Optimizing Static Calendar Queues

K. BRUCE ERICKSON and RICHARD E. LADNER
University of Washington
and
ANTHONY LAMARCA
Yahoo! Inc.

The calendar queue is an important implementation of a priority queue that is particularly useful in discrete event simulators. We investigate the performance of the static calendar queue that maintains N active events. The main contribution of this article is to prove that, under reasonable assumptions and with the proper parameter settings, the calendar queue data structure will have constant (independent of N) expected time per event processed. A simple formula is derived to approximate the expected time per event. The formula can be used to set the parameters of the calendar queue to achieve optimal or near optimal performance. In addition, a technique is given to calibrate a specific calendar queue implementation so that the formula can be applied in a practical setting.

Categories and Subject Descriptors: E.1 [**Data**]: Data Structures; F.2 [**Theory of Computation**]: Analysis of Algorithms and Problem Complexity

General Terms: Algorithms, Performance, Theory, Verification

Additional Key Words and Phrases: Algorithm analysis, calendar queue, data structures, discrete event simulation, Markov chain, optimization, priority queue

1. INTRODUCTION

The *calendar queue* data structure, as described by Brown [1988], is an important implementation of a priority queue that is useful as the event

R. E. Ladner's research was supported by NSF grant CCR-9108314. A. LaMarca's research was done while he was a graduate student at the University of Washington, and was partially supported by an AT&T Fellowship.
A preliminary version of this paper appeared in *Proceedings of the 35th Annual IEEE Symposium on Foundations of Computer Science* (Nov. 20–22, 1994), pp. 732–743.
Authors' addresses: K. B. Erickson, Department of Mathematics, University of Washington, Box 354350, Seattle, WA 98195; email: erickson@math.washington.edu; R. E. Ladner, Department of Computer Science and Engineering, University of Washington, Box 352350, Seattle, WA 98195; email: ladner@cs.washington.edu; A. LaMarca, Yahoo! Inc., 3420 Central Expressway, Santa Clara, CA 95051; email: lamarca@yahoo-inc.com.

queue in a discrete event simulator. At any time in a discrete event simulator there are N active events, where each event e has an associated event time $t(e)$ when it is intended to occur in simulated time. The set of events is stored in the priority queue ordered by their associated event times. A basic simulation step consists of finding an event e_0 that has the smallest $t(e_0)$, removing the event from the priority queue, and processing it. As a result of the processing, new events may be generated. The parameter N can vary if zero or more than one new even is generated. Each new event e has an event time $t(e) > t(e_0)$ and must be inserted in the priority queue accordingly.

The events are stored in buckets in the calendar queue, with each bucket containing events whose times are close to each other. All the events with the smallest times are in the same bucket, so they can be accessed quickly and simulated. Any newly generated event can be quickly put into its bucket. When the events in one bucket are consumed, the next bucket is considered. The details of the algorithm are given later. The calendar queue has several user controllable parameters, the bucket width, and number of buckets, that affect its performance. Brown [1988] provided empirical evidence that the calendar queue, with its parameters properly set, achieves expected constant time per event processed. The goal of this article is to prove the constant time per event of the calendar queue behavior in a reasonable model where, for each new event e, the quantity $t(e) - t(e_0)$ is a nonnegative random variable sampled from some distribution.

Generally, the number of active events may vary over time. An important case is the *static* case, which arises when N is a constant, such as the case of simulating a parallel computer. In this case, each event corresponds to either the execution of a segment of code or an idle period by one of the processors. Thus, if there are N processors, then there are exactly N active events in the priority queue. In this article we focus on the static calendar queue.

Even before Brown's [1988] article, the calendar queue was used in discrete event simulators when the number of events was large. In many of these situations the calendar queue significantly outperforms traditional priority queue data structures [Brown 1978; Francon et al. 1978; Knuth 1973; Sleator and Tarjan 1985; 1986; Vuillemin 1978]. An interesting new development is the employment of a calendar queue-like data structure as part of the queueing mechanism of high-speed network switches and routers [Rexford et al. 1997]. In this case the calendar queue-like data structure is implemented in hardware.

1.1 Organization

In Section 2 we define the calendar queue data structure and the parameters that govern its performance. Section 3 presents the Markov chain model of calendar queue performance. Section 4 presents an expression

Table I

0	1	2	3 ⇓	4	5	6	7	8	9
-	111	128	31	-	54	-	-	85	98
			138		251				

that describes the performance of the calendar queue in the infinite bucket case. The bucket width can be chosen to approximately minimize the expected time per event at a constant. Section 5 describes how to choose the number of buckets without significantly compromising the performance over the infinite bucket calendar queue. Section 6 develops a technique for calibrating a calendar queue implementation and demonstrates the effectiveness of the technique. Section 7 presents our conclusions and the Appendix contains the longer technical proofs.

2. THE CALENDAR QUEUE

A calendar queue has M *buckets* numbered 0 to $M - 1$, a *current bucket* i_0, a *bucket width* δ, and a *current time* t_0. We have the relationship that $i_0 = t_0/\delta \bmod M$. For each event e in the calendar queue, $t(e) \geq t_0$, and event e is located in bucket i if and only if $i \leq t(e)/\delta \bmod M < (i + 1)$. The analogy with a calendar can be stated by the following: there are M days in a year each of duration δ, and today is i_0, which started at absolute time t_0. Each event is found on the calendar on the day it is to occur, regardless of the year.

As an example, choose $N = 8, M = 10, \delta = 10, i_0 = 3$, and $t_0 = 30$ (see Table I). The eight events have times 31, 54, 85, 98, 111, 128, 138, 251. In this example, the next event to be processed has time 31, which is in the current bucket numbered 3. Suppose it is deleted and the new event generated has time 87. Then the new event is placed in bucket 8 next to the event with time 85. Since $138 \geq 40$, it will not be processed until the current bucket has cycled around all the buckets once. Thus, t_0 is increased by δ, and the next bucket to be examined is bucket 4, which happens to be empty. Thus, the processing of the buckets is done in cyclic order and only the events e that are in the current cycle, $t_0 \leq t(e) < t_0 + \delta$, are processed.

A calendar queue is implemented as an array of lists. The current bucket is an index into the array, and the bucket width and current time are either integers, fixed-point, or floating-point numbers. Each bucket can be implemented in a number of ways, most typically as an unordered linked list or as an ordered linked list. In the former case, insertion into a bucket takes constant time and deletion of the minimum from a bucket takes time proportional to the number of events in the bucket. In the latter case, insertion may take time proportional to the number of events in the bucket, but deletion of the minimum takes constant time. The choice of algorithm for managing the individual buckets is called the *bucket discipline*. In this article we focus on the unordered list bucket discipline.

2.1 Calendar Queue Performance

For the calendar queue, the performance measure we are most interested in is the *expected time per event*, that is, the time to delete the event with minimum time and insert the generated new event. There are two key (user controllable) parameters in the implementation of a calendar queue that affect its performance, namely, the *bucket width* δ and the *number of buckets* M. The choice of the best δ and M depends on the number of events N and the process by which $t(e)$ is chosen for a newly generated event e. Assuming M is very large (infinite), if δ is chosen too large, then the current bucket will tend to have many events, which is inefficient. On the other hand if δ is chosen too small, then there will be many empty buckets to traverse before reaching a nonempty bucket, which again is inefficient. Regardless of the choice of δ, if M is chosen too small, then the current bucket will again tend to have too many events in it which are not to be processed until later visits to the same bucket.

In order to analyze the calendar queue, we make some simplifying assumptions on the process by which $t(e)$ is chosen for a new event e. The main assumption we make is that the quantity $t(e) - t(e_0)$, called the *jump*, is a random variable sampled from some distribution that has a mean μ, where e_0 is the event with minimum time $t(e_0)$. We will fully delineate the simplifying assumptions later. The choice of a good δ certainly depends on both μ and N. As μ grows, so should δ. As N grows, δ should decrease. Determining exactly how δ should change as a function of μ and N to achieve optimal performance is a goal of this article.

Assume that we have infinitely many buckets. In addition to the two parameters μ and N, the choice of a good δ also depends on three hidden implementation parameters b, c, and d where b is the incremental time to process an empty bucket, c is the incremental time to traverse a member of a list in search of the minimum in the list, and d is the fixed time to process an event. If m empty buckets are visited before reaching a bucket with n events ($n \geq 1$), then the time to process an event is defined to be

$$bm + cn + d.$$

Define $K_N(\delta)$ to be the stationary expected value of $bm + cn + d$. Then $K_N(\delta)$ is the expected time per event in the infinite bucket calendar queue.

In a real implementation of a calendar queue, the number of buckets M is finite. In this case, it may happen that some events in the bucket have times that are not within δ of the current time and are not processed until much later. Define $K_N^M(\delta)$ to be the expected time per event in the M bucket calendar queue. Generally, $K_N^M(\delta) \geq K_N(\delta)$ because extra time may be spent traversing events in buckets that are not processed until later. Another goal of this article is to determine how to choose M so that $K_N^M(\delta)$ is the same or only slightly larger than $K_N(\delta)$.

Table II. Parameters of the Calendar Queue

N	Number of events	Known parameter
μ	Mean of the jump	Known or estimated parameter
b	Time per empty bucket	Hidden parameter determined by calibration
c	Time per list entry	Hidden parameter determined by calibration
d	Fixed time per event	Hidden parameter determined by calibration
δ	Bucket width	User controlled parameter
M	Number of buckets	User controlled parameter

Table II summarizes the various parameters that affect the performance of the calendar queue.

Figure 1 illustrates the existence of an optimal δ for minimizing the expected time per event. Figure 2 illustrates the effect of selection of M on the expected time per event. The graphs in both figures were generated by simulating the calendar queue with an exponential jump with mean 1 and $b = c = d = 1$. Measurements were taken after a suitably long warm-up period and over a long enough period so that average time per event was very stable. The simulation of Figure 1 uses an infinite number of buckets with 100 events. The simulation of Figure 2 uses the optimal bucket width for $N = 1{,}000$ for the infinite bucket calendar queue, then varying the number of buckets. In choosing $M = 2{,}000$ or $3{,}000$ the performance curve is almost flat approaching the performance with infinitely many buckets.

3. MODELING THE CALENDAR QUEUE PERFORMANCE

To model the calendar queue performance, we begin by specifying the properties of the random variable $\xi = t(e) - t(e_0)$, where e_0 is the event with current minimal time, and e is the newly generated event. We assume that ξ is a random variable with density f defined on $[0, \infty]$, the nonnegative reals. We call f the *jump density* and its random variable simply the *jump*. Successive jumps are assumed to be mutually independent and identically distributed. Let μ be the mean of the jump; that is

$$\mu = \int_0^\infty zf(z)dz = \int_0^\infty \{1 - F(z)\}dz, \tag{1}$$

where F is the distribution function of the ξ: $F(x) = \int_0^x f(z)dz$. We call F the *jump distribution*.

We define the *support* of the jump distribution to be

$$\beta = \sup\{x \geq 0 : F(x) < 1\}$$

The value $\beta = \infty$ is not excluded.

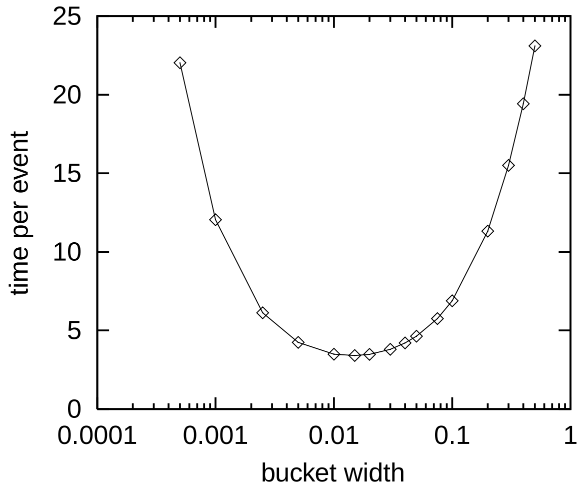

Fig. 1. Graph of bucket width δ vs. expected time per event $K_N(\delta)$ in the simulated infinite bucket calendar queue with 100 events, exponential jump with mean 1, and $b = c = d = 1$.

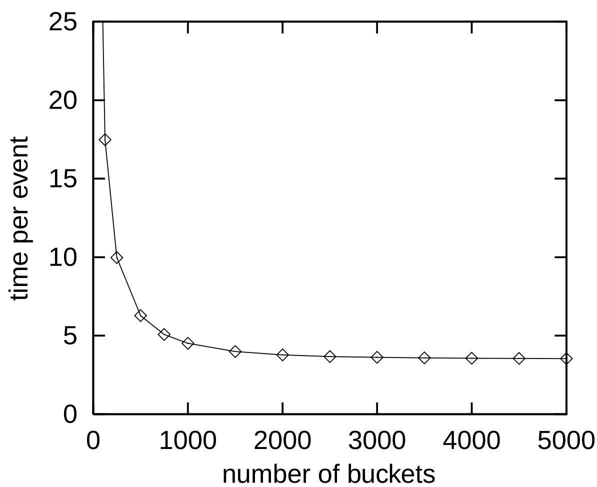

Fig. 2. Graph of number of buckets M vs. expected time per event $K_N^M(\delta)$ for δ chosen optimally in the M bucket simulated calendar queue with 1,000 events, exponential jump with mean 1, and $b = c = d = 1$.

Technical Assumptions About f and F

In order to facilitate the proofs, we make several technical assumptions about f and F that are in force throughout, except as noted:

J1. The density $f(x) > 0$ for all x in the interval $(0, \beta)$.

J2. The mean μ is finite.

J3. There is an $\epsilon_0 > 0$ and c_0 such that $F(x) \leq c_0 x$ for all $x \leq \epsilon_0$.

Assumption *J2* is crucial; it guarantees the existence of a nontrivial "steady state." Note that *J3* holds if the density f is bounded in a neighborhood of 0.

3.2 The Markov Chain

We model the infinite bucket calendar queue as a Markov chain $\hat{\mathbf{X}}$ with state space in $[0, \infty)^N$. For $t = 0, 1, 2, \ldots$, let $(X^1(t), X^2(t), \ldots, X^N(t))$ denote the state of the chain at time t. The state $(X^1(t), X^2(t), \ldots, X^N(t))$ represents the positions, relative to the beginning of the current bucket, of the N events (indexed 1 to N) in the calendar queue at step t. A step of the calendar queue consists of examining the current bucket, and either moving to the next bucket if the current bucket is empty, or removing the event with smallest time from the current bucket and inserting a new event (with the same index) according to the jump distribution. Accordingly, the transitions of $\hat{\mathbf{X}}$ are as follows. Let m be the index such that $X^m(t) = \min\{X^1(t), X^2(t), \ldots, X^N(t)\}$. If $X^m(t) \geq \delta$, then $X^i(t + 1) = X^i(t) - \delta$ for all i. If $X^m(t) < \delta$, then for $i \neq m$, $X^i(t + 1) = X^i(t)$, and $X^m(t + 1) = X^m(t) + \xi_t$ where $\xi_0, \xi_1, \ldots$ are independent nonnegative random variables. It is assumed that these random variables, ξ_t, $t \geq 0$, all have the same probability density f. The parameter δ is a fixed nonnegative real number.

We can think of $X^i(t)$ as the position of the ith particle in an N particle system. If no particle is in the interval $[0, \delta)$, then all particles move δ closer to the origin. Otherwise, the particle closest to the origin in the interval, $[0, \delta)$, jumps a random distance from its current position and the other particles remain stationary. Thus, a particle in the Markov chain $\hat{\mathbf{X}}$ represents an event in the infinite bucket calendar queue where the position of the particle corresponding to an event e is the quantity $t(e) - t_0$. The interval $[0, \delta)$ corresponds to the currently active bucket in the infinite bucket calendar queue.

It is important to note that a step of the Markov chain $\hat{\mathbf{X}}$ does not correspond to the processing of an event in the calendar queue. The processing of an event in the calendar queue corresponds to a number of steps of the Markov chain where the interval $[0, \delta)$ is empty, followed by one step where the interval $[0, \delta)$ is nonempty.

Define q_i to be the limiting probability, as t goes to infinity, that the interval $[0, \delta)$ has exactly i particles in it. Technically, $q_i = q_i(N, \delta)$ is a function of N and δ, but we drop the N and δ to simplify the notation. The quantity q_0 is the probability that the interval $[0, \delta)$ is empty. It is not obvious that q_i exists for $0 \leq i \leq N$, so we prove the following lemma in Appendix A.

LEMMA 3.1 *If the jump density has properties J1, J2, and J3, then the limiting probabilities q_i for $0 \le i \le N$ exist and are independent of the initial state of $\hat{\mathbf{X}}$.*

Let us also define $E_N(\delta)$ to be the limiting expected number of particles in the interval $[0, \delta)$; that is,

$$E_N(\delta) = \sum_{j=1}^{N} j q_j. \tag{2}$$

4. EXPECTED TIME PER EVENT IN INFINITE BUCKET CASE

The expected time to process an event in the infinite bucket calendar queue is closely related to the function $E_N(\delta)$, as we can see from the following lemma.

LEMMA 4.1 *The expected time per event in the infinite bucket calendar queue is*

$$K_N(\delta) = \frac{q_0 b + (1 - q_0)d + E_N(\delta)c}{1 - q_0}. \tag{3}$$

PROOF. The Markov chain $\hat{\mathbf{X}}$ models the calendar queue. Thus q_0 is the portion of buckets visited that are empty and for $j > 0$, q_j is the portion of buckets visited that have j events. Each empty bucket visited, which happens with probability q_0, has cost b, but does not result in finding an event to process. Each bucket visited with $j > 0$ events, which happens with probability q_j, has cost $cj + d$, and results in finding an event to process. Thus, the expected cost per event in the calendar queue is

$$K_N(\delta) = \frac{q_0 b + \sum_{j=1}^{N} q_j(cj + d)}{1 - q_0},$$

which yields Equation (3) using Equation (2). □

Let us define the following important quantity:

$$p = p(\delta) = \frac{1}{\mu}\int_0^{\delta} [1 - F(x)]dx. \tag{4}$$

The second part of Equation (1) implies that $0 \le p \le 1$. Note also that $\delta[1 - F(\delta)]/\mu \le p \le \delta/\mu$.

In order to derive a good approximating formula for $K_N(\delta)$, we first need to find good bounds for the quantities q_i for $0 \le i \le N$. The following technical lemma, proved in the Appendix, Section B, provides those bounds.

LEMMA 4.2 *For $N \geq 2$ and all $\delta > 0$, we have*

$$q_0 = \frac{\mu}{\mu + N\delta}, \tag{5}$$

and for $j = 1, 2, \ldots, N$, we have

$$q_0 B(j) \leq q_j \leq \frac{q_0 B(j)}{1 - F(\delta)} \tag{6}$$

where $B(j)$ is the tail of the binomial distribution for N trials with "success" parameter p:

$$B(j) = \sum_{k=j}^{N} \binom{N}{k} p^k (1-p)^{N-k}. \tag{7}$$

The simple exact formula for q_0 is interesting. It is possible to write down some very complicated integrals that give exact expressions for the other q_j, but these formulas are highly unwieldy and their proofs are not informative (cf., [Erickson 1999]).

It is also interesting to note that our assumption *J1*, requiring the probability density f to be positive on its support can be removed, but the proofs of the theorems become even longer. Without *J1*, if the density f has the property that there is a constant $c > 0$ such that $f(x) = 0$ for $x \leq c$, then if $\delta \leq c$, we have $F(\delta) = 0$, and (6) yields exact expressions for q_j for all $j \geq 1$.

Lemma 4.2 yields the following upper and lower bounds on $K_N(\delta)$.

LEMMA 4.3 *For $N \geq 2$ and all $\delta > 0$,*

$$\frac{\mu b}{N\delta} + \frac{\mu c}{2\delta}[(N-1)p^2 + 2p] \leq K_N(\delta) - d \leq \frac{\mu b}{N\delta} + \frac{\mu c}{2\delta}\frac{[(N-1)p^2 + 2p]}{1 - F(\delta)}. \tag{8}$$

PROOF. From (5) and (6), we have

$$E_N(\delta) = \sum_j j q_j \geq q_0 \sum_j jB(j) = \frac{\mu}{\mu + N\delta} \sum_j jB(j).$$

Define $b(i) = \binom{N}{i} p^i (1-p)^{n-i}$ to be the standard binomial distribution with N trials and success parameter p. Since $b(i) = B(i) - B(i+1)$ and $b(i)$ has mean Np and second moment $(Np)^2 + Np(1-p)$, we sum by parts to derive

$$\sum_{j=1}^{N} jB(j) = \sum_{i=1}^{N} b(i)(\sum_{j=1}^{i} j) = \sum_{i=1}^{N} \frac{1}{2}(i^2 + i)b(i)$$

$$= \frac{1}{2}[(N^2 - N)p^2 + 2Np],$$

which, upon substituting into Equation (3) and doing a little rearranging, yields the left side of (8). The right side of (8) is derived similarly. □

When $\delta = \mathcal{O}(1/N)$, δ gives good calendar queue performance. In this case the bounds of (8) give us a wonderfully simple, and accurate, approximating formula for $K_N(\delta)$.

THEOREM 4.1 *If $\delta = \mathcal{O}(1/N)$, then the expected time per event in the infinite bucket calendar queue with bucket width δ is*

$$K_N(\delta) = d + c + \frac{cN}{2\mu} \cdot \delta + \frac{b\mu}{N} \cdot \frac{1}{\delta} + \mathcal{O}(N^{-1}). \tag{9}$$

In fact, there are numbers τ_1, τ_2 such that for any fixed $\gamma > 0$ the $\mathcal{O}(N^{-1})$ term is bounded by $(\tau_2\gamma^2 + \tau_1\gamma)/N$ uniformly for $0 < \delta \leq \gamma/N$.

The proof of this theorem is an almost immediate consequence of Equations (3), (5), and (8); but is also postponed to the Appendix, Section C. Interestingly, the expected time depends on the mean of the jump, and not on the shape of its probability density.

Note that one immediate consequence of Theorem 4.1 is that if the bucket width is chosen to be $\theta\mu/N$ for θ in a fixed interval, then the infinite bucket calendar queue has constant expected time per event performance. Indeed, a formula for the optimal performance of the calendar queue can be derived, as seen in the following theorem.

THEOREM 4.2 *The expected time per event $K_N(\delta)$ achieves a global minimum in the interval $(0, \infty)$ at δ_{opt} where*

$$\delta_{opt} = \sqrt{\frac{2b\,\mu}{c\,N}} + \mathcal{O}(N^{-3/2}), \tag{10}$$

$$K_N(\delta_{opt}) = d + c + \sqrt{2bc} + \mathcal{O}(N^{-1}). \tag{11}$$

The proof of this theorem is in the Appendix, Section D. Theorem 4.2 shows that the optimal choice of δ depends only on the ratio of b to c, the mean μ of the jump, and N.

5. CHOOSING THE NUMBER OF BUCKETS

Now that we have found how to select δ so as to approximately minimize the expected time per event in the infinite bucket calendar queue, our next goal is to select M, the number of buckets, so that the M bucket calendar queue has the same or similar performance as the infinite bucket calendar queue.

For the case in which the jump distribution has finite support ($\beta < \infty$), there is a natural choice for M, which guarantees that the calendar queue with M buckets has exactly the same performance as the infinite bucket calendar queue. If $M \geq \beta/\delta + 1$, then it is guaranteed that in the long run all the events e in the current bucket will have $t_0 \leq t(e) < t_0 + \delta$. In this case, each event in the current bucket will eventually be processed during the current visit to the bucket, and not postponed until future visits to the bucket. For the case in which the support β of the jump distribution is either infinite or is finite but β/δ is too large to be practical, it will be necessary to choose a number M that gives performance less than that of the infinite bucket calendar queue.

5.1 Expected Time per Event In the Finite Bucket Case

The same Markov chain $\hat{\mathbf{X}}$ can be used to analyze this case. Let $L_N^M(\delta)$ be the (steady state) expected number of particles in the set

$$\Gamma = \bigcup_{j=1}^{\infty} [jM\delta, jM\delta + \delta).$$

In terms of the M bucket calendar queue, if an event e has $t(e) - t_0 \in \Gamma$, then the event is in the current bucket but is not processed. The occurrence of such an event will cause the M bucket calendar queue to run less efficiently than the infinite bucket calendar queue. The following lemma quantifies the difference between the performance of the finite and infinite bucket calendar queues.

LEMMA 5.1 *The expected time per event in the M bucket calendar queue with bucket width δ is*

$$K_N^M(\delta) = K_N(\delta) + \frac{c(\mu + N\delta)}{N\delta} L_N^M(\delta).$$

PROOF. In the Markov chain $\hat{\mathbf{X}}$, let q_{ij} be the limiting probability that there are i particles in the interval $[0, \delta)$ and j particles in Γ. The probabilities q_{ij} can be shown to exist in the same way as the probabilities q_i in Lemma 1 by using Corollary 1 in the Appendix, Section A. In the M bucket calendar queue the cost of visiting a bucket with i events whose times are in the interval $[t_0, t_0 + \delta)$ and j events whose times are in the set $\{t_0 + x : x \in \Gamma\}$ is $c(i + j) + d$ if $i > 0$ and $cj + b$ if $i = 0$. Thus, the expected cost per event $K_N^M(\delta)$ equals

$$\frac{\sum_{i=1}^{N}\sum_{j=0}^{N} q_{ij}(c(i+j)+d)+\sum_{j=0}^{N} q_{0j}(cj+b)}{1-\sum_{j=0}^{N} q_{0j}}.$$

But $\sum_{j=0}^{N} q_{oj} = q_0 = \mu/\mu + N\delta$ (Equation (5) of Lemma 4.2) and $\sum_{j=0}^{N} j \sum_{i=0}^{N} q_{ij} = L_N^M(\delta)$. By using Equation (3) in the proof of Lemma 4.1 we derive the equation for $K_N^M(\delta)$. □

In the Appendix, Section E, we indicate how to derive the following rather horrible looking bounds for $L_N^M(\delta)$.

LEMMA 5.2 *The function $L_N^M(\delta)$ is bounded above by*

$$\frac{\mu N(Np+1)\Pi_1+\mu Np(Np+2-p)\Pi_2}{(\mu+N\delta)[1-F(\delta)]^2}$$

and bounded below by

$$\frac{\mu N(Np+1-p)\Pi_1}{\mu+N\delta},$$

where p and F have the same meaning as before (see Equation (4)) and

$$\Pi_1 = \frac{1}{\mu}\sum_{j=1}^{\infty}\int_{jM\delta}^{jM\delta+\delta}[1-F(x)]dx,$$

$$\Pi_2 = \max_{0\le y\le\delta}\sum_{j=1}^{\infty}[F(jM\delta+\delta-y)-F(jM\delta-y)].$$

Note that under the hypothesis $\mu < \infty$ (*J2*), the above series converge and can be given bounds in terms of μ, δ, and M. However, using the bounds as stated in the lemma, we can derive a more useful asymptotic expression for $L_N^M(\delta)$. The lemma is proven in the Appendix, Section F.

LEMMA 5.3 *If $\delta = x\mu/N$ and $r = M/N$, where r and x are constants, then*[1]

$$L_N^M(\delta) \simeq x\sum_{j=1}^{\infty}[1-F(j\mu xr)].$$

[1]We define $g(N) \simeq h(N)$ if $\lim_{N\to\infty} g(N)/h(N) = 1$.

5.2 Degradation in Performance Due to Finitely Many Buckets

Define ϵ_M to be the performance *degradation* in choosing M buckets instead of infinitely many buckets; that is,

$$\epsilon_M = \frac{K_N^M(\delta) - K_N(\delta)}{K_N(\delta)}.$$

If we choose δ optimally, then Lemma 5.3 and Theorem 4.2 yield the following asymptotic expression for ϵ_M.

THEOREM 5.1 *If M/N is constant and*

$$\delta = \sqrt{\frac{2b}{c} \cdot \frac{\mu}{N}}$$

then

$$\epsilon_M \simeq \frac{c + \sqrt{2bc}}{d + c + \sqrt{2bc}} \sum_{j=1}^{\infty} [1 - F(jM\delta)] \tag{12}$$

The following asymptotic bound is implied by Theorem 5.1.

THEOREM 5.2 *If M/N is constant and*

$$\delta = \sqrt{\frac{2b}{c} \cdot \frac{\mu}{N}}$$

then [2]

$$\epsilon_M \lesssim \frac{c + \sqrt{2bc}}{d + c + \sqrt{2bc}} \cdot \sqrt{\frac{c}{2b} \cdot \frac{N}{M}} \tag{13}$$

PROOF. By Theorem 5.1, it suffices to show that $\sum_{j=1}^{\infty}[1 - F(jD)] \leq \mu/D$ for $D > 0$. To see this, let $k \geq 2$, then

$$\begin{aligned} \frac{\mu}{D} = \frac{1}{D}\int_0^{\infty} xf(x)dx &\geq \sum_{j=0}^{k} \frac{1}{D}\int_{jD}^{jD+D} xf(x)dx \\ &\geq \sum_{j=0}^{k} j[F((j+1)D) - F(jD)] \\ &= \sum_{j=1}^{k} [1 - F(jD)] - k[1 - F((k+1)D)]. \end{aligned}$$

[2] We define $g(N) \lesssim h(N)$ if $\limsup_{N\to\infty} g(N)/h(N) \leq 1$.

The finiteness of μ implies that $x[1 - F(x)] \to 0$ as $x \to \infty$. Therefore, if we let $k \to \infty$, we get $\mu/D \geq \sum_{j=1}^{\infty}[1 - F(jD)]$. $\square$

Equation (13) shows that for a fixed $\epsilon > 0$ (such as .01) M can be chosen to be $\mathcal{O}(N)$, so that $\epsilon_M \leq \epsilon$. In other words, one can always choose the number of buckets M to be a multiple of N and still obtain a performance almost as good as that of the infinite bucket case.

For the interesting case of the exponential jump density

$$f(x) = \frac{1}{\mu} e^{-x/\mu}, x \geq 0,$$

we can calculate the series in Equation (12) exactly:

$$\epsilon_M \simeq \frac{c + \sqrt{2bc}}{d + c + \sqrt{2bc}} \cdot \frac{1}{e^{\sqrt{\frac{2bM}{c\,N}}} - 1}. \tag{14}$$

Let us suppose $b = c = d = 1$. Choosing δ optimally equal to $\sqrt{2}\mu/N$ allows us to solve Equation (16) for M/N when given an acceptable ϵ_M. For example, if we choose $\epsilon_M = .05$, then M/N should be approximately 1.92, and if $\epsilon_M = .01$ then M/N should be approximately 3.02. Figure 3 illustrates that asymptotic Equation (14) provides an excellent choice of M over a wide range of N. Using our simulation of the calendar queue, we plot for a wide range of N the value of ϵ_M for each of $M/N = 1.92$ and $M/N = 3.02$. Again, measurements were taken after a suitably long warm-up period and over a long enough period, so that average time per event was very stable. Both plots are relatively flat near the asymptotic values .05 and .01, respectively. Thus, Equation (14) seems quite accurate. The bound of Theorem 5.2 is not necessarily tight because we are crudely approximating an integral. For example, if we choose $\epsilon_M = .01$, then formula (13) requires M/N to be at least 50.

6. CALIBRATING A CALENDAR QUEUE IMPLEMENTATION

In an actual calendar queue implementation, we would like to find the best bucket width δ and number of buckets M. The preceding theory tells us how to do so if we know the hidden implementation parameters b, c, and d. In this section we give a relatively simple method for estimating these parameters by simply timing executions of the simulation for various values of δ proportional to μ/N. The key to the method is Equation (10) for the expected time per event. We can write $K_N(\delta)$ as a linear function of the unknowns b, c, and d. The general calibration method is as follows: First, estimate M large enough so that the degradation in using M buckets over infinitely many is small. Second, find $K_N^M(\delta)$ for a number of different δ's by

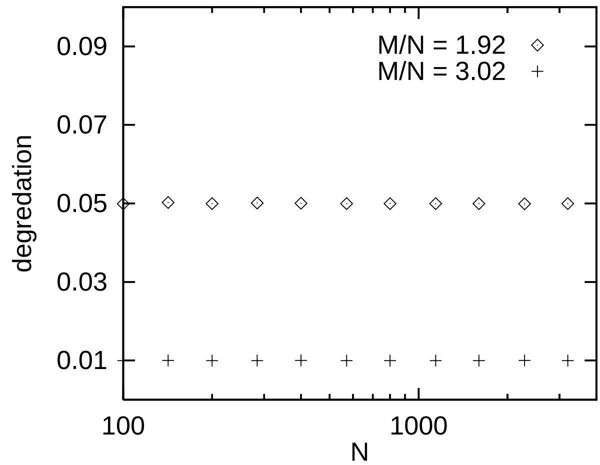

Fig. 3. Graph of N vs. degradation ϵ_M from simulations for $M/N = 1.92$ and $M/N = 3.02$.

timing executions of the implementation. Third, use a linear least squares approximation to find the b, c, and d that best fit the function

$$\frac{\mu}{N\delta} \cdot b + \left(1 + \frac{N\delta}{2\mu}\right) \cdot c + d. \tag{15}$$

We illustrate this method with an example. We developed a calendar queue implementation in C++ and ran it on a *DEC* Alphastation 250. We chose $N = 1{,}000$ and an exponential jump with mean $10{,}000$. Just by examining the code we felt that b, the time to process an empty bucket, was considerably larger than c, the cost of traversing a list entry. We made an educated guess that the optimal δ was certainly greater than $.5\mu/N = 5$. We chose $M = 10 \cdot N = 10{,}000$. Hence, for $\delta = 5$ or larger, there is only a small chance that an event in the current bucket is not processed because its time is too large. We timed the calendar queue for 20 values of δ ranging over several orders of magnitude, namely, $\delta = 5, 15, 25, \ldots, 195$. Using these data, we used linear least squares approximation to compute $b = 837.619$, $c = 44.3039$, and $d = 1439.69$ using Equation (15). Figure 4 shows the curve of Equation (15) using these parameters. The figure also shows the time per event for $\delta = 10, 20, 30, \ldots, 200$. Thus, this method accurately predicts datapoints that were not used in the linear least squares approximation. It is interesting to note that using these values of b, c, and d in Equations (10) and (11), we obtain $\delta_{opt} \approx 61.5$ and $K(\delta_{opt}) \approx 1756.38$. By contrast, the best δ among the 40 executions is $\delta = 50$ with execution time 1754.92.

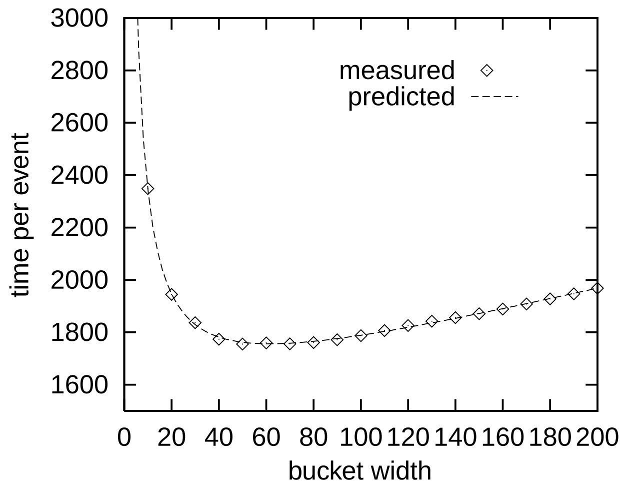

Fig. 4. Measured and predicted expected time per event for a calibrated implementation of a calendar queue.

Care must be taken in applying this calibration method because the hidden parameters b, c, and d are measured indirectly by measuring the expected time per event. For fixed M, N, and δ, the expected time per event can vary over different runs because of interruptions by other processes, page faults, or other effects. However, in our experimental setting we carefully controlled the environment so that our running times varied little for a fixed parameter setting. In a real computing environment which cannot be controlled this calibration method might not yield such good results.

Ideally, using a fixed N, M (large enough), and μ we can estimate the hidden parameters b, c, and d, which could then be used for any other N, M (large enough), μ, and δ. However, due to the cache behavior of modern processors, the values of b, c, and d are not actually constant independent of M, N, and properties of the jump distribution other than its mean. For example, a smaller M might achieve fewer cache misses reducing the running time, and thereby effectively lowering the values of these constants. It may be that in applying the calibration method, δ is chosen so large that the original M chosen is far larger than necessary. In this case, it might be wise to choose a smaller M, then recalibrate the calendar queue starting with a larger δ.

In a real application of the calendar queue, it is unlikely that the jumps are mutually independent, identically distributed random variables, as described in our model. Nonetheless, the mean of the jump can be empirically estimated, the calibration done, and Equation (10) for the optimal δ applied to find a potentially good δ.

7. CONCLUSION

We have shown that there is an expression for the expected time to process an event in the infinite bucket calendar queue and that the bucket width can be chosen optimally. With the bucket width near the optimal bucket width, the calendar queue has expected constant time per event. The optimal bucket width depends on a few parameters only, the incremental time to process an empty bucket (b), the incremental time to traverse a list item (c), the mean of the jump (μ), and the number of events (N). We have shown that the number of buckets M can be chosen to be $\mathcal{O}(N)$, so as to achieve minimal or almost minimal expected time per event. Finally, we have shown that the implementation parameters can be determined by using approximation based on the linear least squares method.

Although the calendar queue runs very fast for certain applications, it has the disadvantage that its performance depends on the choice of parameters δ and M. An interesting problem would be to design a priority queue based on the calendar queue that automatically determines good choices for δ and M. We believe that the calibration method described in this article might give insight into the design of a dynamic calendar queue where N and/or μ can vary over time.

APPENDIX

Section A sets up the notation and concepts that are used throughout the Appendix.

A. INVARIANT DISTRIBUTION, POSITIVITY, AND LIMITS

Consider the Markov chain $\hat{\mathbf{X}}$ described in Section 3. The symbol P^{ν} denotes the probability measure induced on the trajectory space of the chain $\hat{\mathbf{X}}$ when the initial distribution is ν, and $P^{\hat{\mathbf{x}}}$ denotes trajectory space probabilities when the chain starts at the point $\hat{\mathbf{X}}$. (Note: $P^{\nu} = \int P^{\hat{\mathbf{x}}} d\nu(\hat{\mathbf{x}})$, the integration being carried out over the entire state space.) Integration (better known as *expectation*) with respect to P^{ν} and $P^{\hat{\mathbf{x}}}$ is denoted E^{ν} and $E^{\hat{\mathbf{x}}}$, respectively.

Let B_i stand for the set of points $\hat{\mathbf{x}} = (x^1, x^2, \ldots, x^N)$ in $[0, \infty)^N$ such that

$$x^i < x^j, \text{ for all } j \neq i \text{ and } x^i < \delta,$$

and let $A_0 = [\delta, \infty)^N$. For an $\hat{\mathbf{x}}$ in the state space and a (measurable) subset A, the one-step transition probability ($T.\ P.$) that the chain will move from $\hat{x}$ to a point in A is given by

$$P(\hat{\mathbf{x}}, A) = \sum_{i=1}^{N} 1_{B_i}(\hat{\mathbf{x}}) \int_0^{\infty} 1_A(\hat{\mathbf{x}} + z\hat{\mathbf{e}}_i) f(z) dz + 1_{A_0}(\hat{\mathbf{x}}) 1_A(\hat{\mathbf{x}} - \delta \hat{\mathbf{1}}), \tag{16}$$

where $\mathbf{1} = (1, 1, \ldots, 1)$, $\hat{\mathbf{e}}_i =$ the standard i^{th} unit coordinate vector, and $1_A(\hat{x})$ is the function that is 1 for $\hat{x}$ in A and 0 otherwise.[3]

Let Λ be the set $[0, \delta + \beta)^N$. It follows from (16) that

$$P(\hat{\mathbf{x}}, A) = 0 \text{ forall } \hat{\mathbf{x}} \text{ in } \Lambda \text{ and } A \subset \Lambda^C.$$

In other words, Λ is an *absorbing* set for the chain. Moreover, the dynamical description of the chain implies that a particle that starts outside Λ will reach Λ in a finite (but possibly random) number of steps. (One can show, using the method in the proof in B.3, that the number of steps required to eventually enter Λ has a finite expectation.) Thus Λ^C is *transient* for the chain. (Of course, $\Lambda^C = \emptyset$ if $\beta = \infty$.)

A measure m is an *invariant measure* for the chain if m is σ-finite and for every measurable subset A of the state space

$$m\{A\} = \int P(\hat{\mathbf{x}}, A) m\{d\hat{\mathbf{x}}\}. \tag{17}$$

A Markov chain is called a *Harris recurrent chain*, or simply a Harris chain, if there exists a unique, up to positive multiples, invariant measure m such that if A is any Borel subset with $m(A) > 0$, then $P^{\hat{\mathbf{x}}}(\hat{\mathbf{X}}_t \in A \text{ i.o.}) = 1$ for all $\hat{\mathbf{x}}$ in the state space. (The initials *i.o.* stand for "infinitely often.") A Harris chain with an invariant *probability* measure, necessarily unique, is called *positive*.

The state space $\mathcal{D}$ of a Harris chain can be written as a disjoint union: $\mathcal{C}_1 \cup \ldots \cup \mathcal{C}_d \cup \mathcal{D}$, where $m(\mathcal{D}) = 0$, and for $j = 1, \ldots, d$, $P(\hat{\mathbf{x}}, \mathcal{C}_{j+1}) = 1$ for all $\hat{\mathbf{x}} \in \mathcal{C}_j$, $P(\hat{\mathbf{x}}, \mathcal{C}_1) = 1$, for all $\hat{\mathbf{x}} \in \mathcal{C}_d$ (see Revuz [1984, Chap. 6]). The sets $\mathcal{C}_1, \ldots, \mathcal{C}_d$ are known as the *recurrent cyclic classes*. The integer d is finite and if $d = 1$, the chain is called *aperiodic*.

THEOREM A.1. *If the jump density satisfies J1, J2, and J3, then the Markov chain $\hat{\mathbf{X}}$ with T.P. (18) is a positive, aperiodic, recurrent Harris chain. Its invariant probability is concentrated on $\Lambda = [0, \beta + \delta)^N$.*

Corollary A.1. If φ is any bounded measurable function, then for any initial distribution ν, we have

[3]Actually, (16) does not define a proper transition probability at all points! Indeed, if $\hat{\mathbf{x}}$ is any point with two or more coordinates that are equal and strictly less than δ, then $\hat{\mathbf{x}}$ does not lie in A_0 or in B_i for any i. Hence, $P(\hat{\mathbf{x}}, A) = 0$ for all subsets A. However, the jump distribution F has a density, so it has no *atoms* (discrete points of positive probability). This and the dynamical description of the chain in Section 3 imply that two points that start at the same position eventually become and remain separated *w. p.1.* Indeed, this occurs as soon as one of them jumps to the right. Thus, it does no harm if we initially banish such points from the state space. With this understanding, P is indeed a transition probability on its state space.

$$\lim_{t\to\infty} \frac{\sum_{s=1}^{t} \varphi(\hat{\mathbf{X}}_s)}{t} = \lim_{t\to\infty} \frac{\sum_{s=1}^{t} E^{\nu}\{\varphi(\hat{\mathbf{X}}_s)\}}{t}$$

$$= \lim_{t\to\infty} E^{\nu}\varphi(\hat{\mathbf{X}}_t)$$

$$= \int \varphi(\hat{\mathbf{x}})dm(\hat{\mathbf{x}}),\ P^{\nu} - a.s. \tag{18}$$

Note: The leftmost term is a limit of averages of random quantities and the assertion is that the limit exists with P^{ν}-probability 1 and equals the (nonrandom) quantities on the right. If φ is unbounded but integrable with respect to m and if the chain is Harris and positive, then the above limit relations remain valid at least in the case that ν is point mass at some $\hat{\mathbf{x}}$ or that $\nu = m$.

PROOF (COROLLARY A.1). The deterministic limit statements of Corollary A.1 are immediate consequences of Proposition 2.5 in Chap. 6, §2 of Revuz [1984]. The a.s. limit-of-averages assertion is a consequence of the ergodic theorem for Harris chains; see Theorem 4.3, and its companion remark, in Revuz [1984, Chap. 4 §4]. □

Remark. The main significance of aperiodicity is that it justifies the existence of the limit of $E^{\hat{\mathbf{x}}}\{\phi(\widehat{\mathbf{X}}_t)\}$ in (18). The existence of limits of averages does not require aperiodicity.

PROOF (LEMMA 3.1). Let

$$A_j = \{\hat{\mathbf{x}} : \text{ exactly } j \text{ components of } \hat{\mathbf{x}} \text{ lie in } [0, \delta)\}.$$

Then, $Z(t) = \sum_k kI_{A_k}(\hat{\mathbf{X}}_t) =$ the number of particles in interval $[0, \delta)$ at time t, and from (18) it follows immediately that $q_j = \lim_t P^{\nu}\{Z(t) = j\} = m\{A_j\}$. □

The proof of Theorem A.1 is postponed to the very last appendix. It is lengthy and somewhat tedious, but there are some interesting features.

B. PROOF OF LEMMA 4.2

B.1 Computing $\mathbf{q_0 = m\{A_0\}}$

In this section we prove (5) of Lemma 4.2.

The sets B_i, defined in the last section, are disjoint and their union is the complement (in Λ) of A_0. Since m assigns 0 mass to $[0, \infty)^N \cap \Lambda^c$, we have

$$m\{A_0\} = 1 - \sum_{1}^{N} m\{B_i\} = 1 - Nm\{B_1\}. \tag{19}$$

Let ψ be any bounded or positive function on the state space. Equation (19) has an analogue for functions, which reads: $\int \psi(\hat{\mathbf{x}})m\{d\hat{\mathbf{x}}\} = \int m\{d\hat{\mathbf{x}}\}\int \psi(\hat{\mathbf{y}})P(\hat{\mathbf{x}};d\hat{\mathbf{y}})$. Noting that $\int_{A_0} m\{d\hat{\mathbf{x}}\}\int \psi(\hat{\mathbf{y}})P(\hat{\mathbf{x}}, d\hat{\mathbf{y}}) = \int_{A_0}\psi(\hat{\mathbf{x}} - \delta\hat{\mathbf{1}})m\{d\hat{\mathbf{x}}\}$, by (16), and doing a little rearranging, we get

$$\int_{A_0}[\psi(\hat{\mathbf{x}} - \delta\hat{\mathbf{1}}) - \psi(\hat{\mathbf{x}})]m\{d\hat{\mathbf{x}}\} = \sum_{j=1}^{N}\int_{B_j} m\{d\hat{\mathbf{x}}\}[\psi(\hat{\mathbf{x}}) - \int \psi(\hat{\mathbf{y}})P(\hat{\mathbf{x}};d\hat{\mathbf{y}})]. \quad (20)$$

Fix i and let $\psi(\hat{\mathbf{x}}) = \exp(-\alpha x^i)$, where α is any complex number with nonnegative real part. Then, for $\hat{\mathbf{x}}$ in B_i,

$$\int \psi(\hat{\mathbf{y}})P(\hat{\mathbf{x}}, d\hat{\mathbf{y}}) = \int_0^\infty \exp\{-\alpha(x^i + z)\}f(z)\,dz = \exp(-\alpha x^i)\phi(\alpha),$$

where $\phi(\alpha) = \int_0^\infty e^{-\alpha z}f(z)dz$ is the Laplace transform of F. Also, for $\hat{\mathbf{x}}$ in B_j with $j \neq i$,

$$\int \psi(\hat{\mathbf{y}})P(\hat{\mathbf{x}};d\hat{\mathbf{y}}) = \int_0^\infty \exp(-\alpha x^i)f(z)dz = \exp(-\alpha x^i) = \psi(\hat{\mathbf{x}}).$$

All but the i^{th} term on the right side of (20) vanishes, and it becomes

$$[1 - \phi(\alpha)]\int_{B_i} \exp(-\alpha x^i)m\{d\hat{\mathbf{x}}\}.$$

Simplification of the left-hand side of (20) leads to

$$[e^{\alpha\delta} - 1]\int_{A_0} \exp(-\alpha x^i)m\{d\hat{\mathbf{x}}\} = [(1 - \phi(\alpha))]\int_{B_i} \exp(-\alpha x^i)m\{d\hat{\mathbf{x}}\}. \quad (21)$$

Divide (21) by α and make $\alpha \to 0$. The result is $\delta m\{A_0\} = -\phi'(0)m\{B_i\} = \mu m\{B_i\}$ for $i = 1, 2, \ldots, N$. Equation (5) follows immediately from this and (19). $\square$

B.2 The Case $\mathbf{N = 1}$

If we observe the successive positions of a single one of our N particles at only those times at which it actually moves, we get a 1-dimensional version of the N-dimensional chain. Let $u^i(0) = 0$ and for $r = 1, 2, \ldots$, let

$$u^i(r) = \min\{t : t > u^i(r-1) \text{ and } X^i(t) \neq X^i(u^i(r-1))\}.$$

From the description of the chain in terms of the independent random variables ξ, we conclude that

(i) each sequence $\{X^i(u^i(r)) : r = 0, 1, 2, \ldots\}$ is itself a Markov chain on the line;

(ii) The N Markov chains are mutually independent.

If we can find an increasing sequence of times $\{S_k\}$ such that each S_k is a common value of every one of the u^i (i.e., for each k there are numbers $r^i(k)$, not necessarily the same, such that $S_k = u^i(r^i(k))$ for every i), then, given $\hat{\mathbf{X}}(0) = \hat{\mathbf{x}}$, $\hat{\mathbf{X}}(S_k)$ has mutually independent components. Here is such a sequence: let $S_0 = 0$, and, for $k > 0$, let

$$S_k = 1 + \min\{t : t \geq S_{k-1}, \hat{\mathbf{X}}(t) \in A_0\}.$$

The times $T_k = S_k - 1$, $k \geq 1$ are the successive (random) times at which the interval $[0, \delta)$ is empty of particles ($Z(T_k) = 0$). Since $\hat{\mathbf{X}}(T_k)$ is obtained from $\hat{\mathbf{X}}(S_k)$ by adding the deterministic constant δ to each of the components of $\hat{\mathbf{X}}(S_k)$, it follows that the components of $\hat{\mathbf{X}}(T_k)$ are also mutually independent. It is easy to show that the chain induced on A_0 (or *trace chain*), the sequence $\{\hat{\mathbf{X}}(T_k), k \geq 0\}$ is also a Markov chain (see Revuz [1984, Exercise 3.13, p. 27]).

An important point to note is that the special structure of $\hat{\mathbf{X}}$ implies that for each i the chain $\{X^i(T_k), k \geq 0\}$ coincides in law with the trace chain on $[\delta, \infty]$ of an $N = 1$ version of $\hat{\mathbf{X}}$. It is at least intuitively clear that the trace chain $\{\hat{\mathbf{X}}(T_k)\}$ is also positive recurrent and has an invariant probability distribution m_0, say, obtained by renormalizing the distribution m restricted to A_0 (see Revuz [1984, Example 3.13, p. 27, and Proposition 2.9, p. 93] for a formal proof). Thus, for subsets B,

$$m_0(B) = \frac{m\{B \cap A_0\}}{m\{A_0\}} = \left(1 + \frac{N\delta}{\mu}\right) m\{B \cap A_0\}. \tag{22}$$

But because this trace chain also has independent components, it follows that m_0 is a "product measure" built up from the invariant distributions of each of its component chains. These component chains have identical *T.P.*s, so the factors in m_0 are the same. Let us call this common factor distribution m_{10}. Once computed, m_{10} (concentrated on $[\delta, \infty)$) may be used to compute the limit, as $k \to \infty$, of the probability of finding exactly j particles in the interval $[0, \delta)$ at times $S_k = T_k + 1$. By now it should be clear that the limiting distribution of $Z(S_k)$ is a binomial distribution corresponding to N Bernoulli trials with parameter $p = m_{10}\{[\delta, 2\delta)\}$. (However, the limit distribution of $Z(t)$ for t tending to infinity without restriction is *not* a binomial.)

The invariant distribution, let us call it m_1 rather than m, in the case where $N = 1$ of our basic chain can be calculated explicitly and then m_{10} obtained from the special case of (22). The measure m_1 turns out to be uniform on $[0, \delta)$ and coincides with the (δ-translate) of the stationary distribution for the renewal process with interarrival distribution F. This stationary distribution has a *density* equal to the normalized tail-sum $1 - F$ (see Feller [1971, XI.4]). One can give queueing theory arguments for the above description of m_1, but since Equation (21) leads to this result almost immediately, we use that equation to give a quick proof. In the case $N = 1$, $A_0 = [\delta, \infty)$, and $B_1 = [0, \delta)$, Equation (21) simplifies to

$$[e^{\alpha\delta}-1]\int_{\delta}^{\infty} e^{-\alpha x}\, dm_1(x) = [1 - \phi(\alpha)]\int_0^{\delta} e^{-\alpha x}\, dm_1(x),$$

valid for any complex number α, $\mathrm{Re}(\alpha) \geq 0$. If we set $\alpha = -2n\pi i/\delta$, where n is an arbitrary integer, we find that the left-hand side vanishes. The density assumption implies that $\phi(\alpha) \neq 1$ for any $\alpha \neq 0$. Hence, $\int_0^{\delta} \exp(2n\pi i x/\delta)\, dm_1(x) = 0$ for every $n \neq 0$. Standard uniqueness results in Fourier series theory imply that we must have $dm_1(x) = C dx$, $0 \leq x \leq \delta$ for some constant C. From (5) in case $N = 1$, we find $C\delta = 1 - m_1\{[\delta, \infty)\} = \delta/(\mu + \delta)$, so $C = 1/(\mu + \delta)$. For the Laplace transform of m_1 on $[\delta, \infty)$, we get

$$\int_{\delta}^{\infty} e^{-\alpha x}\, dm_1(x) = \frac{1 - \phi(\alpha)}{e^{\alpha\delta}-1}\int_0^{\delta} e^{-\alpha x} C dx = e^{-\alpha\delta}\frac{1 - \phi(\alpha)}{\alpha(\mu + \delta)}.$$

Inverting the Laplace transforms in this equation reveals that the density g_1 of m_1 for $x \geq \delta$ is given by $g_1(x) = [1 - F(x - \delta)]/(\mu + \delta)$. From (22) it is then clear that m_{10} has the density $g_{10}(x) = (1 + \delta/\mu)g_1(x)$ for $x \geq \delta$. $\square$

The upshot of the preceding is that we can now conclude that the limit distribution of $Z(S_k)$ is

$$\lim_k P^v\{Z(S_k) = i\} = \lim_n \frac{\#\{k \leq n : Z(S_k) = i\}}{n} = \binom{N}{i} p^i(1 - p)^{N-i} \quad (23)$$

(w.p. 1) for $i = 0, 1, \ldots, N$, where

$$p = m_{10}\{[\delta, 2\delta)\} = \frac{1}{\mu}\int_0^{\delta}[1 - F(x)]dx = \frac{1}{\mu}\int_0^{\delta}\int_x^{\infty} f(z)dzdx.$$

Remark. It follows from the work of the last two sections that the measure m, *when restricted to* A_0, is a product measure because its restriction to A_0 coincides with $q_0 m_0$. However, the m-measures of subsets of A_0 have no particular interest; it is only the m-measures of the other A_j (defined in the proof of Lemma 3.1) that are required, and these sets are contained in the complement of A_0. But m restricted to the complement of A_0 is *not* a product measure.

B.3 Estimates for $\mathbf{q_j}$

Apart from the explicit representation of m on A_0 discussed in the last section, a simple expression for m on all of $[0, \infty)^N$ for $N > 1$ is not available. This means that, with one exception (the case $F(\delta) = 0$), we do not have *simple* explicit formulae for the values of $q_j = m\{A_j\}$ and must resort to approximations. It turns out, however, that the approximate formulas are quite amenable to analysis, particularly in the region of interest: $\delta = \mathcal{O}(1/N)$.

In this section we finish the proof of the two inequalities of (6), which we henceforth designate (LH-6) for the left side and (RH-6) for the right.

To simplify the notation a little, the starting distribution ν is omitted, if not forgotten, when it is not essential. For this proof we introduce the following objects:

$$n^{\circ}(t) = \max\{k : S_k \leq t,\}$$

$$\#(t, j) = \#\{s : s \leq t \text{ and } Z(s) = j\} = \sum_{s=0}^{t} I_{A_j}(\hat{\mathbf{X}}),$$

$$\#^{*}(n, j+) = \{k : k \leq n \text{ and } Z(S_k) \geq j\} = \sum_{k=1}^{n} I_{A_j^*}(\hat{\mathbf{X}}(S_k)),$$

where $A_j^* = \bigcup_{i=j}^{N} A_i$. The variable $\#(t, 0)$ differs from $n^{\circ}(t)$ by at most 1 because the T_ks are the zeros of Z. Therefore, by the ergodic limit theory for $\hat{\mathbf{X}}(t)$,

$$\lim_{t \to \infty} \frac{n \circ (t)}{t} = q_0. \tag{24}$$

Hence,

$$\lim_{n} \frac{S_n}{n} = \lim_{n} \frac{S_n}{n^{\circ}(S_n)} = \lim_{t} \frac{t}{n^{\circ}(t)} = \frac{1}{q_0},$$

and then,

$$\lim_{n} \frac{\#(S_n, j)}{n} = \left[\lim_{n} \frac{\#(S_n, j)}{S_n}\right]\left[\lim_{n} \frac{S_n}{n}\right] = \frac{q_j}{q_0}. \tag{25}$$

Fix $j > 0$ and define a sequence of random variables $\{V_k\}$ $(k \geq 1)$ by $V_k = 0$ if $Z(S_{k-1}) < j$, and otherwise V_k shall be the total number of times t in the interval $[S_{k-1}, S_k)$ that the counting variable $Z(t)$ has the value j:

$$V_k = \sum_{S_{k-1} \leq t < S_k} I_{A_j}(\hat{\mathbf{X}}).$$

Clearly,

$$\#(S_n, j) = \sum_{k=1}^{n} V_k,\ n \geq 1. \tag{26}$$

The occurrence of the event $V_k > r$ implies that $Z(S_{k-1}) \geq j$ and that there are at least r jumps, counting from the first time in $[S_{k-1}, S_k)$ that $Z(t) = j$, that had magnitudes smaller than δ. Hence,

$$\begin{aligned} P\{V_k > r\} &= P\{V_k > r \mid Z(S_{k-1}) \geq j\} P\{Z(S_{k-1}) \geq j\} \\ &\leq F(\delta)^r P\{Z(S_{k-1}) \geq j\},\ r = 0, 1, 2, \ldots, \end{aligned}$$

and therefore, assuming $F(\delta) < 1$,

$$E\{V_k\} = \sum_{r=0}^{\infty} P\{V_k > r\} \leq \frac{P\{Z(S_{k-1}) \geq j\}}{1 - F(\delta)} < \infty.$$

(Indeed, all moments $E\{(V_k)^\beta\}$, $\beta = 1, 2, \ldots,$ are finite.) This inequality, (23), and the ergodic limit theory now yield

$$\frac{q_j}{q_0} = \lim_n \frac{\#(S_n, j)}{n} = E\left[\lim_{n\to\infty} \frac{\sum_{k=1}^{n} V_k}{n}\right] = \lim_{k\to\infty} E\{V_k\}$$

$$\leq \lim_{k\to\infty} \frac{P\{Z(S_{k-1}) \geq j\}}{1 - F(\delta)} = \frac{\sum_{i=j}^{N} \binom{N}{i} p^i (1-p)^{N-i}}{1 - F(\delta)} = \frac{B(j)}{1 - F(\delta)} \tag{27}$$

which is (RH-6).

As to (LH-6), note first that

$$\#(S_n, j) = \sum_{k=1}^{n} V_k \geq \#^*(n, j+).$$

Hence, (see (25), and (23)),

$$\frac{q_j}{q_0} \geq \lim_n \frac{\#^*(n, j+)}{n} = \lim_n \frac{\#\{k \leq n : Z(S_k) \geq j\}}{n} = B(j)$$

which is (LH-6). □

C. PROOF OF THEOREM 4.1

For the purposes of this proof let us write $x = N\delta/\mu$,

$$D_N(\delta) = \frac{1}{2}\{(N^2 - N)p^2 + 2Np\},$$

$$D(x) = \frac{1}{2}(x^2 + 2x).$$

The conclusion of Theorem 4.1 is equivalent to the assertion that

$$K_N(\delta) - d - \frac{b}{x} - \frac{c}{x}D(x) = \mathcal{O}(1/N), \tag{28}$$

uniformly on bounded x-intervals. We base the proof on (8), which states

$$\frac{c}{x}D_N(\delta) \leq K_N(\delta) - d - \frac{b}{x} \leq \frac{c}{x}D_N(\delta)\{1 - F(\delta)\}^{-1} \tag{29}$$

in the new notation. Fix a number $\gamma > 0$ and confine x to the interval $(0, \gamma/\mu)$, so that $0 < \delta \leq \gamma/N$. Let ϵ_0 and c_0 be the numbers introduced in Assumption J3 in Section 3. Keeping $N > \max\{2c_0\gamma, \gamma/\epsilon_0\}$, which makes $F(\delta) \leq \min\{1/2, c_0\delta\}$, we get

$$\{1 - F(\delta)\}^{-1} \leq 1 + 2F(\delta) \leq 1 + 2c_0\delta.$$

Also, $D_N(\delta) \leq 1/2(N^2p^2 + 2Np) \leq 1/2(\gamma^2/\mu^2 + 2\gamma/\mu)$ because $p \leq \delta/\mu \leq \gamma/N\mu$. Hence,

$$\begin{aligned} 0 &\leq K_N(\delta) - d - \frac{b}{x} - \frac{c}{x}D_N(\delta) \\ &\leq \frac{2c_0c\delta}{x}D_N(\delta) = \frac{2c_0c\mu}{N}D_N(\delta) \\ &\leq \frac{C_1}{N}, \end{aligned}$$

where $C_1 = c_0c(\gamma^2/\mu + 2\gamma)$. From this and a little algebra, it is easily seen that to finish the proof of (33) it suffices to find a number C_2, depending on γ, such that

$$\frac{1}{x}|D_N(\delta) - D(x)| \le \frac{C_2}{N}, \text{ for } 0 < \delta \le \frac{\gamma}{N}. \tag{30}$$

But because $Np \le N\delta/\mu = x$, we have

$$\frac{1}{x}|D(x) - D_N(\delta)| \le \frac{1}{x}\int_{Np}^{x}(u+1)du + \frac{Np^2}{2x} \le \frac{x+1}{x}|x - Np| + \frac{\delta}{2\mu}$$

$$= \frac{N(x+1)}{\mu x}\int_0^{\delta} F(z)dz + \frac{\delta}{2\mu} \le \frac{c_0 N(x+1)\delta^2}{2\mu x} + \frac{\delta}{2\mu}$$

$$\le \frac{C_2}{N}, C_2 = \frac{\gamma\{c_0(\mu+\gamma)+1\}}{2\mu}. \qquad \square$$

The numbers C_1 (multiplied by c) and C_2 yield estimates for τ_1 and τ_2 mentioned in Theorem 4.1: $\tau_2 = 3/2cc_0/\mu$ and $\tau_1 = 1/2c(5c_0\mu+1)/\mu$.

D. PROOF OF THEOREM 4.2

Throughout this section we write $a = b/c$,

$$\gamma_0 = \mu\sqrt{2b/c} = \mu\sqrt{2a}, \ \delta_{o,N} = \gamma_0/N$$

$$K_o = c + \sqrt{2bc}.$$

Moreover, there is no harm in also supposing that $d = 0$.

Step 1. As a function of δ, $K_N(\delta)$ is continuous on $(0, \infty)$. The reader is asked to turn to formula (3). To begin with, the variable q_0 is $\mu/(\mu + N\delta)$, which is obviously continuous. The only possible discontinuous term in the formula for K_N is $E_N(\delta)$. However, the continuity of this function is an immediate consequence of the following exact formula, which is discussed after the proof of the theorem is complete.

LEMMA D.1.

$$E_N(\delta) = \frac{N(N-1)}{\mu(\mu+N\delta)}\int_0^{\delta}(\delta - t)[1 - F(t)]dt + \frac{N\delta}{\mu + N\delta}.$$

Step 2. The function $\delta \to E_N(\delta)$ is nondecreasing. One can prove this by differentiating the expression for $E(\delta)$ of Lemma 8 and checking that the result is nonnegative.

Here is an outline of an alternative, but more intuitive, proof: Consider two chains $\hat{\mathbf{X}}^1$ and $\hat{\mathbf{X}}^2$ with the same N and jump density, but with different bucket sizes $\delta_1 < \delta_2$. If we follow the trajectory of an individual particle in each chain, we find that in the chain with the larger bucket size, δ_2, the particle, on average, gets back to the interval $[0, \delta_2)$ quicker than it would to the interval $[0, \delta_1)$ in the chain with the smaller bucket size δ_1. Since a typical particle of the chain $\hat{\mathbf{X}}^2$ is more often in the interval $[0, \delta_2)$ than in $[0, \delta_1)$ of the chain $\hat{\mathbf{X}}^1$, the average number of particles in $[0, \delta_2]$ of $\hat{\mathbf{X}}^2$ is at least as large as the average number of particles in $[0, \delta_1)$ of $\hat{\mathbf{X}}^1$. That is, $E_N(\delta_1) \leq E_N(\delta_2)$.

Step 3. Next, we establish

$$\lim_{N\to\infty} K_N(\delta) = \infty \text{ uniformly on } [\delta_1, \infty)$$

for each fixed $\delta_1 > 0$. The explicit formula for $q_0(\delta) = \mu/(\mu + N\delta)$ yields that $1 - q_0(\delta)$ is also a nondecreasing function of δ. By Equation (3) and Step 2, we find that the function $(1 - q_0(\delta))K_N(\delta) - q_0(\delta)b$ is also nondecreasing in δ. Hence, for $\delta_1 \leq \delta_2$,

$$K_N(\delta_1) \leq \frac{1 - q_0(\delta_2)}{1 - q_0(\delta_1)} K_N(\delta_2) + \frac{q_0(\delta_1) - q_0(\delta_2)}{1 - q_0(\delta_1)} b.$$

For $N \geq \mu/\delta_1$ and $\delta_1 \leq \delta_2$ we have $1/2 \leq 1 - q_0(\delta_1) \leq 1 - q_0(\delta_2) \leq 1$. From the above inequality, for $\delta_1 \leq \delta_2$ and $N \geq \mu/\delta_1$, we have

$$K_N(\delta_1) \leq 2K_N(\delta_2) + 2b.$$

From this inequality, it follows that $K_N(\delta)$ goes to infinity uniformly in the interval $[\delta_1, \infty)$, provided that $\lim K_N(\delta_1) = \infty$ for each fixed $\delta_1 > 0$. This follows from Equation (8).

Step 4. For $\gamma > \gamma_0 = \mu\sqrt{2a}$, we have $\min_{0<\delta<\gamma/N} K_N(\delta) = K_o + \mathcal{O}(N^{-1})$. Moreover, if δ_1 is the minimizing δ on this interval, then

$$\delta_1 = \delta_{o,N} + \mathcal{O}(N^{-3/2}),$$

where $\delta_{o,N}$ is defined as above. To prove all this, let

$$H_N(\delta) = d + c + \frac{cN}{2\mu}\delta + \frac{b\mu}{N}\frac{1}{\delta}$$

By straightforward calculus, for each N, $H_N(\delta)$ has a global minimum on $(0, \infty)$ at point $\delta_{o,N}$. Let δ_{opt} be a value of δ, which gives the minimum value of $K_N(\delta)$ on the given interval $(0, \gamma/N)$. By Theorem 4.1, on this interval we can find a constant C, depending on γ, such that for all N sufficiently large,

$$H_N(\delta) - C/N < K_N(\delta) < H_N(\delta) + C/N$$

uniformly for $0 < \delta < \gamma/N$. Since $\gamma > \gamma_0$, δ_1 is in the interval $(0, \gamma/N)$. On this interval, $K_N(\delta)$ is sandwiched between the two convex functions, $H_N(\delta) \pm C/N$, both of which have a global minimum at the same point $\delta_{o,N}$ interior to the interval. For a fixed N, δ_{opt} must lie between the two solutions to the equation (in δ) $H_N(\delta) - C/N = H_N(\delta_{o,N}) + C/N$. By a simple calculation, we find that the difference between the two solutions is $\mathcal{O}(N^{-3/2})$, and this yields $\delta_{opt} = \delta_{o,N} + \mathcal{O}(N^{-3/2})$. For any δ between the two solutions, we find that $K_N(\delta) = H_N(\delta) + \mathcal{O}(N^{-1})$.

Step 5. The next step is to show that for any $\gamma > \gamma_0$ there is $\delta_1 > 0$ and N_1 such that for all $N \geq N_1$,

$$\inf_{0<\delta<\delta_1} K_N(\delta) = \inf_{0<\delta<\gamma/N} K_N(\delta). \tag{31}$$

Thus, the minimum exhibited in Step 4 extends to the fixed interval $(0, \delta_1]$. This fact and Step 3 imply that for all N sufficiently large,

$$\min_{0<\delta<\infty} K_N(\delta) = K_o + \mathcal{O}(1/N)$$

completing the proof of the theorem.

For the moment, we fix δ_1 such that $0 < F(\delta_1) < 1$. We choose δ_1 later. By inequality (8) and the fact that $p(\delta) \geq \delta[1 - F(\delta)]/\mu$, we have that for all $\delta < \delta_1$,

$$K_N(\delta) \geq c\,Q\left(a/z + 1 + \frac{1}{2}(1 - 1/N)z\right),$$

where $Q = 1 - F(\delta_1)$ and $z = QN\delta/\mu$. Define

$$\mathcal{L}_N(z) = c\,Q\left(a/z + 1 + \frac{1}{2}(1 - 1/N)z\right).$$

For each N, the horizontal line at height $K_o(1 + C_0/N)$ cuts the graph of the convex function $\mathcal{L}_N$ at two points, the larger of which we call z_N^*. Thus, z_N^* is the larger root of the equation

$$Q(N-1)z^2 - 2\left[(N + C_0)(1 + \sqrt{2a}) - NQ\right]z + 2NQa = 0.$$

As a sequence of N, the values z_N^* converge to a bounded positive limit. Define $\delta_N^* = \mu z_N^* / QN$. Hence, the sequence $N\delta_N^*$ also converges to a limit γ^*. Note that γ^* tends to $\sqrt{2a\mu}$ as Q approaches 1. We choose δ_1 (and hence Q), so that $\gamma^* < \gamma$. Now choose N_1 such that $\delta_N^* < \gamma/N < \delta_1$ for all $N \geq N_1$. Since the minimum of $K_N(\delta)$ is bounded above by $K_o(1 + c_0/N)$ and bounded below by the function $\mathcal{L}_N(QN/(\mu\delta)$, the minimum of $K_N(\delta)$ in the interval $(0,\ \delta_1]$ must already lie in the interval $(0,\ \delta_N^*]$, and hence in the interval $(0,\ \gamma/N]$. □

D.1 A Discussion of the Exact Formula for $E_N(\delta)$

The proof of this result is quite long and is based on some exact, although very complicated, integral formulas for the q_js (see Erickson [1999] for details). The exact formula for $E_N(\delta)$ leads to an exact formula for $K_N(\delta)$, but our work has led us to the conclusion that the excellent asymptotic formulas of Theorems 4.1 and 4.2 (and the simple inequalities of Lemma 4.2 that lead to them) are of much greater practical use, and are certainly easier to prove. For this reason we do not include the long proof of the exact formula.

Our main use of Lemma D.1 was to slightly shorten the proof of the global minimization of K_N. Note that D.1 immediately yields the continuity of K_N as a function of δ. We require continuity in order to speak sensibly of the existence of a minimizing δ. Even without the continuity, however, the basic result of Theorem 4.2 is essentially correct; only the language used to express it needs to be changed. (We must use the term "greatest lower bound" in place of "minimum," and we can only assert that there are points δ at which the greatest lower bound is approximately attained.)

E. PROOF OF LEMMA 5.2

We write L_N^M for $L_N^M(\delta)$, and $Z^A(t)$ for the number of particles in set A at time t. The set Γ (see 5.1) is a subset of $[\delta,\ \infty)$, so

$$Z^{\Gamma+\delta}(T_{j-1}) = Z^{\Gamma}(S_{j-1}) \le Z^{\Gamma}(s) \le Z^{\Gamma}(T_j)$$

for $S_{j-1} = T_{j-1} + 1 \le s \le T_j$. (Recall that T_k, $k \ge 1$ are the successive times at which the interval $[0,\ \delta)$ is empty of particles.) Hence,

$$L_N^M = \lim_t\ (1/t)\sum_{s\le t} Z^{\Gamma}(s) = \lim_k\ (k/T_k)\lim_k\ (1/k)\sum_{j=1}^{k}\left[\sum_{s=S_{j-1}}^{T_j} Z^{\Gamma}(s)\right]$$

$$= q_0 \lim_k E\left[\sum_{s=S_{k-1}}^{T_k} Z^{\Gamma}(s)\right] \le q_0 \limsup_k\ E[Z^{\Gamma}(T_k)\Delta_k], \tag{32}$$

w.p.1, where $\Delta_k = T_k - T_{k-1}$. Suppose that at time T_{k-1} there are $r = Z^{[\delta,\ 2\delta)}(T_{k-1})$ particles at positions $x_1 + \delta, \ldots, x_r + \delta$ in $[\delta,\ 2\delta)$. Then, at S_{k-1} there will be r particles in $[0,\ \delta)$ at positions $x_1, \ldots, x_r$. So

$$Z^{\Gamma}(T_k) = Z^{\Gamma+\delta}(T_{k-1}) + u_1 + \cdots + u_r,$$

where $u_i = 1$ or 0, according to which of the ith of these particles lands in Γ or not, when it is finally removed from the interval $[0, \delta)$. (This removal must occur during $[S_{k-1}, T_k]$.) Writing $u = u_i$, $x = x_i$, then the strong Markov property gives

$$P\{u = 1 | \mathcal{F}_{T_{k-1}}\} = \sum_{j=1}^{\infty} [H_{\delta-x}(jM\delta) - H_{\delta-x}(jM\delta - \delta)],$$

where $\mathcal{F}_{T_{k-1}}$ is the σ-field of the random variables T_{k-1}, $X(s)$, $s = 0, 1, \ldots, T_{k-1}$ and where $H_t(b)$ is the probability that a particle starting at the origin lands in the interval $[t, t + b]$ when it first jumps over t. H satisfies

$$H_t(b) = \int_{0-}^{t} [F(t + b - y) - F(t - y)]U\{dy\},$$

where U is the renewal measure (see Feller [1971, p. 369]).[4]

In general, $U(z) \equiv U\{[0, z]\} \le [1 - F(z)]^{-1}$ for distributions on $[0, \infty)$, so that

$$\begin{aligned} & P\{u = 1 | \mathcal{F}_{T_{k-1}}\} \\ & \le \sup_{x \in [0, \delta)} \int_{0-}^{\delta - x} \sum_{j=1}^{\infty} [F(jM\delta + \delta - x - z) \\ & - F(jM\delta - x - z)]U\{dz\} \\ & \le \Pi_2 / [1 - F(\delta)]. \end{aligned}$$

(Recall the definition of Π_2 in Theorem 5.2.) Calling the right-hand side p^* and noting that, conditional on the σ-field $\mathcal{F}_{T_{k-1}}$, the variables u_i are independent of Δ_k, we have

$$\begin{aligned} E\{Z_k^{\Gamma} \Delta_k | \mathcal{F}_{T_{k-1}}\} & \le [Z_{k-1}^{\Gamma+\delta} + rp^*] E\{\Delta_k | \mathcal{F}_{T_{k-1}}\} \\ & \le [Z_{k-1}^{\Gamma+\delta} + rp^*] \left(\frac{r}{1 - F(\delta)} + 1 \right) \\ & \le [1 - F(\delta)]^{-1} (Z_{k-1}^{\Gamma+\delta} + rp^*)(r + 1), \end{aligned}$$

[4]Feller defines H in terms of the *open* interval (δ, ∞), whereas we use the closed interval. Because F has no atoms, this difference in definition has no consequence.

where $r = Z_{k-1}^{[\delta,\,2\delta)}$, and $Z_j^{\{\cdot\}}$ means $Z^{\{\cdot\}}(T_j)$, $j = k - 1, k$. Now, at times $\{T_j\}$, the particles are independent, so the limiting joint distribution of $Z_{k-1}^{\Gamma+\delta}$ and $Z_{k-1}^{[\delta,\,2\delta)}$ is a trinomial. (Separately, they have binomial limit distributions.) Letting $k \to \infty$,we get $E\{Z_{k-1}^{[\delta,\,2\delta)}\} \to Np$, $E\{Z_{k-1}^{\Gamma+\delta}\} \to N\Pi_1$, $\mathrm{Var}(Z_{k-1}^{[\delta,\,2\delta)}) \to Np(1 - p)$ and

$$E\{Z_{k-1}^{[\delta,\,2\delta)} \cdot Z_{k-1}^{\Gamma+\delta}\} \to N(N-1)p\Pi_1,$$

where p is defined in (4) and $\Pi_1 = G_0(\Gamma + \delta)$. Going back to (32) with these calculations, we obtain

$$\begin{aligned} L_N^M &\le q_0[1 - F(\delta)]^{-1} \lim_k E\{(Z_{k-1}^{\Gamma+\delta} + p^* Z_{k-1}^{[\delta,\,2\delta)})(Z_{k-1}^{[\delta,\,2\delta)} + 1)\} \\ &= \frac{Nq_0\{(Np + 1)\Pi_1 + pp^*[Np + 2 - p]\}}{1 - F(\delta)}. \end{aligned}$$

(Recall the basic property of conditional expectations $E\{E[\cdot \mid \mathcal{F}_{T_{k-1}}]\} = E\{\cdot\}$.) Replacing p^* with $\Pi_2/[1 - F(\delta)]$ and q_0 with $\mu/(\mu + N\delta)$, combining fractions and dropping the factor $1 - F(\delta)$, which will occur in the numerator, we finally obtain the upper bound on L_N^M.

For the lower bound, we have

$$\begin{aligned} L_N^M &\ge q_0 \liminf_k \; E[Z_{k-1}^{\Gamma+\delta}\Delta_k] \ge q_0 \lim_k E\{Z_{k-1}^{\Gamma+\delta}(Z_{k-1}^{[\delta,\,2\delta)}) + 1\}, \\ &= q_0\{N(N-1)p\Pi_1 + N\Pi_1\}, \end{aligned}$$

which evaluates to the lower bound on L_N^M. □

F. PROOF OF LEMMA 5.3

In the following we let lim, lim sup, and lim inf stand for the limits of various quantities as $N \to \infty$, with the other variables constrained to vary as stated in the hypothesis.

First note that $\lim F(\delta) = 0$, $\lim Np = \lim N\delta/\mu = x$, and $\lim M\delta = rx\mu$, $\lim N\delta = x\mu$.

Let $t_j = jM\delta = jrx\mu$. Note that t_j does not vary with N. For all sufficiently large N, δ will be so small that the intervals $(t_j - \delta, t_j + \delta]$ will not overlap. Put $J_N = \cup_j (t_j - \delta, t_j + \delta]$, then, for all N, we have $J_{N+1} \subset J_N$. Moreover, $\cap_N J_N$ is at most a countable discrete set. Hence,

$$\Pi_2 \le \sum_{j=1}^{\infty} \{F(jM\delta + \delta) - F(jM\delta - \delta)\} = F\{J_N\} \to 0, \text{ as } N \to \infty$$

by continuity of F (no atoms). (The letter F stands for both the distribution function and the induced probability measure, as is customary.) From this (and the limits $Np \to x$), we get $\lim Np(Np + 2 - p)\Pi_2 = 0$.

Next,

$$\frac{\delta}{\mu}\sum_{j=1}^{\infty}[1 - F(j\mu rx + \delta)] \leq \Pi_1 \leq \frac{\delta}{\mu}\sum_{j=1}^{\infty}[1 - F(j\mu rx)].$$

As $\delta \to 0$, the first sum on the left goes to $\sum_j[1 - F(jrx\mu)]$. Because $N\delta/\mu \to x$, it follows that

$$\lim N\Pi_1 = x\sum_{j=1}^{\infty}[1 - F(jrx\mu)].$$

Using these limits in the upper bound for $L_N^M(\delta)$, we get

$$\limsup L_N^M(\delta) \leq \lim \frac{\mu\{N(Np + 1)\Pi_1 + Np(Np + 2 - p)\Pi_2\}}{(\mu + N\delta)[1 - F(\delta)]^2}$$

$$= \lim \frac{N\mu(Np + 1)\Pi_1}{\mu + N\delta} = \frac{\mu(x + 1)}{\mu + x\mu}\lim N\Pi_1 = x\sum_{j=1}^{\infty}[1 - F(jrx\mu)].$$

Similarly, from the lower bound for $L_N^M(\delta)$:

$$\liminf L_N^M(\delta) \geq \lim \frac{N\mu(Np + 1 - p)\Pi_1}{\mu + N\delta} = \frac{\mu(x + 1)}{\mu + x\mu}\lim N\Pi_1$$

$$= x\sum_{j=1}^{\infty}[1 - F(jrx\mu)].$$

Thus the $\limsup L_N^M(\delta) = \liminf L_N^M(\delta) = x\sum[1 - F(jrx\mu)]$, so the limit of $L_N^M(\delta)$ exists and its value is as stated.

G. PROOF OF THEOREM A.1

We have seen that for each $k \geq 1$ and, given the initial positions, the N components of $\hat{\mathbf{X}}(S_k)$ are mutually independent random variables. It is also not hard to see that the conditional distribution of $X^i(S_k)$, given $X^i(S_{k-1}) = x$, equals the distribution of the *residual waiting time* at epoch δ of a delayed renewal process, starting at epoch x with interarrival distribution F; see Feller [1971, p. 369] and Erickson [1999]. (We have already seen this distribution in the proof of Lemma 5.2, Section E, although it was described in slightly different language.)

Letting $H_s\{I\}$ denote the probability that the residual waiting time at epoch s lies in I for a pure renewal process starting at 0 ($H_s\{[0, b)\} = H_s(b)$ in Section E), it follows that for any fixed $x > 0$, every integer $k \gg x/\delta$, and any Borel $I \subset [0, \beta]$,

$$P\{X^i(S_k) \in I \mid X^i(0) = x\} = H_{k\delta - x}\{I\}. \tag{33}$$

Using the Markov property, it then follows that for fixed $\hat{\mathbf{x}} = (x^1, \ldots, x^N) \in \Lambda$ and Borel sets $I_i \subset [0, \beta)$,

$$P\{\hat{\mathbf{X}}(S_k) \in I_1 \times I_2 \times \cdots \times I^N \mid \hat{\mathbf{X}}(S_0) = \hat{x}\} = \prod_{i=1}^{N} H_{k\delta - x^i}\{I^i\}. \tag{34}$$

If U denotes the renewal measure, then assumption (a) implies that U has an absolutely continuous part that possesses a strictly positive density on $(0, \infty)$; but see Feller [1971, p. 369],

$$H_s\{I\} = \int_{0-}^{s} F\{I + s - z\} U\{dz\}. \tag{35}$$

Consequently, the measure $I \to H_{k\delta - x}\{I\}$ also has an absolutely continuous part that is strictly positive on $[0, \beta)$. The conclusion we may draw from the preceding is that $\{\hat{\mathbf{X}}(S_k)\}$ is *irreducible* with respect to the measure ℓ_N, the Lebesgue measure in R^N (restricted to $(0, \beta)^N$); see Revuz [1984, Chap. 3 §2]. This implies that the trace chain $\{\hat{\mathbf{X}}(T_k)\}$ is also ℓ_N irreducible on its state space A_0. Together, these two assertions imply that the full chain $\hat{\mathbf{X}} = \hat{\mathbf{X}}(t); t \geq 0$ is also ℓ_N-irreducible. But we can also draw additional useful conclusions from (34) and (35).

From Stone's decomposition theorem Revuz [1984, chap. 5, §5], we can write $U = U_1 + U_2$ where U_2 is a *finite* measure and U_1 is absolutely continuous with a bounded continuous density u such that $\lim_{x \to \infty} u(x) = 1/\mu$. For any Borel $I \subset [0, \infty)$ and $s > z$, $F\{I + s - z\} \leq 1 - F(s - z) \to 0$ as $s \to \infty$. Hence, by dominated convergence,

$$\lim_{k \to \infty} H_{k\delta - x}\{I\} = \lim_{k \to \infty} \int_{0-}^{k\delta - x} F\{I + k\delta - x - z\} U_2\{dz\} +$$

$$+ \lim_{k \to \infty} \int_{0}^{k\delta - x} F\{I + y\} u(k\delta - x - y)\, dy$$

$$= \frac{1}{\mu}\int_0^\infty F\{I + y\}dy = \frac{1}{\mu}\int_I [1 - F(x)]dx.$$

This and the product formula (34) yield that, for any Borel set $A \subset [0, \infty)^N$,

$$\lim_{k\to\infty} P^{\hat{\mathbf{x}}}\{\hat{\mathbf{X}}(S_k) \in A\} = m_0^*(A), \tag{36}$$

where m_0^* is the product measure $F_0 \times F_0 \times \cdots \times F_0$ and F_0 is the probability distribution on $[0, \beta)$ with density $\{1 - F(x)\}/\mu$. Not only does (36) give us one of the limit theorems we used earlier, but it implies that the subchains $\hat{\mathbf{X}}(S_k)$ and $\hat{\mathbf{X}}(T_k)$ are both Harris recurrent (with invariant probabilities $m_0{}^*$ and m_0 (translate by $\delta\mathbf{1}$ of m_0^*), respectively. If $m_0^*(A) > 0$, then (34) implies that, with probability 1, $\hat{\mathbf{X}}(S_k)$ enters A infinitely many times k, whatever the initial position. But $\hat{\mathbf{X}}(T_k)$ is obtained from $\hat{\mathbf{X}}(S_k)$ by adding δ to each component, so the previous assertion is also correct for $\hat{\mathbf{X}}(T_k)$ with respect to its invariant probability (see Revuz [1984, Chap. 2, §3]).

Now consider the full chain $\hat{\mathbf{X}} = \{\hat{X}(t),\ t = 1, 2, \ldots\}$. If $\ell_N(A) > 0$, $A \subset \Lambda$, then the preceding makes it clear that $\hat{\mathbf{X}}$ will hit A with positive probability. This implies that $\hat{\mathbf{X}}$ is ℓ_N-irreducible. According to Revuz [1984, Chap. 2, Theorems 2.3, 2.5, and Definition 2.6], either $\hat{\mathbf{X}}$ is a Harris chain with a (unique up to constant multiples) invariant measure m, or else the potential kernel is proper. The potential kernel K is defined by $K(\hat{\mathbf{x}}, A) = \sum_{t\geq 0} P^{\hat{\mathbf{x}}}\{\hat{\mathbf{X}}_t \in A\}$. If K is proper, then Λ can be written as an increasing sequence of subsets D_n, each of which has bounded potential. But eventually any such sets must have positive Lebesgue measure, and in that case (34) implies $P^{\hat{\mathbf{x}}}\{\hat{\mathbf{X}}(t) \in D_n \text{ for infinitely many } t\} = 1$ for all initial positions $\hat{\mathbf{x}}$. But $K(\hat{x}, D_n) < \infty$ implies that the *expected total number* of hits in D_n is finite, which implies that the number of hits must be finite with probability 1. We thus cannot have a proper potential kernel, and therefore $\hat{\mathbf{X}}$ must be a Harris chain.

Consider next the aperiodicity issue. Seeking a contradiction, let us suppose that $\hat{\mathbf{X}}$ is periodic. Let $\{\mathcal{C}_i\}_{i=1}^d$ be recurrent cyclic classes in the decomposition of the state space. These subsets have a positive Lebesgue measure. Without loss of generality, we may suppose that $m\{\mathcal{C}_1 \cap A_N\} > 0$, $A_N = [0, \delta)^N$. Define σ to be the smallest S_k such that $\hat{\mathbf{X}}(S_k) \in A_N \cap \mathcal{C}_1$. This stopping time is finite on account of (36).

Let ξ_s^j be a doubly-indexed sequence of independent random variables each with distribution F. The earliest possible epoch after σ at which $\hat{\mathbf{X}}$ can arrive in $A_0 = (\delta, \infty)^N$ exceeds $N + \sigma - 1$ because no more than one particle can move at any particular step. For any integer $r \geq 1$ and any Borel rectangle $A = I^1 \times I^2 \cdots \times I^N \subset A_0$, we have

$$P\{\hat{\mathbf{X}}(\sigma + N + r) \in A \,|\, \hat{\mathbf{X}}(\sigma) = (x^1, \ldots, x^N)\}$$

$$\geq \left(\int_0^{\delta - x^1} P\{\delta - x^1 - z \leq \xi^1_{r+1} \in I^1 - x^1\} F^{r^*}\{dz\} \right) \times$$

$$\times \left(\prod_{j=2}^{N} P\{\delta - x^j \leq \xi^j_1 \in I^j - x^j\} \right) \tag{37}$$

Due to hypotheses J1, the distribution F and each of its convolutions F^{r^*} puts positive mass on every subinterval in $(0, \beta)$. Hence the right-hand side of (37) is strictly positive whenever the cylinder set A has a positive Lebesgue measure. Standard measure theory implies that this is also correct for any Borel set $A \subset [\delta, \beta + \delta)$ with a positive measure.

With probability 1 at time σ, $\hat{\mathbf{X}}(\sigma)$ belongs to $\mathcal{C}_1 \cap A_N \subset \mathcal{C}_1$. So if the chain is periodic, then w.p.1 at all future epochs of the form $t = \sigma + nd$ the chain will always be found in $\mathcal{C}_1$. Also, for each $k = 1, 2, \ldots, d - 1$, at times $t - \sigma \equiv k(d) = k \bmod d$, the chain must belong to the set $\mathcal{C}_{1+k(d)}$, which is disjoint from the other classes including $\mathcal{C}_1$. But in (37), for *any* $N + r \geq N$, the right-hand side is strictly positive. By choosing A in (37) to be any set in $\mathcal{C}_k \cap [\delta, \beta + \delta)^N$ with a positive measure and letting r take on different (mod d) values, we get a contradiction to the previous assertion about belonging to disjoint sets. There is no contradiction if $d = 1$, that is, if $\hat{\mathbf{X}}$ is aperiodic.

It remains to show that the invariant measures m (unique up to constant multiples) of the full chain $\hat{\mathbf{X}}$ are *finite*: $m(\Lambda) < \infty$.

The trace chain $\{\hat{\mathbf{X}}(T_k)\}$ is positive recurrent with invariant probability m_0. But m_0 is also a multiple of m restricted to A_0. Hence, $m(A_0) < \infty$. It follows from the Renewal theorem that the mean return time to A_0 is also finite. Let $\phi \geq 0$ be a bounded measurable function *on* A_0. Then ϕ is m_0-(and hence m-) summable, and

$$\lim_{n\to\infty} \frac{\sum_{k=1}^{n} \phi(\hat{\mathbf{X}}(T_k))}{n} = m_0(\phi) \equiv \int_{A_0} \phi(\hat{\mathbf{x}}) dm_0(\hat{\mathbf{x}}) > 0.$$

Let $n^\circ(t)$ denote the number of visits to A_0 by the full chain during $0 \leq s \leq t$. Then, $n^\circ(t) \to \infty$ and

$$\lim_t \frac{\sum_{s=1}^{t} \phi(\hat{\mathbf{X}}_s)}{t} = \left[\lim_t \frac{n^\circ(t)}{t}\right]\left[\lim_n \frac{\sum_{k=1}^{n} \phi\{\hat{\mathbf{X}}(T_k)\}}{n}\right]$$
$$= \{E^{m_0}(T_1)\}^{-1} m_0(\phi) > 0. \tag{38}$$

But if the chain $\hat{\mathbf{X}}$ is *null*, that is, if m has infinite mass, then for any bounded m-summable function ϕ [Revuz 1984, Theorem 2.6, p. 198],

$$\lim_t E^{\hat{x}} \left\{ \left(\frac{1}{t}\right) \sum_{s=1}^{t} \phi(\hat{\mathbf{X}}) \right\} = 0.$$

By Fatou's lemma, one can readily see that this contradicts (38). □

ACKNOWLEDGMENTS

We would like to thank the two referees and the editor for their many valuable suggestions for improving this article.

REFERENCES

BROWN, M. 1978. Implementation and analysis of binomial queue algorithms. *SIAM J. Comput. 7*, 298–319.

BROWN, R. 1988. Calendar queues: a fast 0(1) priority queue implementation for the simulation event set problem. *Commun. ACM 31*, 10 (Oct.), 1220–1227.

ERICKSON, K. 1999. Calendar queue expectations. *Stochastic Models 15*, 617–638.

FELLER, W. 1971. *An Introduction to Probability Theory and its Applications*. John Wiley and Sons, Inc., New York, NY.

FRANCON, J., VIENNOT, G., AND VUILLEMIN, J. 1978. Description and analysis of an efficient priority queue representation. In *Proceedings of the 19th Annual Symposium on Foundations of Computer Science*. 1–7.

KNUTH, D. E. 1973. *The Art of Computer Programming*. Addison-Wesley Longman Publ. Co., Inc., Reading, MA.

REVUZ, D. 1984. *Markov Chains*. Elsevier North-Holland, Inc., New York, NY.

REXFORD, J., BONOMI, F., GREENBERG, A., AND WONG, A. 1997. Scalable architectures for integrated traffic shaping and link scheduling in high-speed atm switches. *IEEE J. Sel. Areas Commun. 7*, 938–950.

SLEATOR, D. D. AND TARJAN, R. E. 1985. Self-adjusting binary search trees. *J. ACM 32*, 3 (July), 652–686.

SLEATOR, D. D. AND TARJAN, R. E. 1986. Self adjusting heaps. *SIAM J. Comput. 15*, 1 (Feb.), 52–69.

VUILLEMIN, J. 1978. A data structure for manipulating priority queues. *Commun. ACM 21*, 4 (Apr.), 309–315.

Received: March 1998; revised: August 1999; accepted: December 1999