2, followed by the final swap producing $A{\sim}D$.

b) Parallel: This scheme leverages the simultaneous swap action defined in Section III-B, which waits for all entangled pairs to be generated and then swaps them simultaneously. A protocol on the chain $A-B-C-D$, distributing the end-to-end entangled pair $A{\sim}D$, is specified as:

$$\mathsf{p} = \big((\mathsf{ge}\langle A{\sim}B\rangle \parallel \mathsf{ge}\langle B{\sim}C\rangle \parallel \mathsf{ge}\langle C{\sim}D\rangle)^{(\nexists_{A\sim B}\vee\nexists_{B\sim C}\vee\nexists_{C\sim D})}\,;$$
$$\mathsf{sw}\langle A{\sim}D \,@\, B, C\rangle\big)^{(\nexists_{A\sim D})}$$

c) Ad hoc: A swap-as-soon-as-possible protocol distributes $A{\sim}D$ along the chain $A-B-C-D$ by concurrently attempting all actions until $A{\sim}D$ is produced. The protocol can be written as follows:

$$\begin{aligned}\mathsf{p} = \big(&\mathsf{ge}\langle A{\sim}B\rangle \parallel \mathsf{ge}\langle B{\sim}C\rangle \parallel \mathsf{ge}\langle C{\sim}D\rangle \parallel \\ &\mathsf{sw}\langle A{\sim}C \circledast B\rangle \parallel \mathsf{sw}\langle B{\sim}D \circledast C\rangle \parallel \\ &\mathsf{sw}\langle A{\sim}D \circledast B\rangle \parallel \mathsf{sw}\langle A{\sim}D \circledast C\rangle\big)^{(\nexists_{A\sim D})}\end{aligned}$$

d) Doubling: In this scheme, nested protocols perform l levels of entanglement swapping on 2^l+1 nodes. A protocol acting on the chain $A-B-C-D-E$ to produce the end-to-end entangled pair $A{\sim}E$ is specified as:

$$\begin{aligned}\mathsf{p1} &= \big(\mathsf{ge}\langle A{\sim}B\rangle \parallel \mathsf{ge}\langle B{\sim}C\rangle\big)^{(\nexists_{A\sim B}\vee\nexists_{B\sim C})}\\ \mathsf{p2} &= \big(\mathsf{ge}\langle C{\sim}D\rangle \parallel \mathsf{ge}\langle D{\sim}E\rangle\big)^{(\nexists_{C\sim D}\vee\nexists_{D\sim E})}\\ \mathsf{p3} &= \big(\mathsf{p1}\,;\, \mathsf{sw}\langle A{\sim}C \circledast B\rangle\big)^{(\nexists_{A\sim C})}\\ \mathsf{p4} &= \big(\mathsf{p2}\,;\, \mathsf{sw}\langle C{\sim}E \circledast D\rangle\big)^{(\nexists_{C\sim E})}\\ \mathsf{p5} &= \big((\mathsf{p3} \parallel \mathsf{p4})\,;\, \mathsf{sw}\langle A{\sim}E \circledast C\rangle\big)^{(\nexists_{A\sim E})}\end{aligned}$$

The first level of entanglement swapping establishes pairs $A{\sim}C$ and $C{\sim}E$, whereas the second level produces $A{\sim}E$.

e) Heuristic: Alternatively, the order of swapping can be arbitrary and decided by a heuristic. Consider a protocol distributing $A{\sim}E$ on the same chain as in the doubling scheme above. For example, the protocol starts swapping $B{\sim}C$ and $C{\sim}D$ at the center node C, then swaps the produced $B{\sim}D$ with $A{\sim}B$, and finally swaps the obtained $A{\sim}D$ with $D{\sim}E$:

$$\begin{aligned}\mathsf{p1} &= (\mathsf{ge}\langle B{\sim}C\rangle \parallel \mathsf{ge}\langle C{\sim}D\rangle)^{(\nexists_{B\sim C}\vee\nexists_{C\sim D})}\\ \mathsf{p2} &= (\mathsf{p1}\ ;\ \mathsf{sw}\langle B{\sim}D \circledast C\rangle)^{(\nexists_{B\sim D})}\\ \mathsf{p3} &= \big((\mathsf{p2}\ \parallel\ \mathsf{ge}\langle A{\sim}B\rangle^{(\nexists_{A\sim B})})\ ;\ \mathsf{sw}\langle A{\sim}D \circledast B\rangle\big)^{(\nexists_{A\sim D})}\\ \mathsf{p4} &= \big((\mathsf{p3}\ \parallel\ \mathsf{ge}\langle D{\sim}E\rangle^{(\nexists_{D\sim E})})\ ;\ \mathsf{sw}\langle A{\sim}E \circledast D\rangle\big)^{(\nexists_{A\sim E})}\end{aligned}$$

V. Numerical Evaluation

First, we validate the correctness of our tool against analytical results from the literature (in Section V-A). To make the approaches comparable, we assume instantaneous entanglement swapping, i.e., we set $t_{\mathrm{c}}^{\mathrm{sw}} = 0$ in our tool. Then, we proceed with the analysis of resource contention in a concurrent protocol acting on a star topology (in Section V-B).

For our computations to be finite, we cap the maximum number of attempts of each action. We use shorthand notation

$$e^{(\alpha,r)} \triangleq \underbrace{(e +_\alpha 1); \ldots ; (e +_\alpha 1)}_{r}$$

for the r-th unrolling of the while loop $e^{(\alpha)}$, i.e., it is attempted at most r times if α is not met earlier. We limit our examples to a handful of nodes, leaving tool optimization and large-scale evaluation for future work. Our tool (with instructions on how to reproduce results) is publicly available [23].

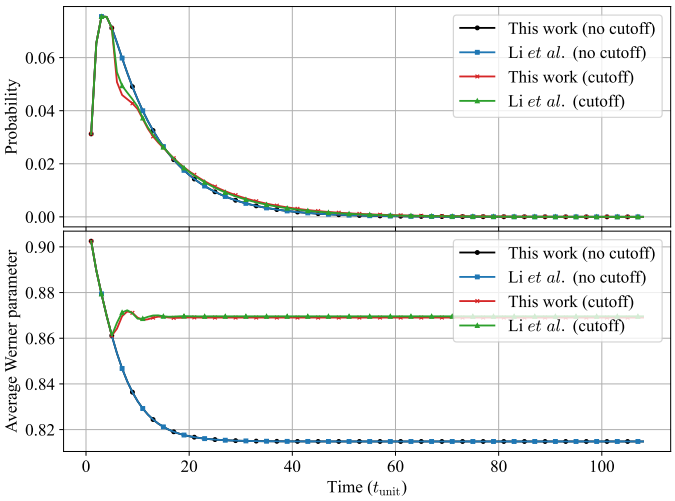


Fig. 4. Probability mass function and average Werner parameters for the distribution of the pair $A{\sim}C$ in the protocols from Section V-A. Comparison of the result obtained from the algorithm by Li *et al.* [10] against our tool.

A. Cutoff Protocol

We consider a time-cutoff condition (see Section III-D) for the following protocol,

$$\begin{aligned}\big((&\mathsf{ge}\langle A{\sim}B\rangle_{p_{\mathrm{ge}},w_0} \parallel \mathsf{ge}\langle B{\sim}C\rangle_{p_{\mathrm{ge}},w_0})\,;\\ &\mathsf{sw}\langle A{\sim}C \circledast B\rangle_{p_{\mathrm{sw}},t_{\mathrm{coh}},t_{\mathrm{coh}},t_{\mathrm{coh}}}\big)^{(\nexists_{A\sim C},r)}\end{aligned}$$

acting on a three-node homogeneous chain with $p_{\mathrm{ge}} = 0.25$, $w_0 = 0.95$, $p_{\mathrm{sw}} = 0.5$, $t_{\mathrm{coh}} = 10^2$.

This protocol generates pairs $A\sim B$ and $B\sim C$ and then swaps them, attempting $r = 108$ times unless $A\sim C$ is produced earlier. The protocol can be evaluated using the algorithm developed by Li *et al.* [10], used here to validate our tool. We showcase two different instances of the protocol: one includes a time-cutoff condition ($\tau = 4$), whereas the other does not ($\tau = +\infty$). Results with and without the cutoff condition are shown in Figure 4. As expected, the time-cutoff leads to a trade-off between distributing entanglement with a higher probability and obtaining a higher-quality pair with a lower probability. The results match the values obtained from the algorithm of Li *et al.* [10] up to a negligible error.

B. Star Network Concurrent Protocol

Finally, we analyze an example that goes beyond linear chains of nodes. In the context of aggregated quantum networks and multi-path distribution, we consider the star topology in Figure 1 and attempt to distribute pairs $A\sim C$ and $B{\sim}C$ concurrently, leveraging the following protocol:

$$\begin{aligned}\big((&\mathsf{ge}\langle A{\sim}H\rangle_{\hat{p}_{\mathrm{ge}},\hat{w}_0} \parallel \mathsf{ge}\langle B{\sim}H\rangle_{p_{\mathrm{ge}},w_0} \parallel \mathsf{ge}\langle C{\sim}H\rangle_{p_{\mathrm{ge}},w_0})\,;\\ &(\mathsf{sw}\langle A{\sim}C \circledast H\rangle_{p_{\mathrm{sw}}^H,t_{\mathrm{coh}}^A,t_{\mathrm{coh}}^C,t_{\mathrm{coh}}^H} \circ\\ &\qquad \mathsf{sw}\langle B{\sim}C \circledast H\rangle_{p_{\mathrm{sw}}^H,t_{\mathrm{coh}}^B,t_{\mathrm{coh}}^C,t_{\mathrm{coh}}^H})\big)^{(\nexists_{A\sim C}\wedge\nexists_{B\sim C},r)}\end{aligned}$$

The protocol attempts to generate $A\sim H$, $B\sim H$, and $C\sim H$. Then, it concurrently attempts to produce end-to-end entangled pairs $A\sim C$ and $B\sim C$ via entanglement

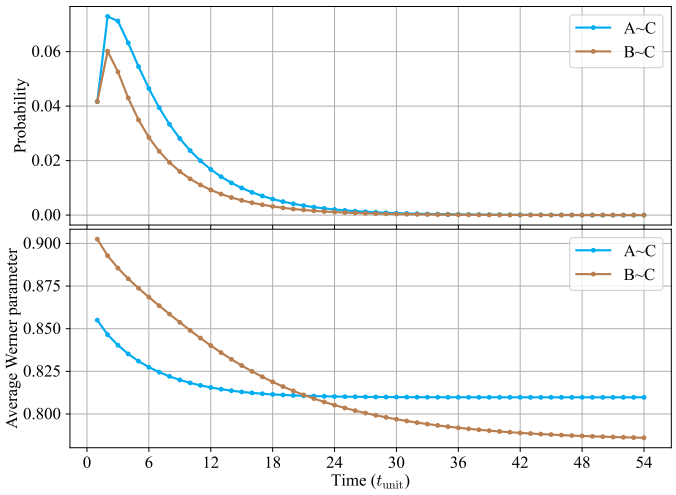


Fig. 5. Probability mass function and average Werner parameters for the distribution of the pairs $A{\sim}C$ (cyan) and $B{\sim}C$ (brown) in the star network concurrent protocol from Section V-B.

swapping. Since both swaps require $C \sim H$ to execute, this leads to resource contention; the specification above exhibits the $\circ$ operator to impose priority on the swap $\mathsf{sw}\langle A{\sim}C \mathbin{@} H\rangle$ over the swap $\mathsf{sw}\langle B \sim C \mathbin{@} H\rangle$. Thus, if all three pairs are successfully generated at the same time, the prioritized swap $\mathsf{sw}\langle A{\sim}C \mathbin{@} H\rangle$ executes, whereas $\mathsf{sw}\langle B{\sim}C \mathbin{@} H\rangle$ is attempted only if $B \sim H$ and $C \sim H$ are the only two pairs available. Notice that the negation of the predicate $\nexists_{A\sim C} \wedge \nexists_{B\sim C}$ is the exit condition for the while loop, exiting once either $A{\sim}C$ or $B \sim C$ is present, bounded by $r = 54$ attempts. We consider a heterogeneous hardware scenario with $\hat{p}_{\text{ge}} = 0.25$ and $\hat{w}_0 = 0.9$ being the hardware parameters for the pair $A{\sim}H$, representing worse conditions than other pairs, having values $p_{\text{ge}} = 0.33$ and $w_0 = 0.95$. The coherence time of the memories at the end nodes is set to $t^A_{\text{coh}} = t^B_{\text{coh}} = t^C_{\text{coh}} = 10^2$, with the hub's memory holding a higher value $t^H_{\text{coh}} = 2 \cdot 10^2$. The probability of success for swapping is $p^H_{\text{sw}} = 0.5$.

Results highlight the impact of the imposed priority, with the probability of distributing the entangled pair $B{\sim}C$ being generally lower than the probability of distributing $A{\sim}C$, as shown in the top plot of Figure 5.

The pair $B \sim H$ is generated with a higher initial quality ($w_0 > \hat{w}_0$). However, it is more likely that the pair $B{\sim}H$ is generated but not swapped immediately, thus decohering for a longer time when stored in memory. This impacts the Werner parameter of the pair $B{\sim}C$ produced from $B{\sim}H$; therefore, on average, pairs $B{\sim}C$ decohere faster than pairs $A{\sim}C$, as shown in the bottom plot of Figure 5.

VI. Conclusion and Future Work

This work provides foundations for formal reasoning about the performance of entanglement distribution in quantum networks. We present a language and tool for specifying and analyzing quantum repeater protocols in terms of rate and quality of the distributed entangled pairs. Our approach captures realistic behaviors of protocols beyond linear topologies, taking into account resource contention arising when multiple network operations require the same resources to execute. The tool is currently capable of reproducing analytical results from the literature and evaluating a protocol concurrently distributing multiple entangled pairs in a 4-node star topology.

Future efforts will focus on optimizing our tool to evaluate larger topologies and more protocol iterations. The improved tool will allow us to compare swapping schemes for heterogeneous networks and find optimal strategies for the scheduling order of actions competing for resources.

Our ultimate goal is to develop a methodology that yields design principles for optimal repeater protocols. With the envisioned improvements, we aim to make our language and tool a scalable framework for design-space exploration and automated optimization of quantum repeater protocols.

References

[1] S. Wehner, D. Elkouss, and R. Hanson, "Quantum Internet: A Vision for the Road Ahead," *Science*, vol. 362, no. 6412, pp. 1–9, 2018.

[2] C. Wang, A. Rahman, R. Li, M. Aelmans, and K. Chakraborty, "RFC 9583: Application Scenarios for the Quantum Internet," 2024.

[3] S. Pirandola, U. L. Andersen, L. Banchi *et al.*, "Advances in Quantum Cryptography," *Advances in Optics and Photonics*, vol. 12, no. 4, pp. 1012–1236, 2020.

[4] A. K. Ekert, "Quantum Cryptography based on Bell's Theorem," *Physical Review Letters*, vol. 67, no. 6, pp. 661–663, 1991.

[5] W. J. Munro, K. Azuma, K. Tamaki, and K. Nemoto, "Inside Quantum Repeaters," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 21, no. 3, pp. 78–90, 2015.

[6] S. Muralidharan, L. Li, J. Kim, N. Lütkenhaus, M. D. Lukin, and L. Jiang, "Optimal Architectures for Long Distance Quantum Communication," *Scientific Reports*, vol. 5, no. 1, 2016.

[7] W. Wootters and W. Zurek, "A single quantum cannot be cloned," *Nature*, vol. 299, pp. 802–803, 1982.

[8] C. H. Bennett, G. Brassard, S. Popescu, B. Schumacher, J. A. Smolin, and W. K. Wootters, "Purification of Noisy Entanglement and Faithful Teleportation via Noisy Channels," *Physical Review Letters*, vol. 76, no. 5, p. 722, 1996.

[9] W. J. Munro, A. M. Stephens, S. J. Devitt, K. A. Harrison, and K. Nemoto, "Quantum communication without the necessity of quantum memories," *Nature Photonics*, vol. 6, no. 11, pp. 777–781, Nov. 2012.

[10] B. Li, T. Coopmans, and D. Elkouss, "Efficient optimization of cutoffs in quantum repeater chains," *IEEE Transactions on Quantum Engineering*, vol. 2, pp. 1–15, 2021.

[11] Á. G. Iñesta, G. Vardoyan, L. Scavuzzo, and S. Wehner, "Optimal entanglement distribution policies in homogeneous repeater chains with cutoffs," *npj Quantum Information*, vol. 9, no. 1, p. 46, 2023.

[12] S. Haldar, P. J. Barge, S. Khatri, and H. Lee, "Fast and reliable entanglement distribution with quantum repeaters: principles for improving protocols using reinforcement learning," *Physical Review Applied*, vol. 21, no. 2, p. 024041, 2024.

[13] T. Coopmans, R. Knegjens, A. Dahlberg *et al.*, "NetSquid, a NETwork Simulator for QUantum Information using Discrete Events," *Communications Physics*, vol. 4, no. 164, pp. 1–15, 2021.

[14] X. Wu, A. Kolar, J. Chung, D. Jin, T. Zhong, R. Kettimuthu, and M. Suchara, "Sequence: a customizable discrete-event simulator of quantum networks," *Quantum Science and Technology*, vol. 6, no. 4, p. 045027, 2021.

[15] L. La Corte, K. Goodenough, A. G. Maity, S. Santra, and D. Elkouss, "Bayesian optimization for repeater protocols," in *2025 International Conference on Quantum Communications, Networking, and Computing (QCNC)*. IEEE, Mar. 2025, pp. 135–142.

[16] K. Goodenough, T. Coopmans, and D. Towsley, "On noise in swap asap repeater chains: exact analytics, distributions and tight approximations," *Quantum*, vol. 9, p. 1744, 2025.

[17] A. Buckley, P. Chuprikov, R. Otoni, R. Soulé, R. Rand, and P. Eugster, "A language for quantifying quantum network behavior," *Proc. ACM Program. Lang.*, vol. 9, no. OOPSLA2, Oct. 2025.

[18] R. van Meter, J. Touch, and C. Horsman, "Recursive quantum repeater networks," *Progress in Informatics*, no. 8, pp. 65–79, 2011.
[19] R. F. Werner, "Quantum states with einstein-podolsky-rosen correlations admitting a hidden-variable model," *Physical Review A*, vol. 40, no. 8, p. 4277, 1989.
[20] S. Brand, T. Coopmans, and D. Elkouss, "Efficient computation of the waiting time and fidelity in quantum repeater chains," in *OSA Quantum 2.0 Conference*. Optica Publishing Group, 2020.
[21] W. Kozlowski, S. Wehner, R. V. Meter, B. Rijsman, A. S. Cacciapuoti, M. Caleffi, and S. Nagayama, "Architectural Principles for a Quantum Internet," RFC 9340 (https://www.rfc-editor.org/info/rfc9340), 2023.
[22] A. Abane, M. Cubeddu, V. S. Mai, and A. Battou, "Entanglement routing in quantum networks: A comprehensive survey," *IEEE Transactions on Quantum Engineering*, vol. 6, pp. 1–39, 2025.
[23] P. Chuprikov and L. La Corte, "Fidelity-BellKAT." [Online]. Available: https://github.com/swystems/prob-bellkat/releases/tag/fidelity-bellkat-artifact