Queues Served by a Rotating Ring

E. G. COFFMAN, Jr., E. N. GILBERT, A. G. GREENBERG
AT&T Bell Laboratories, Murray Hill, New Jersey 07974

F. T. LEIGHTON
Mathematics Department and Laboratory for Computer Science
Massachusetts Institute of Technology, Cambridge, MA 02139

Philippe ROBERT
Institut National de Recherche en
Informatique et en Automatique, 78150 Le Chesnay, France

A. L. STOLYAR*
Motorola, Inc., Arlington Heights, IL 60004

Abstract A ring of N cells rotates in discrete steps past N queues, moving customers from their queues of arrival to randomly chosen destinations. The model has applications in communication systems, processor interconnection networks, and flexible manufacturing. The arrivals to the queues are independent and stochastically identical. The total numbers of arrivals to the system during successive steps are independent, identically distributed random variables with mean λ and finite second and third moments. A *greedy* policy governs the insertion of customers on the ring: A customer waiting at the head of a queue enters the next unoccupied cell to appear at that queue. The customer then remains on the ring a random travel time d, leaves, and frees its cell for another customer.

A necessary condition for the system to be stable is $\lambda E[d] < N$. If no customer travels further than once around the ring ($d \leq N$), $\lambda < 1$ is sufficient for stability. Other results assume d to be stochastically bounded by an exponentially distributed travel time with mean N/μ. Then $\lambda < \mu$ is sufficient for stability. In the limit of large N, stable systems with fixed

*Alexander Stolyar's work was completed in April 1992 while visiting AT&T Bell Labs, Murray Hill.

λ and μ have expected numbers $o(1)$ of waiting customers per queue; then a customer's wait in a queue is usually negligible compared with his travel time. Simulations suggest that the mean number waiting in a queue may even be $O(1/N)$.

Keywords: Processor rings, loop communication systems, cyclic service, stability, asymptotics.

1. INTRODUCTION

A sequence of N queues is arranged around a circular conveyor, or *ring* of cells numbered $1, 2, \ldots, N$, $N \geq 2$. Each queue has access to one cell, but the ring rotates unidirectionally in discrete steps so that the queue at cell i will be at cell $i-1$ on the next step (or at cell N if $i = 1$). The purpose of the ring is to transport customers from one queue location to another; once placed in a cell a customer remains there for a random number d of steps, its *travel time*, and then leaves. In each step, events occur in the following sequence: (i) the ring rotates one position as the queues receive new arrivals, (ii) customers that have reached their destinations leave, and (iii) each nonempty queue places one of its waiting customers on the ring if it is opposite an empty cell.

The placement rule in (iii) is called the *greedy* policy, since an empty cell is supplied with a new customer as soon as that cell encounters a nonempty queue. Making (iii) follow (ii) prevents departing customers from blocking new arrivals to the ring, i.e., a cell can discharge a customer and accept a new one both in the same step.

At each step, the N queues receive i.i.d. numbers of new arrivals with means λ/N. Then λ is the mean total arrival rate to the system. The number of arrivals to a queue at each step is also assumed to have $O(1/N)$ second and third moments; this assumption will be needed for the existence of the first and second moments of the queue-length distribution. Customer travel times are i.i.d. random variables; cases of special interest include the uniform and geometric distributions.

Each cell of the ring acts as a server but the ring is far from the usual queueing system with N parallel servers. Idle servers and occupied queues can coexist if these servers and queues are in different places. The rotating ring creates an interdependence among the N queues that makes a complete probabilistic analysis appear intractable. A principal source of difficulty is a clumping effect that tends to concentrate the occupied cells of the ring into isolated trains of nearby cells.

The lesser goals of stability conditions and asymptotic behavior hold more promise. Section 2 includes simulations of ring systems with d uniform on $\{1, \ldots, N-1\}$. These indicate that the queues are stable if $\lambda < 2$; indeed, the mean queue lengths tend to 0 as $N \to \infty$. No proof of that is given but Section 3 allows any distribution with $d \leq N$ and proves that $\lambda < 1$ is sufficient for stability. Sections 3 and 4 also consider travel times d that are dominated by a geometrically distributed time with mean N/μ; Sections 3 and 4 show that, if $\lambda < \mu$, then the queues are stable and their mean lengths approach 0 as N increases.

A number of systems are modeled by rings. Pierce [21] proposed a communication system featuring interconnected rings with information packets as customers; combinatorial studies of such systems can be found in [4, 7, 21]. In other applications, customers might be automobiles at a traffic circle or production items on a conveyor. Recently, rings have found important applications in local area computer networks [2, 3, 8, 9, 12, 15, 20] and in parallel computers, replacing earlier bus-based interconnects; examples include the Scalable-Coherent-Interface (SCI) ring [5, 14, 22], the Kendall Square Research (KSR-1) system [17], and the Hector system [23]. Arrays of computers were studied in [16, 19] using a probability model very similar to the present one. The model here, like those in [16, 19] is idealized; more realistic models are briefly discussed in Section 5 as directions for further research.

2. RING BEHAVIOR

Figure 1 shows the ring system in operation for one complete revolution. Each column represents one cell, with time increasing downward. Rows describe the cells just as they leave for their next queues. A slash mark \ denotes a full cell at a queue containing one or more waiting customers. A vertical mark | denotes an empty cell opposite a queue that is also empty. All other cells are full and are opposite empty queues. Small circles represent new arrivals to queues. Since Fig. 1 is drawn from the viewpoint of an observer fixed in the ring, the queues appear to be moving to the right.

Figure 1 is a computer simulation of the greedy policy; occupied queues fill empty cells as soon as they arrive. The greedy policy, used throughout this paper, is a natural one. All others let queues waste opportunities to unload customers. However, one can easily construct examples in which, for given arrivals and destinations, the greedy policy causes more total customer delay than some nongreedy policy. (An example appears in Section 4.)

Customers are assumed to leave the ring without blocking waiting customers from entering their emptied cells in the same step. A system with blocking departures would be equivalent to our nonblocking system with all travel times d increased to $d+1$ (the extra step provides the required blocking at step d but leaves the cell available at step $d+1$).

When the queues are stable, an application of Little's theorem to the ring alone shows the mean number of customers on the ring to be $\lambda E[d]$. In Fig. 1 with $\lambda = 1.5$ and $E[d] = N/2$, one quarter of the cells are expected to be empty. Since the ring holds at most N customers, a necessary condition for stability is

$$\lambda < N/E[d] \ . \tag{2.1}$$

In many queueing problems, queues are automatically stable if the traffic arrives slower than the maximum rate of the servers. Here too, (2.1) is perhaps sufficient for stability, but that requires proof. A problem arises because empty cells may remain idle while queues, located elsewhere, are occupied. A trivial case in which (2.1) is sufficient for stability occurs when $d = 1$ for all customers, since in that case each cell is able to accept a new customer at each step.

The uniform distribution of d on $\{1, \ldots, N-1\}$ is an important special case studied in other models of processor interconnection networks (see [16, 19]). The necessary condition for stability in this case is $\lambda < 2$, by (2.1), but whether this condition is sufficient, and whether it implies that expected queue lengths tend to 0 as $N \to \infty$ remain to be proved. Insight into the problem with uniform distributions was provided by simulations. The simulations used $\lambda = 1.0$, 1.9 with Poisson-distributed arrivals to each queue at rate λ/N per step. In tests on rings of reasonable size, $N = 10$, 50, 100, 200, 10^6 consecutive steps of the ring were simulated and 100 snapshots of the state of the system were recorded at intervals of 10^4 steps. Other tests to show asymptotic behavior took $N = 8K$, $16K$, $32K$, $64K$; simulations used 10 snapshots $10N$ steps apart in a total of 100 complete ring rotations. The simulations were costly, requiring about 10^9 random numbers for some values of N. A MasPar MP-1 parallel computer was used.

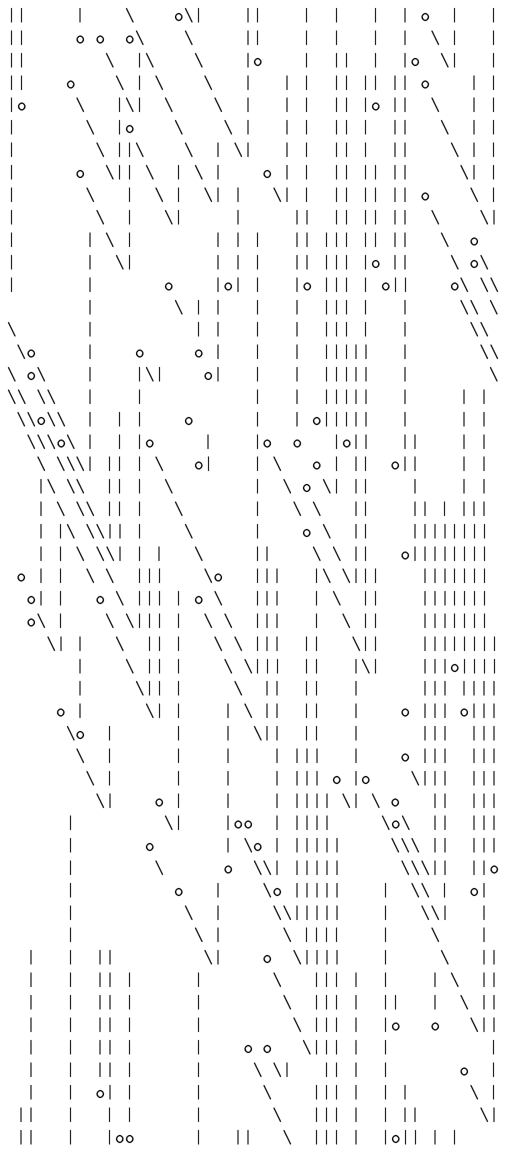

Figure 1: One complete revolution of a ring system with $N = 50$, arrival rate $\lambda = 1.5$, and uniform destinations.

Table 1 shows total numbers of customers in all N queues, averaged over the available snapshots. The accuracy of this average as an estimator for the mean number of customers is indicated by the number in parentheses; this is the standard deviation of the numbers from the snapshots. There was no tendency for queues to continue growing as the numbers of trials increased. This observation strongly suggested that the queues were indeed stable and that $\lambda < 2$ was sufficient as well as necessary for stability. Dividing the average total queue lengths in Table 1 by the appropriate value of λ gives estimates of average waiting times, by Little's theorem.

Table 1 is strong evidence that the expected total number waiting is bounded by a constant as $N \to \infty$. This is powerful support indeed for the conjecture that individual expected queue lengths tend to 0 as $N \to \infty$, and at a rate $O(1/N)$. This result implies that expected waiting times are bounded by a constant as $N \to \infty$. To see this, look upon an individual queue as a single-server system so that a customer's time in the system is his waiting time; then apply Little's theorem.

Table 1: Total Queue Lengths from Simulations

	$\lambda = 1.0$		$\lambda = 1.9$	
N	average	median	average	median
10	1.09(.015)	1	50.7(3.6)	47
50	1.30(.019)	1	81.6(5.6)	79.5
100	1.33(.02)	1	90.1(7.1)	84.5
200	1.35(.02)	1	96.6(7.0)	97
...	...	...	...	...
8192	1.36(.027)	1	101.5(8.6)	118.5
16384	1.36(.017)	1	108.8(6.6)	104
32768	1.35(.023)	1.5	106.9(6.6)	109
65536	1.35(.022)	1.5	105.3(4.5)	98

A striking feature of the simulations was the appearance of long trains of consecutive occupied cells on the ring. These trains are somewhat evident in Fig. 1. The clumping that appears in Fig. 1 is understandable because, when a queue must wait for several steps to discharge a customer, it often adds a customer at the end of the train that caused the delay.

Statistics on train lengths were obtained from the snapshots used for

Table 1; these are summarized in Table 2. Each occupied cell following an empty cell counted as the start of a new train; also a ring with all N cells occupied counted as a train of length N. It is interesting that a mean train length in Table 2 does not differ greatly from the mean length $1/(1-\lambda/2)$ that one can derive for a long row of cells occupied independently with probabilities equal to the fraction of occupied cells, i.e., with probabilities $\lambda E[d]/N = \lambda/2$. However, the long trains are evident from the large maximum and small median train lengths. With independently occupied cells, the median train length would have been about 69% of the mean; by contrast, Table 2 shows that for $\lambda = 1.9$ and large N, the median is less than 38% of the mean. In such cases, train lengths are distributed with long tails; there are a few very long trains but the remaining customers scatter about the ring in a large number of short trains.

Table 2: Train Lengths from Simulations

	$\lambda = 1.0$			$\lambda = 1.9$		
N	average	median	maximum	average	median	maximum
10	2.345	2	10	8.70	10	10
50	2.476	2	25	19.10	14	50
100	2.390	2	18	21.5	12	100
200	2.303	2	21	21.8	10	200
...	...	...	...	...	...	...
8192	2.39	2	28	22.0	8	549
16384	2.38	2	29	23.2	8	1052
32768	2.39	2	35	22.7	8	927
65536	2.39	2	43	21.9	8	973

The independent-cell idea can be incorporated into an approximate ring model that has a $O(1)$ mean total number of waiting customers as $N \to \infty$. Replace the cells of the ring by new cells that are occupied independently at each step, and with the same probability $\lambda E[d]/N$ as the original cells. This simplification removes the tendency to form very long trains and allows the N queues to operate independently of one another. See [22] for similar approximations. New customers arrive to each queue independently at rate λ/N. To keep the model simple, let the number of arrivals to a queue at each step be 1 with probability λ/N and 0 otherwise. To obtain a system resembling one with uniformly distributed destinations, take $\lambda/2$ to be the probability that a queue finds its cell occupied. With these assumptions, the length of a queue increases by 1 with probability $(\lambda/2)(\lambda/N)$ at each step, and the length of an occupied queue decreases by 1 with probability $(1-\lambda/2)(1-\lambda/N)$ at each step. Otherwise queue lengths remains the same. For $\lambda < 2$ the stationary queue-length distribution is then found to be geometric with the mean total number in all N queues approaching $\lambda^2/(2-\lambda)$ as $N \to \infty$. The means approach 1 and 36.1 for $\lambda = 1$ and 1.9. The agreement with the average queue lengths in Table 1 is much better when λ is small. The greater discrepancies at larger λ suggest that the underestimate of the tail of the train-length distribution has a major effect for large λ.

3. STABILITY

Define the system state $Z(t)$, $t = 0, 1, 2, \ldots$, at the beginning of step t as a pair consisting of the *queue state* and the *ring state*. The queue state gives the number in each queue and the ring state indicates for each cell whether it is empty or occupied, and in the latter case how long its occupant has been on the ring. The ring process $\{Z(t)\}$ is clearly a Markov chain; this section proves sufficient conditions for ergodicity in the two cases described in Section 2, beginning with the case $d \leq N$ for all customers.

Theorem 1. *Let d have any distribution on $\{1, \ldots, N\}$. If $\lambda < 1$, then $\{Z(t)\}$ is an ergodic Markov chain. The stationary mean value of the number Q of customers waiting in each queue is bounded by*

$$E[Q] \leq \frac{L\sigma^2}{2N^2(L-1-L\lambda)} + L\lambda \ ,$$

where σ^2 is the variance of the number of arrivals to the system in each step, and $L = \lfloor 1/(1-\lambda) \rfloor + 1$, with $\lfloor x \rfloor$ the integer part of x.

Before proving Theorem 1, we note that a ring with $\lambda < 1$ can easily be stabilized by adopting a nongreedy *gating* policy that forbids customers from entering the ring except at times $0, N, 2N, \ldots$. Since $d \leq N$, all queues will find a cell available once every N steps. Between these times a mean number $\lambda < 1$ of new customers will arrive at each queue. Then an easy analysis of individual queues shows that each is stable. By allowing empty cells to pass occupied queues without accepting new customers, the gating policy appears to be less effective than the greedy policy, which suggests that the ring with the greedy policy is also stable for $\lambda < 1$.

However, the two policies are hard to compare exactly, especially during heavy traffic. For example, consider the test case where $d = N$ for all customers and $\lambda < 1$ is necessary and sufficient for stability under both the greedy and gating policies. In heavy traffic, during periods when all queues are nonempty, the ring under the gating policy acts as N parallel servers, operating independently except for the synchronization created by the ring. Under the same conditions, the behavior of the greedy policy can be quite erratic, since one or more nonempty queues can fill all of the cells, locking out the remaining queues for long periods of time. (This type of behavior is the focus of the combinatorial analysis in [4].) Theorem 1 proves that, in spite of this potential for erratic behavior, the stability condition under the greedy policy remains the same as that under N independent parallel servers.

Finally, note that large-N behavior for fixed $\lambda < 1$ argues strongly in favor of the greedy policy. Section 2 conjectured that the expected waiting time is bounded by a constant as $N \to \infty$. On the other hand, under the gating policy, arrivals to an empty queue must wait about $N/2$ steps on average. It follows easily that the expected waiting time must grow linearly in N, under the gating policy.

Return now to Theorem 1.

Proof. Two lemmas are required.

Lemma 1. *Suppose some queue remains nonempty for kN steps, where k is an integer > 1. During that time, each cell of the ring receives at least $k-1$ new customers.*

Proof. During steps $N+1, \ldots, kN$, each cell visits the nonempty queue $k-1$ times and departs occupied. The occupant of a cell cannot have been there for more than a rotation, so it must have entered after the cell's last visit to the queue. Then each of $k-1$ visits accounts for a different new occupant. ∎

Let $M(t)$ and $A(t)$ denote respectively the number of customers waiting in all queues after step t and the number of arrivals to all queues during steps $1, \ldots, t$.

Lemma 2. *For all $k \geq 1$,*

$$M(kN) \leq A(kN) + \max\{0, M(0) - (k-1)N\} \ .$$

Proof. Write $M(kN) = A(kN) + M(0) - M'$ where M' is the number of customers that left the queues on steps $1, 2, \ldots, kN$. If all $M(0)$ original customers have left, $M' \geq M(0)$. Otherwise, some queue contained an original customer for kN steps and Lemma 1 shows that $M' \geq (k-1)N$. Then $M(0) - M' \leq \max\{0, M(0) - (k-1)N\}$ and the lemma is proved. ∎

To finish the proof of the theorem, consider the sequence $M(0)$, $M(kN)$, $M(2kN), \ldots$ of numbers waiting at times kN steps apart. During steps $nkN+1, \ldots, (n+1)kN$ there are $A((n+1)kN) - A(nkN) = a(n)$ new arrivals to the queues. Define a new sequence $\{W(t)\}$ by

$$\begin{aligned} W(0) &= M(0) \ , \\ W(t+1) &= a(t+1) + \max\{0, W(t) - (k-1)N\} \ , \quad t = 0, 1, 2 \ldots \ . \end{aligned}$$

The inequality $M(tkN) \leq W(t)$, $t \geq 0$, follows from Lemma 2 by induction. Note that $W(t)$ has the same form as the sojourn time of customer t in an ordinary queue where one customer arrives every $(k-1)N$ units of time, service is by order of arrival, and $\{a(t)\}$ is the sequence of service times. If $k > 1/(1-\lambda)$, then

$$(k-1)N - E[a(t)] = (k-1)N - \lambda kN > 0 \ ,$$

i.e., the mean service time is less than the time between arrivals. Standard theory [1] shows that, under this condition, $\{W(t)\}$ is an ergodic Markov chain. This means that $\{W(t)\}$ returns infinitely often to a finite set near the origin and that the time between consecutive visits to this set is integrable. Since $M(tkN) \leq W(t)$, $t \geq 0$, and since there is only a finite number of ring states, it is easy to see that both $\{M(t)\}$ and $\{Z(t)\}$ must also have this property. Then $\{Z(t)\}$ is an ergodic Markov chain.

By the above arguments the probability distributions of $M(t)$ and $W(t)$ converge to stationary distributions of random variables M and W, with M stochastically smaller than W. Now take $k = L = \lfloor 1/(1-\lambda) \rfloor + 1$. Then the stationary mean number waiting in all queues is

$$NE[Q] = E[M] \leq E[W] \leq \frac{L\sigma^2}{2N(L-1-L\lambda)} + LN\lambda \ ,$$

where the bound on $E[W]$ is an immediate consequence of inequality (9) in Kingman [18]. The bound in the theorem statement follows immediately. ■

Define $p_k = P\{d = k | d \geq k\}$, the analog of the failure rate in continuous time. The remainder of this section considers the class of travel-time distributions for which

$$p_k \geq \mu/N, \quad k \geq 1, 2, \ldots \tag{3.1}$$

where $\mu > 0$ is a given parameter, i.e., each customer on the ring at any step has a probability at least μ/N of leaving, regardless of how long it has already been on the ring. Note that the uniform distribution on $\{1, \ldots, N\}$ is covered by this model when $\mu = 1$; the theorem below will give the same sufficient condition ($\lambda < 1$) for stability as in Theorem 1.

Each cell that visits any given queue is either empty or it will become empty with probability at least μ/N. Then the waiting time at the head of the queue is stochastically smaller than a geometric random variable with parameter μ/N, independently of the other queues. This suggests that an individual queue length can be bounded stochastically by that of a single-server discrete-time queue with geometric service. Theorem 2, to follow, proves this when the number Y of arrivals to a queue in one step possesses first three moments satisfying

$$E[Y] = \lambda/N, \quad E[Y^2] = O(1/N), \quad E[Y^3] = O(1/N) \tag{3.2}$$

as $N \to \infty$. Most of the important arrival processes, e.g., those having Y with Bernoulli or Poisson distributions, satisfy (3.2). $E[Y^3]$ can fail to exist only if Y is distributed with a long tail, e.g., if $P\{Y = i\} = 90/[\pi(i+1)]^4$.

Theorem 2. *Suppose $\lambda < \mu$ and (3.1) and (3.2) hold. Then (i) $\{Z(t)\}$ is an ergodic Markov chain, (ii) the stationary queue length Q_i of the*

i^{th} queue is stochastically smaller than a random variable X, with first and second moments bounded uniformly in N, independently of the other queue lengths. That is,

$$P\{Q_1 \leq q | Q_2, \ldots, Q_N\} > P\{X \leq q\} \ . \tag{3.3}$$

Proof. Part (ii) is proved below; part (i) follows by standard arguments which are omitted.

Introduce an artificial mechanism to decide customer departures from the ring in two stages. At each step, cells draw independent binary random numbers G with $P\{G = 1\} = \mu/N$. Consider a customer beginning his k^{th} step on the ring. In stage 1, the customer examines the G in his cell and leaves if $G = 1$. Stage 2 is needed if $G = 0$; the customer now makes a random decision to leave with probability $(p_k - \mu/N)/(1-\mu/N)$, where $p_k = P\{d = k | d \geq k\} \geq \mu/N$. The effect of the two stages is to make the customer leave with probability p_k, as required.

Now replace one of the queues with a new queue that can release a customer only into a cell with $G = 1$. Compare the original system with the new one by starting them in the same state at time $t = 0$ and letting them have the same arrivals and the same sequences of G values in the cells. The times the new queue can release a customer (when $G = 1$) are a subset of the times when the original queue can release a customer (the new queue cannot release a customer when the original queue releases one in stage 2 nor can it release one if the cell arrives empty but has $G = 0$). It follows by induction on t that the number in the original queue remains bounded by the number X in the new queue. Moreover X is not influenced by events occurring at the other $N - 1$ queues.

Note that X is the number in an ordinary queueing system in which the customer at the head of the queue is considered to be receiving service; his service ends and he enters the ring when a cell with $G = 1$ arrives. A routine analysis of this queueing system gives the generating function $X(z) = \sum_{i \geq 0} P\{X = i\} z^i$ for the distribution of X. Introduce the generating function $Y(z)$ of the distribution of the number Y of arrivals to a single queue in one step. The identity relating $X(z)$ to $Y(z)$ is

$$NX(z)D(z) = (\mu - \lambda)(1 - z) \tag{3.4}$$

where

$$D(z) = [z + \mu(1-z)/N]Y(z) - z \ .$$

The moment $E[X^k]$ is obtainable from the first k derivatives of $X(z)$, evaluated at $z = 1$. Differentiate (3.4) $k+1$ times to get an identity for derivatives of $X(z)$. In it, the derivative $k+1$ is multiplied by $D(z)$. Since $D(1) = 0$, the moment identity only relates the first k moments of X to the first $k+1$ moments of Y. It then follows from (3.2) that X has bounded first and second moments. ∎

Theorem 2 will be useful in a study of the limit $N \to \infty$, when the mean and variance of the random variable X remain $O(1)$. The fact that the stochastic bound holds simultaneously for all queues will be crucial.

4. ASYMPTOTICS

This section examines the stationary queue-length distribution under the greedy policy in the limit $N \to \infty$. The arrival and travel-time assumptions are those of Theorem 2. Then the stability condition (2.1) is satisfied and the stochastic dominance (3.3) holds.

Theorem 3. *If $\lambda < \mu$ and (3.1), (3.2) hold, then*

$$\lim_{N\to\infty} E[Q] = 0 \ .$$

Proof. The proof is by contradiction. Assume that $E[Q]$ does not approach 0 as $N \to \infty$, and hence that there is a positive constant ν, e.g., $\nu = \limsup_{N\to\infty} E[Q]/2$, such that

$$E[Q] > \nu \tag{4.1}$$

for infinitely many values of N. For the remainder of the proof, N is restricted to the values for which (4.1) holds. In broad outline, the proof starts by assuming that the ring is in equilibrium at time 0, and then examines what happens in the next time interval of αN steps, where α is a small parameter that remains fixed as $N \to \infty$. The proof shows that, because of (4.1), there are enough admissions to the ring in $[0, \alpha N]$ to produce a net decrease in $E[Q]$; but in equilibrium $E[Q]$ is the same at 0 and αN, a contradiction.

Let C denote the set of all customers in queues during $[0, \alpha N]$; C is the union of C_0, the set of customers waiting in the initial state, and the set C_a of arrivals during $[0, \alpha N]$. A customer ϕ is represented by a pair (s, t), meaning that the customer arrives at time step t to the queue that is then opposite cell s. Customers in C_0 are given time $t = 0$. In diagrams like Fig. 1, customers are points (shown as small circles) in the array $H = \{1, \ldots, N\} \times \{0, \ldots, \alpha N\}$.

Exactly as in the proof of Theorem 2, departures from the ring are decided in two stages using a Bernoulli random variable G. An event $G = 1$ in a cell is called an *enabling* and denoted by ψ. An enabling also has a pair of coordinates, the time t and cell s of the enabling. In diagrams, enablings in H appear as x's. When an enabling occurs in an occupied cell, the cell must evict its occupant. Enablings that occur in empty cells have no effect. Some customers may depart before their cells become enabled, if travel times do not have the geometric distribution with parameter μ/N (as do the times between enablings at a cell). To account for cells that are empty at time 0, let each such cell have an enabling with $t = 0$. The set of all enablings in H is denoted by E.

To estimate the number of admissions to the ring during $[0, \alpha N]$, the greedy policy will be compared with a more tractable one chosen from the class of policies defined below. Lemma 3 proves that no policy in the class has a smaller sum of waiting times

$$S = \sum_{\phi \in C} w(\phi) , \tag{4.2}$$

where $w(\phi)$ is the time ϕ spends waiting in queues during $[0, \alpha N]$.

Consider those policies that place customers only into enabled cells (when a departure from a cell does not coincide with an enabling, the cell must remain vacant until its next enabling). These policies induce one-to-one maps, or *matchings*, between customers admitted to the ring and the enablings that allow the customers to enter cells. Matchings have the properties:

(i) If ϕ is matched to ψ, then the cell enabled by ψ must appear at ϕ's queue no later than ϕ's arrival there, and no later than αN.

(ii) Suppose ϕ, ϕ' are matched to ψ, ψ' in the same cell, with ψ' later than ψ. Then ψ' must occur no earlier than ϕ's departure (which is enabled by ψ' itself or some earlier enabling).

The following property trivially implies (ii) and leads to a useful dominance result.

(ii′) No two enablings matched to customers in $[0, \alpha N]$ occur in the same cell (a cell admits at most one new customer during $[0, \alpha N]$).

Lemma 3. *For given sets C and E, no matching that satisfies* (i) *and* (ii′) *has a smaller sum of waiting times S than the greedy matching.*

Remark. Figure 2 shows that the lemma need not hold for matchings satisfying (i) and (ii) but not (ii′). The nongreedy matching in Fig. 2 gives a smaller sum of waiting times than does the greedy matching, but the nongreedy matching does not satisfy (ii′).

Proof of Lemma 3. Let π be a policy that only generates matchings satisfying (i) and (ii′). If π places a customer ϕ into a cell during $[0, \alpha N]$, let $c(\phi)$ denote that cell. Both the greedy policy and π will be compared to an artificial policy π^* that sometimes returns customers from cells to queues. Under π^*, an occupied queue places a customer into the first available empty cell, just as with the greedy policy. However, if π matches a customer ϕ to an enabling, π^* gives ϕ priority over cell $c(\phi)$; if a queue containing ϕ ever meets $c(\phi)$, then ϕ must enter $c(\phi)$ even if that requires evicting a customer from $c(\phi)$. An evicted customer joins the queue that ϕ just left. Thereafter, ϕ remains in $c(\phi)$ and can not be evicted for another customer. Interchanging customers in this way leaves numbers in the queues and cells the same as for the greedy policy; π^* merely exchanges customer names. But the sum of waiting times in $[0, \alpha N]$ can be expressed as

$$S = \sum_{t=1}^{\alpha N} \sum_{i=1}^{N} Q_i(t) \ , \tag{4.3}$$

where $Q_i(t)$ is the number in queue i after step t. Therefore, π^* and the greedy policy have the same S.

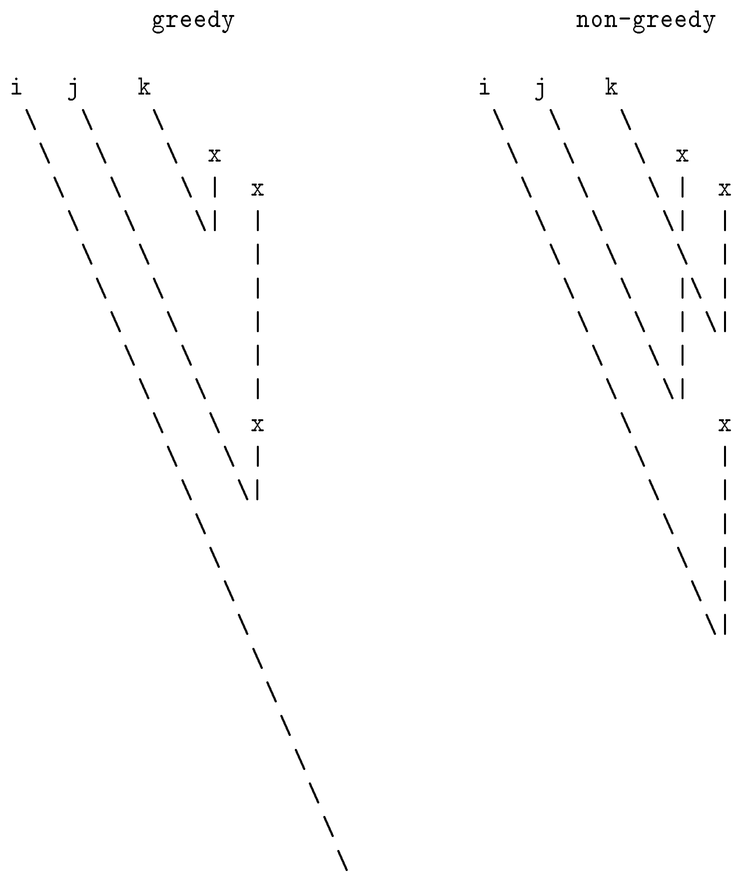

Figure 2: Two policies. In the initial state all cells are full and all queues are empty, except queues i, j, and k, which have one customer each.

It now suffices to show that $w^*(\phi) \leq w(\phi)$ for all ϕ in C, where w^* and w denote the respective times spent waiting in queues during $[0, \alpha N]$ under π^* and π. If π fails to place ϕ into a cell, then $w^*(\phi) \leq \alpha N = w(\phi)$, so suppose ϕ enters a cell under π. Although π^* may put ϕ on and off the ring several times, each step that ϕ waits in a queue brings $c(\phi)$ one step closer. Once ϕ has waited time $w(\phi)$ in queues, π^* places ϕ into $c(\phi)$. Thereafter π^* will never return ϕ to a queue, for (ii$'$) forbids π from matching two customers to enablings in $c(\phi)$. Then $w^*(\phi) \leq w(\phi)$. ∎

Next, define as follows a policy π satisfying properties (i) and (ii$'$). The remainder of the proof of Theorem 3 will show that the sum of waiting times under π, and hence the sum under the greedy policy by Lemma 3, is small enough to contradict (4.1).

Policy π first subdivides H into αK^2 squares with sides N/K; K is an integer that will ultimately be chosen "sufficiently large," but K remains fixed as $N \to \infty$. Figure 3 shows a typical subdivision. The square H_{ij} contains the points (s,t) with

$$s - iN/K = 1, \ldots, N/K \quad \text{and} \quad t - jN/K = 1, \ldots, N/K \ .$$

For simplicity, treat $2K$ as a divisor of N and αN; otherwise, unimportant edge effects become involved. On the top row of H, π likewise subdivides the points with $t = 0$ into K intervals H_{0j} lying above H_{1j}. The initial customers in C_0 appear as x's in the H_{0j} and the arrivals in C_a appear as x's in the H_{ij}, $i \geq 1$.

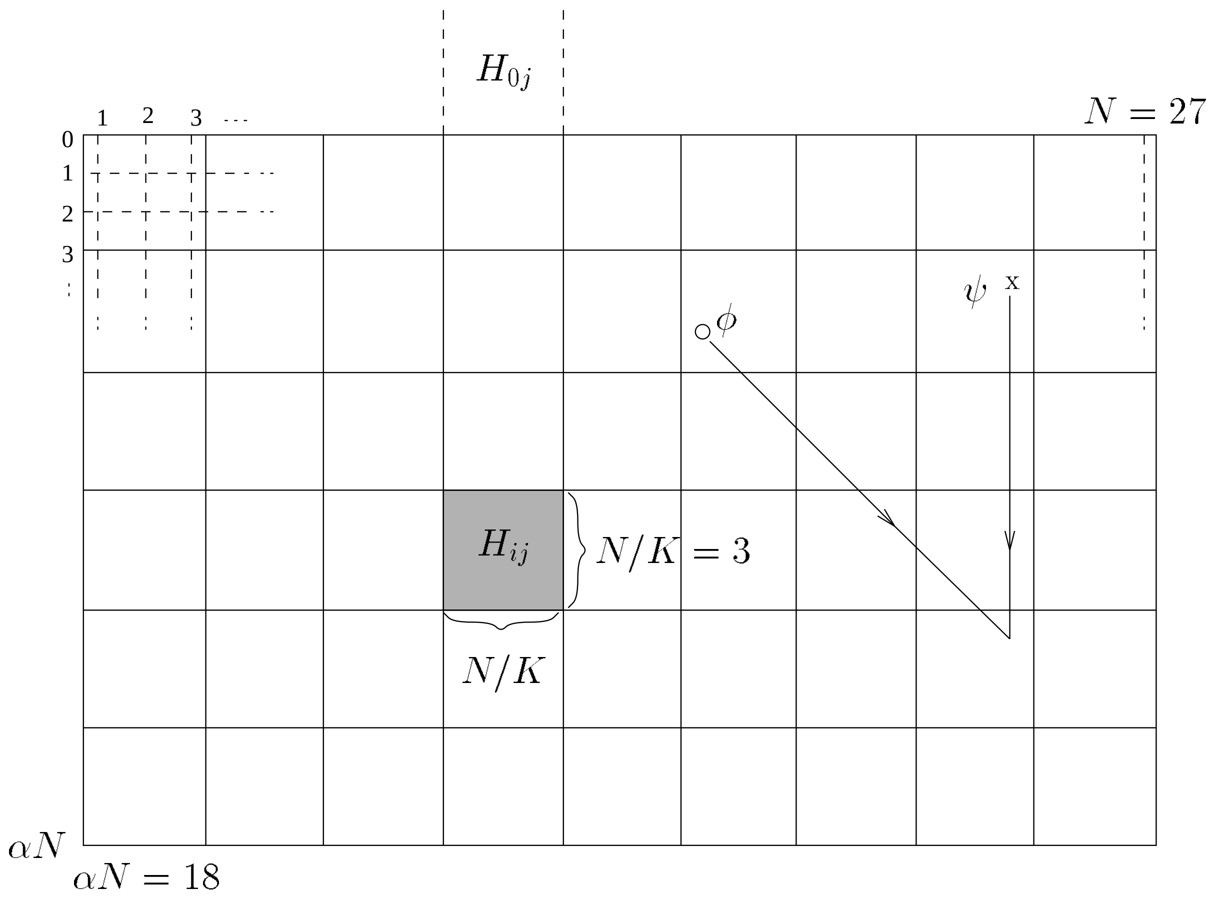

Figure 3: Partition of H ($K = 9$, $N = 27$, $\alpha = 2/3$), showing a customer ϕ in $H_{2,6}$ served by a cell with enabling ψ in $H_{2,8}$.

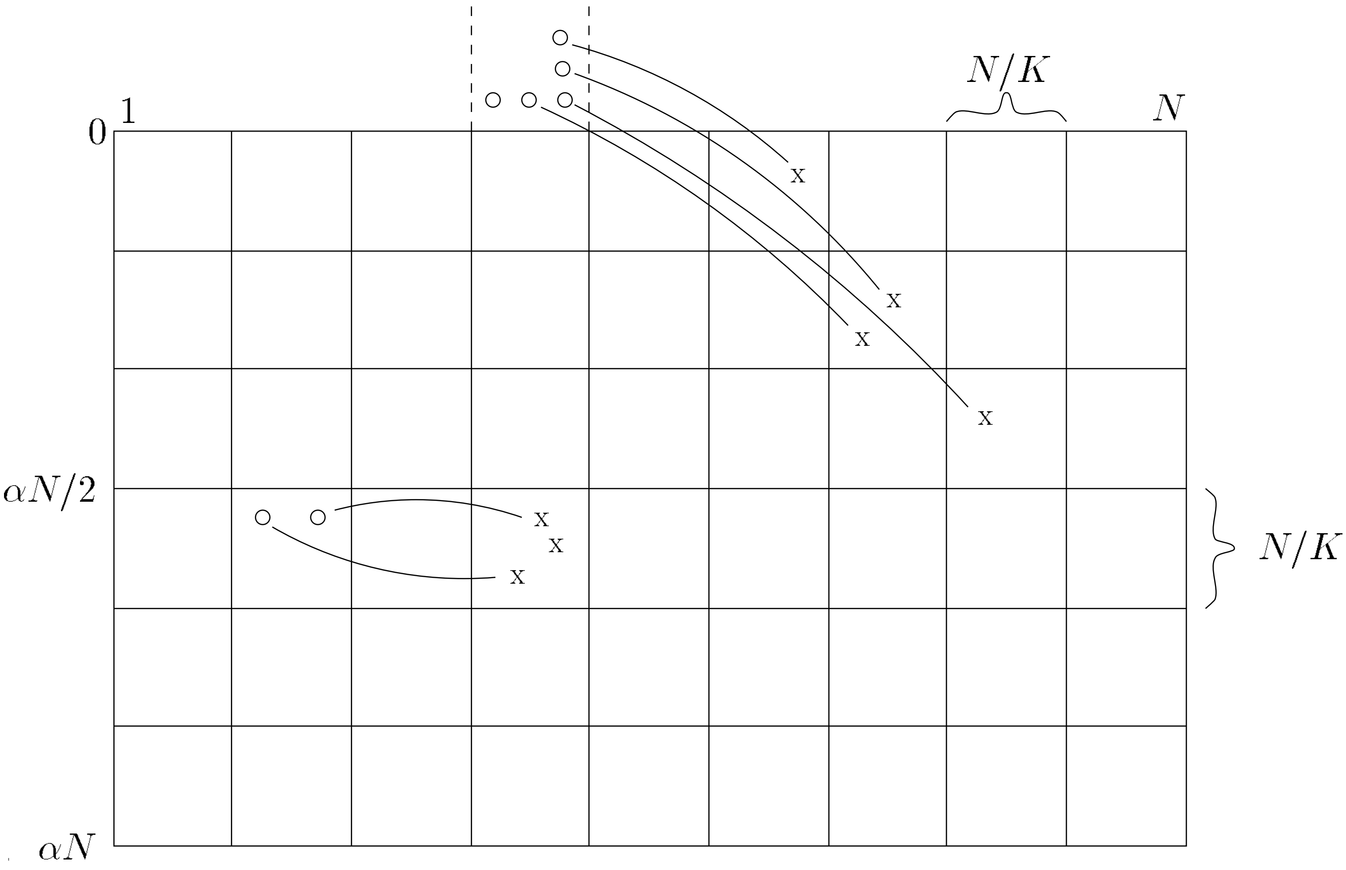

Figure 4: Illustrating steps 2, 3 of π. Matches from H_{42} to H_{44} illustrate step 2, and the others illustrate step 3. The x's shown in H_{16}, H_{27}, H_{38} are among those left over after step 2.

Next, π performs the following three steps (see Figure 4).

1. If a cell has any enablings in $[0, \alpha N]$, π chooses one at random, with all enablings equally likely. The enabling chosen is the only one from that cell that may be matched to a customer during $[0, \alpha N]$; the other enablings are discarded. This step ensures that the matching will satisfy property (ii$'$).

2. Next, π matches the arrivals in $H_{ij} (i \geq 1)$, in any convenient way, to the enablings that remain in $H_{i,j+2}$ (subscripts are added modulo K); π leaves no customer in H_{ij} unmatched if it can be matched to an unused enabling in $H_{i,j+2}$. Customers in H_{ij} all move below $H_{i,j+2}$ (see Fig. 4), and so these matchings have property (i).

3. To match initial customers to enablings that remain from steps 1 and 2, π first matches as many customers in H_{0j} as possible to enablings

left in $H_{1,j+2}$. Next, π matches as many of these customers, that are still unmatched, to enablings left in $H_{2,j+3}$; π continues diagonally downward, but uses only the enablings in $[0, \alpha N/2]$, the top half of H (see Fig. 4). Any customers that remain unmatched after trying $H_{\alpha K/2, j+\alpha K/2+1}$ are unmatched throughout $[0, \alpha N]$. Again, the matching has property (i).

Let R_a and R_0 denote the subsets of customers in C_a and C_0 that are matched in steps 2 and 3. The waiting-time bounds

$$w(\phi) \leq 3N/K, \quad \phi \in R_a \ , \tag{4.4}$$

and

$$w(\phi) \leq \alpha N/2 + 3N/K, \quad \phi \in R_0 \tag{4.5}$$

follow easily from the definition of π (see Fig. 4). All unmatched customers have

$$w(\phi) \leq \alpha N, \quad \phi \in (C_a - R_a) \cup (C_0 - R_0) \ . \tag{4.6}$$

The proof of Theorem 3 will be constructed from these bounds and the following two lemmas.

Lemma 4. *There exists an $\alpha > 0$ independent of both K and N such that the expected number of unmatched arrivals in $[0, \alpha N]$ satisfies*

$$\frac{E|C_a - R_a|}{N} \to 0 \quad as \quad N \to \infty \ .$$

Lemma 5. *Under the assumption $E[Q] > \nu$ in (4.1), there exists a constant $\gamma > 0$ independent of N such that the expected number of matched initial customers satisfies*

$$\liminf \frac{E|R_0|}{N} > \gamma \quad as \quad N \to \infty \ .$$

The analysis below first shows that these lemmas lead to the desired contradiction of (4.1); the proofs of Lemmas 4 and 5 then conclude the proof of Theorem 3.

By Lemma 3 and (4.4)–(4.6), the sum S of waiting times under the greedy policy has an expected value bounded by

$$\begin{aligned} E[S] \ \leq \ & E|R_0|(\alpha N/2 + 3N/K) + E|C_0 - R_0| \cdot \alpha N \\ & + E[R_a| \cdot 3N/K + E|C_a - R_a| \cdot \alpha N \ . \end{aligned}$$

Rearranging and applying Lemma 4 gives

$$E|R_0|(\alpha N/2-3N/K) \leq E|C_0|\cdot\alpha N-E[S]+E|R_a|\cdot 3N/K+o(N^2)\ , \quad (4.7)$$

as $N \to \infty$. Now, $E|R_a| \leq E|C_a| = \lambda\alpha N$ and (4.3) implies $E[S] = \alpha N^2 E[Q] = \alpha N E|C_0|$. Substituting for $E|R_a|$ and $E[S]$ in (4.7), and then dividing both sides by $N^2(\alpha/2-3/K)$ gives

$$\frac{E|R_0|}{N} \leq \frac{6\lambda}{K-6/\alpha} + o(1)\ , \quad (4.8)$$

as $N \to \infty$. But α is chosen independently of K, so the right-hand side of (4.8) can be made as small as desired by taking K and N sufficiently large. This contradicts Lemma 5 and proves the theorem.

Proof of Lemma 4. Let μ' be the effective rate of enablings remaining after step 1 of π. The first requirement is to verify that $\alpha > 0$ can be chosen so that this rate of enablings is at least the rate of new customers, i.e., $\lambda < \mu'$.

Recall that the enabling chosen to remain in a column was selected at random with each one equally likely. Columns with no enabling have probability $(1-\mu/N)^{\alpha N} < e^{-\alpha\mu}$. Then

$$\mu' > \frac{(1-e^{-\alpha\mu})N}{\alpha N} = \mu - \frac{\alpha\mu^2}{2} + O(\alpha^2)\ .$$

If α is small enough, $\lambda < \mu' < \mu$, as desired.

Next, let n_{ij}° denote the number of arrivals (circles) in H_{ij}, and let n_{ij}^{x} denote the number of enablings (x's) in H_{ij} after step 1 of π. The means are $E[n_{ij}^{\circ}] = \lambda N/K^2$ and $E[n_{ij}^{\mathrm{x}}] = \mu' N/K^2$, so for any $\epsilon > 0$ the law of large numbers says that inequalities

$$\frac{n_{ij}^{\circ}}{N} < \frac{\lambda+\epsilon}{K^2}, \quad (4.9)$$

$$\frac{n_{ij}^{\mathrm{x}}}{N} > \frac{\mu'-\epsilon}{K^2}, \quad (4.10)$$

hold with probability $1-o(1)$. As i and j range over $1 \leq i \leq \alpha K$ and $1 \leq k \leq K$, (4.9) and (4.10) give $2\alpha K^2$ inequalities. Boole's inequality [11, p. 23] then shows that all the inequalities hold simultaneously with probability at least $1-2\alpha K^2 o(1) = 1-o(1)$. Now choose α so small that

$\mu' > \lambda$, and take $\epsilon < (\mu' - \lambda)/2$. Then with high probability all of the arrivals in C_a ($|C_a| = \sum_{i,j} n^\circ_{ij}$) become matched, i.e.,

$$Pr\{|C_a - R_a| = 0\} \to 1 \quad \text{as} \quad N \to \infty \ . \tag{4.11}$$

By the assumptions of the arrival process,

$$\frac{E|C_a|}{N} \to \alpha\lambda, \quad \frac{Var|C_a|}{N} = O(1) \ , \quad \text{as} \quad N \to \infty \ ,$$

so the family of random variables $\left\{\frac{|C_a|}{N}\right\}_{N\geq 1}$ is uniformly integrable (see [6], pp. 219–221). Then the family $\left\{\frac{|C_a - R_a|}{N}\right\}_{N\geq 1}$ is also uniformly integrable, which together with (4.11) implies the lemma, i.e., $\frac{E|C_a - R_a|}{N} \to 0$ as $N \to \infty$. ■

Proof of Lemma 5. Write the number of initial waiting customers as a sum $|C_0| = n_1 + \cdots + n_K$, where n_j is the number waiting in the j^{th} block H_{0j} of N/K queues. The proof shows first that three conditions limiting the n_j hold jointly with some positive probability as $N \to \infty$ (see (4.14)–(4.16) below). The remainder of the proof argues that the lemma follows from these conditions and the probability estimate.

The N/K queue lengths of the j^{th} block are not independent, but they are bounded stochastically by independent random variables X_k, each distributed as X in Theorem 2. Then, for each j, the law of large numbers gives

$$Pr\{n_j < 2E[X]N/K\} \ > \ Pr\left\{\sum_{k=1}^{N/K} X_k > 2E[X]N/K\right\} \tag{4.12}$$

$$= \ 1 - o(1) \quad \text{as} \quad N \to \infty \ . \tag{4.13}$$

Because of Theorem 2, bounds (4.13) hold independently for $n_1, \ldots, n_K$, so all K inequalities

$$n_j < 2E[X]N/K, \quad 1 \leq j \leq K \ , \tag{4.14}$$

hold simultaneously with probability $[1 - o(1)]^K = 1 - o(1)$ as $N \to \infty$.

Next, choose $\epsilon = \frac{\mu - \lambda}{4}$ in (4.9), (4.10) and choose α sufficiently small that the rate μ' of enablings after step 1 satisfies $\mu - \epsilon < \mu' < \mu$. Then the overall rate of unmatched enablings after step 2 (i.e., after customers in C_a have been matched) exceeds

$$\mu' - \epsilon - (\lambda + \epsilon) = \mu - \lambda - 3\epsilon = \frac{\mu - \lambda}{4} .$$

Then (4.9) and (4.10) imply that the αK^2 inequalities

$$n^{\mathrm{x}}_{ij} - n^{\mathrm{o}}_{i,j-2} > \frac{N(\mu - \lambda)}{4K^2}, \quad 1 \le i \le \alpha K, \ \ 1 \le j \le K \ , \tag{4.15}$$

for numbers of unmatched enablings in H_{ij}, hold simultaneously with probability $1 - o(1)$.

Finally, by assumption (4.1),

$$\frac{E|C_0|}{N} = E[Q] > \nu > 0 .$$

Now $|C_0| = \sum Q_i(0)$ is stochastically no larger than the sum of N independent random variables, each distributed as X in Theorem 2. By Theorem 2, X has bounded first and second moments, so $\frac{|C_0|}{N}$ is uniformly integrable. Together with $E|C_0|/N > \nu$, this implies that there exist positive numbers ν' and δ such that $|C_0|/N \ge \nu'$, and therefore

$$n_1 + \cdots + n_K \ge \nu' N \ , \tag{4.16}$$

holds with a probability exceeding δ.

Now when all of (4.14)–(4.16) hold there cannot be many small n_j. Indeed, a fraction

$$\beta = \frac{\nu'}{4E[X]} \tag{4.17}$$

or more of them satisfy

$$n_j > \frac{\nu' N}{2K} . \tag{4.18}$$

For, otherwise, at least $(1 - \beta)K$ of the n_j fail to satisfy (4.16) and the remaining n_j are bounded by (4.14); then

$$n_1 + \cdots + n_K < (1 - \beta)\frac{KN\nu'}{2K} + \beta K \frac{2E[X]N}{K} < \nu' N$$

in violation of (4.16).

After step 2, π uses the unmatched enablings from a sequence of $\alpha K/2$ squares for matching initial customers in a block H_{0j}. By (4.15) the number of these customers that are matched is then $\min\{\eta, n_j\}$, where

$$\eta = \frac{\alpha K}{2} \cdot (\mu - \lambda)\frac{N}{4K^2} = (\mu - \lambda)\alpha N/(8K) \ .$$

Then the βK or more blocks that satisfy (4.18) have a total number of matched customers at least

$$\beta K \min\{\eta, \nu' N/(2K)\} = \beta N \min\{\alpha(\mu - \lambda)/8, \nu'/2\} = \Omega(N) \ . \quad (4.19)$$

This lower bound on $|R_0|$ obtains whenever (4.14)–(4.16) hold jointly. But (4.14), (4.15) each hold with probability $1 - o(1)$ and (4.16) holds with probability $\delta > 0$, so $|R_0|$ has the lower bound (4.19) with probability $\delta(1 - o(1))$. Then $\liminf \frac{E|R_0|}{N} > \gamma$ as $N \to \infty$ for any γ satisfying $0 < \gamma < \delta\beta \min\{\alpha(\mu - \lambda)/8,\ \nu'/2\}$, and the lemma is proved. ∎

This completes the proof of Theorem 3. ∎

5. FINAL REMARKS

In a parallel effort, Dai and Weiss [10] have applied different techniques in a proof that $\lambda E[d] < N$ is a sufficient as well as necessary condition for stability for any distribution of d with a finite mean. Note that Theorems 1 and 2 here contain more information than just this stability result; bounds are given in both cases.

Asymptotics in N pose interesting open questions. Does $\lim E[Q] = 0$, $N \to \infty$, when $1 \leq \lambda < 2$ and d is uniformly distributed on $\{1, \ldots, N-1\}$? Recall also that the simulations of Section 2 suggested the even stronger result, $E[Q] = O(1/N)$.

Adding more realistic features to the computer-ring model broadens the scope for further research. For example, the model here makes no allowance for variable-length packets or cache memory coherence [14]. A further complication is that a packet's destination may be too busy to accept it. The ring then carries extra messages of acceptance or rejection as well as retrial packets.

The greedy policy for placing customers on the ring idealizes register insertion or SCI [14] and gives a discipline that some might consider unfair. An analysis of policies with greater fairness (e.g., smaller waiting-time variances) would be interesting; such policies can require empty cells to pass occupied queues and serve other queues where customers have waited longer. However, when comparing such policies with the greedy policy, a trade-off must be expected; fairness may entail longer mean queue lengths.

References

[1] Asmussen, S., *Applied Probability and Queues*, Wiley, New York, 1987.

[2] van Arem, B., "Analysis of a queueing model for the Orwell slotted ring," *Performance '90*, King, Mitrani and Pooley (eds.), Elsevier, 1990, 61–75.

[3] van Arem, B. and van Doorn, E. A., "Analysis of a queueing model for slotted ring networks," *Computer Networks and ISDN Systems*, **20** (1990), 309–314.

[4] Avi-Itzhak, B., "Heavy Traffic Characteristics of a Circular Data Network," *Bell. Sys. Tech. J.*, **50** (1971), 2521–2549.

[5] Barroso, L. A. and Dubois, M., "The Performance of Cache-Coherent Ring-Based Multiprocessors," *Proc. 20th Ann. Internat. Symp. Comput. Arch.*, San Diego, CA (1993), 268-277.

[6] Billingsley, P., *Probability and Measure*, 2nd Edition, Wiley, New York, 1986.

[7] Brandenburg, L. H., Gopinath, B., and Kurshan, R. P., "On the Addressing Problem of Loop Switching," *Bell Sys. Tech. J.*, **51** (1972), 1445–1469.

[8] Bux, W. and Schlatter, M., "An approximate method for performance analysis of buffer insertion rings," *IEEE Trans. on Communications*, **31**, No. 1, 51–55.

[9] Cidon, I. and Ofek, Y., "MetaRing – A Full-Duplex Ring with Fairness and Spatial Reuse," *IEEE Trans. on Comm.*, **COM-41** (1993), 110–120.

[10] Dai, J. G. and Weiss, G., "Stability and Instability of Fluid Models for Certain Re-Entrant Lines," Tech. Rep., Georgia Institute of Technology, Atlanta, GA, Feb. 6, 1994.

[11] Feller, W., *An Introduction to Probability Theory and Its Applications, Vol. 1, 2nd Ed.*, Wiley, NY 1957.

[12] Fuwa, Y. and Tasaka, S., "Register insertion type slotted rings: a performance analysis," *IEEE proceedings* (1991), 191–201.

[13] Graham, R. L. and Pollak, H. O., "On the Addressing Problem for Loop Switching," *Bell Sys. Tech. J.*, **50** (1971), 2495–2519.

[14] Gustavson, D., "The Scalable Coherent Interface and Related Standards Projects," *IEEE Micro*, **12**, (1992), 10–22.

[15] Hopper, A. and Needham, R. M., "The Cambridge fast ring network system," *IEEE Trans. on Computers*, **37**, No. 10 (1988), 1214–1223.

[16] Kahale, N. and Leighton, F. T., "Greedy Routing on Arrays," Laboratory for Computer Science, MIT, Cambridge, MA 02139 (in preparation).

[17] Kendall Square Research, *The KSR1 System – Technical Summary*, Waltham, Mass. (1992).

[18] Kingman, J., "Inequalities in the Theory of Queues," *J. Roy. Stat. Soc.*, B 32 (1970), 102–110.

[19] Leighton, F. T., "Average-Case Analysis of Greedy Routing Algorithms on Arrays," *Proc. Symp. Parallel Alg. & Arch.*, ACM Press, New York, 1990, 2–10.

[20] Ofek, Y., "Overview of the MetaRing Architecture," *Computer Networks and ISDN Systems*, **26**, No. 6–8 (1994), 817–830.

[21] Pierce, J. R., "Network for Block Switching of Data," *Bell Sys. Tech. J.*, **51** (1972), 1133–1145.

[22] Scott, S. L., Goodman, J. R., and Vernon, M. K., "Performance of the SCI Ring," *Proc. 19th Internat. Symp. Comput. Arch.*, (1992), 403–414.

[23] Vranesic, Z. G., Stumm, M., Lewis, D. M., and White R., "Hector: A Hierarchically Structured Shared Memory Multiprocessor," *IEEE Computer*, **24** (1991), 72–78.