

**The Maximum Size of
Dynamic Data Structures**

Claire M. Kenyon-Mathieu and Jeffrey Scott Vitter

Technical Report No. CS-89-22
March 1989

The Maximum Size of Dynamic Data Structures *

Claire M. Kenyon-Mathieu¹ and Jeffrey Scott Vitter²

Abstract. We develop two probabilistic methods that allow us to analyze the maximum data structure size encountered during a sequence of insertions and deletions in data structures such as priority queues, dictionaries, linear lists, and symbol tables, and in sweepline structures for geometry and VLSI applications. The notion of the “maximum” is basic to issues of resource preallocation. We apply our methods to combinatorial models of file histories and probabilistic models, as well as to a non-Markovian process (algorithm) for processing sweepline information in an efficient way, called “hashing with lazy deletion” (HwLD). We derive expressions for the expected maximum data structure size that are asymptotically exact, that is, correct up to lower-order terms; in several cases of interest the expected value of the maximum size is asymptotically equal to the maximum expected size. This solves several open problems, including longstanding questions in queueing theory. Both of our approaches are robust and rely upon novel applications of techniques from the analysis of algorithms. At a high level, our first method isolates the primary contribution to the maximum and bounds the lesser effects. In our second technique we relate the continuous-time probabilistic model to its discrete analog—the maximum slot occupancy in hashing.

1. Introduction

The size attained by data structures is fundamental to issues of resource allocation, yet until recently, little was known about analyzing the *maximum* size attained over a period of time, which is important for preallocating resources. A possible explanation of this deficiency is that classical methods of analysis with generating functions and recurrences cannot be applied readily for the maximum function. In this paper we develop two asymptotic methods to study the distribution of the maximum size of data structures. The methods are robust in that they apply to several different combinatorial and probabilistic models. We also study a non-Markovian process called *hashing with lazy deletion* (HwLD), which corresponds to an efficient way of processing sweepline information in computational geometry and VLSI layout applications [Van Wyk and Vitter, 1986].

* An extended abstract of this work appears in [Kenyon-Mathieu and Vitter, 1989].

¹ Current address: Department of Computer Science, Princeton University, Princeton, NJ 08544. This work was done while the author was at Ecole Normale Supérieure, LIENS, 45, rue d’Ulm, 75230 Paris Cedex 05, France.

² Current address: Department of Computer Science, Brown University, Box 1910, Providence, R. I. 02912. Support was provided in part by an NSF research grant and by an NSF Presidential Young Investigator Award with matching funds from IBM.

One of the motivations for our study is the need to develop and analyze practical space-efficient plane-sweep algorithms. Some work in this area has been done by Van Wyk and Vitter [1986], Morrison, Shepp, and Van Wyk [1987], Mathieu and Vitter [1988], and Ottmann and Wood [1986], but as the latter point out, “Surprisingly there has been little theoretical investigation of space-economical plane-sweep algorithms even though such algorithms have significant practical applications.” Ottmann and Wood [1986] do not investigate the maximum data structure size (that is, the maximum number of items cut by the sweepline); they express the running times of their algorithms in terms of the maximum size. Our approach in this paper is to examine the distribution of the maximum data structure size, based upon several popular input models, and in addition to show that the HwLD algorithm introduced is optimum simultaneously for both average running time and preallocated space.

Data structures process a sequence of items over time; at time t the data structure stores the items that are “living” at time t . Let us think of the i th item as being an interval $[s_i, t_i]$ in the unit interval, containing a unique key k_i of supplementary information. The i th item is “born” at time s_i , “dies” at time t_i , and is “living” when $t \in [s_i, t_i]$. The data structure also handles dynamic queries over time. Let us denote the data structure size at time t by $Size(t)$. If we think of the items as horizontal intervals, then $Size(t)$ is just the number of intervals “cut” by the vertical line at position t . In a typical planesweep application, having to do with VLSI artwork analysis, we might have 10^6 intervals in the time range $[0, 1]$, with $E(Size) = 10^3$; that is, only square root of the total number of items tend to be present at any given time [Szymanski and Van Wyk, 1983]. Thus, for space efficiency, it is important to use a dynamic data structure whose size follows the growth rate of $Size(t)$.

In HwLD, items are stored in a hash table of H buckets, based upon the hash value of the key. The distinguishing feature of HwLD is that an item is not deleted as soon as it dies; the “lazy deletion” strategy deletes a dead item only when a later insertion accesses the same bucket. The number H of buckets is chosen so that the expected number of items per bucket is small. HwLD is thus more time-efficient than doing “vigilant-deletion,” at a cost of storing some dead items.

To model insertions, deletions, and queries, we consider two classes of models: *combinatorial* models and *probabilistic* models. The combinatorial models are the discrete-time models of file histories introduced in [Flajolet, Françon, and Vuillemin, 1979] and [Flajolet, Françon, and Vuillemin, 1980] to model the evolution of several classical types of dynamic data structures, such as priority queues, dictionaries, stacks, and linear lists. The second class of models consists of probabilistic continuous-time models in equilibrium in which the birthtimes of items are independent and form a Poisson process with birth rate λ . Various models of lifetime are considered. Not only does our approach work for these models, but it can also be adapted to handle models in which lifetimes are not independent, such as the M/M/1 probabilistic model and the non-Markovian models corresponding to HwLD.

We denote by $Use(t)$ the number of items stored at time t in the HwLD data structure. The lazy deletion strategy means that $Use(t) \geq Size(t)$. Let t^* be any time t that maximizes $E(Size(t))$. (For the probabilistic models, $E(Size(t))$ is the same for all t .) Van Wyk and Vitter [1986] compute $E(Size(t^*))$ and show via generating function

and approximation techniques that $E(Use(t^*)) \sim E(Size(t^*)) + H$ for the combinatorial model of priority queues and for the M/M/ ∞ probabilistic model. Big-oh bounds on $E(\max_{t \in [0,1]} \{Size(t)\})$ and $E(\max_{t \in [0,1]} \{Use(t)\})$ were only recently obtained by [Mathieu and Vitter, 1988] under certain assumptions for the M/G/ ∞ probabilistic model. Exact formulas were also developed for several combinatorial and probabilistic models that could be used to compute the distribution of $\max_{t \in [0,1]} \{Size(t)\}$ numerically, but they do not seem to give any asymptotic information. However, the fact that the relevant transform in each case was expressed simply as the ratio of consecutive classical orthogonal polynomials gave informal evidence that some common asymptotic method(s) might exist to analyze the different models.

In this paper we develop general asymptotic methods using techniques from analysis of algorithms to settle the open problems posed in [Van Wyk and Vitter, 1986], [Morrison, Shepp, and Van Wyk, 1987], and [Mathieu and Vitter, 1988]. We derive *asymptotically exact* expressions for $E(\max_{t \in [0,1]} \{Size(t)\})$ and $E(\max_{t \in [0,1]} \{Use(t)\})$ for several combinatorial and probabilistic models. In particular we show that *HwLD is asymptotically optimal in terms of preallocated storage*. The gist of our first method is to concentrate on the primary contribution to the maximum and to show via probabilistic techniques that the rest of the contribution is negligible. The hard part is coping with the inherent lack of independence of the size as a function of time. We show exactly when the expected maximum size is asymptotically more than the maximum expected size and when they are equal. The second method we use, for the continuous-time probabilistic models, is a discrete counterpart having to do with the maximum slot occupancy in hashing. This approach offers another illustration of the strong connections between discrete and continuous models in the analysis of algorithms.

2. Analysis of Combinatorial Models

File histories, as introduced in [Flajolet, Françon, and Vuillemin, 1979] and [Flajolet, Françon, and Vuillemin, 1980], model the evolution of several classical types of dynamic data structures, including priority queues (PQ), dictionaries (D), symbol tables (ST), stacks (S), and linear lists (LL). The data structures are treated as combinatorial objects; their performance characteristics are determined by the *relative* order of the elements they contain, not by the actual values of the elements. Thus, we say that there are $k + 1$ ways of inserting a new element into a dictionary of size k , since there are $k + 1$ “gaps” where the new element can fit in, relative to the k elements already present. The evolution of the data structure is represented as a path in $\mathbf{Z}^2$, where the x -coordinate counts the number of operations, whether they be insertions, deletions or queries, and the y -coordinate counts the size. Each step is of the type $(a, b) \rightarrow (a + 1, b \pm 1)$ (insertion or deletion) or $(a, b) \rightarrow (a + 1, b)$ (positive or negative query). To each step we associate a certain choice among the possibilities. For example, in priority queues, deletions can be performed only for the minimum element, so the number of possibilities for a deletion is 1. The probability model is that all possibilities are equally likely, with the constraint that the data structure is empty initially and at the end. The following table summarizes the number of possibilities for each type of

data structure and operation, in terms of the current data structure size k :

	PQ	D	LL	ST	S
insertions	$k + 1$	$k + 1$	$k + 1$	$k + 1$	1
deletions	1	k	k	1	1
positive queries	0	k	0	k	0
negative queries	0	$k + 1$	0	0	0

As an introduction to our first method, let us consider the combinatorial model of file histories corresponding to the size of priority queues (PQ). An equivalent formulation, as considered in [Van Wyk and Vitter, 1986], is to generate the $2n$ birthtimes and deathtimes of the n items as independent uniform random variates in the unit interval $[0, 1]$. The i th item is born at time $\min\{s_i, t_i\}$ and dies at time $\max\{s_i, t_i\}$. The average priority queue size $E(\text{Size}(t)) = 2nt(1 - t)$ varies parabolically in the unit interval and attains its maximum $n/2$ when $t = 1/2$, as shown in Figure 1.

This “peak” in the value of $E(\text{Size}(t))$ suggests that the value $\max_{t \in [0, 1]} \{\text{Size}(t)\}$ should be achieved in a neighborhood of $t = 1/2$ and thus should be $\sim n/2$. This was conjectured in [Van Wyk and Vitter, 1986]. In this section we introduce our first method and use it to prove the conjecture. The method will be developed further in the next section for the probabilistic models, where the expected values in question are flat and have no peaks like the ones considered in this section.

Theorem 2.1. *For priority queue and dictionary file histories of length $2n$, we have*

$$E(\max\{\text{Size}(t)\}) \sim \max_{t \in [0, 1]} \{E(\text{Size}(t))\} = \frac{n}{2}.$$

Theorem 2.2. *For priority queue and dictionary file histories of length $2n$, if $H = o(n)$ we have*

$$E(\max_{t \in [0, 1]} \{\text{Use}(t)\}) \sim \frac{n}{2}.$$

Proof of Theorem 2.1. First we prove Theorem 2.1 for priority queues. Since we have $E(\max_{t \in [0, 1]} \{\text{Size}(t)\}) \geq E(\text{Size}(\frac{1}{2})) = n/2$, our main problem is to show the other direction, namely, that $E(\max_{t \in [0, 1]} \{\text{Size}(t)\}) \lesssim E(\text{Size}(\frac{1}{2}))$.*

We consider the neighborhood $\mathcal{V} = (\frac{1}{2} - \epsilon, \frac{1}{2} + \epsilon)$ of $t = \frac{1}{2}$, where $\epsilon = n^{-1/6}$, as pictured in Figure 1. We shall prove that with high probability the maximum M of $\text{Size}(t)$ is reached for $t \in \mathcal{V}$. Let v denote the number of births and deaths in interval $\mathcal{V}$. If the maximum is reached inside $\mathcal{V}$, then its value is at most $\text{Size}(\frac{1}{2}) + v$; thus we have

$$E(M) \leq E(\text{Size}(\frac{1}{2})) + E(v) + n \cdot \Pr\{M \text{ reached outside } \mathcal{V}\}. \quad (2.1)$$

The first two terms are clearly equal to $n/2 + 4n\epsilon \sim n/2$. All that remains is to prove that $\Pr\{M \text{ reached outside } \mathcal{V}\} = o(1)$. By symmetry, this probability can be bounded by

$$2 \cdot \Pr\{\exists t \in [0, \frac{1}{2} - \epsilon], \text{Size}(t) > \text{Size}(\frac{1}{2})\}. \quad (2.2)$$

* We adopt the notation $f(n) \lesssim g(n)$ (as $n \rightarrow \infty$) if there is a function $h(n)$ such that $h(n) \sim g(n)$ (as $n \rightarrow \infty$) and $f(n) \leq h(n)$ for all $n \geq 1$.




Figure 1. Graph of $E(\text{Size}(t))$, as a function of t , for the combinatorial model of priority queues considered in Section 2. The quantity $\text{Size}(t)$ has a geometric interpretation as the number of items (horizontal intervals) “cut” by a vertical line (sweep line) at position t . (The graph is bell-shaped, and this makes the analysis of $E(\max_{t \in [0,1]} \{\text{Size}(t)\})$ easier; however, this is not the case for the “flat” distributions of the probabilistic models in Section 3.)

The problem is that the values of $\text{Size}(t)$ at two different t are clearly not independent. In order to get around this problem, we divide $[0, \frac{1}{2} - \epsilon]$ into $\ell = n(\frac{1}{2} - \epsilon)$ equal-sized intervals $[t_i, t_{i+1}]$, with endpoints $t_i = i/n$, for $0 \leq i \leq \ell$. Let $a = n/2 - n^{26/50}$ and $a' = n/2 - 2n^{26/50}$. We have

$$\begin{aligned}
 & \Pr\{\exists t \leq \tfrac{1}{2} - \epsilon, \text{Size}(t) > \text{Size}(\tfrac{1}{2})\} \\
 & \leq \Pr\{\text{Size}(\tfrac{1}{2}) < a\} + \Pr\{\exists t \leq \tfrac{1}{2} - \epsilon, \text{Size}(t) > a\} \\
 & \leq \Pr\{\text{Size}(\tfrac{1}{2}) < a\} + \sum_{0 \leq i \leq \ell} \Pr\{\text{Size}(t_i) > a'\} + \sum_{0 \leq i \leq \ell} \Pr\{x_i > a - a'\} \quad (2.3)
 \end{aligned}$$

where x_i is the number of births in interval $[t_i, t_{i+1}]$. Since $\Pr\{\text{Size}(t_i) > a'\}$ is maximized at $t_i = \frac{1}{2} - \epsilon$, and since $x_1, x_2, \dots$ are identically distributed, we get

$$\begin{aligned}
 & \Pr\{\exists t \leq \tfrac{1}{2} - \epsilon, \text{Size}(t) > \text{Size}(\tfrac{1}{2})\} \\
 & \leq \Pr\{\text{Size}(\tfrac{1}{2}) < a\} + \tfrac{n}{2} \cdot \Pr\{\text{Size}(\tfrac{1}{2} - \epsilon) > a'\} + \tfrac{n}{2} \cdot \Pr\{x_0 > a - a'\}. \quad (2.4)
 \end{aligned}$$

We analyze each term separately. For fixed t , the distribution of $\text{Size}(t)$ is well-known. It is binomially distributed with parameter $t(1-t)$:

$$\begin{aligned}\Pr\{\text{Size}(t) = k\} &= \binom{n}{k} (2t(1-t))^k (1-2t(1-t))^{n-k}; \\ \frac{E(\text{Size}(t))}{2} &= \{n\}; \\ \text{Var}(\text{Size}(t)) &= 2nt(1-t)(1-2t(1-t)).\end{aligned}$$

Thus Chebyshev's inequality yields

$$\Pr\{\text{Size}(\tfrac{1}{2}) < \tfrac{n}{2} - 2n^{26/50}\} < \frac{1}{4}n^{-2/50} = o(1).$$

As for the next term, we have

$$\Pr\{\text{Size}(\tfrac{1}{2} - \epsilon) > \tfrac{n}{2} - 2n^{26/50}\} = \sum_{k > n/2 - 2n^{26/50}} \binom{n}{k} (\tfrac{1}{2} - 2\epsilon^2)^k (\tfrac{1}{2} + 2\epsilon^2)^{n-k}.$$

The terms of the sum form a decreasing sequence, and the ratio between two successive terms is at most

$$\left(\frac{n/2 + 2n^{26/50}}{n/2 - 2n^{26/50}}\right) \left(\frac{\frac{1}{2} - 2\epsilon^2}{\frac{1}{2} + 2\epsilon^2}\right) = \left(\frac{\frac{1}{2} + 2n^{-24/50}}{\frac{1}{2} - 2n^{-24/50}}\right) \left(\frac{\frac{1}{2} - 2n^{-1/3}}{\frac{1}{2} + 2n^{-1/3}}\right) < 1.$$

The sum can thus be replaced by a geometric sum, and after computing its asymptotics using Stirling's formula, we find that

$$\frac{n}{2} \cdot \Pr\{\text{Size}(\tfrac{1}{2} - \epsilon) > a'\} \leq n^{5/6} e^{-8n^{1/3}} = o(1).$$

The third term is also computed easily, since

$$\Pr\{x_0 = k\} = \binom{2n}{k} \left(\frac{1}{n}\right)^k \left(1 - \frac{1}{n}\right)^{2n-k}.$$

Using Chebyshev's inequality, we find again that

$$\frac{n}{2} \Pr\{x_0 > n^{26/50}\} = o(1).$$

Thus the theorem is proved for priority queues.

The proof for dictionaries goes along the same lines. Let $H_{k,l,n}$ be the number of histories going from level k to level l in n steps. If we consider the size after m steps of a dictionary history going from level 0 to 0 in n steps, its distribution is given by

$$\Pr\{\text{Size}(m) = i\} = \frac{H_{0,i,m} H_{i,0,n-m}}{H_{0,0,n}}.$$

We know from [Flajolet Francon Vuillemin 79] that

$$\sum_{k,l,n} H_{k,l,n} u^k v^l \frac{z^n}{n!} = \frac{1}{1 - z(1+u)(1+v) - uv}.$$

We deduce that

$$\Pr\{Size(m) = i\} = \frac{\binom{m}{i} \binom{n-m}{i}}{\binom{n}{m}}.$$

The expected value is

$$E(Size(m)) = n \frac{m}{n} \left(1 - \frac{m}{n}\right).$$

Thus the graph of $E(Size(m))$, when m goes from 0 to n , describes a parabola, just as for priority queues. Since we know the distribution of $Size(m)$, we can from then on work out a proof exactly similar to the priority queue case (further details are omitted). ■

Proof of Theorem 2.2. For Theorem 2.2, as in Theorem 2.1, the lower bound $n/2$ applies. To get an upper bound, we use the bound

$$E\left(\max_{t \in [0,1]} \{Use(t)\}\right) \leq H \cdot E\left(\max_{t \in [0,1]} \{Use_1(t)\}\right) = H \cdot E\left(\max_{t \in [0,1]} \{Size_1(t)\}\right), \quad (2.5)$$

where $Size_1(t)$ is the number of living items present in the first bucket of the HwLD table at time t . The inequality in (2.5) follows from the uniformity of the hash function in HwLD, and the subsequent equality follows from the fact that the priority queue is initially empty and thus $Use_1(t)$ and $Size_1(t)$ attain the same maxima. Let n_1 be the number of living items that hash to bucket 1 in HwLD. We have

$$\Pr\{n_1 = k\} = \binom{n}{k} \left(\frac{1}{H}\right)^k \left(1 - \frac{1}{H}\right)^{n-k}. \quad (2.6)$$

We can now use the result of the previous theorem:

$$E\left(\max_{t \in [0,1]} \{Size_1(t)\}\right) \lesssim \sum_{k > n/2H} \left(\frac{k}{2} \cdot \Pr\{n_1 = k\}\right) + \frac{n}{2H} \cdot \Pr\{n_1 \leq \frac{n}{2H}\}, \quad (2.7)$$

from which we find after some elementary calculations that

$$E\left(\max_{t \in [0,1]} \{Size_1(t)\}\right) \lesssim \frac{n}{2H}. \quad (2.8)$$

By summing on the H buckets, we get $E(\max_{t \in [0,1]} \{Use(t)\}) \lesssim n/2$. The proof for dictionaries is similar. ■

These techniques can also be applied to the combinatorial models for other types of file histories, namely linear lists and symbol tables, to get similar results. (Note that the techniques cannot be used for stacks, because the random variable “maximum size” is not concentrated enough around its expected value in that case).

3. Analysis of the Probabilistic Models

In this section we consider continuous-time probabilistic models in equilibrium. The birthtimes of items are independent and form a Poisson process with intensity λ ; the probability of a birth during a time interval of length Δt is $\sim \lambda \Delta t$, as $\Delta t \rightarrow 0$. In the first model we consider, the lifetimes of the items are independently distributed according to an arbitrary distribution with mean $1/\mu$. An important special case is when the lifetimes satisfy the memoryless exponential distribution. This model and its special case are well-known and are referred to as the stationary M/G/ ∞ and M/M/ ∞ models.

The maximum is more difficult to analyze for the probabilistic models, since the expected value of $Size(t)$ is no longer “peaked” around a certain value of t . An easy analysis (see [Feller, 1968], for example) shows that for each t the distribution of $Size(t)$ is Poisson with mean λ/μ . But the maximum value of $Size(t)$ in each case is sufficiently concentrated about its mean so that our method is applicable.

Theorem 3.1. *In the stationary M/M/ ∞ model with birth rate λ and average lifetime $1/\mu$, we have, assuming either that $\mu \rightarrow 0$ or that $\mu = \Omega(1)$ and $\lambda \rightarrow \infty$,*

$$E\left(\max_{t \in [0,1]} \{Size(t)\}\right) \sim \begin{cases} \frac{\lambda}{\mu} & \text{if } f(\lambda, \mu) \rightarrow 0; \\ d \frac{\lambda}{\mu} & \text{if } f(\lambda, \mu) \rightarrow c; \\ \frac{f(\lambda, \mu)}{\ln f(\lambda, \mu)} \frac{\lambda}{\mu} & \text{if } f(\lambda, \mu) \rightarrow \infty, \end{cases}$$

where $f(\lambda, \mu) = (\ln \lambda)/(\lambda/\mu)$ and the constant $d \geq 1$ is defined implicitly from the constant c by $d \ln d - d = c - 1$. In the more general M/G/ ∞ case, in which the lifetime distribution can be arbitrary, the asymptotic upper bounds hold; in the first case $\ln \lambda = o(\lambda/\mu)$, the corresponding lower bound is trivial, so we get asymptotic equality.

Theorem 3.2. *In the stationary M/G/ ∞ model, if $\ln \lambda = o(\lambda/\mu)$, we have*

$$E\left(\max_{t \in [0,1]} \{Use(t)\}\right) \sim \frac{\lambda}{\mu} + H.$$

The condition $\ln \lambda = o(\lambda/\mu)$ is typically met in practice in geometry applications, as in [Szymanski and Van Wyk, 1983]. Similar results for $E(\max_{t \in [0,1]} \{Use(t)\})$ hold as for cases 2 and 3 of Theorem 3.1, except that the conditions are more intricate.

Proof of Theorem 3.1. If $\mu \rightarrow 0$, then $\lambda = o(\lambda/\mu)$, and

$$Size(0) \leq \max_{t \in [0,1]} \{Size(t)\} \leq Size(0) + \#[0,1],$$

where $\#[0,1]$ denotes the number of births in $[0,1]$. Taking expectations we get

$$\frac{\lambda}{\mu} \leq E\left(\max_{t \in [0,1]} \{Size(t)\}\right) \leq \frac{\lambda}{\mu} + \lambda,$$

and hence $E(\max_{t \in [0,1]} \{Size(t)\}) \sim \lambda/\mu$.

From now on, we assume that $\lambda \rightarrow \infty$ and that $\mu \geq \alpha$, for some positive constant α . First we derive the upper bounds for $E(\max_{t \in [0,1]} \{Size(t)\})$. We use the basic identity

$$E(\max_{t \in [0,1]} \{Size(t)\}) = \sum_{k \geq 1} \Pr\{\max_{t \in [0,1]} \{Size(t)\} \geq k\}. \quad (3.1)$$

The probabilities in the sum form a decreasing sequence. We are going to show that $\max_{t \in [0,1]} \{Size(t)\}$ has a distribution concentrated near some value V (to be specified later). We trivially have

$$E(\max_{t \in [0,1]} \{Size(t)\}) \leq V(1 + 2\epsilon) + \sum_{k > V(1+2\epsilon)} \Pr\{\max_{t \in [0,1]} \{Size(t)\} \geq k\}. \quad (3.2)$$

The desired upper bound $E(\max_{t \in [0,1]} \{Size(t)\}) \lesssim V$ follows if we show for an adequate choice of $\epsilon \rightarrow 0$ that $\sum_{k > V(1+2\epsilon)} \Pr\{\max_{t \in [0,1]} \{Size(t)\} \geq k\} = o(1)$.

In order to evaluate the probabilities, we have to deal with the lack of independence of the successive values of $Size(t)$, as t goes from 0 to 1. We partition $[0, 1]$ into N intervals of equal size, $I_0 = [t_0, t_1)$, $I_1 = [t_1, t_2)$, $\dots$, $I_j = [t_j, t_{j+1})$, $\dots$. The number of intervals N will be defined to be λ in case 1 and $(\lambda/\epsilon V)^{1+\epsilon}$ in cases 2 and 3. The key point is that $\max_{t \in [0,1]} \{Size(t)\} > k$, where $k > V(1 + 2\epsilon)$, only if there is an interval endpoint t_i where $Size(t_i) \geq (k + V)/2$ or if one of the intervals has at least $(k - V)/2$ births:

$$\Pr\{\max_{t \in [0,1]} \{Size(t)\} \geq k\} \leq \Pr\{\exists j, Size(t_j) \geq \frac{k+V}{2}\} + \Pr\{\exists j, \#I_j \geq \frac{k-V}{2}\}, \quad (3.3)$$

where $\#I_j$ denotes the number of births during time interval I_j . By (3.2) we get

$$E(\max_{t \in [0,1]} \{Size(t)\}) \leq V(1 + 2\epsilon) + 2N \sum_{k > V(1+\epsilon)} \Pr\{Size(0) \geq k\} + 2 \sum_{k > \epsilon V} \Pr\{\exists j, \#I_j \geq k\}. \quad (3.4)$$

For the $M/M/\infty$ process, the random variables $Size(0)$ and $\#I_j$ are Poisson distributed with means λ/μ and λ/N , respectively. The rest of the proof consists of technical computations and approximations, with adequate choices for the parameters V , N and ϵ .

First we compute

$$S = 2 \sum_{k > \epsilon V} \Pr\{\exists j, \#I_j \geq k\} \leq 2N \sum_{k > \epsilon V} \Pr\{\#I_0 \geq k\}. \quad (3.5)$$

The inequality holds because the random variables $\#I_j$ are identically distributed. The Poisson probability function of $\#I_0$ is

$$\Pr\{\#I_0 = k\} = e^{-\lambda/N} \frac{(\lambda/N)^k}{k!}. \quad (3.6)$$

The ratio between consecutive probabilities is $< \lambda/(N\epsilon V) = o(1)$, for suitable choices of N , ϵ , and V . By (3.6) we get

$$\frac{\Pr\{\#I_0 = k+1\}}{\Pr\{\#I_0 = k\}} = \frac{\lambda}{N(k+1)} < \frac{\lambda}{N\epsilon V} = o(1),$$

for a suitable choice of N and V . Thus we have

$$\Pr\{\#I_0 \geq k\} = \sum_{j \geq k} \Pr\{\#I_0 = j\} \leq \frac{1}{1 - \lambda/(N\epsilon V)} \cdot \Pr\{\#I_0 = k\},$$

and by (3.5),

$$S \leq \frac{2N}{1 - \lambda/(N\epsilon V)} \sum_{k \geq \epsilon V} \Pr\{\#I_0 = k\} \leq \frac{2N}{(1 - \lambda/(N\epsilon V))^2} \cdot \Pr\{\#I_0 = \epsilon V\}.$$

Substituting (3.6), we get for large λ

$$S \leq 3Ne^{-\lambda/N} \frac{(\lambda/N)^{\epsilon V}}{(\epsilon V)!}.$$

If we pick V so that $V \geq \frac{\lambda}{\mu} \rightarrow \infty$, we can choose $\epsilon \rightarrow 0$ so that $\epsilon V \rightarrow \infty$. We now apply Stirling's approximation formula to get

$$S \leq \frac{3N}{\sqrt{\epsilon V}} e^{-\lambda/N} \left(\frac{\lambda e}{N\epsilon V} \right)^{\epsilon V}. \quad (3.7)$$

Our choices of N and V for the three cases of Theorem 3.1 are as follows:

Case 1. Assuming that $\ln \lambda = o(\frac{\lambda}{\mu})$, we fix $N = \lambda$ and $V = \frac{\lambda}{\mu}$. We find if $\epsilon \rightarrow 0$ slowly enough that $S = o(1)$.

Case 2. Assuming that $\ln \lambda = c\frac{\lambda}{\mu}$, where c is a positive constant, we fix

$$N = \left(\frac{\lambda}{\epsilon V} \right)^{1+\epsilon} \quad \text{and} \quad V = d\frac{\lambda}{\mu}(1 + \sqrt{\epsilon}),$$

where d is the solution of $d \ln d - d = c - 1$. For $\epsilon \rightarrow 0$ slowly enough we find that $S = o(1)$.

Case 3. Assuming that $f(\lambda, \mu) = (\ln \lambda)/(\lambda/\mu) \rightarrow \infty$, we fix

$$N = \left(\frac{\lambda}{\epsilon V} \right)^{1+\epsilon} \quad \text{and} \quad V = \frac{f}{\ln f} \frac{\lambda}{\mu}.$$

The analysis works as in case 2 to show that $S = o(1)$.

We now turn our attention to bounding the other sum in (3.4), namely,

$$P = N \sum_{k > V(1+\epsilon)} \Pr\{Size(0) \geq k\}. \quad (3.8)$$

We want to show that $P = o(1)$. The random variable $Size(0)$ has the Poisson distribution:

$$\Pr\{Size(0) = k\} = e^{-\lambda/\mu} \frac{(\lambda/\mu)^k}{k!}. \quad (3.9)$$

The ratio between two successive terms is at most $(\lambda/\mu)/(V(1+\epsilon))$ when $k > V(1+\epsilon)$. Thus we can write

$$P \leq \frac{N}{\left(1 - \frac{\lambda}{\mu}/(V(1+\epsilon))\right)^2} \cdot \Pr\{Size(0) = V(1+\epsilon)\}. \quad (3.10)$$

By using (3.9), the fact that $V \geq \frac{\lambda}{\mu}$ in all cases, and Stirling's formula, we get for large V

$$P \leq \frac{N}{\epsilon^2} e^{-\lambda/\mu} \frac{(\lambda/\mu)^{V(1+\epsilon)}}{(V(1+\epsilon))!} \leq \frac{N}{\epsilon^2} \frac{e^{-\lambda/\mu}}{\sqrt{V}} \left(\frac{e\lambda/\mu}{V(1+\epsilon)} \right)^{V(1+\epsilon)}.$$

We can now evaluate this expression for all three cases of the theorem, and we find that, if ϵ goes to 0 slowly enough, we have in all cases $P = o(1)$, which concludes the proof of the upper bound of the theorem.

For the lower bound, let us for clarity restrict ourselves to case 3 of the theorem, and assume that $\lambda = \mu$ (the proof in case 2 is similar). The above proof yields $E(\max_{t \in [0,1]} \{Size(t)\}) \leq ((\ln \lambda)/\ln \ln \lambda)(1+\epsilon)$, for some positive $\epsilon = o(1)$. We shall now show that the reverse also holds, namely, that $E(\max_{t \in [0,1]} \{Size(t)\}) \geq ((\ln \lambda)/\ln \ln \lambda)(1-\epsilon)$, for some positive $\epsilon = o(1)$. We shall prove this by starting with (3.1) and showing that

$$\Pr\left\{\max_{t \in [0,1]} \{Size(t)\} \geq k\right\} \sim 1, \quad \text{for } k \leq \frac{\ln \lambda}{\ln \ln \lambda}(1-\epsilon). \quad (3.11)$$

We once again partition $[0, 1]$ into $\ell = \lambda^{1-\epsilon}$ equal-sized intervals, with endpoints $t_i = i/\lambda^{1-\epsilon}$, for $0 \leq i \leq \ell$. We have

$$\begin{aligned} \Pr\left\{\max_{t \in [0,1]} \{Size(t)\} < k\right\} &= \sum_{j_1, \dots, j_\ell < k} \Pr\{\forall i, Size(t_i) = j_i\} \\ &= \sum_{j_1, \dots, j_\ell < k} \prod_{1 \leq i \leq \ell} \Pr\{Size(t_i) = j_i \mid Size(t_{i-1}) = j_{i-1}\}. \end{aligned} \quad (3.12)$$

The motivation for our choice of interval size is to have enough births and deaths in each interval so that the values of $Size(t)$ at the endpoints are “sufficiently independent.” Let $P_n(t)$ denote $\Pr\{Size(t) = n \mid Size(0) = j_{i-1}\}$. We define the generating function $P(s, t) = \sum_n P_n(t) s^n$, which is equal to

$$P(s, t) = e^{-(1-s)(1-e^{-\mu t})} (1 - (1-s)e^{-\mu t})^{j_{i-1}} \quad (3.13)$$

(cf. [Feller, 1968]). The conditional probability term in (3.12) is $P_{ji}(1/\lambda^{1-\epsilon}) = \langle s^{ji} \rangle P(s, 1/\lambda^{1-\epsilon})$. By extracting the coefficient of s^{ji} in (3.13) and using asymptotic approximations, we find that there is “sufficient independence”:

$$\Pr\left\{\max_{t \in [0,1]} \{Size(t)\} < k\right\} \sim (\Pr\{Size(0) < k\})^{\lambda^{1-\epsilon}} \sim o(1), \quad (3.14)$$

for $k \leq ((\ln \lambda)/\ln \ln \lambda)(1 - \epsilon)$. By letting $\epsilon \rightarrow 0$ at an appropriate rate, we prove our goal (3.11), which completes the proof. ■

Proof of Theorem 3.2. The lower bound here is easy:

$$E\left(\max_{t \in [0,1]} \{Use(t)\}\right) \geq E(Use(0)) = \frac{\lambda}{\mu} + H.$$

We prove the upper bound when $\ln \lambda = o(\frac{\lambda}{\mu})$. First, we consider the $M/M/\infty$ case. From [Van Wyk and Vitter, 1986], the stationary distribution of $Use(t)$ is

$$\Pr\{Use(t) = k\} = e^{-\lambda/\mu} \frac{(\lambda/\mu)^{k-H}}{(k-H)!}. \quad (3.15)$$

We partition the interval $[0, 1]$ into λ intervals of equal size, $I_j = [t_j, t_{j+1})$. Let $V = \frac{\lambda}{\mu} + H$. If $k > V(1 + 2\epsilon)$, we have, with the same techniques used for the previous theorem

$$\Pr\left\{\max_{t \in [0,1]} \{Use(t)\} \geq k\right\} \leq \lambda \cdot \Pr\left\{Use(0) \geq \frac{k+V}{2}\right\} + \Pr\left\{\exists j, \#I_j \geq \frac{k-V}{2}\right\}. \quad (3.16)$$

And the same approximations as above show that when $\ln \lambda = o(\lambda/\mu)$ we have

$$E\left(\max_{t \in [0,1]} Use(t)\right) \lesssim \frac{\lambda}{\mu} + H.$$

In order for the proof to work in the $M/G/\infty$ model, we need only show that the stationary distribution of $Use(t)$ is the same as in the $M/M/\infty$ model. We shall compute the probability

$$p_{m,n}(t) = \Pr\{Size(t) = m, Use(t) = m + n \mid Size(0) = Use(0) = 0\} \quad (3.17)$$

and let $t \rightarrow \infty$ to get the stationary distribution. We assume for simplicity that $H = 1$.

$$p_{m,n}(t) = \int_0^t \lambda e^{-\lambda t-x} \cdot \Pr\{Size(x) = m + n - 1 \text{ and } n \text{ deaths in } (x, t)\} dx. \quad (3.18)$$

We assume that $x > \sqrt{t}$; the lower part of the integral is negligible. We have

$$\Pr\{Size(x) = m + n - 1\} \sim e^{-\lambda/\mu} \frac{(\lambda/\mu)^{m+n-1}}{(m+n-1)!}. \quad (3.19)$$

Let $B(x)$ be the distribution of the service time. The probability that there are n deaths in (x, t) can be split into two terms, depending on whether the element born at time x dies before time t or not. The probability that a given element, alive at time x , dies before time t is equal to

$$p(x) = \frac{\frac{1}{x} \int_0^x (B(t-u) - B(x-u)) du}{1 - \frac{1}{x} \int_0^x B(x-u) du}. \quad (3.20)$$

Thus we find that $p_{m,n}(t)$ is asymptotically equal to

$$\begin{aligned} \int_{\sqrt{t}}^t \lambda e^{-\lambda t-x} e^{-\lambda/\mu} \frac{(\lambda/\mu)^{m+n-1}}{(m+n-1)!} & \left((1-B(t-x)) \binom{m+n-1}{n} p(x)^n (1-p(x))^{m-1} \right. \\ & \left. + B(t-x) \binom{m+n-1}{n-1} p(x)^{n-1} (1-p(x))^m \right) dx. \end{aligned}$$

With standard asymptotics, we get

$$\Pr\{Use(t) = s\} = \sum_{m+n=s} p_{m,n}(t) \sim \int_{\sqrt{t}}^t \lambda e^{-\lambda t-x} e^{-\lambda/\mu} \frac{(\lambda/\mu)^{s-1}}{(s-1)!} dx \sim e^{-\lambda/\mu} \frac{(\lambda/\mu)^{s-1}}{(s-1)!}. \quad (3.21)$$

The case for general H is similar. Thus the stationary distribution of $Use(t)$ is the same for $M/G/\infty$ and $M/M/\infty$ processes, and Theorem 3.2 is proved. ■

It is worthwhile noting that Theorem 3.1 derives results in queuing theory, using non-queueing theory techniques from the analysis of algorithms. By a simple change of scale, we can extend the range over which we take the maximum from the unit interval $[0, 1]$ to $[0, T]$. The last subcase of Theorem 3.1 says that, if λ and μ are constant, then $E(\max_{t \in [0, T]} \{Size(t)\}) \sim \ln T / \ln \ln T$, which was a longstanding open problem. Our method also applies to the $M/M/1$ model: For constants λ and μ with $\lambda/\mu = c < 1$, we have $E(\max_{t \in [0, T]} \{Size(t)\}) \sim -\ln T / \ln c$, which previously had been proved only by Brownian motion techniques [Fayolle, 1988].

Theorem 3.3. *In the stationary $M/M/1$ model, with birth rate λ and average lifetime $1/\mu$, we have*

$$E\left(\max_{t \in [0, 1]} \{Size(t)\}\right) \sim \begin{cases} 1 & \text{if } \lambda = o(\mu^{2/3}); \\ \frac{\ln \lambda}{-\ln c} & \text{if } \lambda/\mu \rightarrow c < 1, \lambda \rightarrow \infty. \end{cases}$$

Proof of Theorem 3.3. Let us restrict ourselves to case 2 (the more difficult case). The stationary distribution of the size of an $M/M/1$ process is well-known (see [Feller, 1968]).

$$\Pr\{Size(t) = k\} = \frac{\left(\frac{\lambda}{\mu}\right)^k}{1 - \frac{\lambda}{\mu}}.$$

For the upper bound, we divide $[0, 1]$ into λ intervals $I_i = [t_i, t_{i+1}[$, with $t_i = i/\lambda$, and use the same technique as before:

$$\begin{aligned} \Pr\left\{\max_{t \in [0,1]} \{Size(t)\} \geq k\right\} &\leq \lambda \cdot \Pr\{Size(0) \geq \frac{k+k_0}{2}\} \\ &\quad + \lambda \cdot \Pr\{\#I_0 \geq \frac{k-k_0}{2}\}. \end{aligned}$$

Fixing $k_0 = \ln \lambda / (-\ln c)$, we find after some calculations that

$$E\left(\max_{t \in [0,1]} \{Size(t)\}\right) \leq \frac{\ln \lambda}{-\ln c} (1 + o(1)).$$

The lower bound proof follows the proof in the $M/M/\infty$ case. ■

4. Time Hashing: the Discrete Analog

In this section we analyze $\max_{0 \leq t \leq 1} \{Size(t)\}$ by relating the problem to its discrete version—the maximum slot occupancy in hashing. The tricky part is handling the lack of independence of slot occupancies.

Theorem 4.1. *In the stationary $M/G/\infty$ model, with birth rate λ and average lifetime $1/\mu$, we have, assuming either that $\mu \rightarrow 0$ or that $\mu = \Omega(1)$ and $\lambda \rightarrow \infty$,*

$$E\left(\max_{t \in [0,1]} \{Size(t)\}\right) \lesssim \begin{cases} \frac{\lambda}{\mu} & \text{if } f(\lambda, \mu) \rightarrow 0; \\ d \frac{\lambda}{\mu} & \text{if } f(\lambda, \mu) \rightarrow c; \\ \frac{2f(\lambda, \mu)}{\ln f(\lambda, \mu)} \frac{\lambda}{\mu} & \text{if } f(\lambda, \mu) \rightarrow \infty, \end{cases}$$

where $f(\lambda, \mu) = (\ln \lambda)/(\lambda/\mu)$ and the constant $d \geq 1$ is defined implicitly from the constant c by $d \ln d - d = 2c - 1$. When $f(\lambda, \mu) \rightarrow 0$, we have asymptotic equality.

It is interesting to note the close correspondence between the above formulas and the formulas for the maximum bucket occupancy in hashing, given in [Kolchin, Sevast'yanov, and Chistyakov, 1978] (cf. Section 4). The formulas for cases 2 and 3 of Theorem 4.1 are slightly weaker than the corresponding bounds in Theorem 3.1.

Theorem 4.2. *In the stationary $M/G/\infty$ model, with birth rate λ and average lifetime $1/\mu$, we have, assuming either that $\mu \rightarrow 0$ or that $\lambda = o(H)$ or that $\mu = \Omega(1)$ and $\lambda \rightarrow \infty$,*

$$E\left(\max_{t \in [0,1]} \{Use(t)\}\right) \lesssim \begin{cases} \frac{\lambda}{\mu} + H & \text{if } f(\lambda, \mu) \rightarrow 0; \\ d \frac{\lambda}{\mu} + H & \text{if } f(\lambda, \mu) \rightarrow c; \\ \frac{2f(\lambda, \mu)}{\ln f(\lambda, \mu)} \frac{\lambda}{\mu} + H & \text{if } f(\lambda, \mu) \rightarrow \infty, \end{cases}$$

where $f(\lambda, \mu) = (\ln \lambda)/(\lambda/\mu)$ and the constant $d \geq 1$ is defined implicitly from the constant c by $d \ln d - d = 2c - 1$. When $f(\lambda, \mu) \rightarrow 0$, we have asymptotic equality.

Theorem 4.3. *In the stationary $M/G/\infty$ model, with birth rate λ and average lifetime $1/\mu$, assuming that $\ln \lambda = o(H)$, $\mu = \Omega(1)$, and $\lambda \rightarrow \infty$, we have*

$$E\left(\max_{t \in [0,1]} \{Use(t)\} - \max_{t \in [0,1]} \{Size(t)\}\right) \sim H.$$

Results similar to those in cases 2 and 3 of Theorems 4.1 and 4.2 also hold for $E(\max_{t \in [0,1]} \{Use(t)\} - \max_{t \in [0,1]} \{Size(t)\})$, except that the conditions are more complicated.

An approach called “time hashing” was introduced in [Mathieu and Vitter, 1988] to give optimum bounds to within a constant factor for $E(\max_{t \in [0,1]} \{Size(t)\})$ when $f(\lambda, \mu) \rightarrow 0$ and $E(\max_{t \in [0,1]} \{Use(t)\} - \max_{t \in [0,1]} \{Size(t)\})$ when $H \geq (\ln \lambda)^{1+\epsilon}$, for constant $\epsilon > 0$. The approach we use here to show that the constant factors are in fact 1 is the “inverse” of the approach used in [Mathieu and Vitter, 1988], so a brief explanation of the former technique is called for.

For example, in the analysis of $E(\max_{0 \leq t \leq 1} \{Size(t)\})$ when $\ln \lambda = o(\lambda/\mu)$ in [Mathieu and Vitter, 1988], all the items that are alive for at least some time in $[0, 1]$ are considered. There are stages $k = 0, 1, 2, \dots, K$, and each stage has an associated hash table. For $0 \leq k \leq K$, all items (intervals) born in $(-\frac{1}{\mu}2^k, 1]$ with lifespan in the range $(\frac{1}{\mu}2^{k-1}, \frac{1}{\mu}2^k]$ are put into stage k ; in addition, all such items with lifespan $\leq \frac{1}{2\mu}$ are put into stage 0. Each stage consists of a hash table of $\mu 2^{-k} + 1$ slots. The j th slot, for $0 \leq j \leq \mu 2^{-k}$, represents the interval of time $(\frac{1}{\mu}(j-1)2^k, \frac{1}{\mu}j2^k]$. An item in stage k is placed into the slot corresponding to its birthtime. A special stage $K+1$ is constructed to store all the items that do not fit into one of the earlier stages. The parameter K is chosen large enough so that the number of items in stage $K+1$ is $O(\lambda/\mu)$. The important link between this discrete version of the problem and the original continuous one is the following relation:

$$\max_{0 \leq t \leq 1} \{Size(t)\} \leq 2 \sum_{0 \leq k \leq K+1} \max_{0 \leq j \leq \mu 2^{-k}} \{N_k(j)\},$$

where $N_k(j)$ denotes the number of items in the j th slot of stage k . The bound $\max_{0 \leq t \leq 1} \{Size(t)\} = O(\lambda/\mu)$ was proved by showing that $\max_{0 \leq j \leq \mu 2^{-k}} \{N_k(j)\} = O(\lambda/(2^k \mu))$, under the assumption that $\ln \lambda = O(\lambda/\mu)$.

Proof of Theorem 4.1. We shall prove Theorem 4.1 simultaneously for all three cases, $f(\lambda, \mu) \rightarrow 0$, $\rightarrow c$, and $\rightarrow \infty$, by showing that $E(\max_{0 \leq t \leq 1} \{Size(t)\}) \lesssim d\lambda/\mu$, where $d = d(\lambda, \mu)$ is the solution of the equation $d \ln d - d = 2f - 1$. In case 1, for example, we have $d \sim 1$, and in case 3 we have $d \sim 2f/\ln f$.

Note that for case 1 this upper bound will prove the asymptotic result claimed in Theorem 4.1, namely, $E(\max_{0 \leq t \leq 1} \{Size(t)\}) \sim \frac{\lambda}{\mu}$, since $E(\max_{0 \leq t \leq 1} \{Size(t)\}) \geq Size(0) = \lambda/\mu$.

If $\mu \rightarrow 0$, then the result follows immediately, as noted in Section 3. So we assume that $\lambda \rightarrow \infty$ and $\mu = \Omega(1)$. We “invert” the process used in [Mathieu and Vitter, 1988] to prove the big-oh bound, as shown in Figure 2. Instead of letting the stages of time hashing grow coarser and coarser, we consider the limiting case in which the slots represent smaller and smaller units of time relative to the item sizes. We use only one




Figure 2. Typical items that contribute an entry to slot j in the time hashing table.

stage, but the number of slots varies with λ and μ . We use $m = g\mu$ slots, for any $g \rightarrow \infty$. (Without loss of generality we choose g so that $g\mu$ is an integer.) The j th slot, for $1 \leq j \leq g\mu$, represents the time interval $((j-1)/(g\mu), j/(g\mu)]$. For each item we place an entry into each slot whose associated time interval intersects the item's lifetime. If we define $N(j)$ to be the slot occupancy of slot j , it is easy to see that the following upper bound holds:

Lemma 4.1. *We have*

$$\max_{0 \leq t \leq 1} \{Size(t)\} \leq \max_{1 \leq j \leq g\mu} \{N(j)\}.$$

To prove Theorem 4.1 it suffices to show that $E(\max_{1 \leq j \leq g\mu} \{N(j)\}) \lesssim d\lambda/\mu$.

Lemma 4.2. *The slot occupancies $N(j)$ are Poisson distributed with mean n/m , where $n = \lambda(g+1)$ is the average number of items inserted into the hash table of $m = f\mu$ slots.*

Proof. The slot occupancy $N(j)$ of the j th slot is equal to the sum of two independent quantities: the number of items living at time $(j-1)/(g\mu)$ plus the number of items born during the time interval $((j-1)/(g\mu), j/(g\mu)]$. These two quantities are Poisson-distributed with means λ/μ and $\lambda/(g\mu)$, respectively. Hence, $N(j)$ is Poisson-distributed with mean $\frac{\lambda}{\mu}(1 + \frac{1}{g}) = n/m$. ■

The following lemma is useful for studying the maximum slot occupancy in time hashing, because the random variables X_j are not required to be independent.

Lemma 4.3 [Mathieu and Vitter, 1988]. *For random variables $X_1, \dots, X_m$, if we have $\Pr\{X_j > b\} \leq 1/(nm)$, for all $1 \leq j \leq m$, where $n = E(\sum_j X_j)$, then*

$$E(\max_{1 \leq j \leq m} \{X_j\}) \leq b + \frac{1}{n} \cdot E(\max_{1 \leq j \leq m} \{X_j\} \mid \max_{1 \leq j \leq m} \{X_j\} > b).$$

Proof. We condition the expectation based upon if $\max_{1 \leq j \leq m} \{X_j\}$ is $\leq b$ or $> b$:

$$\begin{aligned} E(\max_{1 \leq j \leq m} \{X_j\}) &\leq b \cdot \Pr\{\max_{1 \leq j \leq m} \{X_j\} \leq b\} \\ &\quad + E(\max_{1 \leq j \leq m} \{X_j\} \mid \max_{1 \leq j \leq m} \{X_j\} > b) \cdot \Pr\{\max_{1 \leq j \leq m} \{X_j\} > b\}. \end{aligned}$$

The first probability term is bounded trivially by 1, and the second is bounded by

$$\Pr\{\max_{1 \leq j \leq m} \{X_j\} > b\} \leq \Pr\{X_1 > b\} + \Pr\{X_2 > b\} + \cdots + \Pr\{X_m > b\} \leq m \cdot \frac{1}{nm} = \frac{1}{n}. \quad \blacksquare$$

To apply Lemma 4.3 to prove Theorem 4.1, we use $X_j = N(j)$, $b = d_\mu^\lambda(1 + \frac{1}{g})$, $n = \lambda(g + 1)$, and $m = g\mu$, where $g \rightarrow \infty$ slowly. (In particular, we require that $g \leq (d - 1)\lambda/\mu$.) We have

$$\begin{aligned} nm \cdot \Pr\{N(j) > b\} &= nm e^{-n/m} \sum_{k > b} \frac{(n/m)^k}{k!} \\ &\leq \frac{n^2 e^{-n/m} (n/m)^b}{b! (b + 1 - n/m)} \\ &\leq \frac{g^2}{\sqrt{2\pi(\lambda/\mu)d(1 + 1/g)}(d - 1)_\mu^\lambda(1 + \frac{1}{g})} \cdot \frac{\lambda^2 e^{(d-1)(\lambda/\mu)(1+1/g)}}{d^{d(\lambda/\mu)(1+1/g)}} \\ &\leq \frac{\lambda^2 e^{(d-1)(\lambda/\mu)(1+1/g)}}{d^{d(\lambda/\mu)(1+1/g)}}, \end{aligned} \quad (4.1)$$

for large λ , by Stirling's formula. Taking logarithms of (4.1) and using the definition of d , we get

$$\ln(nm \cdot \Pr\{N(j) > b\}) \leq 2 \ln \lambda - \frac{\lambda}{\mu} \left(1 + \frac{1}{g}\right) (d \ln d - d + 1) = -\frac{2 \ln \lambda}{g} < 0. \quad (4.2)$$

This implies that the left-hand side of (4.1) is ≤ 1 for large λ , and hence the conditions for Lemma 4.3 are satisfied.

Lemma 4.3 gives us

$$\begin{aligned} \max_{1 \leq j \leq f\mu} \{N(j)\} &\leq d_\mu^\lambda \left(1 + \frac{1}{g}\right) \\ &\quad + \frac{1}{n} E(\max_{1 \leq j \leq g\mu} \{N(j)\} \mid \max_{1 \leq j \leq g\mu} \{N(j)\} > d_\mu^\lambda(1 + \frac{1}{g})). \end{aligned} \quad (4.3)$$

The random variables $N(j)$ are not independent, but the conditional expectation on the right-hand side of (4.3) can be bounded by

$$\begin{aligned} E(\max_{1 \leq j \leq g\mu} \{N(j)\} \mid N(1) > d_\mu^\lambda(1 + \frac{1}{g})) \\ \leq n + E(N(1) \mid N(1) > d_\mu^\lambda(1 + \frac{1}{g})) = O(n). \end{aligned} \quad (4.4)$$

Plugging (4.4) back into (4.3) gives us $E(\max_{1 \leq j \leq f\mu} \{N(j)\}) \lesssim d_\mu^\lambda$, which proves Theorem 4.1. $\blacksquare$

Proof of Theorem 4.2. Theorem 4.2 can be proved in an identical way to Theorem 4.1. For each t , $Use(t) - H$ is Poisson distributed with mean λ/μ [Mathieu and Vitter, 1988]. The techniques in the proof can then be applied to $Use(t) - H$ instead of to $Size(t)$.

For example, let us define the time hashing as in the proof of Theorem 4.1, except that we account for lazy deletion. For each item, we place an entry into each slot whose associated time interval intersects the item's presence in the HwLD data structure. We define $N(j)$ to be the slot occupancy of slot j . We have the corresponding versions of Lemmas 4.1 and 4.2:

Lemma 4.4. *We have*

$$\max_{0 \leq t \leq 1} \{Use(t)\} \leq \max_{1 \leq j \leq g\mu} \{N(j)\}.$$

To prove Theorem 4.2 it suffices to show that $E(\max_{1 \leq j \leq g\mu} \{N(j)\}) \lesssim d\lambda/\mu + H$.

Lemma 4.5. *The slot occupancies $N(j) - H$ are Poisson distributed with mean n/m , where $n = \lambda(g + 1)$ is the average number of items inserted into the hash table of $m = f\mu$ slots.*

The rest of the proof proceeds analogously as for Theorem 4.1. ■

The proof of Theorem 4.3 is more complicated and is omitted for brevity. It uses the techniques developed in [Mathieu and Vitter, 1988]. We bound $E(\max_{t \in [0,1]} \{Use(t) - Size(t)\})$ using Lemma 4.3 and an application of Chernoff's bound, which in turn gives us a bound for $E(\max_{t \in [0,1]} \{Use(t)\} - \max_{t \in [0,1]} \{Size(t)\})$, as desired.

5. Conclusions

We have developed two probabilistic methods that are useful for the analysis of the distribution of the maximum size of data structures. We get the asymptotic value of the expected maximum of $Size(t)$ and $Use(t)$ for several different combinatorial and probabilistic models of insertion and deletion. This solves the open problems from [Van Wyk and Vitter, 1986], [Morrison, Shepp, and Van Wyk, 1987], and [Mathieu and Vitter, 1988], as well as a longstanding open problem from queueing theory.

In our first method we isolate the primary contribution to the maximum and bound the lesser effects. Our second technique, which we use for a continuous model, takes advantage of the close connections between the model and its discrete analog, namely, the maximum slot occupancy in hashing. These methods can be used to get estimates of second-order terms and higher moments of the expected maximum, as well as estimates of the shape of the distribution of the maximum. The techniques also appear directly applicable to the study of the maximum size of other dynamic data structures, such as quad trees, k -d trees, and radix-exchange tries.

References

- W. Feller. *An Introduction to Probability Theory and Its Applications*. Volume 1. Wiley, New York (third edition 1968).
- G. Fayolle. Personal communication (1988).
- P. Flajolet. "Analyse d'algorithmes de manipulation d'arbres et de fichiers," *Cahiers du bureau universitaire de recherche opérationnelle*, **34–35** (1981), 1–209.
- P. Flajolet, J. Françon, and J. Vuillemin. "Computing Integrated Costs of Operations with Applications to Dictionaries," *Proceedings of the 11th Annual ACM Symposium on Theory of Computing*, Atlanta (April–May 1979), 49–61.
- P. Flajolet, J. Françon, and J. Vuillemin. "Sequence of Operations Analysis for Dynamic Data Structures," *Journal of Algorithms*, **1**(2) (1980), 111–141. A shortened version appeared in *Proceedings of the 20th Annual IEEE Symposium on Foundations of Computer Science*, Puerto Rico (October 1979), 183–195.
- S. Karlin and J. L. McGregor. "The Differential Equations of Birth-and-Death Processes, and the Stieltjes Moment Problem," *Trans. of the American Mathematical Society*, **85** (1957), 489–546.
- L. Kleinrock. *Queueing Systems*. Volume I: *Theory*. Wiley & Sons, New York (1975).
- V. F. Kolchin, B. A. Sevast'yanov, and V. P. Chistyakov. *Random Allocations*. V. H. Winston & Sons, Washington (1978).
- C. M. Mathieu and J. S. Vitter. "Maximum Queue Size and Hashing with Lazy Deletion," *Proceedings of the 20th Annual Symposium on the Interface of Computing Science and Statistics*, Reston, VA (April 1988).
- C. M. Mathieu and J. S. Vitter. "General Methods for the Analysis of the Maximum Size of Dynamic Data Structures," *Proceedings of the 16th International Colloquium on Automata, Languages and Programming*, Stresa, Italy (July 1989).
- J. Morrison, L. A. Shepp, and C. J. Van Wyk. "A Queueing Analysis of Hashing with Lazy Deletion," *SIAM Journal on Computing* **16**, 6 (December 1987), 1155–1164.
- T. Ottmann and D. Wood. "Space-Economical Plane-Sweep Algorithms," *Computer Vision, Graphics, and Image Processing*, **34** (1986), 35–51.
- T. G. Szymanski and C. J. Van Wyk. "Space-Efficient Algorithms for VLSI Artwork Analysis," *Proceedings of the 20th IEEE Design Automation Conference* (1983), 743–749.
- C. J. Van Wyk and J. S. Vitter. "The Complexity of Hashing with Lazy Deletion," *Algorithmica*, **1**(1) (1986), 17–29.