On the Capacity Region of Bipartite and Tripartite Entanglement Switching

Gayane Vardoyan*, Saikat Guha†, Philippe Nain‡, and Don Towsley*
* College of Information and Computer Sciences, University of Massachusetts, Amherst, {gvardoyan, towsley}@cs.umass.edu
†College of Optical Sciences, The University of Arizona, saikat@optics.arizona.edu
‡ Inria, France, philippe.nain@inria.fr

***Abstract*—We study a quantum switch serving a set of users. The function of the switch is to create bi- or tripartite entangled state among users at the highest possible rates at a fixed ratio. We model a set of randomized switching policies. Discovering that some are better than others, we present analytical results for the case where the switch stores one qubit per user, and find that the best policies outperform a time division multiplexing (TDM) policy for sharing the switch between bipartite and tripartite state generation. This performance improvement decreases as the number of users grows. The model is easily augmented to study the capacity region in the presence of qubit decoherence, obtaining similar results. Moreover, decoherence appears to have little effect on capacity. We also study a smaller class of policies when the switch stores two qubits per user.**

I. Introduction

Multi-qubit entangled states are fundamental building blocks for quantum computation, sensing, and security. Consequently there is a need for a quantum network that can generate such entanglement on demand between pairs and groups of users [1]–[3]. In this paper, we study the performance of the simplest multi-user network, a star-topology quantum switch connecting k users, where each user is connected to the switch via a separate link. Bipartite, two-qubit maximally-entangled states, *i.e.*, Bell pairs (or EPR states) are generated at a constant rate across each link, with the qubits getting stored at local quantum memories at each end of the links. As these link entanglements start appearing, the switch uses two-qubit Bell-state measurement (BSM) between pairs of locally-held qubits and three-qubit Greenberger-Horne-Zeilinger (GHZ) basis measurements between triples of locally-held qubits to provide two-qubit and three-qubit entanglements to pairs and triples of users, respectively [4]. The capacity of such a switch to provide these two types of entanglements to the users depends on the switching mechanism, the number of quantum memories and their decoherence rates, and the number of links.

In this paper, we study the capacity region when the switch can store either $B = 1$ or $B = 2$ qubits for each link at any given time. We consider a simple time division multiplexing (TDM) policy between the two types of entanglements, along with a class of randomized policies. When properly configured, the latter outperform TDM. However the relative difference between the two policies goes to zero as $k \to \infty$. We also observe that increasing the number of memories from one to two increases capacity but that the increase diminishes as k gets large. We also explore the effect that decoherence—the locally stored qubits at each end of the link being subject to a noise process that reduces the entanglement between the two qubits—has on capacity. In the cases of $B = 1$ with and without decoherence, we have simple closed form expressions for capacity whereas for the case of $B = 2$, our results are numerical.

The remainder of this paper is organized as follows: in Section II, we formulate the problem and propose a method of solving it. In Section III, we present the case where the system has a per-link buffer of size one, and provide analytical and numerical results. In Section IV, we present numerical results for the case where the system has a per-link buffer of size two and observe similar behavior to the buffer size one case. In section V, we introduce a simple technique for modeling decoherence and use both the buffer size one and two systems to illustrate it. For the former, we also have analytical results. We make concluding remarks in Section VI.

II. System Description and Assumptions

We consider a switch that connects k users over k separate links. The creation of an end-to-end entanglement requires two steps. First two-qubit Bell states are generated pairwise between a qubit stored locally at the switch and a qubit owned by a user. Once such link-level two-qubit entangled states have been created, the switch performs joint (entangling) measurements (over $j \geq 2$ locally-held qubits that are entangled with qubits held by j distinct users), which, if successful, produces a j-qubit maximally-entangled state between the corresponding j users. Link-level entanglement generation, as well as entangling measurements, when realized with practical systems, are inherently probabilistic [5]. We assume that only two-user (two-qubit) and three-user (three-qubit) entanglements are created, i.e., BSMs and 3-qubit GHZ basis measurements are done at the switch. For simplicity, we will assume that these $j = 2$ or 3 qubit measurements at the switch take negligible time and always succeed.

Each link attempts two-qubit entanglements in each time slot of length τ seconds, and with probability p, establishes one entangled pair successfully. For simplicity, we model the time to successfully create a link entanglement as an exponential random variable with mean $1/\mu = \tau/p$. We assume that each link can store $B = 1, 2, \ldots$ qubits. We also assume

that qubits at the switch can decohere and model decoherence time as an exponential r.v. with mean $1/\alpha$. We assume a step-function decoherence model where the two-qubit entanglement goes from a maximally-entangled qubit pair (one ebit) to zero entanglement. In this paper, we only consider $B = 1, 2$. Last, when a qubit is stored at the switch, with its entangled pair stored at a user, we refer to this as a *stored link entanglement*.

We assume that all possible bipartite and tripartite user entanglements are of interest and consider two classes of probabilistic policies, one for $B = 1$ and the second for $B = 2$, that provide the flexibility to generate both types of entanglements with arbitrary rates. Policies in both classes incorporate the *oldest link entanglement first* (OLEF) rule whereby when a link entanglement is created it is always matched up with stored link entanglements when possible rather than be stored. This has the nice consequence, when coupled with the assumption that links are homogeneous but statistically independent, that the system can be modeled by a continuous time Markov chain (CTMC) where the state simply tracks the number of stored entanglements for two users. The next section describes the class of policies for $B = 1$ and Section IV for the class of $B = 2$ policies.

III. System with Per-Link Buffer Size One

In this section, we assume that each link can store one qubit in the buffer, so that the per-link buffer size $B = 1$. We model this system using a CTMC, and by obtaining its stationary distribution, we are able to compute the capacities of bi- and tripartite entanglements. By studying the entire capacity region of this system, we discover that it is always possible to configure the randomized policy to be better than TDM, although as the number of links grows, the advantage of choosing such a policy diminishes.

A. Description

In a system where the switch can make tripartite measurements, we must keep track of two variables for each state of the CTMC: each representing a link with a stored qubit. Hence, $(1, 1)$ represents the state where two of the k links have a qubit stored, one each. Note that we do not need to keep track of all links individually due to the OLEF rule and link homogeneity assumption. States $(1, 0)$ and $(0, 0)$ represent cases where only one link has a stored qubit or no link has a qubit, respectively.

The system is fully described in Figure 1. For a variable $x \in [0, 1]$, we use the notation $\bar{x} \equiv 1 - x$. When the system is in state $(0, 0)$, new entanglements are generated with rate $k\mu$; this is the rate of transitioning from $(0, 0)$ to $(1, 0)$. When the system is in state $(1, 0)$, any new entanglements generated on the link that already has one stored qubit causes the switch to drop one of the qubits. New entanglements on other links are generated with rate $(k - 1)\mu$, and the switch must decide whether to immediately use the two qubits for a BSM or keep both and wait for a new link entanglement. To generalize the policy as much as possible, we add a policy parameter, $r_1 \in [0, 1]$, that specifies the fraction of time the switch performs a BSM. Note that $r_1 = 1$ corresponds to the policy of always using qubits for BSMs. While this maximizes C_2, it also means that $C_3 = 0$.

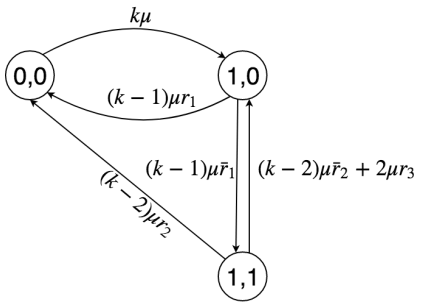

Fig. 1: CTMC for a system with at least three links and buffer size one for each link. k is the number of links, μ is the rate of entanglement generation, and r_1, r_2, and r_3 are parameters that specify the scheduling policy.

Now, suppose that the system is in state $(1, 1)$ and a third link generates an entanglement. This event occurs with rate $(k - 2)\mu$. The switch has two choices: either use all three qubits for a tripartite measurement, or choose two of them for a BSM. We add another policy parameter, $r_2 \in [0, 1]$, that specifies fractions of times the switch performs a BSM and tripartite measurements in the event of three qubits on three different links. Another event that can occur in the $(1, 1)$ state is the generation of an entanglement on either of the two links that already have stored entanglements. This event occurs with rate 2μ. Since $B = 1$, the switch cannot store the new entanglement. A decision must be made: to either discard one of the link entanglements (and remain in state $(1, 1)$) or perform a BSM on two of them and keep the third (and transition to state $(1, 0)$). Since it is not clear which policy is most advantageous, we add another parameter, $r_3 \in [0, 1]$, which specifies the fraction of time that the switch performs a BSM.

B. Numerical Results

We plot the capacity region for the switch with $B = 1$ for all values of r_1, r_2, $r_3 \in [0, 1]$ and compare it against TDM. The entanglement generation rate μ simply scales the capacities, so we set it equal to one. In Figure 2, the number of links is three, and the TDM line is shown in red. Clearly, it is possible to design a policy that yields better performance than TDM: the triangular blue region above TDM represents the maximum capacities of the set of such policies.

Recall that TDM connects points $(0, C_2^*)$ and $(C_3^*, 0)$, where C_2^* and C_3^* are the maximum achievable capacities for bi- and tripartite measurements, respectively. The point farthest from and above TDM (the vertex of the triangular region above the line, shown in green in Figure 2) is achieved by setting $r_1 = 0$ and $r_2 = r_3 = 1$. In other words, the most efficient policy is the following: (i) never perform BSMs in state $(1, 0)$; and (ii) when in state $(1, 1)$ and a third entanglement is generated on a *different* link, always use it in a tripartite measurement, but

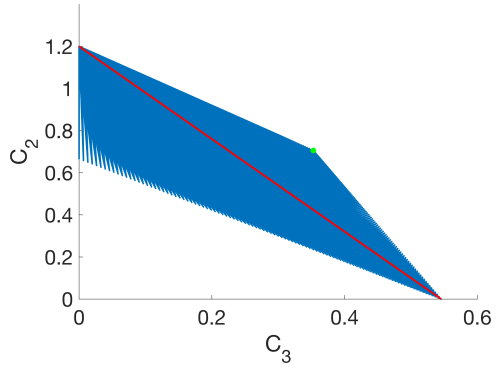

Fig. 2: Capacity region for a system of buffer size one and three links. The red line represents the set of TDM policies.

when a third entanglement is generated on one of the links that already has a stored qubit, *always* perform a BSM. Note that the latter rule directs the switch to not waste an entanglement whenever it is possible to use it in a measurement.

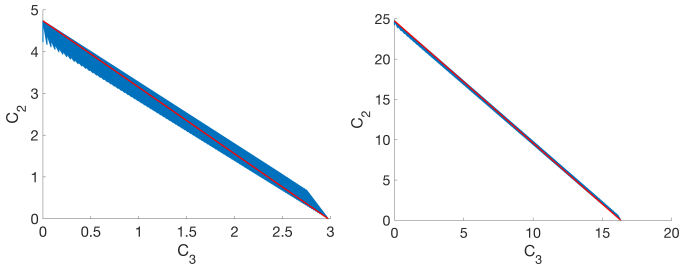

(a) $k = 10$ (b) $k = 50$

Fig. 3: Capacity region for a system of buffer size one and varying number of links. The red line represents the set of TDM policies.

The capacity regions for $k = 10$ and 50 are shown in Figure 3. Note that as the number of links increases, the differences between TDM and the more efficient random policies diminish. In the next section, we provide an analytical proof of this phenomenon.

C. Analysis

Let $\pi(0,0)$, $\pi(1,0)$, and $\pi(1,1)$ represent the stationary distribution of the CTMC in Figure 1. The balance equations (excluding μ, as it cancels out due to every transition rate being its multiple), are:

$$\pi(0,0)k = \pi(1,0)(k-1)r_1 + \pi(1,1)(k-2)r_2,$$
$$\pi(1,1)((k-2)r_2 + (k-2)\bar{r}_2 + 2r_3) = \pi(1,0)(k-1)\bar{r}_1,$$
$$\pi(0,0) + \pi(1,0) + \pi(1,1) = 1.$$

Solving these equations, we obtain

$$\pi(1,1) = \frac{k(k-1)\bar{r}_1}{D},$$
$$\pi(1,0) = \frac{k(k-2+2r_3)}{D}, \text{ where}$$
$$D = (k-2+2r_3)((k-1)r_1 + k) + (k-1)\bar{r}_1((k-2)r_2 + k).$$

Then the bi- and tripartite capacities for this system, $C_2 \equiv C_2(r_1, r_2, r_3)$ and $C_3 \equiv C_3(r_1, r_2, r_3)$, are

$$\begin{aligned} C_2 &= \pi(1,0)(k-1)\mu r_1 + \pi(1,1)((k-2)\mu\bar{r}_2 + 2\mu r_3) \\ &= \frac{k(k-1)\mu(k-2+2r_3-(k-2)r_2\bar{r}_1)}{(k-2+2r_3)((k-1)r_1+k)+(k-1)\bar{r}_1((k-2)r_2+k)}, \\ C_3 &= \pi(1,1)(k-2)\mu r_2 \\ &= \frac{k(k-1)(k-2)\mu r_2\bar{r}_1}{(k-2+2r_3)((k-1)r_1+k)+(k-1)\bar{r}_1((k-2)r_2+k)}. \end{aligned}$$

Claim 1: $C_2(r_1, 0, 1) = C_2(1, 0, r_3)$ and $C_3(0, 1, 0)$ are the maxima for C_2 and C_3, *i.e.* they are C_2^* and C_3^*.

Proof: We start by proving this for C_2. First, note that to maximize C_2's numerator and minimize its denominator, r_2 must be set to 0. This yields

$$\begin{aligned} C_2(r_1, 0, r_3) &= \frac{k(k-1)\mu(k-2+2r_3)}{(k-2+2r_3)((k-1)r_1+k)+k(k-1)\bar{r}_1} \\ &= \frac{k(k-1)\mu}{(k-1)r_1 + k + \frac{k(k-1)\bar{r}_1}{k-2+2r_3}}. \end{aligned}$$

To minimize the denominator above, we should set $r_3 = 1$, which yields

$$C_2(r_1, 0, 1) = \frac{k(k-1)\mu}{(k-1)r_1 + k + (k-1)\bar{r}_1} = \frac{k(k-1)\mu}{2k-1}.$$

Note that $C_2(1, 0, r_3)$ yields the same expression as $C_2(r_1, 0, 1)$. Next, consider the expression for C_3. To minimize the denominator, we should set $r_3 = 0$. This yields

$$\begin{aligned} C_3(r_1, r_2, 0) &= \frac{k(k-1)(k-2)\mu r_2\bar{r}_1}{(k-2)((k-1)r_1+k)+(k-1)\bar{r}_1((k-2)r_2+k)} \\ &= \frac{k(k-1)(k-2)\mu r_2\bar{r}_1}{(k-1)r_1((k-2)\bar{r}_2-k)+k(k-2)+(k-1)((k-2)r_2+k)}. \end{aligned}$$

It is clear that r_1 must be 0, which yields

$$\begin{aligned} C_3(0, r_2, 0) &= \frac{k(k-1)(k-2)\mu r_2}{k(k-2)+(k-1)((k-2)r_2+k)} \\ &= \frac{k(k-1)(k-2)\mu r_2}{k(2k-3)+(k-1)(k-2)r_2} \\ &= \frac{k(k-1)(k-2)\mu}{\frac{k(2k-3)}{r_2}+(k-1)(k-2)}. \end{aligned}$$

From above, we can see that r_2 must be 1, so the maximum is at $C_3(0, 1, 0)$. □

For brevity, let $(C_3(0,1,1), C_2(0,1,1)) \equiv (\hat{C}_3, \hat{C}_2)$; this is the point farthest above the TDM line within the achievable capacity region (*e.g.* the green point in Figure 2). We prove this as part of the proof of the claim below.

Claim 2: Any point (C_3, C_2) in the achievable capacity region satisfies the following constraints:

$$C_2 \le -\frac{3k-2}{2k-1}C_3 + \frac{\mu k(k-1)}{2k-1} \text{ and} \tag{1}$$
$$C_2 \le -\frac{k(k-2)+2(k-1)^2}{k(k-2)}C_3 + \mu(k-1). \tag{2}$$

Moreover (1) and (2) define a tight upper bound on the achievable capacity region.

Proof: First, we must show that the point $(\hat{C}_3, \hat{C}_2)$ is indeed the farthest from the TDM line. To do so, let us find a point (C_3, C_2) on the plane such that the (negative) slope of the line that passes through it and $(0, C_2^*)$ is maximized. This is equivalent to minimizing the quantity

$$\frac{C_2^* - \hat{C}_2}{\hat{C}_3} = \frac{(3k-2)(k-2)r_2 + 2(k-1)(1-r_3)}{r_2(2k-1)(k-2)}.$$

To do so, we must set $r_3 = 1$. Next, note that the TDM line is given by the equation $f(x,y) = y - C_2^*(1 - x/C_3^*)$, and the distance between it and any point (C_3, C_2) is given by $|f(C_3, C_2)|/\sqrt{1 + (C_2^*/C_3^*)^2}$. Hence, it is sufficient to maximize $|f(C_3(r_1, r_2, 1), C_2(r_1, r_2, 1))|$, given by

$$\frac{2\mu k(k-1)}{(2k-1)\left(k - 2 + \frac{2k^2-k}{(k-1)r_2(1-r_1)}\right)}.$$

It is clear that we must set $r_2 = 1$ and $r_1 = 0$, yielding $(\hat{C}_3, \hat{C}_2)$ as the point farthest from the TDM line, as expected.

Next, consider the line passing through $(0, C_2^*)$ and $(\hat{C}_3, \hat{C}_2)$:

$$y_1 = -\frac{3k-2}{2k-1}x_1 + \frac{\mu k(k-1)}{2k-1}, \tag{3}$$

and the line passing through $(\hat{C}_3, \hat{C}_2)$ and $(C_3^*, 0)$:

$$y_2 = -\frac{k(k-2) + 2(k-1)^2}{k(k-2)}x_2 + \mu(k-1). \tag{4}$$

It is not hard to show that for any point (C_3, C_2), (1) and (2) are satisfied. In other words, all points in the achievable capacity region fall on or below these two lines. To prove that this upper bound is tight, it remains to show that all points on lines (3) and (4) are achievable. To see this, let $r_1 = 0$ and $r_3 = 1$:

$$C_2(0, r_2, 1) = \frac{(k - (k-2)r_2)k(k-1)\mu}{k^2 + (k-1)(k + (k-2)r_2)},$$
$$C_3(0, r_2, 1) = \frac{(k-2)r_2 k(k-1)\mu}{k^2 + (k-1)(k + (k-2)r_2)}.$$

Note that any point $(C_3(0, r_2, 1), C_2(0, r_2, 1))$ is *on* line (3), and these two functions are continuous in $r_2 \in [0, 1]$. Similarly, letting $r_1 = 0$ and $r_2 = 1$, we have

$$C_2(0, 1, r_3) = \frac{2r_3 k(k-1)\mu}{k(k-2+2r_3) + 2(k-1)^2},$$
$$C_3(0, 1, r_3) = \frac{k(k-1)(k-2)\mu}{k(k-2+2r_3) + 2(k-1)^2}.$$

Any point $(C_3(0, 1, r_3), C_2(0, 1, r_3))$ is *on* line (4), and these two functions are continuous in $r_3 \in [0, 1]$. Using these facts, we conclude that *all* points on (3) and (4) are achievable. □

Claim 3: As $k \to \infty$, the benefit of using an alternate policy (one that lies above TDM) diminishes.

Proof: We prove this by showing that as $k \to \infty$, the ratio of the achievable area above the TDM line, which we call $A_\triangle$ (because this area has the shape of a triangle) to the total area below the capacity region, which we call A_T, goes to zero. For $A_\triangle$, the length of the base of the triangle is simply the distance between the points $(0, C_2^*)$ and $(C_3^*, 0)$, or $\sqrt{(C_2^*)^2 + (C_3^*)^2}$. The height is given by $|f(\hat{C}_3, \hat{C}_2)|/\sqrt{1 + (C_2^*/C_3^*)^2}$. Then

$$A_\triangle = \frac{|f(\hat{C}_3, \hat{C}_2)|C_3^*}{2}.$$

Then, the area below the TDM line is given by

$$A_{TDM} = \frac{(C_2^*)^2 + (C_3^*)^2}{4}, \text{ so the total area is}$$
$$A_T = A_\triangle + A_{TDM} = \frac{2|f(\hat{C}_3, \hat{C}_2)|C_3^* + (C_2^*)^2 + (C_3^*)^2}{4}.$$

Then the ratio of the area above the TDM to the total area is

$$\frac{A_\triangle}{A_T} = \frac{2|f(\hat{C}_3, \hat{C}_2)|C_3^*}{2|f(\hat{C}_3, \hat{C}_2)|C_3^* + (C_2^*)^2 + (C_3^*)^2} = \frac{1}{1 + \frac{(C_2^*)^2 + (C_3^*)^2}{2|f(\hat{C}_3, \hat{C}_2)|C_3^*}}.$$

To prove that this ratio goes to zero with k, it suffices to show that the second term in the denominator goes to ∞. It can be shown that $((C_2^*)^2 + (C_3^*)^2)/2|f(\hat{C}_3, \hat{C}_2)|C_3^*$ simplifies to

$$\frac{39k^6 - 220k^5 + 493k^4 - 568k^3 + 362k^2 - 120k + 16}{4(6k^5 - 33k^4 + 67k^3 - 62k^2 + 26k - 4)} \overset{k\to\infty}{\longrightarrow} \infty.$$

□

IV. System with Per-Link Buffer Size Two

In a system with per-link buffer size two, there are three additional states, as shown in Figure 4. The goal of this part of the study is to show the existence of better policies than TDM, rather than to find the optimal policy. Hence, the design in Figure 4 does not encapsulate all possible policies: for instance, there is no r_1 parameter here, since our exhaustive search over the entire parameter space for the system with $B = 1$ revealed that r_1 is best set to zero. In addition, note that if the system is in state $(1, 1)$ and another entanglement is generated on one of the links that already has a stored qubit, the system is not allowed to use two of the qubits for a BSM. The reasoning is that since $B = 2$, there is enough space to keep the new qubit. Similarly, when the system is in state $(2, 1)$ a BSM is only allowed if (i) another entanglement is generated on one of the $k - 2$ links that does not have a stored qubit, or (ii) another entanglement is generated on the link that already has two qubits stored. In the latter scenario, not performing a BSM would cause a qubit to be discarded. While this design does not grant the switch access to the full range of policies, it does enable us to find a class of policies that are more efficient than TDM.

Figure 5 shows capacity regions for $B = 2$ with number of links $k = 3$ and 10. We observe that policies more efficient than TDM can be found, but as the number of links grows, the advantage of such policies relative to TDM diminishes. This phenomenon mimics that of the $B = 1$ switch. Figure 6 shows a comparison of $B = 1$ and $B = 2$ switches for three and ten links. We observe that while there is a clear benefit to extra buffer space for a small number of users,

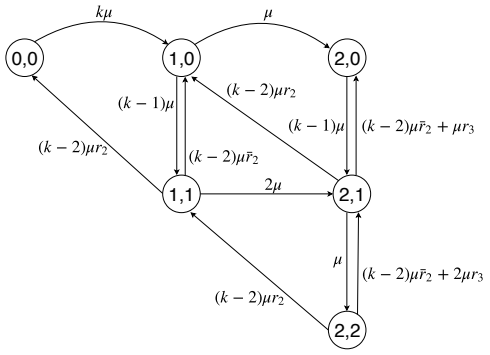

Fig. 4: CTMC for a system with at least three links and buffer size two for each link. k is the number of links, μ is the rate of entanglement generation, and r_2 and r_3 are parameters that specify the scheduling policy.

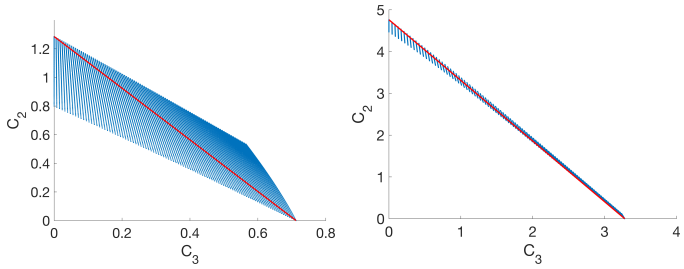

(a) $k = 3$ (b) $k = 10$

Fig. 5: Capacity region for per-link buffer size $B = 2$, for $k = 3$, 10 links. The red line represents the set of TDM policies.

the advantage becomes less apparent as the number of users grows. In addition, it appears that C_3 benefits more from the extra buffer space than C_2.

V. Modeling Decoherence

In this section, we present a simple way to augment the model of Section III to account for the decoherence of quantum states. We also present numerical results of adding decoherence to the system from Section IV. Our decoherence model is described in Section II. For $B = 1$, the resulting CTMC is illustrated in Figure 7.

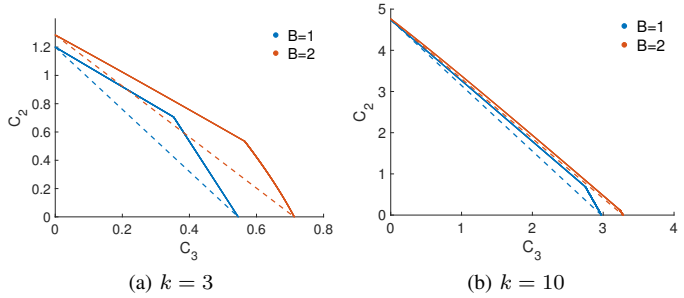

(a) $k = 3$ (b) $k = 10$

Fig. 6: Comparison of capacity regions for systems of buffer sizes one and two with varying number of links k, and entanglement generation rate $\mu = 1$.

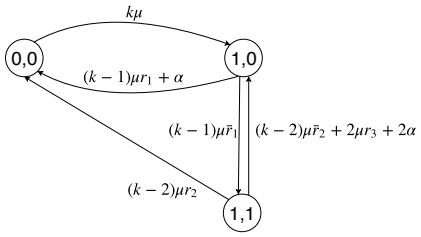

Fig. 7: CTMC for a system with at least three links and buffer size one. k is the number of links, μ is the rate of entanglement generation, α is the decoherence rate, and r_1, r_2, r_3 are parameters that specify the scheduling policy.

The analysis of this model is almost identical to that of the $B = 1$ system without decoherence. As with the latter, the capacity region is bounded above by two lines:

$$y_1 = -\frac{\mu(3k-2)(\alpha+(k-2)\mu)+2\alpha^2}{\mu(k-2)((2k-1)\mu+\alpha)}x_1 + \frac{k(k-1)\mu^2}{(2k-1)\mu+\alpha},$$
$$y_2 = -\frac{2(k-1)^2\mu^2+(k\mu+\alpha)((k-2)\mu+2\alpha)}{\mu(k-2)(k\mu+\alpha)}x_2 + \frac{k(k-1)\mu^2}{k\mu+\alpha}.$$

To avoid ambiguity, let C_2' and C_3' denote the bi- and tripartite capacities of a system with decoherence. As with the previous model, C_2' is maximized at $r_1 = 1$, $r_2 = r_3 = 0$; C_3' is maximized at $r_1 = r_3 = 0$, $r_2 = 1$, and the point farthest from TDM is obtained by setting $r_1 = 0$, $r_2 = r_3 = 1$. The first bounding line passes through the points $(0, C_2'(1,0,0))$ and $(C_3'(0,1,1), C_2'(0,1,1))$; and the second line passes through $(C_3'(0,1,1), C_2'(0,1,1))$ and $(C_3'(0,1,0), 0)$. Moreover, all points *on* the bounding lines are achievable, indicating that the bound is tight.

The capacities are given by

$$C_2' = (k(k-1)\mu^2\,(2(\alpha r_1+\mu r_3)+(k-2)\mu(1-r_2\bar{r}_1)))/D,$$
$$C_3' = (k\mu^3(k-1)(k-2)\bar{r}_1 r_2)/D, \text{ where}$$
$$D = (k-1)\mu\bar{r}_1((k-2)\mu r_2+k\mu) + (k\mu+(k-1)\mu r_1+\alpha)((k-2+2r_3)\mu+2\alpha).$$

Note that the denominator is quadratic in α. This causes both C_2' and C_3' to tend to zero as $\alpha \to \infty$.

Figure 8 shows a comparison of the capacity regions for systems with $B = 1$, for three and ten number of links and different decoherence rates. For all cases, μ is set to one: for qualitative results, we only need to concern ourselves with the value of α *relative* to μ. In real scenarios, we expect α to be at least one order of magnitude less than μ. From numerical results, we observe that the effect of decoherence on the capacity region is not significant, especially as the number of links grows. Analysis supports this observation, since we can show that

$$\lim_{k\to\infty}\frac{C_2'}{C_2} = 1 \text{ and } \lim_{k\to\infty}\frac{C_3'}{C_3} = 1.$$

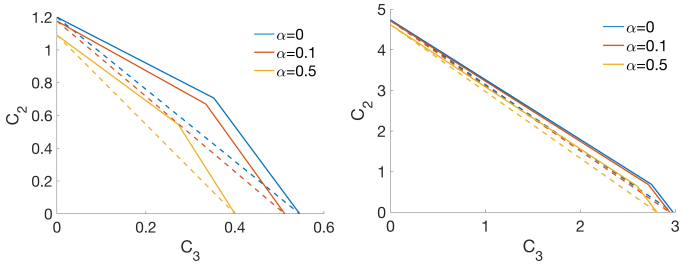

(a) $k = 3$ (b) $k = 10$

Fig. 8: Capacity region for a system of buffer size one and varying number of links k, decoherence rates α, and entanglement generation rate $\mu = 1$. The solid lines are the upper boundaries of the capacity region, and the dashed are TDM lines.

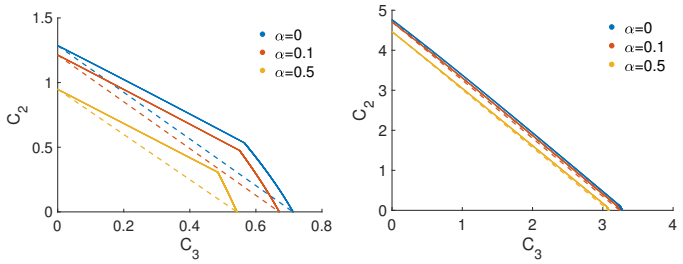

(a) $k = 3$ (b) $k = 10$

Fig. 9: Capacity region for a system of buffer size two and varying number of links k, decoherence rates α, and entanglement generation rate $\mu = 1$. The solid lines are the upper boundaries of the capacity region, and the dashed are TDM lines.

Figure 9 shows a comparison for systems with $B = 2$ and varying number of links and decoherence rates. Results are consistent with that of the $B = 1$ case: the effects of decoherence on capacity are less apparent for larger k values.

VI. Conclusion

In this work, we explore a set of policies for a quantum switch that can store up to two qubits per link and whose objective is to perform bi- and tripartite joint measurements to distribute two and three qubit entanglement to pairs and triples of users. We present analytical results for the case where the per-link buffer has size one. By comparing against TDM policies, we discover that better policies in terms of achievable bi- and tripartite capacities exist, but that as the number of links grows, the advantage of using such policies diminishes. We also compare the capacity regions for systems with different per-link buffer sizes and observe that systems with fewer links benefit more from the extra storage space than systems with a larger number of links. Finally, we model decoherence for both types of systems and present analytical results for the case with per-link buffer size one. Observations and analysis show that as the number of links increases, the effects of decoherence become less apparent on systems.

References

[1] M. Pant, H. Krovi, D. Towsley, L. Tassiulas, L. Jiang, P. Basu, D. Englund, and S. Guha, "Routing Entanglement in the Quantum Internet," 2019.
[2] E. Schoute, L. Mancinska, T. Islam, I. Kerenidis, and S. Wehner, "Shortcuts to quantum network routing," Oct. 2016.
[3] R. Van Meter, *Quantum networking*. John Wiley & Sons, 2014.
[4] M. A. Nielsen and I. Chuang, "Quantum Computation and Quantum Information," 2002.
[5] S. Guha, H. Krovi, C. A. Fuchs, Z. Dutton, J. A. Slater *et al.*, "Rate-loss Analysis of an Efficient Quantum Repeater Architecture," *Phys. Rev. A*.