On the guaranteed throughput and efficiency of closed re-entrant lines*†

James R. Morrison‡and P. R. Kumar§
Dept. of Electrical and Computer Engineering, and
Coordinated Science Laboratory
University of Illinois, Urbana–Champaign

Abstract

A closed network is said to be "guaranteed efficient" if the throughput converges under all non-idling policies to the capacity of the bottlenecks in the network, as the number of trapped customers increases to infinity.

We obtain a necessary condition for guaranteed efficiency of closed re-entrant lines. For balanced two-station systems, this necessary condition is almost sufficient, differing from it only by the strictness of an inequality.

This near characterization is obtained by studying a special type of virtual station called "alternating visit virtual station." These special virtual stations allow us to relate the necessary condition to certain indices arising in heavy traffic studies using a Brownian network approximation, as well as to certain policies proposed as being extremal with respect to the asymptotic loss in the throughput.

Using the near characterization of guaranteed efficiency we also answer the often pondered question of whether an open network or its closed counterpart has greater throughput – the answer is that neither can assure a greater guaranteed throughput.

*Please address all correspondence to the second author at University of Illinois, CSL, 1308 West Main Street, Urbana, IL 61801, USA.

†The research reported here has been supported in part by the U.S. Army Research Office under Contract No. DAAH-04-95-1-0090, the National Science Foundation under Grant No. ECS-94-03571, and the Joint Services Electronics Program under Contract No. N00014-96-1-0129.

‡Email: morrison@decision.csl.uiuc.edu

§Email: prkumar@decision.csl.uiuc.edu. WWW: http://black.csl.uiuc.edu/ ~prkumar

1 Introduction

We consider the throughput of closed re-entrant lines. A scheduling policy u is said to be *efficient* if the throughput λ_N^u converges to the bottleneck capacity λ^* as the number of trapped customers N increases to infinity. The network itself is said to be *guaranteed efficient* if it is efficient under *all* non-idling scheduling policies.

The notion of efficiency is the analog for closed systems of the concept of "stability" for open systems. In [1] and [2] it was shown that under some scheduling policies the two open networks in Figure 1 are unstable when $m_2 + m_4 > 1$, even though $m_1 + m_4 < 1$ and $m_2 + m_3 < 1$. This instability is caused by an unstable eigenvalue $\frac{m_4}{(1-m_2)} > 1$, and leads to a throughput no more than $\frac{1}{m_2+m_4}$.

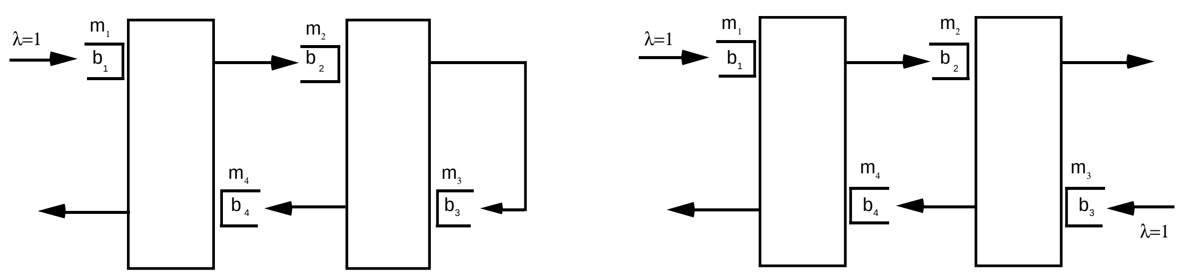

Figure 1: Unstable open networks. The service time for buffer b_i is m_i. The arrival rate is denoted λ.

In [3], the issue of efficiency was raised for closed systems, and a sufficient condition based on the value of a linear program was provided both for guaranteed efficiency as well as efficiency of certain policies. In [4], Harrison and Nguyen (see also Dumas [5]) proved that the closed network in Figure 2, the closed counterpart of the unstable open networks in Figure 1, was indeed inefficient. This was accomplished by noting that under certain scheduling policies, buffers b_2 and b_4 could not simultaneously be in service.

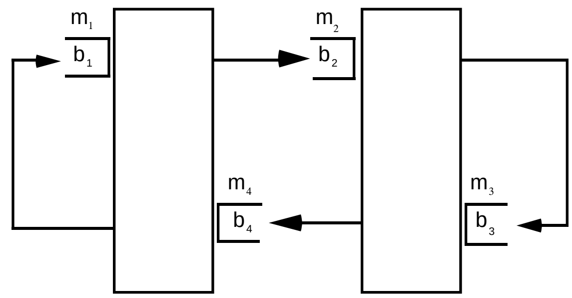

Figure 2: An inefficient closed network.

This observation of "virtual stations" was further developed for open networks by Dai and Vande Vate [6] who identified necessary and nearly sufficient conditions for guaranteed stability under all non-idling policies of two-station open networks. These necessary conditions were extended by Hasenbein [7] to open networks with more than two stations.

Meanwhile, pursuing a formal Brownian network approximation of two-station closed networks, Harrison and Wein [8] developed a certain buffer priority policy (HW-policy) as a candidate for having nearly optimal throughput as the number of trapped customers increases. The priorities were expressed in terms of certain buffer indices, where the index η_i of buffer b_i is the difference in the mean work remaining to be done by the two stations, for a customer starting in buffer b_i, prior to that customer's next exit from some fixed buffer b_L. The performance measure $\lim_{N\to\infty} \frac{N(\lambda^* - \lambda_N^u)}{\lambda^*}$ was dubbed the "asymptotic loss" in [9]. It was established there that the HW-policy was efficient and that no policy could have lesser asymptotic loss than that conjectured for the HW-policy. Further, and very closely related to the present paper, its exact opposite, the Anti-HW policy, was also studied. A sufficient condition was obtained for the efficiency of all non-idling policies, and indeed a finite upper bound was provided for the asymptotic loss of all non-idling policies under this sufficient condition.

In this paper, we combine all these separately developing strands of work concerning the throughput of closed re-entrant lines. We show that except for the sharpness of the inequality, the condition identified in [9] is also necessary for the guaranteed efficiency of two-station systems. This is done by exploring the notion of virtual multiserver stations for closed re-entrant lines, and determining a necessary condition for guaranteed efficiency of multi-station closed re-entrant lines. For two-station systems, a special type of a virtual station, called an alternating visit virtual station, is identified, which serves to nearly characterize guaranteed efficiency. The HW indices are shown to be intimately related to the workloads of these alternating visit virtual stations. The Anti-HW policy is related to the most constrictive of these alternating visit virtual stations.

2 System description

Consider a system as shown in Figure 3, consisting of M stations, labeled $\sigma_1, \sigma_2, \ldots, \sigma_M$, and L buffers, labeled $b_1, b_2, \ldots, b_L$. Buffer b_i is served by station $\sigma(i) \in \{\sigma_1, \sigma_2, \ldots, \sigma_M\}$. There is a population of N trapped customers which follow a closed deterministic route. After completing service at b_i, a customer moves next to buffer b_{i+1}, except that after completing service at b_L it moves next to buffer b_1; equivalently, the buffer visited after b_i is $b_{(i+1) \bmod L}$. The service time required for customers in b_i is exponentially distributed with a mean value of $m_i > 0$. All service times are independent of each other. Each station σ can only serve one customer at a time from the buffers $B(\sigma) := \{b_i : \sigma(i) = \sigma\}$ that it caters to. To avoid trivialities, we shall assume that $M \geq 2$ and that $B(\sigma)$ is nonempty for every $1 \leq \sigma \leq M$. Such a system is called a *closed re-entrant line*.

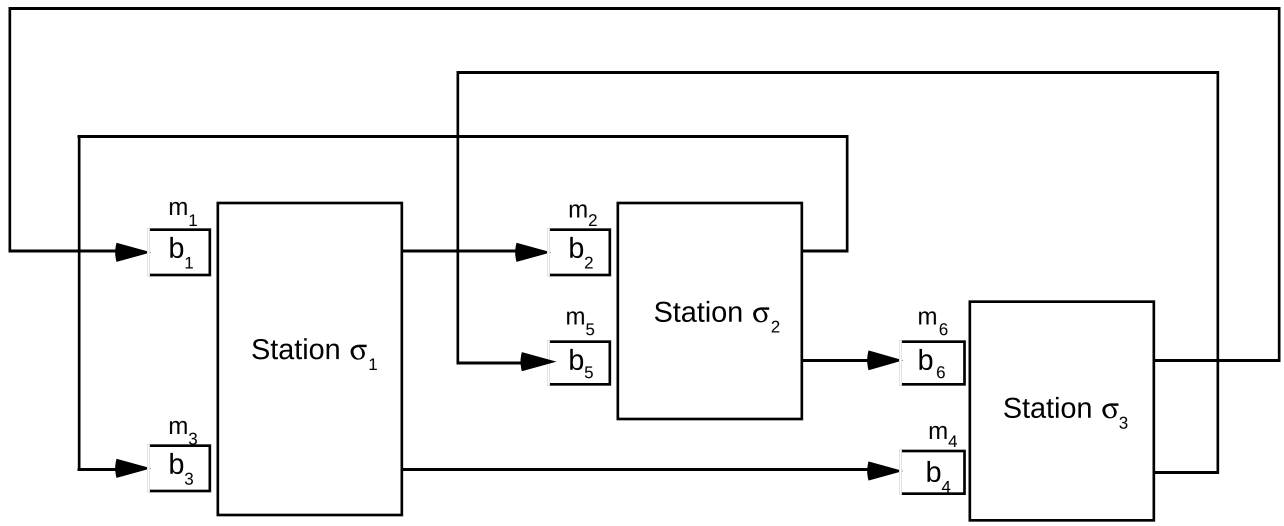

Figure 3: A closed re-entrant line.

The choice of a customer to serve is governed by a scheduling policy, which is assumed to be non-anticipative and *non-idling*. By "non-idling" it is meant that a station σ can only remain idle when all its buffers $B(\sigma)$ are empty of customers; otherwise it has necessarily to serve some one of the customers waiting for its attention. We shall denote by U the set of all non-anticipative and non-idling scheduling policies.

A particular class of non-idling scheduling policies which play an important role in this paper are *buffer priority* policies. These are policies specified by a permutation $\pi = (\pi(1), \pi(2), \ldots, \pi(L))$ of the set $\{1, 2, \ldots, L\}$, with the requirement that a station σ can only serve a customer $b_{\pi(i)} \in B(\sigma)$ only if all the buffers $\{b_{\pi(j)} \in B(\sigma) : j < i\}$ are empty of customers. The priority is implemented in a pre-emptive resume fashion.

Let $x_i(t)$ denote the number of customers in buffer b_i at time t, including any in service. Let $x(t) = (x_1(t), \ldots, x_N(t))$. Since the system is closed, one has $\sum_{i=1}^{L} x_i(t) = N$ for all $t \geq 0$.

Given a scheduling policy $u \in U$, and an initial condition for the system describing the positions of the N trapped customers, we call the random variable

$$\lambda_N^u := \liminf_{T \to +\infty} \frac{1}{T}(\text{Number of departures from } b_L \text{ in the time interval } [0,T]) \qquad (1)$$

the *throughput* of the network.

Let

$$\lambda^* \ := \ \underset{\sigma \in \{\sigma_1,\ldots,\sigma_M\}}{\text{Min}} \frac{1}{\sum_{i \in \sigma} m_i}. \qquad (2)$$

(For notational simplicity we write "$i \in \sigma$" for $b_i \in B(\sigma)$). It is easy to see that the throughput of the system under any scheduling policy $u \in U$ can never exceed λ^*, i.e.,

$$\lambda_N^u \ \leq \ \lambda^* \text{ a.s. for all } u \in U \text{ and all } N \geq 1.$$

We call λ^* the *throughput capacity* of the system. Any station σ which attains the "Min" in (2) is a *bottleneck* station. A system is called *balanced* if all stations are bottlenecks.

Given a scheduling policy[1] $u \in U$, we shall say that it is *efficient* if

$$\lim_{N \to +\infty} \lambda_N^u \ = \ \lambda^* \text{ a.s.}$$

The issue examined in this paper is this: For what closed re-entrant lines is it true that *all* non-idling scheduling policies are efficient? When this is the case we shall say that the closed re-entrant line is *guaranteed efficient.*

If there is a number $\underline{\lambda}$ such that

$$\liminf_{N \to +\infty} \lambda_N^u \ \geq \ \underline{\lambda} \text{ a.s. for all } u \in U,$$

[1]To be more precise, one needs to consider a family of scheduling policies, one for each population level N, since a scheduling policy may depend for its definition on the value of N. This is accommodated by simply expanding the definition of a scheduling policy to depend on N. Similarly one should understand that for each N there is a different initial condition.

we shall say that $\underline{\lambda}$ is a *guaranteed throughput.*

The first main result in this paper (Theorem 1) is an upper bound on the guaranteed throughput. Employing it, we show that for all closed balanced re-entrant lines with $M = 2$, i.e., two stations, a necessary condition for guaranteed efficiency is that

$$\operatorname*{Min}_{i\in\sigma_1} \eta_i \quad \geq \quad \operatorname*{Max}_{j\in\sigma_2} \eta_j. \tag{3}$$

Above,

$$\eta_i \quad := \sum_{\{j:j\geq i \text{ and } b_j\in B(\sigma_1)\}} m_j - \sum_{\{k:k\geq i \text{ and } b_k\in B(\sigma_2)\}} m_k$$

is the difference between the mean works to be done by the two stations σ_1 and σ_2 on a customer situated in b_i, prior to that customer's next departure from b_L. These *indices* $\{\eta_i\}$ were first defined in Harrison and Wein [8], and we shall call them the HW-indices.

Except for the strictness of the inequality, this coincides with the *sufficient* condition for guaranteed efficiency

$$\operatorname*{Min}_{i\in\sigma_1} \eta_i \quad > \quad \operatorname*{Max}_{j\in\sigma_2} \eta_j \tag{4}$$

obtained in Jin, Ou and Kumar [9].

Thus, except for the strictness of the inequality, we have characterized those two station closed re-entrant lines which are guaranteed efficient.

It is worth noting that given the topology of a two-station closed re-entrant line, if we view $\{m_i : 1 \leq i \leq L\}$ as the parameters describing the system, then the region of guaranteed efficiency is convex, in fact a polytope, except for some indeterminacy concerning the boundary points.

The issue of strictness of the inequality can perhaps be explained by the following observation. When the inequality is strict, then it is proved in [9] that not only are all non-idling policies efficient, but that they also have bounded *asymptotic loss*, i.e.,

$$\limsup_{N\to\infty} \frac{N(\lambda^* - \lambda^*_N)}{\lambda^*} \le \bar{J} < +\infty.$$

Thus if one only has "$\geq$" rather than "$>$" in (4), then perhaps one has efficiency but infinite asymptotic loss (for example, if the convergence of λ^u_N to λ^* is logarithmic rather than polynomial). This is akin to a similar issue which arises in the study of open networks where one may have null recurrence rather than positive recurrence on the boundary of the stability region. Of course, given our current knowledge it is also possible that in the indeterminate region one may have inefficiency, or at the other extreme, even bounded asymptotic loss.

The reader should note that the assumption that the service times are exponential is made only for simplicity of exposition. The results continue to hold if the service times at b_i obey the law of large of numbers, with empirical mean converging to m_i, and if the probability of two or more service completions at the same time is zero.

Notation

Throughout this paper we shall use the following notations. If $C \subseteq \{b_1, \ldots, b_L\}$ is a subset of buffers, we shall denote by $S(C)$ the set of stations serving any of the buffers in C, i.e.,

$$\begin{aligned} S(C) &:= \{\sigma : \exists\ i \in \sigma \text{ with } b_i \in C\} \\ &= \{\sigma(i) : b_i \in C\}. \end{aligned}$$

Conversely, we shall denote by $C(\sigma)$ the set of buffers in C which are served by σ, i.e.,

$$\begin{aligned} C(\sigma) &:= \{b_i \in C : i \in \sigma\} \\ &= B(\sigma) \cap C. \end{aligned}$$

Also, given any set A, we shall denote its cardinality by $|A|$. Thus, $|S(C)|$ will denote the number of stations which serve any of the buffers in C.

We use the notation "$k \in C$" as a shorthand for "$b_k \in C$." We shall also denote by $W_C := \sum_{k \in C} m_k$, the sum of the mean service times of the buffers in C. With this notation, note that

$$\lambda^* \;=\; \operatorname*{Min}_{\sigma} \frac{1}{W_{B(\sigma)}}.$$

We shall also interpret all buffer indices "mod L". Thus whenever, we write b_{i+j}, we shall understand it as $b_{(i+j) \bmod L}$. Similarly, by $\sigma(i+j)$ we mean $\sigma((i+j) \bmod L)$.

3 Virtual multiserver station

In Section 4, we will show that under certain buffer priority policies some sets of buffers behave like a station served by multiple servers. In this section we define the notion of such a "virtual multiserver station." Later in Section 4 we shall examine its properties. This definition is an appropriate modification of that in [6] for open systems.

Definition: Virtual multiserver station. *We call a nonempty subset C of buffers a* virtual multiserver station *if it satisfies the following two properties:*

(i) *Whenever $b_k \in C$ and $\sigma(k-1) = \sigma(k)$, then $b_{k-1} \in C$.*

(ii) *Whenever $b_k \in C$ and $\sigma(k-1) \neq \sigma(k)$, then $b_{k-1} \notin C$ and $\sigma(k-1) \in S(C)$.*

In words, if a buffer is in C, and the previous buffer visited by customers is located at the *same* station, then the previous buffer is necessarily to be in C. On the other hand, if a buffer is in C and the previous buffer is located at a *different* station, then the previous buffer should *not* be in C; however, some other buffer at the station serving the previous buffer should be in C.

The following example clarifies the concept.

Example 1: A Virtual multiserver station

Consider the closed re-entrant line shown in Figure 4.

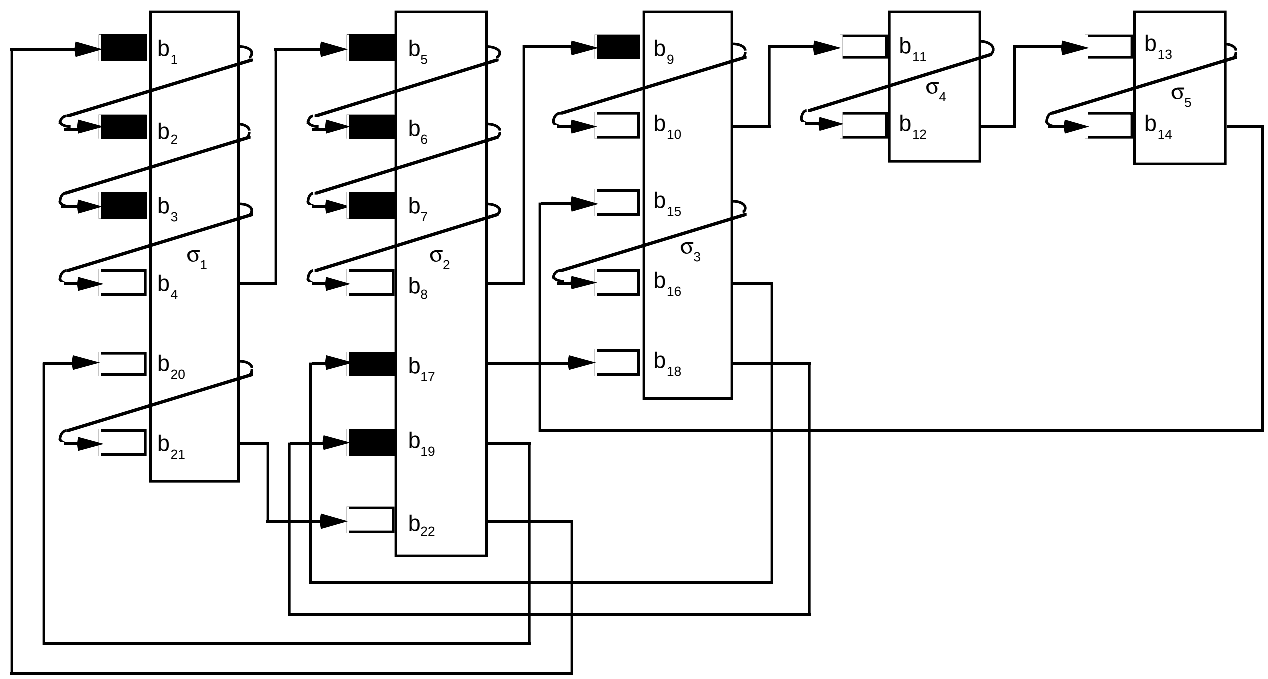

Figure 4: A virtual multiserver station. The buffers in C are colored solid black.

The set $C = \{b_1, b_2, b_3, b_5, b_6, b_7, b_9, b_{17}, b_{19}\}$ satisfies the definition of a virtual multiserver station. Consider, for example, buffer b_6 served by σ_2 which is in C. Its predecessor is buffer b_5 which is also served by σ_2. By condition (i), b_5 is required to be in C. To illustrate (ii), consider b_{19} served by σ_2 which is in C. Its predecessor is b_{18} which is served by σ_3. Hence b_{18} is not allowed to be in C by (ii). Moreover, by (ii), σ_3 is required to have at least one of its buffers in C. This requirement is met by b_9. □

Corollary 1. *If C is a virtual multiserver station then $|S(C)| \geq 2$. Also $|C| < L$.*

Proof. Recall our assumption in Section 2 that there are at least two stations in the system, i.e., $M \geq 2$, and that $B(\sigma) \neq \phi$ for every $1 \leq \sigma \leq M$. Since C is nonempty, there is some buffer $b_i \in C$ which has $\sigma(i) \neq \sigma(i-1)$. Thus by (ii), $S(C)$ contains at least the two stations $\sigma(i-1)$ and $\sigma(i)$. Moreover, since $b_{i-1} \notin C$ it follows that $|C| \leq L-1$. □

Definition: A visit to a station. *We shall say that* $V = \{b_i, b_{i+1}, b_{i+2}, \ldots, b_{i+k}\}$ *is a* visit to the station $\sigma(i)$ *(or simply* visit*) if:*

(i) $\sigma(i) = \sigma(i+1) = \sigma(i+1) = \cdots = \sigma(i+k)$

(ii) $\sigma(i-1) \neq \sigma(i)$

(iii) $\sigma(i+k+1) \neq \sigma(i)$.

We call b_i *the* first buffer *in the visit, and* b_{i+k} *the* last buffer *in the visit.*

Example 1 revisited

As an example, in Figure 4, $\{b_1, b_2, b_3, b_4\}$ is a visit to σ_1. Also $\{b_{17}\}$ is a visit to σ_2.

Definition: Preceding and succeeding visits. *Let* b_j *be the last buffer in a visit* V_1*, and* b_k *the first buffer in a visit* V_2*. If* $k = j+1$*, then we say that* V_2 succeeds V_1*, and that* V_1 precedes V_2. □

Example 1 re-revisited

The visit $V_1 = \{b_1, b_2, b_3, b_4\}$ precedes $V_2 = \{b_5, b_6, b_7, b_8\}$ (which, of course, succeeds it). □

Definition: Predecessors of a buffer in a visit. *Let* $V = \{b_i, b_{i+1}, b_{i+2}, \dots, b_{i+k}\}$ *be a visit. Consider* $1 \le j \le k$. *Then* $b_i, b_{i+1}, \dots, b_{i+j-1}$ *are called the* predecessors *of* b_{i+j} *in the visit.*

Example 1 re-revisited again

In the visit $\{b_9, b_{10}\}$, b_9 is the only predecessor of b_{10}. □

With this terminology in hand, one has the following equivalent definition of a virtual multiserver station.

Corollary 2. *A set of buffers* C *is a virtual multiserver station if and only if the following two conditions are satisfied:*

(i) *If* $b_i \in C$, *then all predecessors of* b_i *in the same visit are also in* C.

(ii) *If* $b_i \in C$, *then the last buffer in the preceding visit is not in* C, *but the station serving the last buffer in the preceding visit has some other buffer in* C.

4 A C-exciting buffer priority policy

For every virtual multiserver station we now define a class of buffer priority policies which will be of special interest in the remainder of this paper.

Definition: A C-exciting buffer priority policy. *Consider a virtual multiserver station* C. *We shall say that a given buffer priority policy* π *is a* C-exciting *policy if it satisfies the property that if* $b_k \notin C$ *but* $b_{k+1} \in C$, *then* b_k *has lower priority than every other buffer in* C *at the same station* $\sigma(k)$ *as it.*

Example 1 again

For the virtual multiserver station $C = \{b_1, b_2, b_3, b_5, b_6, b_7, b_9, b_{17}, b_{19}\}$, the buffer priority policy

$$\pi = \{b_{21}, b_1, b_3, b_2, b_{20}, b_4, b_6, b_5, b_{17}, b_{19}, b_7, b_8, b_{22}, b_9, b_{18}, b_{16}, b_{15}, b_{10}, b_{12}, b_{11}, b_{14}, b_{13}\}$$

is a C-exciting policy. Note that the buffers b_4, b_8, b_{16}, b_{18} and b_{22} are not in C, but their successor buffers b_5, b_9, b_{17}, b_{19} and b_1 are in C. Since buffer b_4 has lower priority than b_1, b_2 and b_3 at σ_1, buffers b_8 and b_{22} have lowest priority at σ_2, and buffers b_{16} and b_{18} have lower priority than b_9 at σ_3, it follows that π is a C-exciting policy.

One should note that a given virtual multiserver station may have more than one C-exciting policy. For example,

$$\pi' = \{b_{20}, b_{21}, b_1, b_2, b_3, b_4, b_5, b_6, b_7, b_{19}, b_{17}, b_{22}, b_8, b_9, b_{10}, b_{18}, b_{15}, b_{16}, b_{12}, b_{11}, b_{13}, b_{14}\}$$

is also a C-exciting policy. $\square$

5 An upper bound on the throughput of a C-exciting buffer priority policy

Consider a virtual multiserver station C, and a C-exciting buffer priority policy π. In this section, similar to [6] for open systems, we exhibit the key property that when policy π is used, then at every time instant t, at least one of the stations σ in $S(C)$ has all buffers in $C(\sigma)$ empty. This has the consequence that at most $S(C) - 1$ stations can be working at any time on the buffers in C. Hence we can imagine C as indeed a virtual station served at any given time by at most $|S(C)| - 1$ servers, which is the reason for calling it a virtual multiserver station. This fact then provides an upper bound on the throughput of the closed re-entrant line under policy π.

Theorem 1: An upper bound on the throughput under a C-exciting policy. *Consider a virtual multiserver station C, and a C-exciting buffer priority policy π. If policy π is used, then the following two facts are true:*

(i) *At any given time t, at least one of the stations $\sigma \in S(C)$ has all buffers in $C(\sigma)$ empty, i.e.,*

$$\prod_{\sigma \in S(C)} \sum_{k \in C(\sigma)} x_k(t) \quad = \quad 0 \text{ a.s.}$$

(ii) *The throughput λ_N^π under π is upper bounded by*

$$\lambda_N^\pi \quad \leq \quad \frac{|S(C)| - 1}{\sum\limits_{k \in C} m_k} \text{ a.s. for every population size } N \geq 1,$$

i.e.,

$$\lambda_N^\pi \quad \leq \quad \frac{S(C) - 1}{W_C} \text{ a.s.}$$

Proof. (i) Fix a population size N. First we note that, as proved in Jin, Ou and Kumar [9], under any buffer priority policy, the finite state Markov chain $\{x(t)\}$ has a single closed communicating class. This follows from the fact that if $b_{\bar{k}}$ has the lowest priority under π at a station $\bar{\sigma}$, then the state $x = (x_1, x_2, \ldots, x_N)$ with $x_{\bar{k}} = N$ (and all other $x_i = 0$ for $i \neq \bar{k}$) is reachable from every other state.

Thus, the "lim inf" in (1) is actually a limit and the throughput λ_N^π is a constant which does not depend on the initial condition of the system.

Consider the initial condition $x_{\bar{k}}(0) = N$ (and $x_i = 0$ for all other $i \neq \bar{k}$), where $b_{\bar{k}}$ is the lowest priority buffer at some station $\bar{\sigma}$, as above. Since all the customers are initially located at just this one station $\bar{\sigma}$, and $S(C) \geq 2$, it follows that there is some other station in $S(C)$ which has all its buffers in C initially empty. Thus, at time $t = 0$, condition (i) of the claim is satisfied.

Suppose now that the result is false. Then there exists a first stopping time $\tau > 0$ at which every station σ in $S(C)$ has a buffer in $C(\sigma)$ which is nonempty. That is, for every $\sigma \in S(C)$, there is a buffer $b_{i(\sigma)} \in C(\sigma)$ with $x_{i(\sigma)}(\tau) \geq 1$.

Since τ is the first such time, it is a jump time of the Markov chain. Let $\tau' < \tau$ denote the previous jump time (or 0 if there is no previous jump time) of the Markov chain. Since τ is the first such time, it must be the case that there is some station $\sigma' \in S(C)$ with all buffers in $C(\sigma')$ empty throughout $[\tau', \tau)$, there was a transition of a customer into a buffer $b_{i'} \in C(\sigma')$ at time τ from some other station, and all other stations $\sigma \neq \sigma'$ with $\sigma \in S(C)$ had at least one nonempty buffer in $C(\sigma)$. (The reason the customer entering $b_{i'}$ had to have come from some other station different from σ' is that if the predecessor buffer $b_{i'-1}$ is also at σ', then by the defining property of a virtual multiserver station it has to be in C and hence in $C(\sigma')$. This however is impossible since $C(\sigma')$ is assumed empty). In particular, it should be noted that for every $\sigma \neq \sigma'$, $\sigma \in S(C)$, there is a buffer $b_{i(\sigma)} \in C$ with $x_{i(\sigma)}(t) \geq 1$ for $\tau' \leq t < \tau$.

As noted above, the customer entering $b_{i'}$ at time τ had a service completion from $b_{i'-1}$ at $\sigma(i'-1) \neq \sigma(i')$ at time τ. By the definition of a virtual multiserver station, since $b_{i'-1} \notin C$ it follows that $\sigma(i'-1) \in S(C)$.

From the definition of a C-exciting policy, $b_{i'-1}$ has lower priority at $\sigma(i'-1)$ than every other buffer in $C(\sigma(i'-1))$. In particular, $b_{i'-1}$ has lower priority than another nonempty buffer in C, say b_n, located at the same station as $b_{i'}$, i.e., $b_n \in C(\sigma(i'-1))$. However since $x_n(t) \geq 1$ for $\tau' \leq t < \tau$, and since b_n has higher priority than $b_{i'}$, the buffer priority policy π could not have served a customer in $b_{i'}$. This leads to a contradiction, proving the result.

(ii) From (i) we see that at any given time, at least one of the servers σ in $S(C)$ cannot be actively serving customers in $C(\sigma)$. Hence at most $|S(C)| - 1$ servers can be simultaneously busy serving customers in C. This leads to the given throughput bound. □

6 A necessary condition for efficiency of a C-exciting buffer priority policy

Recall the throughput capacity of the closed re-entrant line defined in (2). Clearly, if the upper bound on the throughput given in Theorem 1 is strictly less than (2), then the C-exciting buffer priority policy cannot be efficient. This leads to the following necessary condition for efficiency of a C-exciting buffer priority policy.

Theorem 2. *Consider a virtual multiserver station C, and a C-exciting buffer priority policy π. A necessary condition for efficiency of π is*

$$\frac{|S(C)|-1}{\sum\limits_{k\in C} m_k} \;\geq\; \operatorname*{Min}_{\sigma} \frac{1}{\sum\limits_{i\in\sigma} m_i},$$

i.e.,

$$\frac{S(C)-1}{W_C} \;\geq\; \operatorname*{Min}_{\sigma} \frac{1}{W_{B(\sigma)}}.$$

7 Two–station systems

In the remainder of this paper we focus on systems with just two stations, i.e., $M = 2$. We denote the two stations as σ_1 and σ_2.

We note from Corollary 1 that for two station systems one always has $S(C) = 2$ for any virtual multiserver station C. Hence $S(C) - 1 = 1$, and so C is effectively served by at most one server at any given time t. Hence we shall refer to C as just a *virtual station*, dropping the qualifier "multiserver."

The necessary condition of Theorem 2 can therefore be rephrased as follows:

Theorem 3. *A necessary condition for a buffer priority policy* π *exciting a virtual station* C *to be efficient is*

$$W_C \leq \text{Max}\{W_{B(\sigma_1)}, W_{B(\sigma_2)}\}.$$

8 Alternating visit virtual stations

In this section we define a special type of virtual station for two station systems. These special virtual stations are formed by alternating visits to stations as defined below.

Definition: The set $V(i,j)$**.** *Consider two buffers* $b_i \in B(\sigma_1)$ *and* $b_j \in B(\sigma_2)$, $1 \leq i, j \leq L$. *For the two possibilities* $i > j$ *or* $i < j$, *we define a subset* $V(i,j)$ *of buffers as follows:*

Case i: $i > j$. Let $V(i,j)$ consist of all the buffers b_k with $j+1 \leq k \leq i-1$ which are located at σ_1, and all the buffers b_n with $1 \leq n \leq j-1$ or $i+1 \leq n \leq L$ which are at σ_2. That is

$$\begin{aligned} V(i,j) \;:=\; & (\{b_k : j+1 \leq k \leq i-1\} \cap B(\sigma_1)) \\ & \cup(\{b_n : 1 \leq n \leq j-1 \text{ or } i+1 \leq n \leq L\} \cap B(\sigma_2)). \end{aligned}$$

Case ii: $i < j$. Let $V(i,j)$ consist of all the buffers b_k with $1 \leq k \leq i-1$ or $j+1 \leq k \leq L$ which are located at σ_1, and all the buffers b_n with $i+1 \leq n \leq j-1$ which are located at σ_2. That is,

$$\begin{aligned} V(i,j) \;:=\; & (\{b_k : 1 \leq k \leq i-1 \text{ or } j+1 \leq k \leq L\} \cap B(\sigma_1)) \\ & \cup(\{b_n : i+1 \leq n \leq j-1\} \cap B(\sigma_2)). \end{aligned}$$

Example 2

Consider $V(14, 5)$.

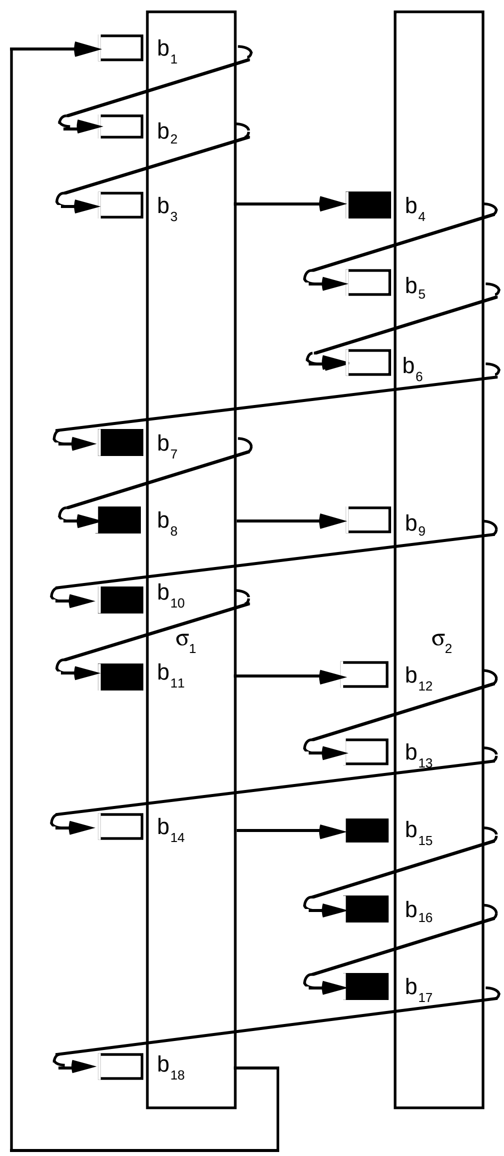

Figure 5: $V(14, 5) = \{b_7, b_8, b_{10}, b_{11}, b_{15}, b_{16}, b_{17}, b_4\}$.

Here $b_{14} \in B(\sigma_1)$ and $b_5 \in B(\sigma_2)$. Since $i = 14$ is greater than $j = 5$, $V(i, j) \cap B(\sigma_1)$ consists of the buffers $\{b_7, b_8, b_{10}, b_{11}\}$ between $j + 1 = 6$ and $i - 1 = 13$ which are located at σ_1. Also $V(i,j) \cap B(\sigma_2)$ consists of the buffers $\{b_{15}, b_{16}, b_{17}, b_4\}$ which lie between $i+1 = 15$ and $L = 18$, or between 1 and $j-1 = 4$, and are located at σ_2. Thus $V(i,j) = \{b_7, b_8, b_{10}, b_{11}, b_{15}, b_{16}, b_{17}, b_4\}$. □

Example 2 revisited

Note that for $i = 14$, and $j = 12$,

$$V(14,12) \quad = \quad \{b_{15}, b_{16}, b_{17}, b_4, b_5, b_6, b_9\}$$

does *not* have *any* buffers at σ_1. □

Example 3

Consider $V(2,13)$. Note that since $i = 2 < j = 13$, $V(2,13) \cap B(\sigma_2)$ consists of the buffers $\{b_4, b_5, b_6, b_9, b_{12}\}$ which lie between $i+1 = 3$ and $j-1 = 12$ and are located at $B(\sigma_2)$. Also, $V(2,13) \cap B(\sigma_1)$ consists of the buffers $\{b_{14}, b_{18}, b_1\}$ which lie between $j+1 = 14$ and $L = 18$, or between 1 and $i-1 = 1$, which are located at $B(\sigma_1)$. □

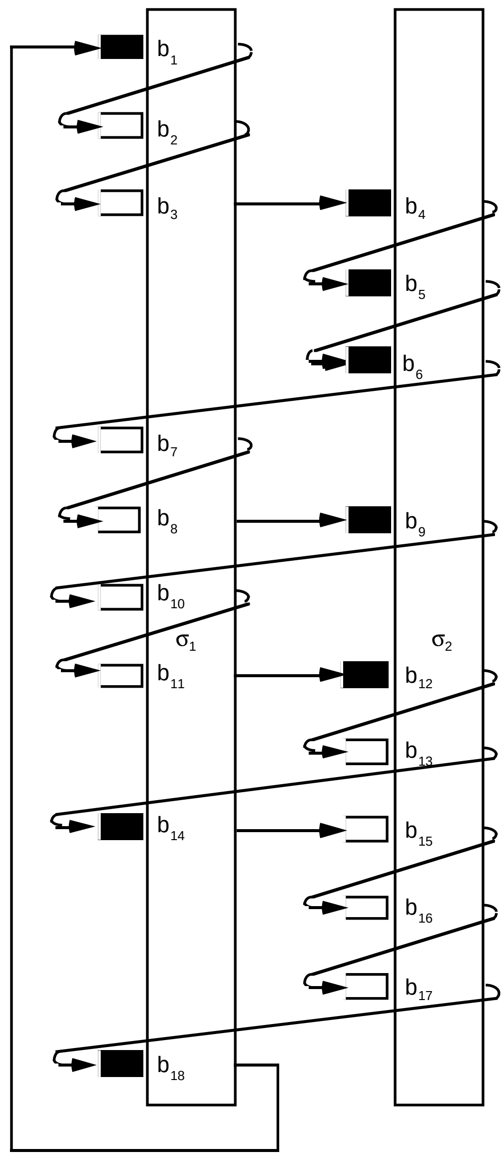

Figure 6: $V(2,13) = \{b_4, b_5, b_6, b_9, b_{12}, b_{14}, b_{18}, b_1\}$.

The next lemma shows that the sets $V(i,j)$'s are indeed virtual stations *if* they contain at least one buffer from each station.

Lemma 1. *$V(i,j)$ is a virtual station if and only if $V(i,j) \cap B(\sigma_1) \neq \phi$ and $V(i,j) \cap B(\sigma_2) \neq \phi$.*

Proof. The necessity is obvious since a virtual station is required to contain buffers at both stations. Turning to sufficiency, suppose for specificity that $i > j$; the proof is similar when $i < j$.

Then

$$V(i,j) \cap B(\sigma_1) \quad = \quad \{b_k : j+1 \leq k \leq i-1\} \cap B(\sigma_1)$$

and

$$V(i,j) \cap B(\sigma_2) \quad = \quad \{b_n : 1 \leq n \leq j-1 \text{ or } i+1 \leq n \leq L\} \cap B(\sigma_2),$$

and both are nonempty by assumption.

Consider a buffer $b_k \in V(i,j) \cap B(\sigma_1)$. Then $k-1 \geq j$. Either $k-1 > j$ or $k-1 = j$. In the former case, $k-1 \geq j+1$. Hence if $b_{k-1} \in B(\sigma_1)$ then necessarily $b_{k-1} \in V(i,j) \cap B(\sigma_1)$, while if $b_{k-1} \in B(\sigma_2)$, then $b_{k-1} \notin V(i,j) \cap B(\sigma_2)$ since $j < k-1 < k \leq i-1$. If however $k-1 = j$, then $k-1 \notin [1, j-1]$. Also, $k-1 \notin [i+1, L]$ since $k \leq i-1$. Hence $b_{k-1} \notin V(i,j) \cap B(\sigma_2)$ and $b_{k-1} = b_j \in B(\sigma_2)$. Thus we see that in either case if $\sigma(k-1) = \sigma(k)$ then $b_{k-1} \in V(i,j)$, while if $\sigma(k-1) \neq \sigma(k)$ then $b_{k-1} \notin V(i,j)$. Thus b_k satisfies the requirement contained in the definition of a virtual multiserver station.

Now consider a buffer $b_n \in V(i,j) \cap B(\sigma_2)$ with $1 \leq n \leq j-1$. Then either $n-1 = 0$ or $n-1 \in [1, j-1]$. If $n-1 = 0$, then $(n-1) \bmod L = L$, and so if $b_L \in B(\sigma_2)$, then $b_L \in V(i,j) \cap B(\sigma_2)$, while if $b_L \in B(\sigma_1)$ then $b_L \notin V(i,j) \cap B(\sigma_1)$. Also if $n-1 \in [1, j-1]$, then again if $b_{n-1} \in B(\sigma_2)$ it lies in $V(i,j) \cap B(\sigma_2)$, while if $b_{n-1} \in B(\sigma_1)$ then $b_{n-1} \notin V(i,j) \cap B(\sigma_1)$.

Now consider a buffer $b_n \in V(i,j) \cap B(\sigma_2)$ with $i+1 \leq n \leq L$. If $n-1 = i$, then $b_{n-1} = b_i \in B(\sigma_1)$ but $b_{n-1} \notin V(i,j) \cap B(\sigma_1)$. However if $n-1 > i$, then $n-1 \in [i+1, L]$. Hence for such an $n-1$, if $b_{n-1} \in B(\sigma_2)$, then $b_{n-1} \in V(i,j) \cap B(\sigma_2)$, while if $b_{n-1} \in B(\sigma_1)$, then $b_{n-1} \notin V(i,j) \cap B(\sigma_1)$.

Thus we see that if $b_n \in V(i,j) \cap B(\sigma_2)$ then $b_{n-1} \in V(i,j)$ if $b_{n-1} \in B(\sigma_2)$, while if $b_{n-1} \notin B(\sigma_2)$, then $b_{n-1} \notin V(i,j)$. Thus b_n satisfies the requirement of the definition of a virtual multiserver station.

This proves the Lemma. □

This motivates the following definition of a special type of virtual station.

Definition. *A set $V(i,j)$ which has buffers at both stations is called an* alternating visit virtual station.

9 The connection between the HW-indices and alternating visit virtual stations

In this section we exhibit the connection between certain indices introduced by Harrison and Wein [8] and the sets $V(i,j)$.

Definition: The HW-index η_i. *For any buffer b_i, the index η_i is defined as*

$$\eta_i \quad := \sum_{\{k: i \le k \le L \text{ and } b_k \in B(\sigma_1)\}} m_k - \sum_{\{k: i \le k \le L \text{ and } b_k \in B(\sigma_2)\}} m_k.$$

The HW-index η_i is simply the difference in the mean works to be done on a customer in buffer b_i by the two stations σ_1 and σ_2, before that customer next exits from buffer b_L.

The following lemma relates the HW-indices to the sets $V(i,j)$.

Lemma 2. *Consider $b_i \in B(\sigma_1)$ and $b_j \in B(\sigma_2)$. If $i > j$, then*

$$\eta_i - \eta_j \quad = \quad W_{B(\sigma_2)} - W_{V(i,j)}.$$

If $i < j$, then

$$\eta_i - \eta_j \quad = \quad W_{B(\sigma_1)} - W_{V(i,j)}.$$

Proof. Consider first the case $i > j$. Then

$$\begin{aligned}
\eta_i - \eta_j &= -\sum_{\{k:j\leq k\leq i-1 \text{ and } k\in\sigma_1\}} m_k + \sum_{\{k:j\leq k\leq i-1 \text{ and } k\in\sigma_2\}} m_k \\
&= -\sum_{\{k:j<k\leq i-1 \text{ and } k\in\sigma_1\}} m_k + \sum_{\{k:j\leq k\leq i-1 \text{ and } k\in\sigma_2\}} m_k \qquad (\text{since } j \notin \sigma_1) \\
&= -W_{V(i,j)} + \sum_{\{k:1\leq k\leq j-1 \text{ or } i+1\leq k\leq L \text{ and } k\in\sigma_2\}} m_k + \sum_{\{k:j\leq k\leq i-1 \text{ and } k\in\sigma_2\}} m_k \\
&= -W_{V(i,j)} + W_{B(\sigma_2)} \qquad (\text{since } i \notin \sigma_2).
\end{aligned}$$

Similarly, if $i < j$, then

$$\begin{aligned}
\eta_i - \eta_j &= \sum_{\{k:i\leq k\leq j-1 \text{ and } k\in\sigma_1\}} m_k - \sum_{\{k:i\leq k\leq j-1 \text{ and } b_k\in\sigma_2\}} m_k \\
&= \sum_{\{k:i\leq k\leq j-1 \text{ and } k\in\sigma_1\}} m_k - W_{V(i,j)} + \sum_{\{k:1\leq k\leq i-1 \text{ or } j+1\leq k\leq L \text{ and } k\in\sigma_1\}} m_k \\
&= W_{B(\sigma_1)} - W_{V(i,j)}.
\end{aligned}$$

thus proving the claim. □

10 Necessary and almost sufficient condition for guaranteed efficiency of balanced two-station closed re-entrant lines

In the remainder of this paper, we restrict attention to *balanced* two station closed re-entrant lines.

Definition. *A two station closed re-entrant line is* balanced *if* $W_{B(\sigma_1)} = W_{B(\sigma_2)}$.

The following theorem provides a necessary and almost sufficient condition for guaranteed efficiency in terms of the HW-indices.

Theorem 4: Necessary and almost sufficient condition for guaranteed efficiency.

(i) *If*

$$\operatorname*{Min}_{i \in \sigma_1} \eta_i \quad > \quad \operatorname*{Max}_{j \in \sigma_2} \eta_j,$$

then all non-idling policies are efficient.

(ii) *If all non-idling policies are efficient, then*

$$\operatorname*{Min}_{i \in \sigma_1} \eta_i \quad \geq \quad \operatorname*{Max}_{j \in \sigma_2} \eta_j.$$

Proof. The sufficient condition (i) is established in Jin, Ou, and Kumar [9]. Turning to (ii), consider $i \in \sigma_1$ and $j \in \sigma_2$. There are two cases; either $V(i,j)$ is a virtual station, or it is not.

If it is *not* a virtual station then from Lemma 1 either $V(i,j) \cap B(\sigma_1)$ or $V(i,j) \cap B(\sigma_2)$ is empty, and so $V(i,j)$ is either contained in $B(\sigma_2)$ or $B(\sigma_1)$, yielding $W_{B(\sigma_2)} - W_{V(i,j)} = W_{B(\sigma_1)} - W_{V(i,j)} \geq 0$. Thus, from Lemma 2, one has

$$\eta_i \quad \geq \quad \eta_j.$$

On the other hand, if $V(i,j)$ is a virtual station, then from Theorem 3, we have

$$W_{V(i,j)} \quad \leq \quad W_{B(\sigma_1)} = W_{B(\sigma_2)}.$$

Hence from Lemma 2, one again has

$$\eta_i - \eta_j \quad \geq \quad 0. \qquad \square$$

In fact, it is shown in [9] that under the sufficient condition (i) of the above Theorem one has the stronger result that all non-idling policies have *finite* asymptotic loss. In particular,

$$\limsup_{N\to\infty} \frac{N(\lambda^* - \lambda_N^u)}{\lambda^*} \leq \frac{\lambda^* \sum_{i=1}^{L} \eta_i m_i}{\underset{i\in\sigma_1}{\text{Min}}\ \eta_i - \underset{j\in\sigma_2}{\text{Max}}\ \eta_j} \quad \text{for all } u \in U.$$

11 Neither closed nor open configurations always dominate the other

Since we now have necessary and almost sufficient conditions for the throughput of open and closed networks, the former due to Dai and Vande Vate [6], we can answer the age old question whether open or closed configurations have better guaranteed throughput. The answer is that neither is always better. We provide two examples which prove this.

Example 4: An open network with lower guaranteed throughput than its closed version.

Consider the system shown in Figure 7.

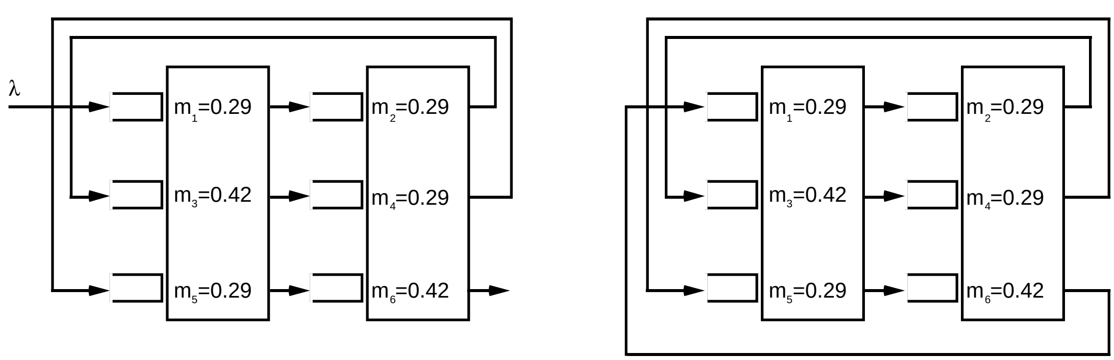

Figure 7: An open network with smaller guaranteed throughput than its closed counterpart.

From Dai and Vande Vate it follows that the open system can only be guaranteed stable under all non-idling policies for arrival rates λ which satisfy,

$$\frac{\lambda m_3}{(1-\lambda m_1)} + \lambda m_6 \leq 1,$$

i.e., for $\lambda \leq 0.990761$. On the other hand, the closed counterpart has $\lambda^* = 1$, and is efficient for all non-idling policies, since

$$-0.13 = \operatorname*{Min}_{i \in \sigma_1} \eta_j > \operatorname*{Max}_{j \in \sigma_2} \eta_j = -0.29. \qquad \square$$

Example 5: An open network with higher guaranteed throughput than its closed version.

Consider the system shown in Figure 8.

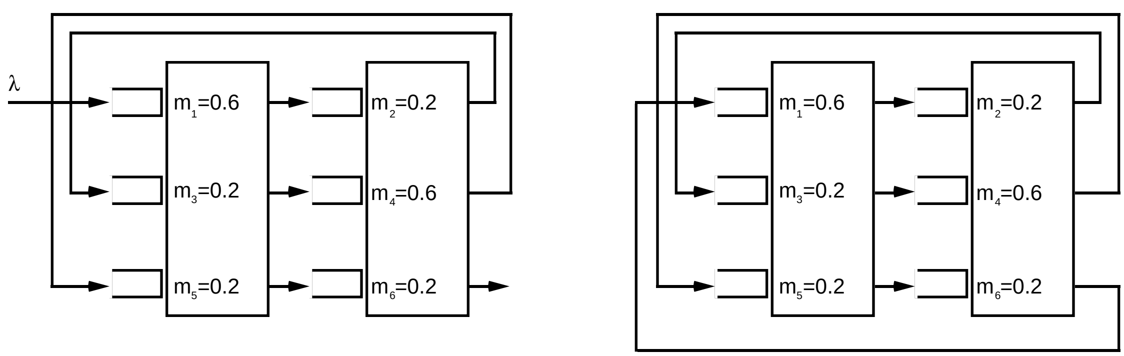

Figure 8: An open network with greater guaranteed throughput than its closed version.

From Dai and Vande Vate [6] it follows that the open network is guaranteed to be stable for all arrival rates $\lambda < 1$. On the other hand, for the closed counterpart, the virtual station $C = \{1, 4\}$ has $W_C = 1.2$. Hence no throughput greater than $\frac{1}{1.2} \leq 0.833$ can be guaranteed.

$\square$

12 Alternating visit virtual stations need not be maximally constrictive for inefficient systems

It is important to note the subtle fact that while alternating visit virtual stations almost characterize inefficient systems, they need *not* be the maximally constrictive virtual stations. Lest the reader wonder how this could be, it suffices to note that our result of Theorem 4 can be construed to read that if there is any virtual station that is maximally constrictive, then there exists an alternating visit virtual station which constricts at least as much as any of the two *real* stations σ_1 and σ_2.

Thus, if one considers the linear program with decision variables λ^{-1} and $\{m_k \text{ for } 1 \leq k \leq L\}$:

$$\text{Min } \lambda^{-1}$$

subject to:

$$\lambda^{-1} \geq \sum_{k \in C} m_k \text{ for every virtual station } C,$$

$$\sum_{k \in V(i,j)} m_k \leq 1 \text{ for every alternating visit virtual station } V(i,j),$$

$$\sum_{k \in \sigma_i} m_k = 1 \text{ for } i = 1, 2,$$

$$m_k \geq 0 \text{ for } 1 \leq k \leq L,$$

then our theory asserts that its value $\lambda_{opt}^{-1} \leq 1$.

In the following example there is a virtual station that is more constrictive than all alternating visit virtual stations.

Example 6: A virtual station that is more constrictive than all alternating visit virtual stations.

Consider the system of Figure 9.

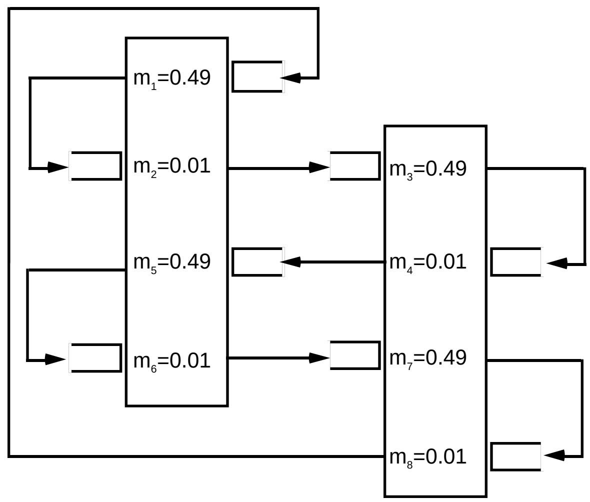

Figure 9: The virtual station $C = \{b_1, b_3, b_5, b_7\}$ is more constrictive of the throughput than all the alternating visit virtual stations $V(1,4)$, $V(2,4)$, $V(2,7)$, $V(2,8)$, $V(5,8)$, $V(6,3)$, $V(6,4)$ and $V(6,8)$.

The virtual station $C = \{b_1, b_3, b_5, b_7\}$ constricts the guaranteed throughput to be no more than $(m_1 + m_3 + m_5 + m_7)^{-1} = \frac{1}{1.96} = 0.5102$. This is more than the upper bounds on throughput $(m_3 + m_5 + m_6)^{-1}$, $(m_1 + m_3 + m_5 + m_6)^{-1}$, $(m_3 + m_4 + m_1)^{-1}$, $(m_3 + m_4 + m_7 + m_1)^{-1}$, $(m_7 + m_1 + m_2)^{-1}$, $(m_5 + m_7 + m_8)^{-1}$, $(m_5 + m_7 + m_8 + m_3)^{-1}$ and $(m_7 + m_1 + m_2 + m_5)^{-1}$, imposed by $V(1,4)$, $V(2,4)$, $V(2,7)$, $V(2,8)$, $V(5,8)$, $V(6,3)$, $V(6,4)$ and $V(6,8)$, respectively. □

13 The Anti-HW policy and maximally constrictive alternating visit virtual stations

The example of the preceding section naturally raises the question of what is the significance of the maximally constrictive alternating visit virtual stations.

Since these maximally constrictive alternating visit virtual stations do capture guaranteed efficiency, it is natural to see if they bear some relationship to a policy one expects be bad – the *Anti-HW* policy. This gives priority at σ_1 to buffers b_i which have higher values of η_i, and at σ_2 to buffers b_j which have lower values of η_j. This policy was studied in Kumar and Kumar [3] and explicitly analyzed in Jin, Ou, and Kumar [9].

The lowest priority buffer b_{i*} at σ_1 under the Anti-HW policy is one for which

$$\eta_{i*} = \operatorname*{Min}_{i \in \sigma_1} \eta_i,$$

and the lowest priority buffer b_{j*} at σ_2 is the one for which,

$$\eta_{j*} = \operatorname*{Max}_{j \in \sigma_2} \eta_j.$$

Thus from Lemma 2 it is immediately clear that among all alternating visit virtual stations, the most constrictive is $V(i^*, j^*)$, i.e.,

$$V(i^*, j^*) = \operatorname*{Max}_{i \in \sigma_1, j \in \sigma_2} V(i, j).$$

However, we now present examples to show that $V(i^*, j^*) = \operatorname*{Max}_{i \in \sigma_1, j \in \sigma_2} V(i, j)$ may or may not be excited by the Anti-HW *policy.* It is worth noting that our current intuition on this topic, as for example provided by heavy traffic Brownian networks, suggests that what are important for the performance of a buffer priority policy are only the identities of the *lowest* priority buffers, and not the precise orders of the more higher priority buffers.

In the following example, all Anti-HW policies (there may be more than one, since there may be buffers at σ_1 (or σ_2) with equal values of η_i (or η_j)) excite all the maximally constrictive alternating virtual stations.

Example 7

Consider the system shown in Figure 10.

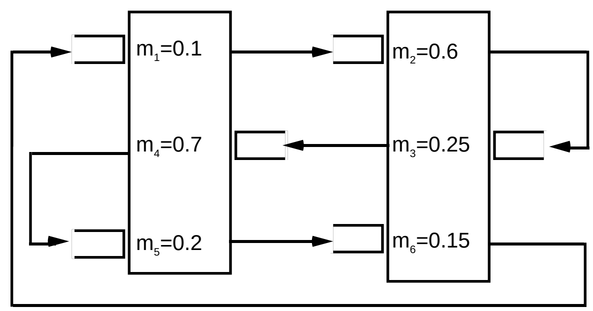

Figure 10: The Anti-HW policy excites the maximally constrictive alternating visit virtual station $V(1,3) = \{b_2, b_4, b_5\}$.

The indices at σ_1 are:

$$\eta_1 = 0, \quad \eta_4 = 0.75, \quad \eta_5 = 0.05,$$

and the indices at σ_2 are:

$$\eta_2 = -0.1, \quad \eta_3 = 0.5, \quad \eta_6 = -0.15.$$

Since no two η_i's are equal, there is only one Anti-HW policy. It gives priority at σ_1 according to the order $\{b_4, b_5, b_1\}$, while at σ_2 it gives priority according to the order $\{b_6, b_2, b_3\}$.

There are three alternating visit virtual stations. They are $V(1,3) = \{b_2, b_4, b_5\}$, $V(5,2) = \{b_4, b_6\}$ and $V(5,3) = \{b_4, b_2, b_6\}$.

The most constrictive of these is $V(1,3)$, restricting the guaranteed throughput to be no more than

$$(m_2 + m_4 + m_5)^{-1} \quad = \frac{1}{1.5} = \quad 0.667.$$

The buffers not in $V(1,3)$ preceding buffers in $V(1,3)$ are b_1 at σ_1 and b_3 at σ_2. To excite $V(1,3)$ a policy needs to give lower priority to b_1 than b_4 or b_5, and lower priority to b_3 than b_2.

The Anti-HW policy does this, and so it does excite the most constrictive alternating visit virtual station. □

However, in the following example, no Anti-HW policy excites any maximally constrictive alternating visit virtual station.

Example 8

Consider the system shown in Figure 11.

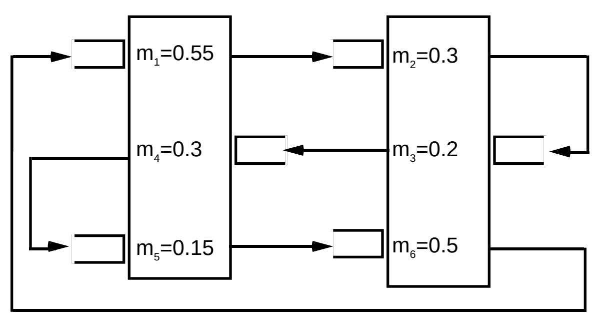

Figure 11: The Anti-HW policy does *not* excite the most constrictive alternating visit virtual station $V(3,5) = \{b_4, b_2, b_6\}$.

At σ_1 the indices are:

$$\eta_1 = 0, \quad \eta_4 = -0.05, \quad \eta_5 = -0.35,$$

while at σ_2 the indices are

$$\eta_2 = -0.55, \quad \eta_3 = -0.25, \quad \eta_6 = -0.5.$$

Thus there is only one Anti-HW policy. At σ_1, it gives priority according to the order $\{b_1, b_4, b_5\}$, while at σ_2 it gives priority according to the order $\{b_2, b_6, b_3\}$.

There are three alternating visit virtual stations. They are $V(1,3) = \{b_2, b_4, b_5\}$, $V(5,2) = \{b_4, b_6\}$, and $V(5,3) = \{b_4, b_2, b_6\}$. The most constrictive of these is $V(5,3)$ which restricts the guaranteed throughput to be no more than

$$(m_4 + m_2 + m_6)^{-1} = \frac{1}{1.1}.$$

The buffers not in $V(5,3)$ which precede buffers in $V(5,3)$ are b_3 at σ_2 and $\{b_1, b_5\}$ at σ_1.

Since the Anti-HW policy does *not* accord lower priority to b_1 than b_4, it is *not* a $V(3,5)$-exciting policy. However, it should be noted that the buffer b_1 is not the *lowest* priority buffer at σ_1, and as commented earlier the Brownian network analysis suggests that what is important for the performance are only the identities of the lowest priority buffers at the two stations and not the higher priority ones. □

In fact, we see that since Examples 7 and 8 have the same topology, whether or not there exists an Anti-HW policy exciting the most constrictive alternating visit virtual station is not purely a function of the topology either.

14 Concluding remarks

The vein of work beginning with [1] and continuing into the present has shown that instability in networks, whether stochastic or deterministic, can have a certain structural character. This lends hope to further studies in this direction, and is clearly worth further pursuit.

From [2] through the present, the body of work is also showing that buffer priority policies have extremal properties with respect to certain carefully defined criteria. For example, for both open and closed two-station networks, the guaranteed throughput is characterized just by examining buffer priority policies. Even more, in the closed case, the HW-policy, a buffer priority policy, may well be asymptotically optimal. This line of investigation into the extremal aspects also needs further pursuit, especially as we delve deeper into issues such as higher order pole multiplicities identified in [10], and other complications, in pursuit of our goal of better understanding multi-station networks and how to schedule them.

References

[1] P. R. Kumar and T. I. Seidman, "Dynamic instabilities and stabilization methods in distributed real-time scheduling of manufacturing systems," *IEEE Transactions on Automatic Control*, vol. AC-35, pp. 289–298, March 1990.

[2] S. H. Lu and P. R. Kumar, "Distributed scheduling based on due dates and buffer priorities," *IEEE Transactions on Automatic Control*, vol. AC-36, pp. 1406–1416, December 1991.

[3] S. Kumar and P. R. Kumar, "Performance bounds for queueing networks and scheduling policies," *IEEE Transactions on Automatic Control*, vol. AC-39, pp. 1600–1611, August 1994.

[4] J. M. Harrison and V. Nguyen, "Some badly behaved closed queueing networks," in *Stochastic Networks* (F. P. Kelly and R. Williams, eds.), vol. 71, pp. 117–124, New York, NY: Springer–Verlag, 1995.

[5] V. Dumas, "A multiclass network with non-linear, non-convex, non-monotonic stability conditions." Preprint, 1995.

[6] J. Dai and J. H. V. Vate, "Global stability of two–station queueing networks." Preprint, 1996. School of Industrial and Systems Engineering, Georgia Institute of Technology.

[7] J. J. Hasenbein, "Necessary conditions for global stability of multiclass queueing networks." Industrial and Systems Engineering Report Series J-96-01, May 1996. Georgia Institute of Technology.

[8] J. M. Harrison and L. M. Wein, "Scheduling networks of queues: Heavy traffic analysis of a two-station closed network," *Operations Research*, vol. 38, no. 6, pp. 1052–1064, 1990.

[9] H. Jin, J. Ou, and P. R. Kumar, "The throughput of closed queueing networks–functional bounds, asymptotic loss, efficiency, and the Harrison-Wein conjectures." Submitted to *Mathematics of Operations Research*, October 1994.

[10] C. Humes, Jr., J. Ou, and P. R. Kumar, "The delay of open Markovian queueing networks: Uniform functional bounds, heavy traffic pole multiplicities, and stability." Submitted to Mathematics of Operations Research, July 1995.