Optimizing Static Calendar Queues*

K. Bruce Erickson Richard E. Ladner Anthony LaMarca
University of Washington
Seattle, WA 98195

September 13, 1994

Abstract

The calendar queue is an important implementation of a priority queue which is particularly useful in discrete event simulators. In this paper we present an analysis of the static calendar queue which maintains N active events. A step of the discrete event simulator removes and processes the event with the smallest associated time and inserts a new event whose associated time is the time of the removed event plus a random increment with mean μ. We demonstrate that for the infinite bucket calendar queue the optimal bucket width is approximately

$$\delta_{opt} = \sqrt{\frac{2b}{c}}\frac{\mu}{N}$$

where b is the time to process an empty bucket and c the incremental time to process a list element. With bucket width chosen to be δ_{opt}, the expected time to process an event is approximately minimized at the constant $c + \sqrt{2bc} + d$, where d is the fixed time to process an event. We show that choosing the number of buckets to be $O(N)$ yields a calendar queue with performance equal to or almost equal to the performance of the infinite bucket calendar queue.

1 Introduction

The *calendar queue* data structure, as described by Brown [2], is an important implementation of a priority queue which is useful as the event 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 intended to occur in simulated time. The set of events are stored in the priority queue ordered by their associated event times. A basic simulation step consists of finding an event e_0 which 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. Each new event e has an event time $t(e) > t(e_0)$ and must be inserted in the priority queue accordingly. The quantity $t(e) - t(e_0)$ can often be modeled as a non-negative random variable defined by some distribution, such as an exponential or uniform 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 execution of a segment of code by one of the processors. Thus, if there are N processors, then there are exactly N active events in the priority queue.

In many situations the calendar queue significantly outperforms traditional priority queue data structures. Brown [2] has given empirical evidence that the calendar queue, with its parameters properly set, achieves expected constant time per event processed.

The main contribution of this paper is to prove that, under reasonable assumptions, the optimally performing calendar queue data structure has constant (i.e., independent of N) expected time per event processed. In addition, simple formulas are derived for setting the parameters of the calendar queue to achieve optimal or near optimal performance.

1.1 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

*This paper appears in the Proceedings of the 35th Annual IEEE Symposium on Foundations of Computer Science, Nov. 20-22, 1994. This paper is University of Washington, Dept. of Computer Science and Engineering Technical Report No. 94-09-02. Erickson's address is Department of Mathematics. Ladner's research was supported by NSF, CCR-9108314. LaMarca's research was supported by an A.T.&T. Fellowship.

if $i \leq t(e)/\delta \bmod M < (i+1)$. The analogy with a calendar can be stated by: 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 $t_0 = 30$. The 8 events have times 31, 54, 85, 98, 111, 128, 138, 251.

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

In this example the next event to process 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 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 which 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, 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*.

1.2 Results

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 parameters in the implementation of a calendar queue which effect 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, the bucket discipline, and the process by which $t(e)$ is chosen for a newly generated event e. Assuming M is very large, 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 too many empty buckets to traverse before reaching a non-empty bucket, which again is inefficient. Assuming δ is chosen well for very large M, if M is chosen too small the current bucket will again tend to have too many events in it, particularly events which are not to be processed until later visits to the same bucket.

The problem we address is first, how to choose δ optimally when M is infinite. Although the case when M is infinite is not directly implementable on a computer, it has bearing on the case where M is finite. The second problem we address is how to choose M large enough so that the calendar queue behaves almost as if M were infinite.

We begin by focusing on the infinite bucket calendar queue with the unordered list implementation as the bucket discipline. In this discipline the time to process an event can be divided into three parts. The first part d is the fixed cost to process an event, which includes the time to process the event and insert a new event into a bucket. The final two parts are the variable costs to process a event. If m empty buckets are visited before reaching a bucket with n events ($n \geq 1$) then the variable cost is $bm + cn$. Thus, b is the incremental time to process an empty bucket and c is the incremental time to traverse a member of a list in search of the minimum in the list. The time to process an event is defined to be:

$$bm + cn + d.$$

If we let $E(\delta)$ be the expected value of $bm + cn + d$, then our goal are first, to find the value of δ which minimizes $E(\delta)$ when M is infinite, then, second, to select a good finite M which does not compromise the minimum very much.

As mentioned earlier, if event e_0 generates event e then $t(e) - t(e_0)$ may be modeled as a random variable chosen according to some distribution. We assume that we are given a probability density f defined on $[0, \infty)$ and that $t(e) - t(e_0)$ is the random variable defined by f. That is, $f(x)$ is non-negative for $x \geq 0$, $\int_0^\infty f(x)dx = 1$, and the probability that $t(e) - t(e_0) \leq t$ is exactly $\int_0^t f(x)dx$. We call f the *jump density* and its random variable simply the *jump*. Let μ be the mean of the jump, that is:

$$\mu = \int_0^\infty x f(x) dx.$$

For example, the exponential jump with mean μ has jump density $f(x) = \frac{1}{\mu} e^{-x/\mu}$ and the uniform jump

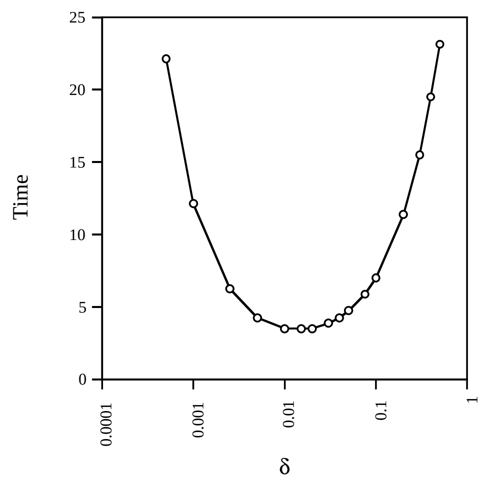

Figure 1: Graph of δ vs. expected time per event ($E(\delta)$) for δ near δ_{opt} in the infinite bucket calendar queue with 100 events.

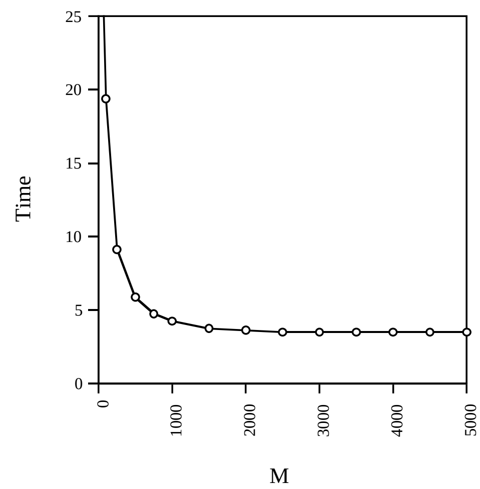

Figure 2: Graph of M vs. expected time per event ($E_M(\delta_{opt})$) in the M bucket calendar queue with 1,000 events.

with mean μ has jump density $f(x) = \frac{1}{2\mu}$ for $0 \leq x \leq 2\mu$ and $x = 0$ for $x > 2\mu$.

Our main result is that if M is infinite, N is large, and δ is chosen to be approximately equal to δ_{opt} where

$$\delta_{opt} = \sqrt{\frac{2b}{c}}\frac{\mu}{N},$$

then the calendar queue performs optimally. Furthermore, when the bucket width is chosen to be δ_{opt} the expected time per event $E(\delta_{opt})$ is approximately

$$E_{min} = c + \sqrt{2bc} + d$$

which is a constant. These results depend on using the unordered list bucket discipline, but do not depend on any specific characteristics of the jump random variable except its mean. The proof of these results comes from the analysis of a continuous state Markov chain which models the calendar queue data structure. Interestingly, a key ingredient in the proof is the analysis of a Markov chain which models a static calendar queue with just one event.

If M is infinite we know how to select the best δ. The next task is to choose M as small as possible yet maintain the minimal, or near minimal, expected time per event of the infinite bucket calendar queue. We say that the jump has *finite support* β if its jump density f satisfies $f(x) = 0$ for all $x > \beta$. The uniform jump of mean μ has support 2μ while the exponential jump does not have finite support. If the jump has finite support β then the choice of any $M \geq \beta/\delta + 1$ will guarantee no loss of performance over infinitely many buckets. If the jump does not have finite support then we present an asymptotic expression for the degradation in performance of choosing an M bucket calendar queue instead of the infinite bucket calendar queue. We show that it is possible to select a constant r such that $M = rN$ buckets is sufficient to achieve a small desired degradation in performance with M buckets over infinitely many buckets.

Figure 1 demonstrates the existence of an optimal δ for minimizing the expected time per event. Figure 2 demonstrates 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 $b = c = d = 1$. 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.

1.3 Related Results

There are a large number of implementations of priorities queues [1, 5, 7, 9, 10, 11]. An interesting empirical comparison of eleven popular implementations is provided by Jones [6]. Most of the popular implementations have expected time per event performance $O(\log N)$, which can be excessive for discrete event simulation with very large N and a large number of simulation steps. As mentioned earlier Brown did an empirical study of the dynamic calendar queue [2]. In the empirical study, Brown's main emphasis was on techniques for changing the bucket width and number of buckets dynamically as N changed. He did not seek to find the optimal bucket width for a given N, just one which would give good performance. The calendar queue provides an attractive alternative priority queue in any discrete event simulation where there is some predictability of the event time of a newly generated event.

1.4 Organization

In section 2 we present our Markov chain model which models the calendar queue and present our main results concerning choosing the optimal δ for M infinite. In section 3 we describe how to choose M finite without significantly compromising the performance of the infinite bucket calendar queue. In section 4 we describe some of our experiences in implementing the calendar queue and give our conclusions.

2 Optimizing the Bucket Width

In this section we show how to choose the optimal bucket width δ_{opt} when the number of buckets is infinite. As mentioned in the introduction we assume we have N events and a jump with probability density f with mean μ. We model the infinite bucket calendar queue as a Markov chain $\widehat{X}$ 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 transitions of $\widehat{X}$ are as follows: Let i be such that $X^i(t) = \min\{X^1(t), X^2(t), \ldots, X^N(t)\}$. If $X^i(t) \geq \delta$ then $X^j(t+1) = X^j(t) - \delta$ for all j. If $X^i(t) < \delta$ then for $j \neq i$, $X^j(t+1) = X^j(t)$, and $X^i(t+1) = X^i(t) + \xi_t$ where $\xi_0, \xi_1, \ldots$ are independent non-negative random variables. It is assumed that these random variables $\xi_t, t \geq 0$, all have the same probability density f. The parameter δ is a fixed non-negative real number. We can think of $X^i(t)$ as the position of the i-th 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 jumps a random distance from where it is while the other particles remain stationary. Thus, a particle in the Markov chain $\widehat{X}$ represents an event in the 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 current bucket in the calendar queue.

Define q_i to be the limiting probability that the interval $[0,\delta)$ has exactly i particles in it. Under mild hypotheses this probability is well defined, depends on δ (and N), and is independent of the of the initial state of $\widehat{X}$. Suppose $a > 0$ is a given constant and that the cost of a step of the Markov chain is $1 + aj$ if there are j particles in the interval $[0,\delta)$. If the expected cost of finding the minimum particle and moving it in the Markov chain is $K(\delta)$, then we have:

$$K(\delta) = \frac{1 + a\sum_{j=1}^{N} jq_j}{1 - q_0}.$$

The relationship between the expected cost per particle in the Markov Chain $\widehat{X}$ and the expected time to process an event in the infinite bucket calendar queue are related in the following lemma.

Lemma 2.1 *Consider an infinite bucket calendar queue with bucket width δ, the unordered list bucket discipline, and parameters b, c, and d. If $E(\delta)$ is the the expected time to process an event and $a = c/b$ then*

$$E(\delta) = bK(\delta) - b + d.$$

Proof: The Markov chain $\widehat{X}$ models the calendar queue, where q_0 is the portion of buckets visited which are empty and for $j > 0$, q_j is the portion of buckets visited which have j events. Each empty bucket visited does not result in finding an event to process, while each non-empty bucket visited will result in finding an event to process. Thus, the expected cost per event in the calendar queue is

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

This quantity equals $bK(\delta) - b + d$. ∎

Lemma 2.1 implies that finding the δ to minimize the expected time per event in the calendar queue is the same δ which minimizes $K(\delta)$.

Theorem 2.1 *Assume that the density $f(x)$ is a bounded function of x for x near 0 and that the mean $\mu = \int_0^\infty xf(x)dx$ is finite. Then the function $K(\delta)$ is minimized at approximately*

$$\delta_{opt} = \sqrt{\frac{2}{a}}\frac{\mu}{N}.$$

The value of $K(\delta_{opt})$ is approximately

$$K_{min} = 1 + a + \sqrt{2a}.$$

The error in δ_{opt} is $O(N^{-\frac{3}{2}})$ and the error in K_{min} is $O(N^{-1})$.

The proof of this theorem is long and requires background in the analysis of Markov chains with a continuous state space. The proof can be found in the Appendix, section A.

As an immediate consequence of lemma 2.1 and theorem 2.1 we have our main theorem.

Theorem 2.2 *Consider an infinite bucket calendar queue with bucket width δ, the unordered list bucket discipline and with parameters b, c, and d. The expected time per event, $E(\delta)$, is minimized at approximately*

$$\delta_{opt} = \sqrt{\frac{2b}{c}\frac{\mu}{N}}.$$

The value of $E(\delta_{opt})$ is approximately

$$E_{min} = c + \sqrt{2bc} + d.$$

The error in δ_{opt} is $O(N^{-\frac{3}{2}})$ and the error in E_{min} is $O(N^{-1})$.

Thus, with the proper choice of δ the infinite bucket calendar queue has a constant expected time per event. Most interestingly, the choice of the optimal bucket width depends only on the mean of the jump and not on the shape of its probability density.

In the derivation of the optimal δ we observed that

$$q_0 = \frac{\mu}{\mu + N\delta}$$

(see equation 8 in section A.2). If δ is chosen to be δ_{opt} then $q_0 = 1/(1+\sqrt{2b/c})$. Thus, the current bucket is empty a significant portion of the time.

3 Choosing the Number of Buckets

Now that we have found how to select δ so as to minimize the expected time per event in the infinite bucket calendar queue our 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.

Recall, that the jump has finite support β if its jump density f satisfies $f(x) = 0$ for all $x > \beta$. In the case when the jump has finite support, 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 $t_0 \leq t(e) < t_0 + \delta$. Thus, we have:

Theorem 3.1 *If the jump has finite support β and $M \geq \beta/\delta + 1$ then the M bucket calendar queue and the infinite bucket calendar queue with bucket width δ have the same expected time per event.*

In the case when the jump does not have finite support or its finite support β would cause β/δ to be too large to be practical, then a M will have to be chosen which gives performance less than that of the infinite bucket calendar queue. The same Markov chain $\widehat{X}$ can be use to analyze this case. Let $L_M(\delta)$ 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 3.1 *Consider an M bucket calendar queue with bucket width δ, the unordered list bucket discipline, and parameters b, c, and d. If $E_M(\delta)$ is the expected time to process an event then*

$$E_M(\delta) = E(\delta) + \frac{c(\mu + N\delta)}{N\delta} L_M(\delta).$$

Proof: In the Markov chain $\widehat{X}$, let q_{ij} be the limiting probability that the there are i particles in the interval $[0, \delta)$ and j particles in Γ. 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 $E_M(\delta)$ equals

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

We have $\sum_{j=0}^{N} q_{0j} = q_0 = \mu/(\mu + N\delta)$ (equation 8 in section A.2) and $\sum_{j=0}^{N} j \sum_{i=0}^{N} q_{ij} = L_M(\delta)$. By using equation 1 in the proof of lemma 2.1 we derive the equation for $E_M(\delta)$. ∎

Let F be the cumulative probability distribution defined by the jump density f, that is,

$$F(x) = \int_0^x f(y)dy.$$

In the Appendix, section B, we indicate briefly how to derive the following rather horrible looking bounds for $L_M(\delta)$.

Theorem 3.2 *$L_M(\delta)$ is bounded above by*

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

and bounded below by

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

where

$$\begin{aligned} p &= \frac{1}{\mu}\int_0^\delta [1-F(x)]dx, \\ \Pi_1 &= \frac{1}{\mu}\sum_{j=1}^{\infty}\int_{jM\delta}^{jM\delta+\delta} [1-F(x)]\,dx, \\ \Pi_2 &= \max_{0\le y\le\delta}\sum_{j=1}^{\infty}[F(jM\delta+\delta-y)-F(jM\delta-y)]. \end{aligned}$$

It should be noted that under the hypothesis $\mu < \infty$, the above series converge and can be given bounds in terms of μ, δ, and M. However, using the bounds as stated in the theorem, we can derive a more useful asymptotic expression for $L_M(\delta)$.

Corollary 3.1 *If $\delta = \lambda\mu/N$ and $r = M/N$ where r and λ are constants, then*

$$L_M(\delta) \simeq \lambda\sum_{j=1}^{\infty}[1-F(j\mu\lambda r)].^1$$

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

$$\epsilon_M = \frac{E_M(\delta)-E(\delta)}{E(\delta)}.$$

If we choose $\delta = \delta_{opt} = \sqrt{\frac{2b}{c}\frac{\mu}{N}}$, then corollary 3.1 and theorem 2.2 yield the following asymptotic expression for ϵ_M

Corollary 3.2 *If M/N is constant and $\delta = \delta_{opt}$, then*

$$\epsilon_M \simeq \frac{c+\sqrt{2bc}}{c+\sqrt{2bc}+d}\sum_{j=1}^{\infty}[1-F(j\mu\sqrt{\frac{2b}{c}}\frac{M}{N})]. \qquad (2)$$

The following asymptotic bound is implied by corollary 3.2.

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

Corollary 3.3 *If M/N is constant and $\delta = \delta_{opt}$, then*

$$\epsilon_M \preceq \frac{c+\sqrt{2bc}}{c+\sqrt{2bc}+d}\cdot\sqrt{\frac{c}{2b}}\cdot\frac{N}{M}.^2$$

Proof: By corollary 3.2 it suffices to show that $\sum_{j=1}^{\infty}[1-F(jMD)] \le \mu/D$ for $D > 0$. To see this let $k \ge 2$, then

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

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 $\mu/D \ge \sum_{j=1}^{\infty}[1-F(jD)]$. ∎

It follows from corollary 3.3 that one can always choose the number of buckets M to be a (moderate) 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 $f(x) = \frac{1}{\mu}e^{-x/\mu}$, $x \ge 0$, we can calculate the the series in (2) exactly:

$$\epsilon_M \simeq \frac{c+\sqrt{2bc}}{c+\sqrt{2bc}+d}\cdot\frac{1}{e^{\sqrt{\frac{2b}{c}}\frac{M}{N}}-1}. \qquad (3)$$

Let us suppose $b = c = d = 1$. If $\delta = \delta_{opt} = \sqrt{2}\mu/N$, then equation (3) allows us to solve for M/N given an acceptable degradation in performance of the M bucket calendar queue over the the infinite bucket calendar queue. 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 (3) 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 $M/N = 3.02$. Both plots are relatively flat near the asymptotic values .01 and .05, respectively.

4 Conclusion

We have shown that bucket width of the infinite bucket static calendar queue can be chosen to minimize the expected time per event. With the optimal

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

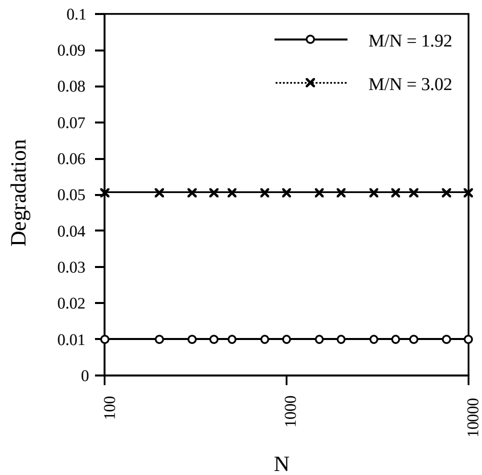

Figure 3: Graph of N vs. degradation (ϵ_M) for $M/N = 1.92$ and $M/N = 3.02$.

bucket width the calendar queue has expected constant time per event. The optimal bucket width depends only on a few parameters, 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). In addition, we have shown that the number of buckets M can be chosen to be $O(N)$ so as to achieve minimal or almost minimal expected time per event.

We implemented our calendar queue in C^{++} for execution on a DEC alpha [3]. Because the DEC alpha has a two-level cache architecture the values of the "constants" b, c, and d actually depend on N, M and even on properties the jump probability density f other than its mean. As N grows the cache miss ratio increases, in effect, causing b, c, and d to grow. Eventually, when N gets large enough the growth in b, c, and d will stop. In the formula for δ_{opt} it is the ratio of b to c that is important. One would hope that this ratio of b to c would not change as N grows, but it does. The value of b depends of the time to process empty buckets, which are in an array, and consequently are spacially correlated. The value of c depends on the time to process list entries, which are not in general spacially correlated. Some jump densities will lead to more temporal correlation between memory accesses than others. For example, the temporal correlation of list items will be higher with the exponential jump than with the uniform jump. Both spacial and temporal correlation of memory accesses effect the cache miss ratio, thereby effecting the values of b, c, and d. We are currently studying methodologies for setting the parameters of the calendar queue on modern computers with multi-level memory hierarchies.

Our analysis applies to the unordered list bucket discipline for the calendar queue. An interesting question is to analyze the calendar queue with the ordered list bucket discipline.

References

[1] M.R. Brown. Implementation and analysis of binomial queue algorithms. *SIAM Journal on Computing*, Vol. 7, pp. 298-319, 1978.

[2] R. Brown. Calendar Queues: A fast $O(1)$ priority queue implementation for the simulation event set problem. *Communications of the ACM*, Vol. 31, pp. 1120-1227, 1988.

[3] DECchip 21064-AA Microprocessor, Hardware Reference Manual. Digital Equipment Corporation. Order Number: EC-N0079-72. 1992.

[4] W. Feller. *An Introduction to Probability Theory and Its Applications, Vol. II.* John Wiley and Sons, New York, 1971.

[5] J. Francon, G Viennot, and J. Vuillemin. Description and Analysis of an efficient priority queue representation. Proceedings of the 19th Annual Symposium on Foundations of Computer Science, pp. 1-7, 1978.

[6] D.S. Jones. An empirical comparison of priority-queue and event-set implementations. *Communications of the ACM*, Vol. 29, pp. 300-310, 1986.

[7] D.E. Knuth. *The Art of Computer Programming, Vol. 3, Sorting and Searching.* Addison-Wesley, Reading, MA, 1973.

[8] D. Revuz. *Markov Chains.* North-Holland, New York. 1984.

[9] D.D. Sleator and R.E. Tarjan. Self-adjusting binary search trees. *Journal of the ACM*, Vol. 32, pp. 652-686. 1985.

[10] D.D. Sleator and R.E. Tarjan. Self-adjusting heaps. *SIAM Journal on Computing*, Vol. 15, pp. 52-69, 1986.

[11] J. Vuillemin. A data structure for manipulating priority queues. *Communications of the ACM*, Vol. 21, pp. 309-315, 1978.

A Proof of Theorem 2.1

A.1 Invariant Distribution, Positive Recurrence, and Limits.

Let F denote the common distribution function of the random variables ξ. Since we are assuming the existence of a density this means that

$$F(x) = \int_0^x f(z)dz, \ 0 \leq x < \infty.$$

(More compactly, $dF(z) = f(z)dz$.) For the most part we work with F rather than f. Let Λ denote the set of points $\widehat{x}$ in $[0,\infty)^N$ with no repeating coordinates. In what follows the reader should understand all set operations taken relative to Λ. The density assumption implies that the probability is zero that any two particles which start at different positions will ever be found at the same position at some time in the future, so the chain stays in Λ once started there. Indeed, Λ is an absorbing set for the chain under the density assumption. The true state space of the chain is, therefore, a subset of Λ but one may certainly refer to $[0,\infty)^N$ as the state space without harm. Let B_i be the set of $\widehat{x} = (x^1, x^2, \ldots, x^N)$ in Λ such that $x^i = \min\{x^1, x^2, \ldots, x^N\} < \delta$, and let $A_0 = [\delta, \infty)^N$ For arbitrary $\widehat{x}$ in the state space and arbitrary (measurable) subset A, the one-step transition probability (T.P.) that the chain will move from $\widehat{x}$ to a state in A is given by

$$\begin{aligned} P(\widehat{x}, A) &= \sum_{i=1}^{N} 1_{B_i}(\widehat{x}) \int_0^\infty 1_A(\widehat{x} + z\widehat{e}_i)dF(z) \\ &\quad + 1_{A_0}(\widehat{x})1_A(\widehat{x} - \delta\widehat{1}), \end{aligned} \tag{4}$$

where, $\widehat{1} = (1, 1, \ldots, 1)$, $\widehat{e}_i$ = the i-th unit coordinate vector with all of its components 0 except the i-th which is 1, and $1_A(\widehat{x})$ is the function which is 1 for $\widehat{x}$ in A and 0 otherwise. In Revuz [8], Chapter 1, may be found the basic theory of Markov chains. We use P^ν to denote the probability measure induced on the "trajectory space" of the chain when the initial distribution is ν, and $P^{\widehat{x}}$ when the chain is started at $\widehat{x}$ (i.e., ν is unit point mass at $\widehat{x}$). Integration of random variables (called "expectation") with respect to these measures is denoted E^ν and $E^{\widehat{x}}$, respectively. A distribution m is called an *invariant probability distribution* for the T.P. P if m has total mass 1 and for every subset A of the state space

$$m(A) = \int P(\widehat{x}, A)dm(\widehat{x}). \tag{5}$$

(The integration is over the entire state space.) For the T.P. defined by (4) it can be established that if the number μ is finite, then the Markov chain is *positive recurrent.* That means there is a unique invariant probability distribution m and if A is any set with $m(A) > 0$, then, with probability 1, $\widehat{X}(t)$ will be in A for an infinite (but usually random) sequence of t's tending to infinity. Under any of various mild assumptions on the support of F (e.g., if it contains at least an interval of length strictly $> \delta$, or if it contains points arbitrarily close to 0), then the chain is also *aperiodic.* This is a property of the chain whose only significance for us is that its presence justifies the first limit in (6) below and later on a similar limit for a related embedded chain. (The existence of limits of averages does not require aperiodicity, however.) Since it is convenient to do so and results in only a minor loss of generality we will assume that F satisfies at least one of the above restrictions making the chain aperiodic.

The limit theory for positive chains (ergodic theory) is comprehensive. See [8], Chapter 6. In our case if A is any set of states and ν any initial distribution

$$\begin{aligned} \lim_t P^\nu\{\widehat{X}(t) \in A\} &= \lim_t \frac{1}{t}\sum_{s=1}^{t} P^\nu\{\widehat{X}(s) \in A\} \\ &= m(A) \qquad (6) \\ &= \lim_t \frac{\#\{s : s \leq t, \widehat{X}(s) \in A\}}{t}. \end{aligned}$$

(Note that the last limit is a limit of random quantities and the assertion is that the limit (as $t \to \infty$) exists with probability 1 (w.p.1) and equals the non-random quantity $m(A)$. Note also that $P^m\{\widehat{X}(0) \in A\} = m(A)$ in our notation.) In particular, if

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

and $Z(t)$ = the number of particles in interval $[0, \delta)$ at time t, then $q_j \equiv \lim_t P^\nu\{Z(t) = j\} = m(A_j)$ and

$$K(\delta) = \lim_t \frac{1 + aE^\nu\{Z(t)\}}{P^\nu\{Z(t) > 0\}} = \frac{1 + a\sum_1^N jq_j}{1 - q_0}. \tag{7}$$

A.2 The computation of $\mathbf{q_0 = m(A_0)}$.

Recall that $A_0 = [\delta, \infty)^N$. In this section we establish the formula

$$m(A_0) = q_0 = \frac{\mu}{\mu + N\delta}, \tag{8}$$

provided μ is finite.

The sets B_i, defined in the last section, are mutually 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

$$m(A_0) = 1 - \sum_1^N m(B_i) \ . \qquad (9)$$

Let ψ be any bounded or positive function on A_0. Equation (5) has a valid analogue for functions ψ which reads: $\int \psi(\widehat{x})dm(\widehat{x}) = \int dm(\widehat{x}) \int \psi(\widehat{y})P(\widehat{x};d\widehat{y})$. The integrals with respect to m are sums of the integrals over the above mentioned sets. Noting that $\int_{A_0} dm(\widehat{x}) \int \psi(\widehat{y})P(\widehat{x},d\widehat{y}) = \int_{A_0} \psi(\widehat{x}-\delta\widehat{1})dm(\widehat{x})$, by (4), and doing a little rearranging, we get

$$\int_{A_0} [\psi(\widehat{x}-\delta\widehat{1}) - \psi(\widehat{x})]dm(\widehat{x}) = \sum_1^N \int_{B_j} dm(\widehat{x})[\psi(\widehat{x}) - \int \psi(\widehat{y})P(\widehat{x};d\widehat{y})]. \ (10)$$

Let us fix i and for ψ take $\psi(\widehat{x}) = \exp(-\alpha x^i)$ where α is any complex number with nonnegative real part independent of $\widehat{x}$. Then, from the above considerations and (4), we find that for $\widehat{x}$ in B_i,

$$\begin{aligned}\int \psi(\widehat{y})P(\widehat{x},d\widehat{y}) &= \int_0^\infty \exp\{-\alpha(x^i+z)\}dF(z) \\ &= \exp(-\alpha x^i)\phi(\alpha),\end{aligned}$$

where $\phi(\alpha) = \int_0^\infty e^{-\alpha z}dF(z)$ is the Laplace transform of F. On the other hand, for $\widehat{x}$ in B_j with $j \neq i$,

$$\begin{aligned}\int \psi(\widehat{y})P(\widehat{x};d\widehat{y}) &= \int_0^\infty \exp(-\alpha x^i)dF(z) \\ &= \exp(-\alpha x^i) = \psi(\widehat{x}).\end{aligned}$$

Thus all but the i^{th} term on the right-hand side of (10) vanishes and it becomes

$$[1-\psi(\alpha)]\int_{B_i} \exp(-\alpha x^i)dm(\widehat{x}).$$

The left-hand side of (10), also simplifies easily and one eventually arrives at

$$[e^{\alpha\delta}-1]\int_{A_0} \exp(-\alpha x^i)dm(\widehat{x}) = [1-\phi(\alpha)]\int_{B_i} \exp(-\alpha x^i)dm(\widehat{x}) \ . \qquad (11)$$

Divide both sides by α and let α tend to 0. Noting that the integrands tend to 1, $[1-\phi(\alpha)]/\alpha$ tends to $-\phi'(0) = \mu$, and $(e^{\alpha\delta}-1)/\alpha$ tends to δ, we get that $\delta m(A_0) = \mu m(B_i)$, for each $i = 1,\ldots,N$. Equation (8) now follows immediately from this and (9).

A.3 The Case 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. For $i = 1,\ldots,N$, let $u^i(0) = 0$, and for $r = 1,2,\ldots$, let

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

From the description of the chain in terms of the independent random variables ξ, one concludes: (I) Each sequence $\{X^i(u^i(r)) : \ r = 0,1,2,\ldots\}$ is itself a Markov chain on the line; (II) These N Markov chains are mutually independent. If one 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 the same, such that $S_k = u^i(r^i(k))$, for every i), then, given $\widehat{X}(0) = \widehat{x}$, $\widehat{X}(S_k)$ has, for each $k > 0$, mutually independent components. Here is such a sequence: let $S_0 = 0$, and for $k > 0$, let $S_k = 1 + \min\{t : \ t \geq S_{k-1} \ \& \ \widehat{X}(t) \in A_0\}$. The times $T_k = S_k - 1$ are the successive (random) times at which the interval $[0,\delta)$ is empty of particles $(Z(T_k) = 0)$. Since $\widehat{X}(T_k)$ is obtained from $\widehat{X}(S_k)$ by adding the deterministic constant δ to each of the components of $\widehat{X}(S_k)$, it follows that the components of $\widehat{X}(T_k)$ are also mutually independent. The sequence $\{\widehat{X}(T_k), k \geq 0\}$ also forms a Markov chain called the *trace chain on* A_0. An important point to note is that the special structure of $\widehat{X}$ implies that for each $i = 1,\ldots,N$ the chain $\{X^i(T_k), k \geq 0\}$ itself coincides, in law, with the trace chain on $[\delta,\infty)$ of an $N = 1$ version of $\widehat{X}$. The general theory of trace chains (cf. [8], Ex.3.13, p. 27 & Prop.2.9, p. 93) implies, reasonably enough, that the trace chain $\{\widehat{X}(T_k)\}$ is also positive recurrent with an invariant probability distribution m_0, say, obtained by renormalizing the distribution m restricted to A_0. That is, for subsets B,

$$m_0(B) = \frac{m(B\cap A_0)}{m(A_0)} = (1+\frac{N\delta}{\mu})m(B\cap A_0). \quad (12)$$

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 distribution G_0. Once computed, G_0 (concentrated on $[\delta,\infty)$) is used to compute the limit, as $k \to \infty$, of the probability of finding exactly j particles in the interval $[0,\delta)$ at the 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 "success" parameter $p = G_0\{[\delta, 2\delta)\}$. (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 G rather than m, in the case $N = 1$ of our basic chain can be calculated explicitly and then G_0 obtained from the special case of (12). In the case $N = 1$, $A_0 = [\delta, \infty)$ and $B_1 = [0, \delta)$ and equation (11), simplifies to

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

which is 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, and $i^2 = -1$, we find that the left-hand side vanishes. The density assumption implies that $\phi(\alpha) \neq 1$ for any $\alpha \neq 0$. Hence $\int_0^\delta \exp(2n\pi ix/\delta) dG(x) = 0$ for every $n \neq 0$. Standard uniqueness results in the theory of Fourier series implies that we must have $dG(x) = C dx$, $0 \leq x \leq \delta$, for some constant C. From (8) in the case $N = 1$ we find $C\delta = 1 - G\{[\delta, \infty)\} = \delta/(\mu + \delta)$, so $C = 1/(\mu + \delta)$. For the Laplace transform of G on $[\delta, \infty)$, we get

$$\begin{aligned} \int_\delta^\infty e^{-\alpha x} dG(x) &= \frac{1 - \phi(\alpha)}{e^{\alpha\delta} - 1} \int_0^\delta e^{-\alpha x} C dx \\ &= e^{-\alpha\delta} \frac{1 - \phi(\alpha)}{\alpha(\mu + \delta)}. \end{aligned}$$

Inverting the Laplace transforms in this equation reveals that the density function g of G on $[\delta, \infty)$ exists and for $x \geq \delta$, is given by $g(x) = [1 - F(x - \delta)]/(\mu + \delta)$, and from (12) it is then clear that G_0 has the density $g_0(x) = (1 + \delta/\mu) g(x)$ for $x \geq \delta$. Thus, the limit distribution of $Z(S_k)$ as $k \to \infty$ is:

$$\begin{aligned} b(i) &= \lim_k P^\nu\{Z(S_k) = i\} = \binom{N}{i} p^i (1 - p)^{N-i} \\ &= \lim_n (1/n) \#\{k : k \leq n, Z(S_k) = i\} \end{aligned} \qquad (13)$$

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

$$\begin{aligned} p &= G_0\{[\delta, 2\delta)\} = \frac{1}{\mu} \int_0^\delta [1 - F(x)] dx \\ &= \frac{1}{\mu} \int_0^\delta \int_x^\infty f(z) dz dx \end{aligned} \qquad (14)$$

A.4 Estimates for the $\mathbf{q_j}$.

Apart from the explicit representation of m on A_0 discussed in the last section, a simple direct expression for m on all of $[0, \infty)^N$ for $N > 1$ is not available. (In particular it appears that m is generally not a product measure on all of $[0, \infty)^N$.) This means that, with one exception noted below, for $j > 0$ we do not have explicit formulas for the values of $q_j = m(A_j)$ and must resort to approximations. In this section we establish two inequalities:

$$q_j \geq q_0 \sum_j^N b(i) , \qquad (15)$$

$$q_j \leq q_0 [1 - F(\delta)]^{-1} \sum_j^N b(i) , \qquad (16)$$

for $j = 1, 2, \ldots, N$. (The one case where we get equality instead of inequality is when the distribution F assigns no mass to $[0, \delta)$, so that $F(\delta) = 0$.)

To simplify the notation a little, the starting distribution ν will be omitted, if not forgotten, when it is not essential. Finally, we put

$$\begin{aligned} D(j) &= \textstyle\sum_{i=j}^N b(i), \\ n^\circ(t) &= \max\{k : S_k \leq t\}, \\ \#(t, j) &= \#\{s : \ s \leq t \ \& \ Z(s) = j\}, \\ \#^*(n, j+) &= \#\{k : \ k \leq n \ \& \ Z(S_k) \geq j\}. \end{aligned}$$

Consider one of the stopping times S_k. If $Z(S_k) \geq j$, then as t goes from S_k to $T_{k+1} = S_{k+1} - 1$, the values of $Z(t)$ must decrease through each value of i from j to 0, possibly pausing for a time at one or more of these values. It follows that for every t

$$\#(t, j) \geq \#^*(n^\circ(t), j+).$$

Now $\#(t, 0)$ differs from $n^\circ(t)$ by at most 1, since the T_k's are the zeros of Z. Therefore, by the law of large numbers for $\widehat{X}(t)$,

$$\lim_{t \to \infty} \frac{n^\circ(t)}{t} = q_0, \qquad (17)$$

and, see (13),

$$\begin{aligned} q_j = \lim_t \frac{\#(t, j)}{t} &\geq \left[\lim_t \frac{n^\circ(t)}{t}\right] \left[\lim_n \frac{\#^*(n, j+)}{n}\right] \\ &= q_0 D(j) \ . \end{aligned}$$

This proves (15).

To prove equation (16) note first that equation (17) entails $\lim_n S_n/n = \lim_n S_n/n^\circ(S_n) = \lim_t [t/n^\circ(t)] = 1/q_0$, and then

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

Fix $j > 0$ and let us define a sequence of random variables $\{V_k\}$ by

$$\begin{aligned} V_k &= \max\{r : Z(s) = Z(s+1) = \cdots Z(s+r) = j \\ &\qquad \text{for some } s \text{ in } [S_{k-1}, S_k)\}. \end{aligned}$$

Note that $V_k = 0$ if $Z(S_{k-1}) < j$ or if $Z(s) = j$ only once during $[S_{k-1}, S_k)$. Clearly

$$\#(S_n, j) = \#^*(n, j+) + \sum_1^n V_k \ .$$

Therefore, by Fatou's Lemma and the dominated convergence theorem,

$$\begin{aligned} \frac{q_j}{q_0} &= E\{\lim_n \frac{\#(S_n, j)}{n}\} \le \liminf_n E\{\frac{\#(S_n, j)}{n}\} \\ &\le \lim_n E\{\frac{\#^*(n, j+)}{n}\} + \limsup_n E\{\frac{1}{n}\sum_1^n V_k\} \\ &\le D(j) + \limsup_k E\{V_k\}. \qquad (18) \end{aligned}$$

To get a bound for the constant $\limsup E\{V_k\}$ we proceed as follows. Let R be the first $s \ge S_{k-1}$ such that $Z(s) = j$. Then, by the Markov property,

$$\begin{aligned} E\{V_k\} &= E\{V_k : Z(S_{k-1}) \ge j\} \\ &= E\{E\{V_k | \mathcal{F}_R\}; Z(S_{k-1}) \ge j\} \\ &= E\{E^{\widehat{X}(R)}\{V_1\}; Z(S_{k-1}) \ge j\} \\ &\le B\, P\{Z(S_{k-1}) \ge j\}, \qquad (19) \end{aligned}$$

where $B = \sup\{E^{\widehat{x}}\{V_1\} : \widehat{x} \in A_j\}$. (The vertical bar with the inner expectation after the second equality above denotes conditional expectation given $\mathcal{F}_R$, the σ-field generated by $\{\widehat{X}(t) : t \le R\}$. The shift back to $k = 1$ is a consequence of the time homogeneity of the chain $\{\widehat{X}(S_k)\}$.)

Note that when the initial value of Z is j, the event $\{V_1 > r\}$ implies that at least $r + 1$ independent random variables ξ (remember the ξ's?) had values not larger than δ. Hence, for any $\widehat{x}$ in A_j, we must have $P^{\widehat{x}}[V_1 > r] \le F(\delta)^{r+1}$. Therefore

$$\begin{aligned} E^{\widehat{x}}\{V_1\} &= \sum_{r=1}^{\infty} r P^{\widehat{x}}\{V_1 = r\} \\ &= \sum_{r=0}^{\infty} P^{\widehat{x}}\{V_1 > r\} \le \frac{F(\delta)}{1 - F(\delta)}, \end{aligned}$$

for all $\widehat{x}$ in A_j. Returning to (19) with this bound for B and then to (18) and recalling (13) again, we get:

$$\begin{aligned} \frac{q_j}{q_0} &\le D(j) + \frac{F(\delta)}{1 - F(\delta)} \lim_k P^{\nu}\{Z(S_{k-1}) \ge j\} \\ &= (1 + \frac{F(\delta)}{1 - F(\delta)}) D(j) = \frac{D(j)}{1 - F(\delta)}, \end{aligned}$$

which proves (16).

A.5 The Minimization Problem.

Throughout this section we assume that $N \ge 2$.

The binomial distribution $\{b(i)\}$ for N trials with success parameter p has mean Np and second moment $(Np)^2 + Np(1-p)$. Hence,

$$\begin{aligned} \textstyle\sum_1^N j \left(\sum_j^N b(i)\right) &= \textstyle\sum_1^N \frac{1}{2}(i^2 + i) b(i) \\ &= \tfrac{1}{2}[(N^2 - N)p^2 + 2Np]. \end{aligned}$$

Now $p \ge \delta\,[1 - F(\delta)]/\mu$ and $(1 - q_0)^{-1} = 1 + \mu/N\delta$. (See (8) and (14.) These considerations and formulas (7), (15), and (16), yield

$$\begin{aligned} K(\delta) &\ge \frac{1 + \frac{1}{2} a q_0 [(N^2 - N)p^2 + 2Np]}{1 - q_0} \\ &\ge [1 - F(\delta)]^2 h(x) \ge [1 - 2F(\delta)] h(x), \end{aligned}$$

where $x = N\delta/\mu$ and

$$h(x) = 1 + a + \frac{1}{x} + \frac{1}{2} a (1 - 1/N) x.$$

On the other hand note that $p \le \delta/\mu$, so if we take $\delta \le \delta'$, where δ' is chosen so that $F(\delta') \le 1/2$, then

$$\begin{aligned} K(\delta) &\le \frac{1 + \frac{1}{2} a q_0 [1 - F(\delta)]^{-1} [(N^2 - N)p^2 + 2Np]}{1 - q_0} \\ &\le [1 - F(\delta)]^{-1} h(x) \le [1 + 2F(\delta)] h(x). \end{aligned}$$

Since the density f is bounded near 0, there is a constant C_0 such that for small $\delta_0 > 0$ (in particular we want $\delta_0 \le \min\{\delta', \frac{1}{2}\}$), $F(\delta) \le C_0 \delta \le C_0 \delta_0 < \frac{1}{2}$ for all $0 < \delta \le \delta_0$, and then

$$(1 - 2C_0\delta_0) h(x) \le K(\delta) \le (1 + 2C_0\delta_0) h(x).$$

Now the extreme members of this last inequality have absolute minimum values on $(0, \infty)$ at the point $x_0 = \sqrt{2N/a(N-1)} = \sqrt{2/a}(1 + \epsilon/N)$, where ϵ is in $(0, 1)$. It follows (as the reader can easily convince himself by sketching some graphs) that the GLB (greatest lower bound) of $K(x\mu/N)$ can be calculated by examining its values for x between the solutions, z_1 and z_2, of the equation $(1 - 2C_0\delta_0) h(z) = (1 + 2C_0\delta_0) h(x_0)$. It can be shown that $|z_2 - z_1| = O(\sqrt{\delta_0})$ for $\delta_0 \simeq 0$. Furthermore, x_0 lies between z_1 and z_2. Hence, if x is such that $K(x\mu/N)$ is near its GLB, then $|x - x_0| \le C_1\sqrt{\delta_0}$, with C_1 a constant which can be chosen independent of δ_0. Therefore the minimizing $\delta (= \mu x/N)$ is bounded by a constant C_2, say, times $1/N$. Replacing δ_0 by

C_2/N, where $N \gg C_2/\delta_0$, and repeating the previous argument, shows that we can replace $C_1\sqrt{\delta_0}$ in the inequality for $|x - x_0|$ by $C_3/\sqrt{N}$ with C_3 independent of N. Multiplying this formally improved inequality by μ/N gives $|\delta_{opt} - \sqrt{2/a}\mu/N| = O(N^{-3/2})$. This estimate and our inequalities for the function K now yield that $K_{min} = 1 + a + \sqrt{2a}$ with an error $O(\delta) = O(N^{-1})$.

B Proof of Theorem 3.2

We continue the notation of the preceding sections of the appendix. In addition we will write L_M for $L_M(\delta)$ and $Z^A(t)$ for the number of particles in set A at time t. Since $\Gamma = [M\delta, M\delta+\delta)\cup[2M\delta, 2M\delta+\delta)\cup\cdots$ is a subset of $[\delta,\infty)$ we have

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

for $S_{j-1} = T_{j-1} + 1 \le s \le T_j$. Hence,

$$\begin{aligned} L_M &= \lim_t (1/t)\sum_{s\le t} Z^\Gamma(s) \\ &= \lim_k (k/T_k)\lim_k (1/k)\sum_{j=1}^k [\sum_{s=S_{j-1}}^{T_j} Z^\Gamma(s)] \\ &= q_0 \lim_k E[\sum_{s=S_{k-1}}^{T_k} Z^\Gamma(s)] \\ &\le q_0 \limsup_k E[Z^\Gamma(T_k)(\Delta_k)], \end{aligned} \tag{20}$$

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 as the i-th of these particles lands in Γ or not when it is finally removed from the interval $[0,\delta)$. (Which removal must occur during $[S_{k-1}, T_k]$.) Writing $u = u_i$, $x = x_i$, then

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

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

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

where U is the renewal function. (See Feller [4], page 369. The function defined by Feller is the hitting probability of $(t, t+b]$; this coincides with our H on account of the density assumption.) In general $U(z) \le [1-F(z)]^{-1}$ for distributions on $[0,\infty)$ so that

$$\begin{aligned} p_x &\equiv \int_{0-}^{\delta-x}\sum_{j=1}^\infty [F(jM\delta + \delta - x - z) \\ &\quad - F(jM\delta - x - z)]\, dU(z) \\ &\le \Pi_2/[1-F(\delta)] \end{aligned}$$

(Recall the definition of Π_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

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

where $r = Z_{k-1}^{[\delta,2\delta)}$, and $Z_j^{\{\cdot\}}$ means $Z^{\{\cdot\}}(T_j)$, $j = k-1, k$. Now at the 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

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

where p is defined at (14) and $\Pi_1 = G_0(\Gamma+\delta)$. Going back to (20) with these calculations we obtain

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

(Recall that $E\{E[\,\cdot\,|\mathcal{F}_{T_{k-1}}]\} = E\{\,\cdot\,\}$.) Replacing p^* with $\Pi_2/[1-F(\delta)]$ and q_0 with $\mu/(\mu+N\delta)$ (equation (8)) we obtain the upper bound on L_M.

For the lower bound we have

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

which evaluates to the lower bound on L_M.