

A Continuous Variable Quantum Switch

Ian Tillman

*College of Optical Sciences
University of Arizona
Tucson, USA
ijtillman@arizona.edu*

Thirupathaiah Vasantam

*Department of Computer Science
Durham University
Durham, UK
thirupathaiah.vasantam@durham.ac.uk*

Kaushik P. Seshadreesan

*Department of Informatics & Networked Systems
School of Computing & Information
University of Pittsburgh, Pittsburgh, USA
kausesesh@pitt.edu*

Abstract—The continuous quadratures of a single mode of the light field present a promising avenue to encode quantum information. By virtue of the infinite dimensionality of the associated Hilbert space, quantum states of these continuous variables (CV) can enable higher communication rates compared to single photon-based qubit encodings. Quantum repeater protocols that are essential to extend the range of quantum communications at enhanced rates over direct transmission have also been recently proposed for CV quantum encodings. Here we present a quantum repeating switch for CV quantum encodings that caters to multiple communication flows. The architecture of the switch is based on quantum light sources, detectors, memories, and switching fabric, and the routing protocol is based on a Max-Weight scheduling policy that is throughput optimal. We present numerical results on an achievable bipartite entanglement request rate region for multiple CV entanglement flows that can be stably supported through the switch. We elucidate our results with the help of exemplary 3-flow networks.

Index Terms—quantum continuous variables, quantum repeater, quantum switch, maximum weight scheduling, entanglement distribution

I. INTRODUCTION

The world of “classical” computing over the past many decades has been on a fascinating journey from individual mainframes and PCs to the internet age of globally networked computers, mobile phones, IoT devices and beyond. Interconnecting computing devices proved to be a game-changer leading to the explosion of information technologies and applications that encompass our lives today, many of which at the outset were hardly thought to be possible, e.g., internet banking and social media. In a similar vein, the quantum 2.0 technologies [1] such as quantum computers, quantum sensors, quantum simulators, and other devices that are currently being developed promise to give rise to novel applications and to outperform classical technologies in certain types of tasks in the future. Arriving at this future largely hinges on interconnecting such quantum 2.0 devices to form quantum networks [2].

The ultimate goal of quantum networking is to enable reliable exchange of quantum information between quantum 2.0 devices at high rates both locally between quantum device modules at quantum data centers as well as at global-scale between distant quantum nodes. This includes faithful transmission of quantum bits (qubits) and distribution of

quantum entanglement, which is the special type of non-classical correlations that can exist between qubits and is believed to be the workhorse behind a large number of the quantum 2.0 technologies. However, quantum networking is not possible over the “classical” internet as it does not support quantum information exchange. It requires new networking infrastructure consisting of novel transceivers that can generate and detect quantum states, quantum memories that can store quantum information, and frequency converters that can transduce photonic quantum information into the telecom band for efficient transmission over optical fibers. Additionally, quantum networking requires novel quantum repeaters [3] that are not just optical amplifiers but special-purpose quantum processors to periodically re-generate quantum information in long range communications and support meaningful end-to-end rates. Furthermore, fast and high efficiency switches are required in conjunction with repeaters to route multiple communication flows through arbitrary network topology.

Quantum information can be broadly classified into discrete and continuous variable (DV and CV) encodings. DV quantum information refers to the information encoded in states of discrete degrees of freedom of quantum systems, e.g., the spin of an electron or nucleus, photon number of an optical mode, and polarization, path, time-bin, or frequency-bin of a photon. On the other hand, CV quantum information refers to quantum states of the continuous quadrature degrees of freedom of a spatial, spectral, or temporal mode of the electromagnetic field. Examples of CV quantum states include coherent states, thermal states, squeezed states, and bosonic cat states. A host of repeater architectures and protocols for quantum networking have been proposed and analyzed extensively for DV encodings since the 1990s based on spin qubits (e.g. in trapped ions [4] and color-centers in diamond [5]) and single photon polarization qubits. Quantum switches capable of routing multiple DV quantum entanglement have also been studied [6], [7]. A switch typically operates in conjunction with repeater elements that help establish and store entanglement between the switch and the different nodes connected to its ports, followed by entanglement swapping across different sets of switch ports to support different entanglement flows. It was proved recently that a Max-Weight scheduling policy is throughput optimal for a switch with quantum memories that have a lifetime of one time step that caters to several DV entanglement flows, where throughput optimality refers

to achieving stability of request queues associated with flows, for all feasible request rates [7].

Quantum repeaters have also been subsequently proposed for CV encodings. A prominent architecture for a CV repeater involves generating entangled two-mode quadrature squeezed states of light from Type-II spontaneous parametric downconversion (SPDC) sources, CV non-Gaussian quantum error correction by injecting and detecting single photons [8], storing heralded link-level entanglement in quantum memories, and entanglement swapping by coherent detection [9]–[11] or by other non-Gaussian methods [11], [12]. For the repeater architecture with coherent detection, the end-to-end entanglement rates in a hub-and-spoke CV network architecture for different relative placements of the central repeater hub were characterized and placements that enhance the rate of each individual flow over its corresponding direct transmission capacity were identified [13]. In the present work, we model the central repeater node as a quantum repeating CV switch that creates CV entangled links with the end nodes connected to its ports and implements routing based on a Max-Weight scheduling policy to cater to requests for the different entanglement flows between the end nodes by suitably performing entanglement swapping across its ports. We present numerical results on an achievable request rate region for the different end-to-end CV bipartite entanglement flows that can be supported by the switch stably. The achievable request rate region translates to an entanglement request rate region when end-to-end entanglement distillation based on the hashing protocol is considered [14].

The manuscript is organized as follows: In Section II, we describe the CV repeater architecture and protocol used in [13], which is also what we will consider in the present work. Subsequently, in Section III, we describe the model for a CV quantum repeating switch. Section IV describes the Max Weight schedule policy used by the switch. This is followed by results in Section V. We conclude with a discussion in Section VI.

II. CONTINUOUS VARIABLE QUANTUM REPEATER ARCHITECTURE

Quantum continuous variables of a bosonic mode [15] are described by the Hermitian quadrature operators $\hat{q} = (\hat{a} + \hat{a}^\dagger)/\sqrt{2}$, $\hat{p} = (\hat{a} - \hat{a}^\dagger)/(\sqrt{2}i)$, where \hat{a}^\dagger, \hat{a} denote the creation and annihilation operators of the mode and we have chosen $\hbar = 1$. The quintessential entangled state in CV is the two-mode squeezed vacuum state (TMSV) $|\chi\rangle_{AB}$, which is expressed in the photon number basis as $|\chi\rangle_{AB} = \sqrt{1 - \chi^2} \sum_{n=0}^{\infty} \chi^n |n, n\rangle_{AB}$, where $0 \leq \chi < 1$. The TMSV state exhibits a reduced, or ‘squeezed’, variance of either the difference of the position quadratures of the two modes or the sum of their momentum quadratures at the expense of an increased variance of the other. When one of the two modes of a TMSV state is transmitted through a lossy channel of transmissivity η , in the limit $\chi \rightarrow 1$ (which is the infinite energy limit) the distillable entanglement of the end-to-end

state coincides with the entanglement distribution capacity of the channel [16] given by

$$C_{\text{direct}}(\eta) = -\log_2(1 - \eta) \text{ ebits/mode}, \quad (1)$$

which in the limit $\eta \ll 1$ is $C(\eta) \propto \eta$. (Here an ebit refers to a maximally entangled pair of qubits.)




Fig. 1. A CV quantum repeater protocol for entanglement distribution. The repeaters rely on multiplexed distribution of two-mode squeezed vacuum (TMSV) states across each link (i.e., lossy links between the end nodes Alice and Bob, and the repeater node Charlie), quantum scissor (QS)-based probabilistic heralded noiseless linear amplification to mitigate the effects of photon loss modeled by finite transmissivities $\eta_i, i \in \{1, 2\}$ (on either side of Charlie), and dual homodyne detection based entanglement swapping at Charlie to connect the heralded and successfully error corrected entangled links on either side of Charlie. The end-to-end entanglement is then stored in quantum memories (QM). Notice the Orange arrows (background) that are aligned in the same direction (left to right) indicating the relative orientation of the links (TMSV tail and QS head) being connected at Charlie (head to tail).

A recently proposed and widely studied repeater architecture [9] for CV that can help surpass the direct transmission capacity of (1) for entanglement distribution between end nodes is shown in Fig. 1. It also involves distributing TMSV states across each individual link in a repeater network. The repeater employs so-called non-deterministic noiseless linear amplification (NLA) [17] (with an associated gain parameter g) at the repeaters to perform non-Gaussian CV quantum error correction on the lossy TMSV states across individual links [18]. The concept of NLA comes under the class of ‘immaculate’ quantum amplifiers [19] that can amplify CV states without adding any noise to the system, but only succeed in doing so probabilistically. In order to boost the success probability of entanglement generation in each time step, the repeaters multiplex their entanglement generation attempts (M fold multiplexing) across each link. In the limit of $\eta \ll 1$, the NLA can be approximately realized using an optical state engineering primitive known as the ‘quantum scissor’ [20], which involves single photon injection and detection. The repeater relies on coherent dual homodyne detection [21], which is a Gaussian measurement, for entanglement swapping at the repeaters to connect adjacent links and have long range end-to-end entanglement, which can then be stored in the quantum memories at the end nodes.

The above multiplexed CV repeater architecture based on TMSV and quantum scissor NLA in the limit $\eta \ll 1$ can support an end-to-end entanglement rate that scales as $\propto \eta^{3/4}$ [13], thus surpassing the direct transmission capacity

of (1). It is important to note that the said rate-loss scaling is achieved when the NLA-end of one link is connected to the TMSV-end of the other, as indicated by the aligned orange shaded arrows in Fig. 1, while the other relative orientations such as connecting 2 NLA ends or 2 TMSV ends result in strictly worse end-to-end entanglement rates. For the different orientations of two adjacent repeater links and for different transmissivities η_i , $i \in \{1, 2\}$ in the two links, the end-to-end entanglement rates were determined in [13] by evaluating an information theoretic lower bound on the asymptotic distillable entanglement of the end-to-end state, known as the reverse coherent information. Using the entanglement rates as a function of the different transmissivities [13] further gave the end-to-end entanglement rates in a hub-and-spoke quantum network with a quantum repeater as the hub node for different placements of the hub relative to the end nodes, thus identifying placements that help surpass the direct transmission capacity of (1) for each individual entanglement flow through the repeater hub.

III. CONTINUOUS VARIABLE QUANTUM SWITCH MODEL

Here, we describe our model for a quantum repeating switch that can route multiple CV entanglement flows. For simplicity, we restrict ourselves to bipartite entanglement flows. The model is based on a hub-and-spoke network topology with the switch as the hub node and the end nodes as the spokes as described in Fig. 2. In every time step, the central repeater hub node attempts CV entanglement generation with each end node along the spokes, or ‘links’, using TMSV states and QS NLA-based error correction over M multiplexed channels. Notably, the entanglement generation across each link is attempted along both link orientations, i.e., half the multiplexed channels use TMSV sources at the end node and the QSs at the switch port, while the other half use TMSV sources at the switch and the QSs at the end node. The rationale behind this is to have a successfully generated entanglement across a link be equally likely to be either possible orientation. This allows the switch to connect aligned links capable of providing larger entanglement rates when possible, and do so impartially between the different entanglement flows through the switch. By suitably tuning the gain parameter of the QS-based NLA the probability of heralding a successfully error corrected CV entangled state across a link, P_{NLA} , can be made to scale as $\eta^{1/4}$ in the limit of $\eta \ll 1$. By attempting entanglement generation across M multiplexed channels, the probability of successfully heralding at least 1 error corrected CV entanglement state across a link is given by $p = 1 - (1 - P_{\text{NLA}})^M$. The switch then picks a request from the queue that can be serviced with the links that were able to generate entanglement in that particular time step and performs a dual homodyne detection based entanglement swap to entangle the two end nodes. The choice of which ports to connect via entanglement swapping in any given time step is driven by a throughput optimal routing policy described in Section IV. A point to note about the dual homodyne detection based CV entanglement swap operation is that it is deterministic unless we choose to perform post-selection to

optimize the end-to-end entangled state, which for simplicity we do not do here. In other words, the entanglement swap success probability of the CV quantum switch is taken to be $q = 1$.

Note that our model assumes that both the end nodes and the switch have unconstrained quantum resources (quantum sources, detectors and memories) and share a synchronized clock. We assume that the lifetime of the quantum memories at the switch (surpressed in Fig. 1) equals the duration of a time step of operating the switch, whereas the lifetime of the memories at the end users are assumed to be several time steps long, enabling the hashing protocol. This simplifies our model and multiple entangled states heralded during one step do not get carried over to subsequent time steps.




Fig. 2. A CV quantum repeating switch architecture. Each end node that is connected to the switch attempts heralded CV entanglement generation with the switch at each time step using multiplexed two-mode squeezed vacuum (TMSV) sources and quantum scissors (QS) depicted in Fig. 1. The attempts happen in both orientations, i.e., TMSV at the end node and QS at the switch and vice versa, as indicated by the pair of Orange arrows. (a), (b), (c) describe the 3 possible scenarios associated with 3 bipartite entanglement flows (denoted by Red, Blue and Purple lines) at the switch. We expand on these possibilities in sec. V

IV. THROUGHPUT OPTIMAL ROUTING

We assume that requests of each entanglement flow type, arrive at the switch according to an i.i.d. process with λ_i denoting the average number of type i requests arriving in each time step. A type i request demands the switch to create an entanglement shared by a set of users, U_i . We make an assumption that requests of each type are processed in the order they arrive at the switch and there is a buffer for each type of requests with infinite classical memory to store unserved requests. Let $Q_i(n)$ be the number of requests of type i waiting in the queue corresponding to type i requests at the beginning of time step n . We use $T_j(n)$ to indicate whether an elementary entanglement is available or not on link j in time step n ; if an elementary entanglement is present then $T_j(n) = 1$ and if no elementary entanglement is present then $T_j(n) = 0$.

Since we pick at most one elementary entanglement per time step along each link, if there are requests of multiple types waiting for service then there exists a competition among requests to use available elementary entanglements. The switch should make scheduling decisions to decide which requests to be processed using available elementary entanglements in such a way that requests are processed as quickly as possible. As in [7], the switch is considered to process requests in each time

step based on a selected matching $\mathbf{r} = [r_1, r_2, \dots, r_k]$ which indicates that the switch attempts to serve r_i requests of type i by performing entanglement swapping operations. The term matching is defined as follows:

Definition 1: Matching: A vector $\mathbf{r} = [r_1, r_2, \dots, r_k]$ with $r_i \in \{0, 1\}$ for $1 \leq i \leq K$ is called matching if for every user j , we have

$$\sum_{m=1, j \in U_m}^K r_m \leq 1. \quad (2)$$

The above equation implies that the number of requested elementary entanglements as per the matching \mathbf{r} , that are shared between user j and the switch should be at most one.

We use the Max-Weight scheduling policy of [7] to make scheduling decisions at the switch that decides which matching should be selected from the set of all feasible matchings. We formally define the scheduling policy below:

Definition 2: Max-Weight scheduling policy: According to the max-weight policy, a matching $\mathbf{W}(n) = [W_1(n), \dots, W_K(n)]$ is selected in time step n if

$$\mathbf{W}(n) = \arg \max_{\mathbf{r} \in \mathcal{M}} \sum_i r_i q_i Q_i(n) I_{\{T_j(n) > 0, \forall j \in U_i\}}, \quad (3)$$

where \mathcal{M} denotes the set of all matchings and q_i is the success probability of entanglement swapping operations of type i requests.

In [7], the capacity region of a switch was found, defined as the set of request rates $\boldsymbol{\lambda} = [\lambda_1, \dots, \lambda_K]$ for which there can exist a scheduling policy that stabilizes the switch, guaranteeing finite average waiting times for requests. If $\boldsymbol{\lambda}$ is not in the capacity region, then no scheduling policy can stabilize the switch. It was also proved in [7] that the Max-Weight policy defined in (3) is throughput optimal, that is, it stabilizes the switch for all request arrival rates that belong to the capacity region. If the switch is stable under a scheduling policy for a rate vector $\boldsymbol{\lambda}$, then in the stationary regime, the departure rate of requests leaving the switch after getting service coincides with $\boldsymbol{\lambda}$. One challenge is that, it is difficult to verify whether a given rate vector $\boldsymbol{\lambda}$ satisfies the necessary conditions of [7] or not to judge stability of the switch. In the next section, for a switch catering to three entanglement flows, we present results that are useful to conclude whether a switch will be stable or not for given $\boldsymbol{\lambda}$.

V. RESULTS

Here we present our results on the bipartite entanglement rate region for multiple CV entanglement flows through the CV quantum switch that can be stably supported. The results are shown in terms of the rate region for the requests corresponding to the different entanglement flows types that can be stably met by the Max-Weight scheduling policy based routing. An achievable end-to-end entanglement rate follows from an achievable request rate given as the latter scaled by the reverse coherent information of the end-to-end entanglement state, which is a lower bound on the distillable entanglement

of the state in the asymptotic limit of many copies of the state [14].

We present our results for the example of 3-flow networks with numerical plots of the request rate regions that are stably supported by the Max-Weight scheduling policy-based routing. Consider the three distinct scenarios that are possible in 3-flow networks as described in Fig. 2, namely (a) where all 3 flows contend with each other in each time step, (b) where 2 out of the 3 flows (Blue and Purple) are disjoint and do not contend with each other, and (c) where all 3 flows are disjoint and thus there is no contention. The request rate regions for each of these scenarios are functions of the entanglement swapping success probabilities q_j , where $j \in \{1, 2, 3\}$, and the success probabilities of heralding entanglement along each link, p_i , where $i \in \{1, 2, 3\}$ for case (a), $i \in \{1, 2, 3, 4\}$ for case (b), and $i \in \{1, 2, 3, 4, 5, 6\}$ in case (c). The probabilities p_i are functions of the link transmissivities η_i , the squeezing parameters χ_i , and the multiplexing parameters M_i . For the sake of simplicity in depicting the results we will use $M_i = 1/P_{\text{NLA},i}$, which in the limit $P_{\text{NLA},i} \rightarrow 0$ (i.e. when $\eta_i \rightarrow 0$) we have $p_i = 1 - (1 - P_{\text{NLA},i})^{M_i} = 1 - (1 - P_{\text{NLA},i})^{(1/P_{\text{NLA},i})} \rightarrow 1 - 1/e \approx 0.632$, approaching from above. This gives a universal and scale-independent minimum value for the link generation success probabilities p_i that also gives a near-optimal entanglement bitrate per mode [13]. Along with our assumption of entanglement swapping probabilities $q_j = 1$ for all flows $j \in \{1, 2, 3\}$, this allows for a greatly simplified discussion.

For the scenario in Fig. 2 (a) where all the 3 flows contend with each other, ignoring the orientation of the successfully generated entanglements along each link, a request rate region that can be stably supported by the Max-Weight scheduling policy based routing at the switch is given by:

$$\lambda_1 + \lambda_2 + \lambda_3 \leq p^3 + 3(1-p)p^2 \quad (4)$$

$$\lambda_i + \lambda_j \leq p^3 + 2(1-p)p^2, \quad i, j \in \{1, 2, 3\}, \quad i \neq j \quad (5)$$

$$\lambda_i \leq p^3 + (1-p)p^2 = p^2, \quad i \in \{1, 2, 3\}, \quad (6)$$

where $\lambda_1, \lambda_2, \lambda_3$ are the request rates corresponding to the 3 contending flows. These bounds can be understood from a combinatorial perspective, as in order for the switch to be stable it must be able to service requests at least as fast as they come in. When considering a set of $1 \leq N \leq 3$ flows the sum of their request rates must be less than the probability that at least one of them can be serviced, which occurs when all elementary links are generated, which has probability p^3 , or exactly two elementary links are formed, which has probability $N(1-p)p^2$. A stable request rate region for the scenario in Fig. 2 (b) is similar to the above with the exception of the constraints on the sum of the two disjoint flows (Blue and Purple) λ_1 and λ_3 now being omitted. Finally, the rate region for the scenario in Fig. 2 (c) is simply given by $\lambda_i \leq p^2, i \in \{1, 2, 3\}$.

These rate regions for $p = 0.632$ are plotted in Figs. 3-5 along with simulated rate regions. The simulations were done

for $0 \leq \lambda_i \leq 1$, $i \in \{1, 2, 3\}$ with step sizes of $\Delta\lambda = 0.005$ and 10^5 time steps each. A rate vector is considered 'stable' if a linear regression of the queue size over all time steps has a slope less than some threshold, here taken to be 10^{-4} . Because these simulations were done with a finite number of time steps we should expect some of the points just outside the actual rate region to be mislabeled as stable due to the average queue size diverging too slowly in that particular simulation.




Fig. 3. Request rate region for 3 overlapping flows as described in Fig. 2 (a) that can be stably supported by the CV repeating switch with link generation probabilities $p = 0.632$ and deterministic entanglement swapping.




Fig. 4. Request rate region for the 3-flow network scenario described in Fig. 2 (b) where 2 out of the 3 flows are disjoint (Blue and Purple flows, with the associated request rates being λ_1 and λ_3 , respectively), that can be stably supported by the CV repeating switch with link generation probabilities $p = 0.632$ and deterministic entanglement swapping.

Since connecting two links that are aligned as shown in Fig. 1 results in an end-to-end entangled state of a higher reverse coherent information than any other orientation, we




Fig. 5. Request rate region for 3 disjoint flows as described in Fig. 2 (c) that can be stably supported by the CV repeating switch with link generation probabilities $p = 0.632$ and deterministic entanglement swapping.




Fig. 6. Request rate region for 3 overlapping flows as described in Fig. 2 (a) that can be stably supported by the CV repeating switch with link generation probabilities $p = 0.632$ and deterministic entanglement swapping, where the relative orientations of the successfully heralded links are taken into consideration and aligned links alone are connected by the switch.

can additionally instruct the switch to take the orientations of the links into consideration in addressing routing requests. For example, in the case of the scenario in Fig. 2 (a) with contending flows and with the orientation of the successfully generated entanglements taken into account by the scheduling

policy, the new request rate region is given by

$$\lambda_1 + \lambda_2 + \lambda_3 \leq 3p^2/2 - 3p^3/4 \quad (7)$$

$$\lambda_i + \lambda_j \leq p^2 - p^3/4, \quad i, j \in \{1, 2, 3\}, \quad i \neq j \quad (8)$$

$$\lambda_i \leq p^2/2, \quad i \in \{1, 2, 3\}. \quad (9)$$

This rate region is plotted in Fig. 6. These are derived in the same way as equations 4-6, except we also take into account that two links must have opposite parity to service a flow. For example the rate of any single flow, λ_i , must not exceed the probability of the two associated elementary links succeeding, p^2 , multiplied by $1/2$ to account for the probability that the two links have opposite parity. When looking at the sum of two flows, $\lambda_i + \lambda_j$, $i \neq j$, this will succeed when all links succeed they are not all the same parity ($\frac{3}{4}p^3$) or when one flow's links succeed with opposite parity and the third link fails ($\frac{1}{2}p^2(1-p)$), which can happen in two ways. Therefore $\lambda_i + \lambda_j \leq \frac{3}{4}p^3 + 2p^2(1-p)\frac{1}{2} = p^2 - p^3/4$. Finally, the sum of all three rates can be analyzed in the same way but there are now three ways in which a single flow's links succeed with opposite parity while the third link fails. So $\lambda_1 + \lambda_2 + \lambda_3 \leq \frac{3}{4}p^3 + 3p^2(1-p)\frac{1}{2} = 3p^2/2 - 3p^3/4$. Though the request rates supported come down after we require opposite parity, the higher reverse coherent information of the end-to-end entangled states may result in higher entanglement rates. Similar request regions with the relative orientations being taking in account for the links can also be arrived at for the other two scenarios described in Fig. 2.

This combinatorial approach of defining a boundary plane for each subset of the request rates can be done for any switch scenario. This includes scenarios where the link generation probabilities p_i are not identical, scenarios where the entanglement fusion probabilities q_i are not unity, or scenarios where the switch offers to service requests for multipartite entanglement. The plane associated with each subset of request rates are calculated by bounding the probability that at least one of those rates can be serviced, which is the same procedure as outlined in the two examples above.

VI. DISCUSSION

Our results can be improved a few obvious ways. First of all, for practical relevance, constrained operation of CV switches with limited resources such as TMSV sources, single photon sources and detectors, homodyne detectors and quantum memories need to be considered. In the resource-constrained case, it should also be interesting to consider spatial vs time multiplexed operation of the CV repeater links. This more practical analysis should also take into account the fidelity decay inherent to the quantum memories being used, particularly when it comes to end nodes storing the final entangled states if the application in question requires entanglement distillation. Secondly, the end-to-end entanglement rates could be further enhanced by relaxing the constraint that unused entangled states successful heralded along links be dropped after each time step and instead by modeling finite but longer lifetime for the quantum memories than the time step of the switch.

Thirdly, it should be interesting to extend the utility of the switch to go beyond just bipartite entanglement flows and to cater to multipartite flows. This would require the switch to perform multimode entangled Gaussian fusion operations on the QS-NLA-enhanced entangled states along different links, which generalize the dual homodyne detection [21]. It should also be interesting to consider more general non-Gaussian fusion operations directly operating on the lossy TMSV state-based links, e.g., generalizing the recently proposed CV repeater protocol that uses a generalized QS to accomplish both error correction and entanglement swapping [22]. We leave these tasks for future work.

In summary, we presented a model for a continuous variable quantum switch where multiple end nodes share CV entanglement links based on TMSV and QS-NLA-based error corrected with the switch, and the switch routes the multiple entanglement flows between the end nodes using a Max-Weight scheduling-based routing protocol that has been shown to be throughput optimal [7]. We presented numerical results on request rate regions for 3-flow networks that can be stably supported by the CV switch. Our work paves the way to determine the capabilities of general CV quantum networks of arbitrary topology, where the switches implement a routing protocol based on a well-defined scheduling policy.

REFERENCES

- [1] J. P. Dowling and G. J. Milburn, "Quantum technology: the second quantum revolution," *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, vol. 361, no. 1809, pp. 1655–1674, 2003.
- [2] S. Wehner, D. Elkouss, and R. Hanson, "Quantum internet: A vision for the road ahead," *Science*, vol. 362, no. 6412, p. eaam9288, 2018. [Online]. Available: <https://www.science.org/doi/abs/10.1126/science.aam9288>
- [3] 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.
- [4] P. Dhara, N. M. Linke, E. Waks, S. Guha, and K. P. Seshadreesan, "Multiplexed quantum repeaters based on dual-species trapped-ion systems," *Phys. Rev. A*, vol. 105, p. 022623, Feb 2022. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevA.105.022623>
- [5] F. Rozpedek, R. Yehia, K. Goodenough, M. Ruf, P. C. Humphreys, R. Hanson, S. Wehner, and D. Elkouss, "Near-term quantum-repeater experiments with nitrogen-vacancy centers: Overcoming the limitations of direct transmission," *Phys. Rev. A*, vol. 99, p. 052330, May 2019. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevA.99.052330>
- [6] G. Vardoyan, S. Guha, P. Nain, and D. Towsley, "On the stochastic analysis of a quantum entanglement distribution switch," *IEEE Transactions on Quantum Engineering*, vol. 2, pp. 1–16, 2021.
- [7] T. Vasantam and D. Towsley, "A throughput optimal scheduling policy for a quantum switch," in *Quantum Computing, Communication, and Simulation II*, P. R. Hemmer and A. L. Migdall, Eds., vol. 12015, International Society for Optics and Photonics. SPIE, 2022, pp. 14 – 23. [Online]. Available: <https://doi.org/10.1117/12.2616950>
- [8] J. Dias and T. C. Ralph, "Quantum error correction of continuous-variable states with realistic resources," *Physical Review A*, vol. 97, no. 3, p. 032335, 2018.
- [9] J. Dias, M. S. Winnel, N. Hosseini-dehaj, and T. C. Ralph, "Quantum repeater for continuous-variable entanglement distribution," *Physical Review A*, vol. 102, no. 5, p. 052425, 2020.
- [10] M. Ghalaii and S. Pirandola, "Capacity-approaching quantum repeaters for quantum communications," *Physical Review A*, vol. 102, no. 6, p. 062412, 2020.
- [11] F. Furrer and W. J. Munro, "Repeaters for continuous-variable quantum communication," *Physical Review A*, vol. 98, no. 3, p. 032335, 2018.

- [12] K. P. Seshadreesan, H. Krovi, and S. Guha, “Continuous-variable quantum repeater based on quantum scissors and mode multiplexing,” *Physical Review Research*, vol. 2, no. 1, p. 013310, 2020.
- [13] I. J. Tillman, A. Rubenok, S. Guha, and K. P. Seshadreesan, “Supporting multiple entanglement flows through a continuous-variable quantum repeater,” 2022. [Online]. Available: <https://arxiv.org/abs/2203.07965>
- [14] I. Devetak and A. Winter, “Distillation of secret key and entanglement from quantum states,” *Proceedings of the Royal Society A: Mathematical, Physical and engineering sciences*, vol. 461, no. 2053, pp. 207–235, 2005.
- [15] C. Weedbrook, S. Pirandola, R. García-Patrón, N. J. Cerf, T. C. Ralph, J. H. Shapiro, and S. Lloyd, “Gaussian quantum information,” *Reviews of Modern Physics*, vol. 84, no. 2, p. 621, 2012.
- [16] S. Pirandola, R. Laurenza, C. Ottaviani, and L. Banchi, “Fundamental limits of repeaterless quantum communications,” *Nature communications*, vol. 8, no. 1, pp. 1–15, 2017.
- [17] T. C. Ralph and A. Lund, “Nondeterministic noiseless linear amplification of quantum systems,” in *AIP Conference Proceedings*, vol. 1110, no. 1. American Institute of Physics, 2009, pp. 155–160.
- [18] J. Niset, J. Fiurášek, and N. J. Cerf, “No-go theorem for gaussian quantum error correction,” *Phys. Rev. Lett.*, vol. 102, p. 120501, Mar 2009. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.102.120501>
- [19] S. Pandey, Z. Jiang, J. Combes, and C. M. Caves, “Quantum limits on probabilistic amplifiers,” *Phys. Rev. A*, vol. 88, p. 033852, Sep 2013. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevA.88.033852>
- [20] D. T. Pegg, L. S. Phillips, and S. M. Barnett, “Optical state truncation by projection synthesis,” *Physical review letters*, vol. 81, no. 8, p. 1604, 1998.
- [21] G. Spedalieri, C. Ottaviani, and S. Pirandola, “Covariance matrices under bell-like detections,” *Open Systems & Information Dynamics*, vol. 20, no. 02, p. 1350011, 2013.
- [22] M. S. Winnel, J. J. Guanzon, N. Hosseini-dehaj, and T. C. Ralph, “Overcoming the repeaterless bound in continuous-variable quantum communication without quantum memories,” 2021. [Online]. Available: <https://arxiv.org/abs/2105.03586>