Gayane Vardoyan

Research Statement

My research spans problems in *classical and quantum communication systems*, with an emphasis on modeling, analysis, and performance evaluation in general. On the classical networking front, I study congestion control algorithms, focusing on ones that are commonly-used in high-speed environments and whose behaviors had not been carefully studied. In one part of this work, with collaborators, I developed a generalized modeling framework for the class of congestion controllers that adjust their sending rate by responding to loss information received from the network [18]. In another part of this work, with collaborators I developed a novel fluid model and applied it to TCP CUBIC (the default TCP variant in Linux kernels) [17]. This model allowed us to complete the first detailed stability analysis of CUBIC and discover that under a certain loss probability model, the algorithm is locally uniformly asymptotically stable.

Currently, I am interested in formulating and analyzing problems in the domain of quantum networking. Some of my recent work is discussed below. In the future, I would like to direct my research toward exploring the full potential of quantum networks, especially when it comes to supporting important applications such as quantum key distribution. Since quantum networks will operate alongside their classical counterparts, I expect that my background in classical networking will prove useful in approaching and solving problems in quantum communication.

Throughout my research, my aim is to effectively use analytical and modeling techniques to accurately represent communication systems and characterize their behavior. Examples of questions that I may ask are "Under what conditions is this protocol stable?" and "How much memory does this quantum device require?" The ultimate goal is to use the knowledge obtained from the performance evaluation process to

1. be able to perform *fair comparisons* of system variants,
2. *detect flaws* or *strengths* within a system's design, and
3. *improve* the design of existing systems and protocols, and *guide* the design of new/future ones.

Background

In recent years, quantum communication technology has seen rapid advances, bringing the vision of a quantum Internet closer to reality. Quantum networks have a variety of applications, ranging from cryptography to sensing and distributed quantum computing. One of the strongest motivations is *quantum key distribution (QKD)*. A unique advantage of QKD over classical key distribution is that eavesdropping is easily detected using simple statistical methods. *Entanglement*[1] is an essential component of QKD, quantum computation, information, and communication. In addition to QKD, a number of other quantum distributed applications rely on entanglement to meet their objectives; examples are quantum error correction [4] and ensemble sensing (*e.g.*, multipartite entanglement for quantum metrology [6] and spectroscopy [9]). These applications drive the increasing need for a quantum network that can supply end-to-end entanglements to geographically diverse groups of endpoints that request them [10, 12, 14]. To this end, it is prudent to model and analyze quantum networks, with the goal of illuminating the challenges of deploying such systems on large scales and for long-term operation, as well as to discover novel solutions to existing problems.

[1] Two quanta are said to be entangled when operations on one affect the other.

The Issue of Limited Distance

One of the major challenges of implementation of distributed tasks in quantum networks is the difficulty of safely transmitting a quantum state across a large distance. For optical fiber, channel transmissivity is $\eta = e^{-\gamma L}$, where L is the length of the link and γ the fiber attenuation coefficient. The probability of successful entanglement generation p on a link is proportional to its transmissivity η. Transmission through free space poses its own challenges, such as photon loss and phase changes due to scattering [14]. Non-entanglement-based protocols, such as BB84 [1], also suffer from limited distance for the same reason: the likelihood of losing a quantum state in transit grows exponentially with distance, while the no-cloning theorem [19] prevents one from making an independent copy of an unknown quantum state, thereby rendering losses irrecoverable. A remedy for the issue of limited distance is the use of quantum repeaters [5] coupled with the process of teleportation [2]. In its basic form, teleportation works by allowing Alice to transport a qubit[2] to Bob using a shared Bell pair[3], local operations, and classical communication.

Note that performing one teleportation consumes exactly one entanglement. To accomplish this process across a larger distance, a quantum repeater positioned between Alice and Bob is needed, with a quantum channel connecting each user to the repeater. Then, link-level entanglements are created: one between Alice and the repeater, and another between Bob and the repeater. The repeater then performs a measurement in the Bell basis[4] on the two locally-held qubits. The result is an end-to-end entanglement between Alice and Bob. Now, Alice can teleport a qubit using this new, longer-distance entanglement. To extend the distance even further, more repeaters can be added, and end-to-end entanglement is created using several link-level entanglements and Bell state measurements (BSMs). This process is what allows QKD protocols and other distributed quantum algorithms to be of practical use.

Prior Work

My work focuses on a quantum switch, which is a device equipped with the functionalities of a quantum repeater, except that it also includes entanglement switching logic. A quantum switch will also have some quantum memory. Moreover, memory coherence times may be heterogeneous: some memories may serve well for long-term qubit storage, while others are better suited for staging purposes. Figure 1 shows an example of what a quantum repeater or switch may look like. In [16], we explore the limits of a star-topology entanglement switching network (similar to the Los Alamos quantum network [7]) and introduce methods to model the process of entanglement generation, memory constraints, link heterogeneity, and decoherence for a switch

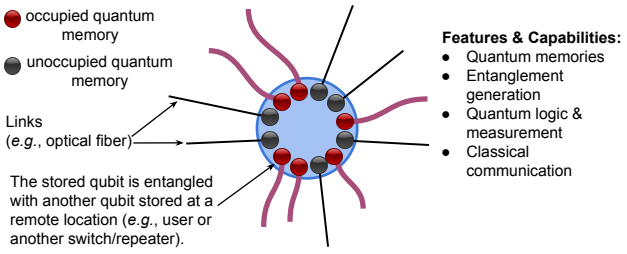

Figure 1: A quantum switch implementation.

[2] A qubit is the quantum analogue of a bit and is represented by a two-level system (*e.g.* spin of an electron – up or down, or polarization of a photon – horizontal or vertical). An important distinction is that a qubit can be in superposition. In Dirac notation, this is represented as $|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle \equiv \alpha\,[1\ 0]^T + \beta\,[0\ 1]^T$, where α and β are complex and $|\alpha|^2 + |\beta|^2 = 1$. The probabilistic interpretation is that if we prepare many states $|\psi\rangle$ and measure them, then over time, $P(|0\rangle) = |\alpha|^2$ and $P(|1\rangle) = |\beta|^2$.

[3] A Bell pair consists of two entangled qubits, *e.g.* $|\Phi^+\rangle = (|0\rangle_A \otimes |0\rangle_B + |1\rangle_A \otimes |1\rangle_B)/\sqrt{2}$, where the subscripts A and B signify the two qubits (possibly separated by distance) and $\otimes$ is a tensor product. For Bell pairs, the probability of an outcome is always 1/2. Since the qubits are entangled, measuring one of them tells us with certainty the state of the other qubit.

[4] A measurement in the Bell basis is an operation that takes as an input a Bell pair and outputs two classical bits. The bits are then used to decode the Bell state.

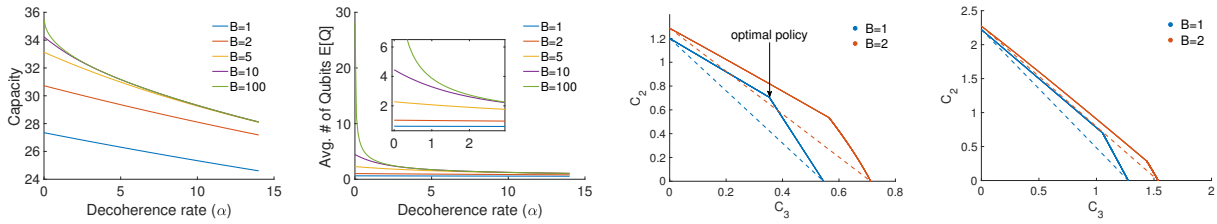

(a) (b) (c) $k = 3$ users/links (d) $k = 5$ users/links

Figure 2: Panels (a) and (b): effect of decoherence on capacity (Mega-ebits/sec) and expected number of stored qubits $E[Q]$ for five links and varying buffer sizes B. The links' entanglement generation rates are 35, 15, 15, 3, and 3 Mega-ebits/sec. For both plots, $B = 100$ curves behave equivalently to $B = \infty$. Panels (c) and (d): comparison of capacity regions for systems of buffer sizes one and two per link with varying number of links k. C_2 and C_3 are the bi- and tripartite capacities, respectively, and B is the number of memories the switch allocates to each link. The dashed lines represent the set of time-division multiplexing policies, while the solid lines profile the achievable capacity regions.

that can serve only bipartite (and in some cases, only tripartite) entanglements. For each set of assumptions, we compute the maximum achievable capacity C of the switch and the expected number of qubits $E[Q]$ stored in memory at the switch. *These metrics can serve as a useful comparison basis to assess the performance of future entanglement switching protocols.*

For bipartite entanglement switching, we observe that in most cases, little memory is required to achieve the performance of an infinite-memory system. We also find that for homogeneous-link systems, decoherence has little effect on performance metrics. In contrast, decoherence can have more significant consequences in heterogeneous-link systems that operate close to their stability constraints. Panels 2a and 2b provide an example of such a system; here, capacity degrades by 7.35 Mega-ebits/sec as the decoherence rate is varied from 0 to 14 Mega-ebits/sec (the latter is the average rate of entanglement generation at the link level for this system). Note that memory occupancy is relatively high when there is no decoherence – this is not something that we observe in similar homogeneous-link systems, for which $E[Q]$ is consistently low (less than two qubits in all our numerical observations). We presented these results at MAMA 2019. In [15], we consider a switch that can store one or two qubits per link and can serve both bipartite and tripartite entanglements. We discover that randomized policies allow the switch to achieve a better capacity than time-division multiplexing between bi- and tripartite entanglements, but the advantage decreases as the number of links grows. Panels 2c and 2d illustrate this result. We observe similar capacity region profiles in the presence of state decoherence. This work was presented at QCrypt 2019 as a contributed talk. These results highlight the importance of modeling quantum networks: *their analysis is vital in order to better conceptualize their operation, as well as to discover and address challenges involved in actualizing them.*

A Vision Toward the Future

In future work, I aim to *expand the analysis from a single quantum switch to a network of switches, focusing specifically on achievability of rates for quantum communication and optimal strategies for end-to-end entanglement preparation between nodes that request them.* Some literature already exists in this broad area, *e.g.*, [8], [10], [11], and [13] to name a few. Nevertheless, much of the existing work relies on simplifying assumptions, such as considering simple topologies (*e.g.*, line, grid, or tree networks) or assuming that measurements at switches succeed deterministically. In contrast, my goal is to consider a network with an arbitrary topology, whose links *and* relays (*e.g.*, switches or repeaters) are both heterogeneous, and where

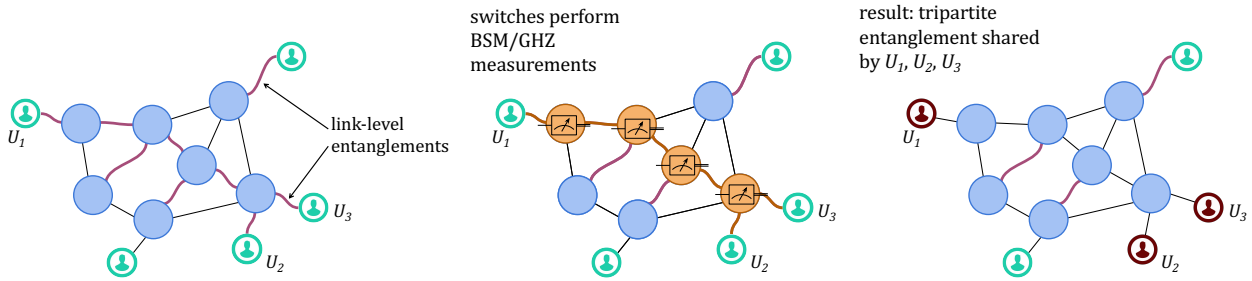

(a) (b) (c)

Figure 3: Entanglement routing in a quantum network. In (a), some of the links have succeeded in creating Bell pairs. In (b), quantum switches perform measurements on locally-held, entangled qubits. (c) is the resulting state of the network: three users share an entangled state and the only entanglement remaining is that of unused Bell pairs.

measurements may fail probabilistically and decoherence may play a significant role. Analytical results for this general setting would be highly applicable to practical scenarios, going beyond the scope of, for example, determining the achievable key generation rates for QKD between users in a quantum network.

Of particular interest for future quantum networks will be the *construction of multipartite entanglement within a network*. These states can be prepared in several ways, but two main classes of strategies exist. In the first, a single node prepares the entanglement and distributes the entangled qubits to nodes that wish to share this state. In the second set of strategies, the network begins by generating link-level entanglements where necessary, and switches perform measurements that eventually result in end-to-end entanglements. An example of this method is shown in Figure 3. Regardless of the method, this process must be accomplished efficiently and reliably. The network must be equipped with ways to counteract the effects of decoherence, errors (*e.g.*, bit or phase flips), and measurement failures. In the first method, for example, qubits may be lost during the distribution phase, and the process may need to restart from the very beginning. Quantum error correction can be implemented to combat this problem. In the second method, in addition to measurement failures, link-level entanglement generation may also fail, and states decohere over time. Purification [3] can be used to combat the latter. While the existence of these remedies is promising for the future of quantum networks, it is yet unclear which workflows and entanglement distribution protocols will produce optimal network performance. I hope to shed light on some of these questions using both analytical tools as well as by further exploring protocol design and making comparisons to existing proposals.

Quantum networks will likely operate alongside classical networks, since many quantum operations, such as teleportation, require classical message exchanges. *For this reason, nodes in the quantum network will need to effectively leverage the classical network to accomplish their tasks.* The exact design and implementation of such hybrid quantum-classical protocols is under active research and discussion, and there is much opportunity for useful and novel contributions. I hope to draw from my prior work on performance and stability analysis of the Transmission Control Protocol, [18] and [17], to gain a unique perspective for the future operation of quantum networks. Due to the uniqueness of the problems encountered in the domain of quantum networking, there arise a number of problem formulations that are of independent theoretical interest, and which are not amenable to standard or existing analytic techniques. A thorough understanding of these problems may allow us to design distributed switching algorithms that efficiently utilize resources (*e.g.*, entanglement) in a quantum network. This is a solid first step toward the design and subsequently, realization of a quantum Internet.

References

[1] C. H. Bennett and G. Brassard. Quantum Cryptography: Public Key Distribution and Coin Tossing. *Theor. Comput. Sci.*, 560(12):7–11, 2014.

[2] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters. Teleporting an Unknown Quantum State Via Dual Classical and Einstein-Podolsky-Rosen Channels. *Physical review letters*, 70(13):1895, 1993.

[3] 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*, 1996.

[4] C. H. Bennett, D. P. DiVincenzo, J. A. Smolin, and W. K. Wootters. Mixed-state Entanglement and Quantum Error Correction. *Physical Review A*, 54(5):3824, 1996.

[5] H.-J. Briegel, W. Dür, J. I. Cirac, and P. Zoller. Quantum Repeaters: the Role of Imperfect Local Operations in Quantum Communication. *Physical Review Letters*, 81(26):5932, 1998.

[6] V. Giovannetti, S. Lloyd, and L. Maccone. Advances in Quantum Metrology. *Nature photonics*, 2011.

[7] R. J. Hughes, J. E. Nordholt, K. P. McCabe, R. T. Newell, C. G. Peterson, and R. D. Somma. Network-Centric Quantum Communications with Application to Critical Infrastructure Protection. *arXiv preprint arXiv:1305.0305*, 2013.

[8] L. Jiang, J. M. Taylor, N. Khaneja, and M. D. Lukin. Optimal Approach to Quantum Communication Using Dynamic Programming. *Proceedings of the National Academy of Sciences*, 2007.

[9] D. Leibfried, M. D. Barrett, T. Schaetz, J. Britton, J. Chiaverini, W. M. Itano, J. D. Jost, C. Langer, and D. J. Wineland. Toward Heisenberg-Limited Spectroscopy with Multiparticle Entangled States. *Science*, 304(5676):1476–1478, 2004.

[10] M. Pant, H. Krovi, D. Towsley, L. Tassiulas, L. Jiang, P. Basu, D. Englund, and S. Guha. Routing Entanglement in the Quantum Internet. 2019.

[11] S. Pirandola. Capacities of Repeater-Assisted Quantum Communications. *arXiv preprint arXiv:1601.00966*, 2016.

[12] E. Schoute, L. Mancinska, T. Islam, I. Kerenidis, and S. Wehner. Shortcuts to Quantum Network Routing. Oct. 2016.

[13] E. Shchukin, F. Schmidt, and P. van Loock. On the Waiting Time in Quantum Repeaters with Probabilistic Entanglement Swapping. *arXiv preprint arXiv:1710.06214*, 2017.

[14] R. Van Meter. *Quantum Networking*. John Wiley & Sons, 2014.

[15] G. Vardoyan, S. Guha, P. Nain, and D. Towsley. On the Capacity Region of Bipartite and Tripartite Entanglement Switching. *arXiv preprint arXiv:1901.06786*, 2019.

[16] G. Vardoyan, S. Guha, P. Nain, and D. Towsley. Performance Evaluation of a Quantum Entanglement Switch. *arXiv preprint arXiv:1903.04420*, 2019.

[17] G. Vardoyan, C. Hollot, and D. Towsley. Towards Stability Analysis of Data Transport Mechanisms: a Fluid Model and an Application. In *IEEE INFOCOM Conference on Computer Communications*, 2018.

[18] G. Vardoyan, N. S. Rao, and D. Towsley. Models of TCP in High-BDP Environments and Their Experimental Validation. In *IEEE 24th International Conference on Network Protocols (ICNP)*, 2016.

[19] W. K. Wootters and W. H. Zurek. A Single Quantum Cannot Be Cloned. *Nature*, 299(5886):802, 1982.