Performance evaluation of fragmented structures: A theoretical study

Ed Coffman [a,b], Robert Margolies [b,*], Peter Winkler [c], Gil Zussman [b]

[a] *Computer Science Department, Columbia University, New York, NY, USA*
[b] *Electrical Engineering Department, Columbia University, New York, NY, USA*
[c] *Mathematics and Computer Science, Dartmouth College, Hanover, NH, USA*

ARTICLE INFO

ABSTRACT

Modern computers with hard-disk storage and networks with dynamic spectrum access illustrate systems having a resource $\mathcal{R}$ that allows fragmented allocations. We model $\mathcal{R}$ as a sequence of $M > 1$ units for which requests are made by items in a FIFO queue; each request is for a desired amount of $\mathcal{R}$, the item size, and the residence time during which it is needed. So long as there are enough currently available units of $\mathcal{R}$, an item at the head of the queue can be divided into fragments accommodated by the gaps in $\mathcal{R}$ formed by these units. Under the key assumption that allocations given to items cannot be changed prior to their departures, fragmentation in the form of alternating gaps and allocated resource builds up randomly as items come and go. The improvements in resource utilization created by fragmentation are acquired at a processing cost, so how fragmentation evolves is an important performance issue.

We define a probability model specifying distributions for item sizes and residence times, and then analyze the system operating at capacity. If M is much larger than the maximum item size, then as the fragmentation process approaches equilibrium, nearly all of the allocated items are completely fragmented, i.e., the occupied units are mutually disjoint. In a suite of four theorems, we show how this result specializes for certain classes of item-size distributions. However, as a counterpoint to these rather intimidating results, we present the findings of extensive experiments which show that the delays in reaching the inefficient states of nearly complete fragmentation can be surprisingly long, and hence that even moderately fragmented states are usually of no concern.

© 2014 Elsevier B.V. All rights reserved.

1. Introduction

Traditionally, fragmentation in one-dimensional dynamic resource allocation has referred to segments of a resource assigned to items alternating with gaps of unassigned, potentially wasted resource. A classic example occurs in dynamic storage allocation [1]; the resource is computer memory and item requests are for the storage of files. Files must be stored in *consecutive* units/slots of memory. Depending on the application, these could be a sequence of bytes, words, or pages. Fragmentation is created solely by the random arrival and departure of files of varying sizes, and by the requirement that
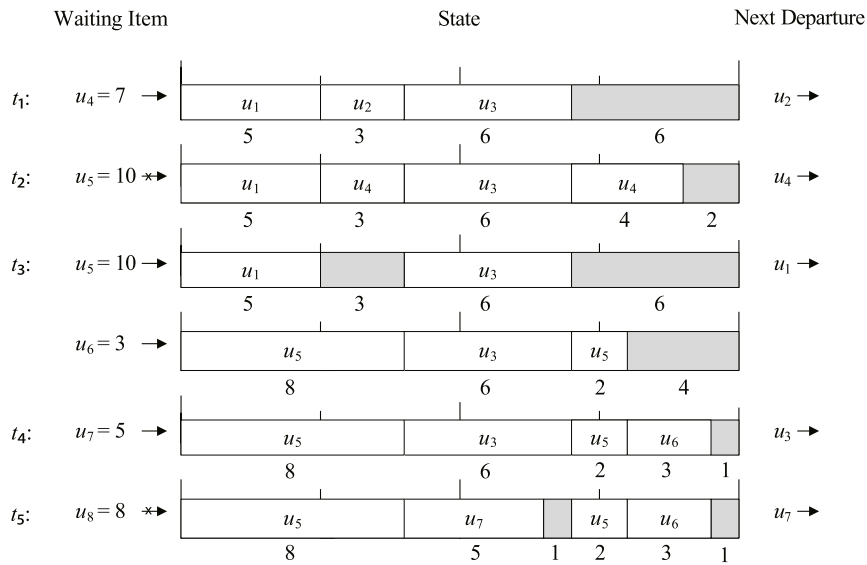
Fig. 1. An example of admission and departure processes and the resulting fragmentation.

files not be moved once they are allocated. *The fragmentation of interest here is fundamentally more general*; it occurs when resources are such that, when no individual gap is sufficiently large, an item's allocation can be broken down into fragments that fit into multiple gaps. A small example will help fix ideas. In what follows, the notation u_i will often be used to denote both the name and size of the ith item; context will always make clear which meaning applies.

In Fig. 1, the resource is 20 slots wide and the items, u_i, are for at most 10 slots. Successive departure times are denoted by the t_i. Allocations for the first 8 items $u_1, \ldots, u_8$ are shown using a simple left-to-right first-fit rule. Assuming the resource is not in use at time 0, items u_1, u_2, and u_3 are the first to be allocated; u_4 must wait for a departure, since the first 4 item sizes sum to more than 20. The first occurrence of fragmentation occurs at the departure of u_2 and the resulting admission of u_4; an initial fragment of u_4 is placed in the gap left by u_2 and a final fragment is placed after u_3. Also, even after u_2 and u_4 have departed, there is still not enough of the resource available for u_5. After the additional departure of u_1, both u_5 and u_6, but not u_7, can be allocated their requested numbers of slots.

On a much larger and more realistic scale, the fragmentation illustrated in Fig. 1 poses important performance trade-offs in modern, hard-disk computing platforms, where items are files that can be split up into fragments occupying gaps of free storage. We assume that files cannot be moved once allocated, so fragmentation allows efficient utilization of storage. But accessing a fragmented file requires a separate access for each of its fragments, so the mechanism implementing file access can be costly in time and space. When performance degrades sufficiently, defragmentation software is typically available, again at a cost of time and space, to shift fragments as needed to consolidate them into entire, unfragmented files. The literature on fragmentation in file-storage platforms appears to be mostly limited to the software engineering community. Yet, the importance of the performance concepts is clear: in the relatively recent past a cottage industry of defragmentation software has emerged [2].

Another application in prospect appears in the design of Dynamic Spectrum Access Networks [3] (also called Cognitive Radio Networks), where items are frequency bands. A portion of the wireless spectrum not licensed to any particular user is reserved for "opportunistic" use. Competing users have to transmit and receive data, and accordingly have specific bandwidth and residence-time requirements. The data transmission can take place over a channel consisting of a number of non-contiguous sub-channels. Once a transmission terminates, the corresponding sub-channel fragments are vacated and made available to other users. As in the storage application, gaps develop randomly in both size and position. The resulting fragmented spectrum also introduces several new problems [4–7] in addition to the fragmentation problems studied here. Further work on resource allocation polices for these networks has been proposed in [8,9].

As indicated above, the extent of fragmentation can have a significant impact on the efficiency of real systems and networks, and therefore, modeling and analysis can provide important insights. Intriguing problems of dynamic storage allocation have been around for over 40 years [1,10], and they continue to resist a stochastic analysis leading to insightful results. This paper studies fragmentation within a probability model of the combinatorial structures illustrated in Fig. 1. Although our model does not refer to any particular application, it is applicable to many practical systems such as those listed above. Two intriguing questions are center stage: How bad can fragmentation become and at what rate does it develop? For example, can fragmentation progress to a point where nearly all items are *completely fragmented*, i.e., for every item size $j > 1$, are nearly all size-j items allocated j mutually disjoint slots? Answers to these questions give insights into those circumstances where defragmentation is worth the cost it incurs. We will discuss algorithmic issues, but only for illustrative purposes

in our experimental study of convergence, since our main results on stochastic fragmentation hold true for all allocation algorithms that one can expect to find in practice.

PLAN OF THE PAPER: Section 2 describes our stochastic model of fragmentation in detail beginning with a definition of the state space of the fragmentation process and the constraints of transition functions imposed on allocation algorithms. The section then introduces a general class of allocation algorithms, describing two in detail that are used in the remainder of the paper when needed for experiments. The background in Section 3 briefly covers earlier results germane to our model. The results in Section 4, the centerpiece of the paper, apply simultaneously to all algorithms in the class defined in Section 2. For given item-size distributions, the results of Section 4 describe the asymptotic fragmentation behavior in statistical equilibrium for large resource sizes. The most general of these results is a theorem showing that the fraction of partially fragmented items tends to 0 as the resource size tends to infinity. In the case where size-1 items have positive probability, the result is strengthened considerably: The expected number of partially fragmented items is shown to be bounded by a constant independent of the resource size. The case where there are only items of sizes 1 and 2 is treated separately, as a tight bound can be found for the constant in this case. Section 5 investigates the convergence times to highly fragmented states, relying primarily on simulation experiments. Section 6 concludes the paper by suggesting areas that merit further research, with an analysis of algorithms being a prime example.

2. The model

The first of the two subsections below introduces in broad terms the fragmentation process at the heart of the paper; the class of allocation algorithms that completes the model definition is covered in the second subsection.

2.1. The fragmentation process

A queue of items waits to use a resource $\mathcal{R}$ consisting of a sequence of M slots.[1] The resource is made available to items on a first-come-first-served (FCFS) basis, each item's request specifying the number of slots that it needs (i.e., the item's size), and the amount of time it needs the slots. The sizes are i.i.d. random variables governed by the distribution $\mathbf{q} = \{q_1, \ldots, q_K\}$. To avoid trivialities we assume that at least two item sizes have positive probability, so that $q_K > 0$ for the maximum item size K. An item waits until it arrives at the head of the queue, at which point it is allocated its requested amount of $\mathcal{R}$ as soon as the total size of the gaps in $\mathcal{R}$ becomes large enough. The allocation of an item of size u, say, may be fragmented as desired to exploit gaps in resource usage, each of which is smaller than u, but whose cumulative size is at least u. Various algorithms other than the first-fit algorithm of Fig. 1 can be adopted to satisfy the requests of new items; we return to further discussion of allocation policies after a few more preliminaries.

The item at the head of the queue is usually called the HOL (head-of-the-line) item. In reference to allocations, the term *fragment* always refers to a piece of some item occupying a maximal sequence of consecutive slots, i.e., a slot adjacent to the sequence must either be empty or occupied by a different item.

We focus on the performance of systems operating at capacity, when throughput is at its maximum. Thus, the queue of waiting items never empties. Unless stated otherwise, the resource $\mathcal{R}$ is initialized, as in the example of Fig. 1, with a sequence of unfragmented allocations to items $u_1, u_2, \ldots$ such that

$$u_1 + \cdots + u_{i-1} \leq M < u_1 + \cdots + u_i.$$

At that point, all $i - 1$ of the items begin independent, identically distributed residence times, or delays. Subsequent state transitions take place at departure epochs when the delays of currently allocated items expire. At such epochs, all fragments of $\mathcal{R}$ in the allocation of a departing item are released, thus enlarging existing gaps and/or creating new ones. Suppose all items up to u_j have received their allocations and an item u_i, $i \leq j$, departs. Then items $u_{j+1}, u_{j+2}, \ldots$ are allocated to gaps until, once again, an item is encountered that requests more of $\mathcal{R}$ than is available. All allocated items then begin or continue their residence times until the next departure. We emphasize that, once an allocation is made to an item, it remains fixed. Our probability model stipulates that item residence times are i.i.d. exponentially distributed random variables independent of item sizes. For convenience, the expected residence time is taken as the time unit. In summary then, for a given gap allocation policy, M and $\mathbf{q}$ are the two parameters of the model.

More formally, a state x in the state space $\mathcal{S}$ of a fragmentation process is given by

$$x = (L_1, \ldots, L_N; u) \tag{1}$$

where u is the size of the item waiting at the head of the queue (the HOL item), $N \geq 1$ is the number of currently allocated items, and L_i, $1 \leq i \leq N$, is the list of disjoint subsequences of $(1, \ldots, M)$ occupied by the fragments of L_i. For x to be admissible, the subsequences of L_i must be disjoint from those of L_j for all $i \neq j$, and for it to be *stable* u must exceed the resource available, i.e., $u > M - \sum_i \|L_i\|$, where $\|L_i\|$ is the size of the item (sum of the fragment sizes) in L_i. The fragmentation process is a Markov process for any allocation algorithm likely to be used in practice, and in particular, any algorithm

---

[1] A summary of the notation appears in Appendix A.

in the class defined in the next section. Asymptotic analysis and the elementary theory of finite Markov chains provide the tools used in the analysis found in Section 3.

2.2. Allocation algorithms

Our results characterize algorithms within a large class $\mathcal{A}$ defined as follows. Let x be the state just after a departure and before new allocations are made and let A be the allocation algorithm. If x is not stable, then A computes an ordering of the gaps in x and allocates gaps to fragments of the HOL item u in that order. These gaps are allocated in their entirety to u, except possibly for the last gap. The fragment in the last gap is left-justified. If the new state x' is still not stable, allocations are repeated as above in state x' for the new HOL item. The process continues until a stable state is reached. An algorithm is said to have the *best-fit* property if the last gap to be allocated is the smallest remaining gap big enough to accommodate the last fragment.

The principal contribution of the paper is the theory of the next section, which *applies to any allocation policy in $\mathcal{A}$—even one that computes a gap ordering based on the sizes and residence times of all future arrivals and departures.* However, we need algorithms for experiments targeting more precise behavior that we have been unable to analyze, a primary example being the convergence rates covered in Section 5. We introduce for illustrative purposes a fast scanning algorithm, Next-Fit (NF), and then a second, complementary algorithm, Best-Fit Decreasing (BFD), which is slower but which creates significantly less fragmentation. Both algorithms treat the resource as a homogeneous ring of M slots; slot 1 follows slot M and items and gaps seamlessly cross the boundary between these two slots. This assumption is largely for convenience and will have a negligible effect on results, since our primary interest is in cases where M is much larger than K.

NF is like the first-fit algorithm of Fig. 1 adapted to the ring, but it starts each scan where the preceding one left off. Specifically, a scan starts at the slot following the end of the last slot filled in making the preceding item allocation. NF is not a best-fit policy like BFD, which allocates an item to the gaps in decreasing order of size until the last fragment of the item (which could be the entire item) is smaller than at least one gap. At that point, the last, perhaps only, fragment of the item is allocated best-fit, i.e., to the smallest gap no smaller than the fragment. To be specific, we say that BFD breaks all ties in favor of the leftmost (lowest indexed) gap providing a best fit for the final fragment.

3. Background

There are two processes closely related to our fragmentation process which have been studied in the past. Although the original model of item fragmentation was the one given here, the first results to appear were for the continuous limit of our model normalized on [0,1], which was studied in [11]. In normalizing our model, the slot size is set to $1/M$ and the support of the item-size distribution is $\{1/M, \ldots, K/M\}$. The continuous limit is reached as $K, M \to \infty$ subject to $K/M \to \alpha$, with $\alpha \in (0, 1]$ the single parameter of the model apart from the item-size distribution. The property of complete fragmentation in our model is analogous to behavior as $\alpha \to 0$ in the continuous model. Experiments have supported the claim that the number of fragments increases as $\alpha \to 0$ in the continuous limit, but this remains an open question. It is rather striking that the corresponding result is indeed provable in our original model, as shown in Section 4. The result gains in significance from experiments that have shown the continuous model to be distinctly limited in its approximation of the original, more realistic discrete scaling of this paper. For example, under a uniform distribution on item sizes, one needs to assume that K is larger than about $M/10$ if the continuous limit with $\alpha = K/M$ is to give measures of fragmentation reasonably close to those of the model analyzed here. This is commonly an unacceptable assumption in practice.

The second related process is defined on the states $\hat{x} = (\|L_1\|, \ldots, \|L_N\|; u)$ where fragment configurations are suppressed. The corresponding process $(X_A(t))$ remains Markovian for any algorithm $A \in \mathcal{A}$. Kipnis and Robert [12] generalized $(X_A(t))$ to a standard queueing process with stochastic arrivals. Their model was motivated by systems where fragmentation is not allowed; allocations are shifted as needed to keep available resource in one block.

The result to follow applies to an asymptotic, large-M analysis of the *gap-wait* process (G_t, H_t) at departure epochs, where G_t and H_t are respectively, at time t, the total gap size and the size of the HOL item. The result will clearly have no direct bearing on fragment configurations. Our focus is on behavior at departure epochs in statistical equilibrium, so we will usually drop the dependence on t. We include the result in this section, as it applies a key observation made in [12] that yields the asymptotics of the equilibrium probability distributions for H and for G, both individually and jointly. We give our own proof as it does not require the considerably more extensive analysis of the general queueing model studied in [12].

Theorem 1 (*Kipnis and Robert [12]*)**.** *At departure epochs, the following asymptotic equilibrium probability distributions apply. As $M \to \infty$, the distribution f_H for the HOL item size H tends to*

$$f_H(h) \sim \frac{h f_U(h)}{\mathbb{E}[U]}, \quad 1 \le h \le K, \tag{2}$$

the total gap-size distribution f_G tends to

$$f_G(g) \sim \frac{1 - F_U(g)}{\mathbb{E}[U]}, \quad 0 \le g \le K - 1, \tag{3}$$

and finally, the distribution π of the gap-wait state tends to

$$\pi_{gh} \sim \frac{f_U(h)}{\mathbb{E}[U]} \tag{4}$$

for all g, h satisfying $0 \le g < h \le K$. ∎

Examples: From (2), we have that

$$\mathbb{E}[H] \sim \frac{\mathbb{E}[U^2]}{\mathbb{E}[U]} \tag{5}$$

so with U uniformly distributed on $\{1, \ldots, K\}$, calculations show that, as $M \to \infty$,

$$\mathbb{E}[H] \sim \frac{2K+1}{3}, \qquad \mathbb{E}[G] \sim \frac{K-1}{3}, \tag{6}$$

and that the asymptotic gap-wait distribution is also uniform:

$$\pi_{gh} \sim \frac{1}{K(K+1)/2}, \quad 0 \le g < h \le K.$$

As another example, one to be used later, consider a general distribution for $K = 2$, with $q_1 := \Pr\{U = 1\}$. We find that $\mathbb{E}[U] \sim 2 - q_1$, and

$$\pi_{01} \sim \frac{q_1}{2-q_1}, \qquad \pi_{02} = \pi_{12} \sim \frac{1-q_1}{2-q_1} \tag{7}$$

as $M \to \infty$.

Proof of Theorem 1. As $M \to \infty$ for fixed K, the frequency of size-i items currently allocated in $\mathcal{R}$ tends to $f_U(i)$ for all $i \in \{1, \ldots, K\}$ by the laws of large numbers. Now let (U_i) be an equilibrium renewal process with intervals U_i between successive renewals having the item-size distribution F_U. Then the size of the interval covering the point M is, asymptotically for large M, equal in distribution to the size of the interval covering a specific point in the renewal process (U_i). For convenience, we adopt the convention that, if M coincides with the end of one interval and the start of the next one, it is the latter interval that covers M. The key observation, extending an argument in [12], is that the duration of the interval covering M is equal in distribution to the size H of the HOL item and, asymptotically, is a random variable with the familiar distribution given in (2).

To calculate the corresponding, asymptotic total-gap size distribution, we note first that, in the asymptotic regime, for an HOL item size h, the total gap size has a uniform distribution on the set $\{0, \ldots, h-1\}$, so

$$f_G(g) \sim \sum_{g+1 \le h \le K} \Pr(G = g | H = h) f_H(h)$$

$$\sim \sum_{g+1 \le h \le K} \frac{1}{h} f_H(h), \quad 0 \le g \le K-1$$

whereupon substitution of (2) yields (3). As noted above, the h gap sizes that can match up with a given HOL item size h are equally likely, so (4) follows directly from $\pi_{gh} \sim f_H(h)/h$ and (2). ∎

4. Asymptotic theory of fragmentation

The general objective of this section is an estimate of the extent of fragmentation in statistical equilibrium. Specifically, in the state notation of Section 2, the goal is a large-M estimate of the equilibrium expected total number of fragments, $|L_1| + \cdots + |L_N|$, for fixed $K \ge 2$. Recall that an item of size $j > 1$ is said to be *completely fragmented* when it is allocated a completely fragmented part of $\mathcal{R}$—no two slots of its allocation are adjacent.

In this section, we prove limit laws which show that for large M nearly all of the items receive completely-fragmented allocations. Before getting into the theory, we look very briefly at some experiments that brought out initial insights (experiments not oriented to asymptotic behavior have been reported in the past (e.g., in [11])). The analysis that follows begins with the case $K = 2$ and accommodates any item-size distribution on $\{1, 2\}$. The next subsection analyzes the case for general $K \ge 2$ and $q_1 > 0$, and the final subsection covers the general case $K \ge 2$ without the constraint that q_1 be positive.

The analysis in this section applies to the fragmentation process in statistical equilibrium: the focus is on large-M asymptotics with $\{q_1, \ldots, q_K\}$ fixed. To improve readability, we will occasionally omit "$M \to \infty$" when the limit is clear in context.

4.1. Preliminaries

For the experimental results in this subsection, the simulation of 10^8 departures produced in all cases a behavior that was indistinguishable from stationarity. Additional details regarding the simulation tool are provided in Appendix B.

Table 1
Simulation results using the NF and BFD policies with $K = 2$, item size probabilities $q_1 = q_2 = 1/2$, and varying M: steady-state averages of the number of unfragmented size-2 items (N_{21}) compared to averages of the total number of items (N), the number of size-1 items (N_1), and the number of fragmented size-2 items (N_{22}).

M	NF				BFD			
	10	100	1000	10000	10	100	1000	10000
$\overline{N}$	6.56829	66.5382	666.552	6666.56	6.56283	66.5806	666.659	6666.32
$\overline{N}_1$	3.47090	33.4098	333.441	3330.45	3.46030	33.4949	333.651	3332.98
$\overline{N}_{21}$	0.79380	0.88536	0.91478	0.91675	1.56803	1.95471	2.05519	1.94872
$\overline{N}_{22}$	2.30359	32.2430	332.197	3332.19	1.53450	31.1310	330.953	3331.40

Table 2
Simulation results using the NF and BFD policies with item sizes uniformly distributed and varying K, M: steady-state probability that an item of size-K is completely fragmented, $\rho_{K,M}$.

M	NF			BFD		
	$\overline{\rho}_{2,M}$	$\overline{\rho}_{5,M}$	$\overline{\rho}_{10,M}$	$\overline{\rho}_{2,M}$	$\overline{\rho}_{5,M}$	$\overline{\rho}_{10,M}$
1000	0.99698	0.94927	0.68151	0.993957	0.825870	0.362162
2000	0.99846	0.97422	0.82478	0.997017	0.902945	0.522318
3000	0.99893	0.98248	0.87911	0.998075	0.933195	0.607607
4000	0.99917	0.98676	0.90714	0.998416	0.948368	0.663502
5000	0.99932	0.98927	0.92475	0.998898	0.958309	0.704494
6000	0.99940	0.99096	0.93668	0.999000	0.964694	0.734542
7000	0.99946	0.99188	0.94572	0.999180	0.969729	0.757376
8000	0.99949	0.99284	0.95187	0.999202	0.973263	0.776753
9000	0.99952	0.99340	0.95694	0.999324	0.976383	0.794857

Table 1 gives experimental results for $K = 2$ that describe behavior at departure epochs in the case where size-1 and size-2 items are equally probable. The values shown in the table for the quantities N_1, N_{21}, and N_{22} are the respective experimental results for the numbers of size-1 items, unsplit size-2 items, and size-2 items that split into 2 fragments placed into non-adjacent slots. It can be seen that, while the average numbers of items, $\overline{N}$, size-1 items, $\overline{N}_1$, and split size-2 items, $\overline{N}_{22}$, grow linearly in M, no sustained growth at all is present in the average number $\overline{N}_{21}$ of unsplit size-2 items. The approximate asymptotic value of 2 for $\overline{N}_{21}$ under BFD is in fact exact; a more general result is proved below.

Detailed complete-fragmentation behavior was also studied in experiments for larger K with a focus on the uniform distribution on $\{1, \ldots, K\}$. The results for a number of values of K are illustrated in Table 2 for NF and BFD, where $\overline{\rho}_{K,M}$ denotes the experimental estimate of the steady-state probability that size-K items are given a completely fragmented allocation of K non-contiguous slots. The trend toward 1 is monotone increasing toward complete fragmentation in all three cases, $K = 2, 5, 10$. Below, we prove an even stronger result than convergence in probability, but precise formulas have been out of reach so far. Note the comparison in the tables between NF and BFD whereby the items under the NF policy fragment significantly faster than under the BFD policy.

4.2. $K = 2$

A precise asymptotic formula is available for $K = 2$ and it requires only that the two item-size probabilities be positive. In deriving this formula, which measures the asymptotic expected number of unsplit 2's at departure times, our approach balances the equilibrium rate at which the number of unsplit 2's increases at departure times against the equilibrium rate at which it decreases at these times. Choose a departure at random in statistical equilibrium, and suppose it occurs at time t. We continue with the notation for N_{ij} from the previous subsection (also described in Appendix A). Recall that the stable gap-wait states for $K = 2$ are $(0,1)$, $(0,2)$, and $(1,2)$, and note immediately that, in the gap-wait state transition at time t, $N_2(t)$ can change by at most 1.

Observation 1. *Suppose first that the departure occurs in a $(0, 2)$ state. If the departure is a 1, then the state transitions to a $(1, 2)$-state and there is no change in $N_{21}(t)$. If the departure is a 2, then the HOL 2 will receive a split or unsplit allocation according as the departing 2 had a split or unsplit allocation, respectively. Again, there is no change in $N_{21}(t)$, so we may restrict the analysis to transitions out of states $(0, 1)$ and $(1, 2)$.*

Observation 2. *In a $(0, 1)$ state, there can be no increase in $N_{21}(t)$ in any of the 3 possible gap-wait state transitions. There can be a decrease, but only by 1, and this happens if and only if $N_{21}(t) > 0$ and the departure is an unsplit 2. Then, since the rate parameter of the exponential residence times is 1, the rate of decrease in the number of unsplit 2's in transitions out of $(0, 1)$ states is simply $N_{21}(t)$. It follows that the average rate at which $N_{21}(t)$ decreases from $(0, 1)$ states is $\mathbb{E}'[N_{21}]$, where the prime indicates that the average is taken in equilibrium over all $(0, 1)$ states.*

Observation 3. *No decrease in $N_{21}(t)$ need occur in $(1, 2)$ states, since, as in the case of $(0, 2)$ states, a departing unsplit 2 can (and will, if a largest-gap-first allocation policy is being used) be replaced by an unsplit 2. For an increase in $N_{21}(t)$ to occur, the next departure must be one of the at most two items using a slot adjacent to the empty slot. (Conceivably, there is a split 2 whose two slots are on either side of the gap, so that there is only one item whose departure can increment $N_{21}(t)$.) Thus $N_{21}(t)$ is incremented at rate at most 2.*

Since in the large-*M* limit the ratio of the equilibrium frequencies of $(0, 1)$-states to $(1, 2)$-states is $q_1/(1-q_1)$, equilibrium demands that $\mathbb{E}'[N_{21}]q_1 \le 2(1-q_1)$ giving $\mathbb{E}'[N_{21}] \le 2(1-q_1)/q_1$. Having observed that changes in N_{21} are rare, in the sense that they occur from any state with probability at most $2/M \to 0$, we conclude that the asymptotic equilibrium expected value $\mathbb{E}[N_{12}]$ of N_{12} over all states is equal to our $\mathbb{E}'[N_{21}]$. Hence, we have the following:

Theorem 2. *In the $K = 2$ case with positive item probabilities, the equilibrium expected value of the number of unsplit items of size 2 satisfies*

$$\mathbb{E}[N_{21}] \le 2(1 - q_1)/q_1$$

asymptotically as $M \to \infty$, regardless of the allocation policy.

Note that, within the constraints of our allocation policies, not even a clairvoyant policy, i.e., one that can see the queue of arrivals and departures for the infinite future, can escape the extreme fragmentation guaranteed by Theorem 2.

If a policy scans the gaps in decreasing order of size, then an argument similar to that given in the proof of Theorem 2 can be used to bound by a constant the expected number of split 2's whose fragments are separated by only a single unoccupied slot (cf. Observation 3). From this there follows:

Corollary 1. *For any largest-gap-first policy in the $K = 2$ case, the asymptotic equilibrium rate at which $N_{21}(t)$ is incremented is precisely 2, and thus, the bound in Theorem 2 becomes*

$$\mathbb{E}[N_{21}] \sim 2(1 - q_1)/q_1, \quad M \to \infty.$$

Although the bound of Theorem 1 applies to the NF algorithm, Corollary 1 does not, since NF is not a largest-gap-first policy (there are transitions when an HOL 2 is not put into a gap of size 2, even when one exists). The problem of finding a tight bound for NF is much more difficult, even in the large-*M* asymptotic regime. The difficulty arises primarily in characterizing the behavior of transitions into and out of state $(1,2)$. This can be expected because a new variable enters the analysis: the index of the position where the scan is to recommence relative to the positions of vacant slots.

4.3. $K \ge 2$ and $0 < q_1 < 1$

For this case, our argument is based on the notion of *bonds*, with a bond being any pair of adjacent slots that are either both empty or both occupied by the same item, i.e., a pair of adjacent slots within a gap or an item fragment. $B(t)$ denotes the total number of distinct bonds at time t. Note that in the $K = 2$ case, $B(t) = N_{21}(t)$.

We will show that, as long as there are some items of size 1, the expected number of bonds in equilibrium is bounded by a constant independent of *M*. To do this we will show, much as in the $K = 2$ case, that bonds are destroyed at a rate proportional to $B(t)$, but created at only a constant rate.

Observation 4. *Bonds can be created only when a departing item vacates one or more slots adjacent to an already present gap. In the worst case, a departure is about to take place while there are already $K - 1$ single-slot gaps. In that case every pair of adjacent slots, exactly one of which is currently unoccupied, will become a bond at rate 1, so the expected increase in the number of bonds is $2K - 2$. For example, if the gaps and a certain item of size K happen to be interleaved to create a giant interval of size $2K - 1$, then that item's departure is the only one to create bonds, but it creates $2K - 2$ of them. If, more typically, the gaps are spread out and every one of the $2K - 2$ slots adjacent to a gap belongs to a different item, then only one new bond will be created but at rate $2K - 2$. In summary, bonds are created at an average rate of at most $2K - 2$.*

Theorem 3. *Suppose the probability of size-1 items satisfies $0 < q_1 < 1$. Then in equilibrium, the number of items only partially fragmented is bounded by a constant independent of the number M of slots, regardless of the allocation policy.*

Proof. In view of Observation 4, we need a suitable bound on the rate of bond destruction to complete the proof. Bonds are destroyed (but only one at a time, assuming a best-fit allocation policy) whenever a departing item has a fragment which is partly but not completely filled by an incoming item. Since no item can contain more than $K - 1$ bonds, $B(t)$ is at most $K - 1$ times the number of items that contain one or more fragments of size at least 2. Suppose that at a moment when there are no gaps an item of size j departs, leaving at least one gap of size $i > 1$. If the next $j - i + 1$ items are all ones, at least one of them will be placed in that gap, destroying a bond. In the worst case, $j = K$ and $i = 2$, but even then the rate of bond destruction is at least the product of (1) the probability of $K - 1$ consecutive 1's, (2) the lower bound $B(t)/(K - 1)$ on

the number of items with at least one fragment of size 2 or greater, and 3) the probability of a full state (i.e., a state with no open slots). That is, we can write as a lower bound to the destruction rate:

$$q_1^{K-1} \times \frac{B(t)}{K-1} \times f_G(0)$$

where, from Eq. (3), $f_G(0) \geq 1/K$ is the limiting probability of a full state. Thus, the mean bond destruction rate is at least $CB(t)$ for some constant $C > 0$ independent of M.

We now finish the argument as before. Since at equilibrium the average rates of increase and decrease of $B(t)$ must be the same, we have that, from Observation 4 and the above argument, $2K - 2 \geq C\mathbb{E}[B]$ and thus $\mathbb{E}[B] \leq (2K - 2)/C$. ■

A note of warning may be in order here. In the proof, the constant provided by the theorem may be exponential in K with base $1/q_1$, suggesting that in practice M may need to be *much* larger than K to get close to complete fragmentation. This is indeed the case, but we defer illustrations along with further discussion to the next section on convergence issues.

Note that, if the item sizes with positive probability have a common divisor $d > 1$, then a best-fit allocation policy will treat the resource as if it were partitioned into unbreakable intervals of d slots each, with (perhaps) some slots left over that are never used. So no fragmentation beyond size d can be expected.

4.4. $K \geq 2$

The result in this section is completely general, but weaker than each of the preceding theorems. This is because of the case yet to be covered: the item sizes having positive probability do not include size 1 and are not all divisible by some $d > 1$. The analysis of this case follows and will be seen to be significantly more subtle.

It remains true that for large M, when an HOL item requesting j slots is waiting, states with 0, 1, 2, . . . , up to $j - 1$ unoccupied slots are equally likely. Moreover, the argument of Observation 4 remains in effect: the average rate at which bonds are created cannot exceed $2K - 2$. It is no longer generally the case, however, that bonds are destroyed at a rate proportional to $B(t)$. Thus, we cannot show $\mathbb{E}[B]$ is bounded by a constant. However, it is still possible, as shown below, to prove nearly complete fragmentation in the sense that $\lim_{M \to \infty} \mathbb{E}[B]/M = 0$.

The difficulty is that even when a gap of size 2 or more is present, it may be that no sequence of items can force the destruction of a bond. For example, suppose item sizes are 2 and 3, all items are currently split, but items of size 3 are split into only two pieces. Then a departing 3 will cause a gap of size 2 to appear, but a best-fit allocation algorithm will fill it immediately using the HOL item, so no bond is destroyed.

Indeed, it is not obvious that bonds will ever get destroyed in this situation; can it be that no further fragmentation will occur? Not quite—the reason is that this situation, with all bonds inside items of size 3, is unstable. Bonds will migrate to items of size 2 and eventually two whole items of size 2 will depart with an HOL 3 waiting, causing the destruction of a bond.

To argue that such an event happens in general, we say that a currently active item is of *type* $\langle i, j \rangle$ if it is for i slots and contains exactly j bonds. A type $\langle i, j \rangle$ is said to be *rife* if the expected equilibrium number of items of type $\langle i, j \rangle$ is bounded below by CM, for some constant factor $C > 0$ and infinitely many values of M. The idea will be to show, by contradiction, that no type $\langle i, j \rangle$ with $j > 0$ can be rife.

To view the dynamic resource allocation process as a continuous-time Markov chain, we imagine that there is a cupboard of M clocks, which "ring" independently and periodically after exponential mean-1 waiting times. When an item is allocated it takes a clock from the cupboard and when the clock rings, the item departs, putting the clock back in the cupboard. When a clock in the cupboard rings, nothing happens. We will make use of the fact that if $\langle i, j \rangle$ is rife, with constant factor C, then for any fixed n and infinitely many M the probability σ that, at a given time t in equilibrium, the next n rings will all signal departures of items of type $\langle i, j \rangle$ is bounded below by $\frac{1}{2}C^n$. To see that this is so, let p_k, $k = 0, 1, 2, . . .$, denote the probability that, in equilibrium, exactly k items of type $\langle i, j \rangle$ are active. Then,

$$\sigma \geq \sum_{k=0}^{M} p_k \prod_{h=0}^{n-1} (k-h)/M \sim \sum_{k=0}^{M} p_k (k/M)^n$$

and so the result follows from

$$\sum_{k=0}^{M} p_k (k/M)^n \geq \left(\sum_{k=0}^{M} p_k \frac{k}{M} \right)^n \geq C^n.$$

Observation 5. *Suppose that some state $\langle i, j \rangle$ with $j > 0$ is rife; we assert that there is another rife state $\langle i', j' \rangle$ with $i'/j' < i/j$. (Of course, since there are only finitely may pairs $\langle i, j \rangle$ with $0 < i \leq K$ and $0 < j < K$, this state of affairs is not possible.) To prove the assertion, note that the fraction i/j exceeds 1 and thus cannot be a common divisor of every item size. Let i' be an item size (with $q_{i'} > 0$) such that $i'/(i/j) = i'j/i$ is not an integer. Suppose that at time t all slots are occupied while i items of size i' are waiting; and at the next i' rings, all signal departures are of items of type $\langle i, j \rangle$. The probability of such an event is at least $q_{j'}^i C^{i'}/(2KM)$, for infinitely many M, assuming C is as in the definition of "rife" for type $\langle i, j \rangle$. If half or more of the times this happens some bond is destroyed, the bond destruction rate would exceed the constant $2K - 2$ bond creation rate, a contradiction.*

It follows that most of the time all $i'j$ bonds in the departing items are distributed among the i incoming items of size i'. Suppose that on this occasion b_k bonds end up in the kth item. If $i'/b_i > i/j$ for each k, $1 \leq k \leq i$, then

$$i'j = \sum_{k=1}^{i} b_k = i' \sum_{k=1}^{i} b_k/i' < i'i(j/i) = i'j,$$

a contradiction. But i'/b_k cannot be equal to i/j because then $i' = b_k i/j$, an integer multiple of i/j. We conclude that for some k, $i'/b_i < i/j$. Since there are fewer than K possible values for this b_k, items of some type $\langle i', j' \rangle$ with $i'/j' < i/j$ are being created with probability at least

$$q_{j'}^i C^{i'}/(4K^3)$$

and therefore at rate M times this constant. So this new type is rife, proving the assertion.

We have shown that there are no rife types $\langle i, j \rangle$ with $j > 0$. If the equilibrium expected fraction of bonds $\mathbb{E}[B]/M$ were not bounded by some constant C, some type $\langle i, j \rangle$ with $j > 0$ would have to have equilibrium expected frequency at least CM/K^2 and thus be rife. We have proved the following by contradiction.

Theorem 4. *Suppose the item size probabilities are such that the values j for which $q_j > 0$ have no nontrivial common divisor. Then in equilibrium, the fraction of items partially fragmented tends to 0 as $M \to \infty$, regardless of allocation policy.*

The caveat following Theorem 3 applies even more strongly in this case. Bounds are not likely to be useful in a real-life problem—the time needed for nearly complete fragmentation is going to be even more astronomical than before.

5. Convergence

So far, we have studied the extent of fragmentation in statistical equilibrium for large M/K. We have yet to study the transient behavior of the fragmentation process and to get a more quantitative appreciation of the properties of this process as a function of the parameters. Unsurprisingly, an analysis of convergence rates is difficult, so we are confined primarily to experiments and suggestive examples in what follows. Our space constraints prevent us from full coverage of our extensive experiments, so we limit ourselves to properties closely related to the theory of the preceding section.

We begin with a return to the red flag raised by the potentially huge constant derived in the proof of Theorem 3. We give an example in which the parameters are far from extreme, yet it illustrates the extremely slow convergence to statistical equilibrium that can occur in practice. As the example makes clear, the problem is that M has to be taken enormously large (far beyond realizable values) to ensure convergence to states where nearly all items are completely fragmented. We follow this example in the second subsection with a more general assessment of convergence to statistical equilibrium, beginning as earlier with the simplest non-trivial case $K = 2$. We then illustrate our study of cases $K \geq 2$ by specializing to uniform item-size distributions on $\{1, \ldots, K\}$. In the third subsection, we direct our attention to the constant bounds on partially fragmented items guaranteed by Theorem 3. These bounds increase with K, and, considering Example 1 below, it is of immediate interest to get a good idea of the rate of increase.

5.1. An instructive example

Consider the implication of the proof of Theorem 3 that there exists a constant C such that, for all M sufficiently large, the number of partially fragmented items is at most C; more specifically, consider estimates of "sufficiently large". The following example shows that they can be egregiously large.

Example 1. Suppose that items are of size 1 and 9 with probabilities .1 and .9, respectively. Suppose further that we are close to complete fragmentation in that it has developed to a point where 90% of the items of size 9 have been completely fragmented, the rest having 7 fragments of size one and one fragment of size 2. Then there will typically be about 4 singleton gaps in play, and bonds will be created (assuming a best-fit allocation policy) at rate about 7. To destroy a bond, however, we need at least 8 size-1 items in a row following the departure of a partially-fragmented 9; this happens at rate about 10^{-8} times the number of such 9s, so if we want the fragmentation to continue forward, we need $10^{-8} \times (M/9)/10$ to be greater than 7, forcing $M > 10^{10}$. Change K to 99 or q_1 to .01, and M will need to exceed the estimated number of elementary particles in the universe. (Note that even the time to reach the initial condition of the example is very likely to be huge.)

Fig. 2 shows the simulation results that estimate the steady-state fraction of size-9 items that are completely fragmented under the BFD algorithm as a function of M for several choices of q_1. The values of M required for nearly complete fragmentation increase by roughly an order of magnitude as q_1 decreases from one value to the next in the sequence 0.9, 0.7, 0.5, 0.3, 0.1. However, consistent with the example above, the required values of M for nearly complete fragmentation grow much faster as q_1 decreases below 0.1. Convergence to complete fragmentation for smaller q_1 appears to require exponentially larger M. In particular, the curve for $q_1 = 0.1$ is still substantially short of complete fragmentation and was found to remain so until $M \approx 10^{10}$. Similar results will hold for any *best-fit* algorithm.

Table 3
Accuracy of approximate convergence times for $K = 2$.

M	$q_1 = 0.1$			$q_1 = 0.9$		
	$\mathbb{E}[T_\gamma]$	$\bar{T_\gamma}$	Error (%)	$\mathbb{E}[T_\gamma]$	$\bar{T_\gamma}$	Error (%)
1,000	1090.70	1221.20	10.69	209.33	221.40	5.45
10,000	10906.98	11115.90	1.88	2093.26	2054.30	1.90
100,000	109069.82	109091.00	0.02	20932.59	20942.20	0.05

5.2. Convergence to nearly complete fragmentation

We begin as in the previous section with $K = 2$, the simplest non-trivial case. We then illustrate our experimental results for general K with an example describing the convergence to complete fragmentation with item size a parameter.

5.2.1. $K = 2$

For this case, Theorem 2 and Corollary 1 give explicit asymptotic results for the constant bound on the number of unsplit 2's in statistical equilibrium. Might we hope to be equally successful in an analysis of convergence rates from an unfragmented initial state? For the moment, no, but a rough calculation is available which leads to an interesting formula that is very accurate when q_1 and hence q_2 are not too close to 0 or 1, and M is large.

We adopt the BFD algorithm for this case, and consider the number N_{22} of split 2's in the chain embedded at the epochs where a 2 departs in gap-wait state $(0,1)$. These are the only epochs of the fragmentation process where N_{22} can increase; an HOL unsplit 2 becomes a split 2 with a probability proportional to the number of unsplit 2's. We ignore increases in the number of unsplit 2's since these events, only in gap-wait state $(1,2)$, occur with probability $O(1/M)$. The quantity to be estimated is $E[T_\gamma]$, $\gamma \in (0, 1)$, which tells how long it takes on average for the embedded fragmentation process to increase the number of split 2's to a fraction γ of the original number of 2's, all of which were unsplit. Note that the embedded process is an occupancy type process. The calculations estimating $E[T_\gamma]$, $\gamma \in (0, 1)$, are routine and can be found in Appendix C. The result is,

$$\mathbb{E}[T_\gamma] \approx M \frac{1 - q_1}{2 - q_1} \ln \frac{1}{1 - \gamma}.$$

The accuracy of this formula is illustrated in Table 3.

Out of a total of D events, the average number of events in the embedded process is asymptotically the total number of events times the probability $\pi_{01} = q_1/(2 - q_1)$ of the $(0,1)$ state (cf. Theorem 1) times the probability $1 - q_1$ of a 2 departure. Then $\mathbb{E}[T_\gamma]/[(1 - q_1)\pi_{01}]$ approximates the asymptotic expected total number $\mathbb{E}[D_\gamma]$ of departures needed to fragment a fraction γ of the unsplit 2's of the initial state. This yields

$$\mathbb{E}[D_\gamma] \approx \frac{M}{q_1} \ln \frac{1}{1 - \gamma},$$

which has an accuracy comparable to that for $\mathbb{E}[T_\gamma]$.

5.2.2. General $K \geq 2$

Example 1 above suggests that questions with realistic answers concern the time to approach states that are just moderately close to complete fragmentation. To get a better handle on the term "moderately close", we resorted to experiments, in which the item-size distribution of choice was the uniform law on $\{1, \ldots, K\}$. The results of a typical experiment are shown in Fig. 3 where the curves represent the growth in average fragmentation normalized by the number of fragments for complete fragmentation, i.e., the item size. The region of near linear growth in the scaling of Fig. 3 shows that even the peak rate of fragmentation is quite slow, in particular approximately logarithmic in the number of departures. Intuitively, the larger item sizes can be expected to approach high levels of fragmentation more slowly the smaller size items. But until the fragmentation of the larger items becomes well advanced, the smaller items may well be inhibited from fragmenting by the relatively large gaps made available by departing, coarsely fragmented large items. Accordingly, the difference in convergence rates across the item sizes is not large.

5.3. Estimates of the number of partially fragmented items from Theorem 3

Recall that, for $q_1 > 0$, Theorem 3 guarantees that, in statistical equilibrium, there is at most a constant number – independent of M – of items which are partially fragmented (i.e., containing bonds). To understand the scaling of this constant as a function of the item size distribution, we resorted to numerical experiments. Fig. 4 presents the steady-state number of bonds using BFD when item sizes are uniformly distributed on $\{1, \ldots, K\}$. A ratio M/K of resource size to maximum item size was selected so as to ensure that the numbers of bonds had leveled off for all K; there was no discernible trend when the simulations terminated.
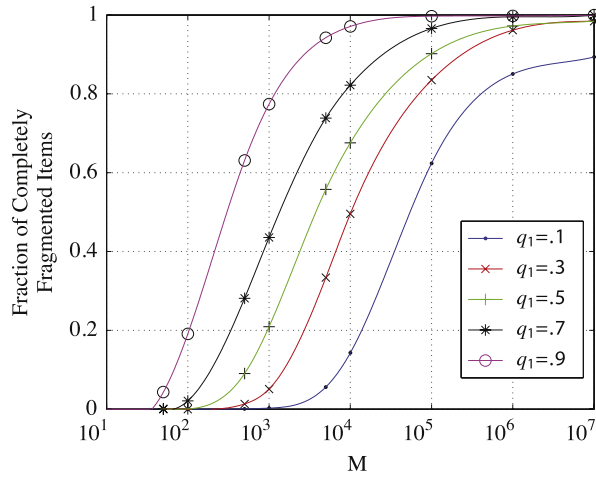
Fig. 2. Steady-state fraction of completely fragmented items as a function of M under the BFD algorithm, $K = 9$, and a parameter $q_1 = 1 - q_9$.
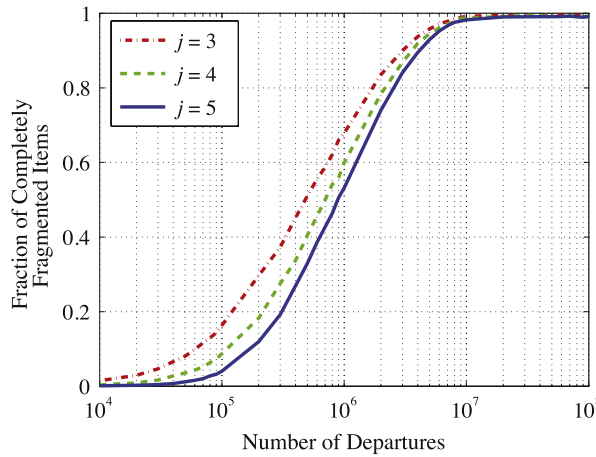

Fig. 3. Transient of the normalized average number of fragments per item for size-j requests under BFD (log scale): $K = 5$ (uniform distribution), $M = 100,000$.
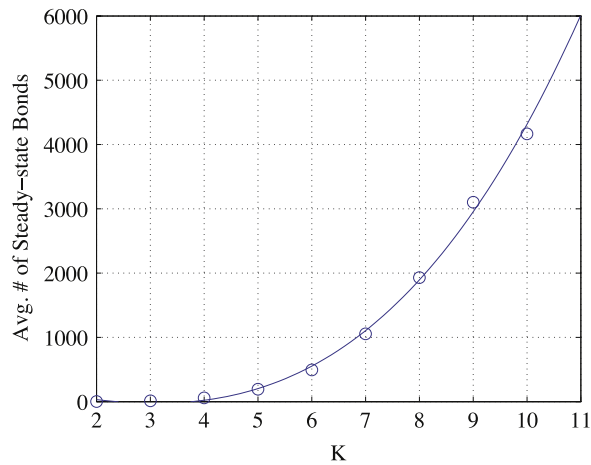

Fig. 4. Average number of bonds in equilibrium under BFD, with $M/K = 10^5$, item sizes uniformly distributed on $\{1, \ldots, K\}$.

The experiments were quite demanding in the time required as K was increased. The figure shows results only for the constants corresponding to $K \leq 10$. The data were fit to a cubic polynomial. Echoing Example 1, if K were increased by any significant amount, the constant bound was effectively unreachable with any value of M that would pass a reality check.

6. Final remarks

File allocation in disk-based storage and dynamic spectrum access are but two applications in which requests for a resource are allowed to be fragmented in the interest of more efficient utilization. The main contribution of this work is the analysis of the fragmentation process in a resource operating at capacity. Despite over four decades of research in the area of dynamic storage allocation, this is the first attempt to study this problem.

Theorems 2–4 constitute our contributions to the theory of fragmentation and show that for general item size distributions with large values of the resource size, nearly all items become *completely fragmented* in statistical equilibrium. In addition, we find a precise asymptotic formula for the number of partially fragmented items in the baseline case where requests are limited to size-1 or size-2. However, precise asymptotic formulas for the number of non-completely fragmented items have been out of our reach so far for arbitrary values of K. We supplement the analysis through examples and experimental results which show that the rates of convergence can be surprisingly slow; even the time required to become *moderately close* to states with nearly all items completely fragmented can be extremely long.

This leaves a number of interesting open problems, among them being more extensive studies of algorithms and of fragment configurations. Needless to say, the methodology of such studies will likely be limited to experimental approaches. An alternative source of open problems arises from generalizations of our assumptions, all of which are unlikely candidates for an analysis leading to explicit formulas. For example, the exponential assumption gives critical simplifications to analysis, but changes in behavior resulting from other residence-time distributions are worth investigating. In addition, instead of a system operating at capacity, in which there is always an item waiting, one could adopt an underlying, fully stochastic model of demand, e.g., a Poisson arrival process. Finally, more realistic models would relax the independence assumptions. A prime example appropriate for Dynamic Spectrum Access applications would allow residence times to depend on fragmentation. One may have to assume that the greater the fragmentation of an item, the longer its residence time.

Appendix A. Notation

$\mathcal{R}$	Resource
M	Size of $\mathcal{R}$ (in slots)
K	Maximum request size (in slots)
U	Random variable denoting the request size
q_j	Equivalent to $f_U(j)$, the request size pmf
u_i	ith realization of U
$H(t)$	The HOL request size at time t
$G(t)$	The total number of available slots at time t
N	Random variable denoting the number of requests being served by $\mathcal{R}$
N_i	Random variable denoting the number of requests of size i
N_{ij}	Random variable denoting the number of requests of size i split into j fragments
N_g	Random variable denoting the number of distinct gaps in $\mathcal{R}$
$\rho_{K,M}$	Steady-state probability of total-fragmentation for resource size M, max request size K
$B(t)$	Total number of distinct bonds at time t
α	The parameter of the continuous model. $K, M \to \infty$ subject to $K/M \to \alpha$, with $\alpha \in (0, 1]$
T_γ	The number of departure epochs required to increase N_{22} to γN_2

Appendix B. Simulation tool

The numerical experiments reported in this paper were conducted using a discrete-event simulator developed in C++. The simulations were executed until the system was reasonably close to steady-state at which point we randomly sampled for the performance measure of interest. The excellent accuracy of the tool was established in tests against the continuous

version of this model studied in [11]. We also tested the accuracy of the tool against results in [12] where the model did not allow fragmentation (cf. Theorem 1). We found that our simulation results were within 1% of the analytical results.

Appendix C. Rough estimate of $\mathbb{E}[T_\gamma]$

Recall that the epochs of the embedded fragmentation process occur in the $(0, 1)$ gap-wait state when a 2 departs. These are the epochs where the number of split 2's can increase. T_γ is the number of such epochs needed to change a fraction γ of the unsplit 2's to split 2's, assuming that all 2's are unsplit in the initial state. To carry the analysis forward, we reduce this occupancy-type process to its elementary form by taking N_2 to be constant at its mean $\beta \equiv E[N_2]$, which will give a useful asymptotic approximation in that region of the parameter space where the variation of N_2 around its mean is relatively small. Within the simpler model, the time to split a fraction γ of the original β unsplit 2's will be the sum of $\lfloor \gamma\beta \rfloor$ independent geometric random variables T_i with parameters $(\beta - i + 1)/\beta$, $i = 1, 2, \ldots, \lfloor \gamma\beta \rfloor$, and mean values $\beta/(\beta - i + 1)$, where T_i is the number of random draws from the set of 2's that have to be made before getting one that is not yet split, given that there are exactly $i - 1$ that have already been split. For simplicity, let $\gamma\beta$ be an integer in what follows. We have, when the number β of size-2 items is large,

$$\sum_{1 \leq i \leq \gamma\beta} E[T_i] = \sum_{1 \leq i \leq \gamma\beta} \frac{\beta}{\beta - i + 1}$$

$$= \beta \sum_{k=\beta(1-\gamma+1)}^{\beta} \frac{1}{k} \approx \beta \ln \frac{1}{1 - \gamma}.$$

A routine calculation shows that $\beta \sim M(1 - q_1)/(2 - q_1)$, and so finally

$$E[T_\gamma] \approx M \frac{1 - q_1}{2 - q_1} \ln \frac{1}{1 - \gamma}$$

References

[1] D.E. Knuth, The Art of Computer Programming, Vol. 1—Fundamental Algorithms, third ed., Addison Wesley Longman Publishing Co, Redwood City, CA, USA, 1997.
[2] M. Paul, 2014 compare best disk defragmenter software, http://disk-defragmenter-software-review.toptenreviews.com/ (2014).
[3] I.F. Akyildiz, W.-Y. Lee, K.R. Chowdhury, CRAHNs: Cognitive radio ad hoc networks, Ad Hoc Netw. 7 (5) (2009) 810–836.
[4] J. Jia, Q. Zhang, X. Shen, HC-MAC: a hardware-constrained cognitive MAC for efficient spectrum management, IEEE J. Sel. Areas Commun. 26 (1) (2008) 106–117.
[5] Y. Yuan, P. Bahl, R. Chandra, T. Moscibroda, Y. Wu, Allocating dynamic time-spectrum blocks in cognitive radio networks, in: Proc. ACM MobiHoc'07, 2007.
[6] A. Shukla, B. Willamson, J. Burns, E. Burbidge, A. Taylor, D. Robinson, A study for the provision of aggregation of frequency to provide wider bandwidth services, Tech. rep., QinetiQ (Aug. 2006).
[7] K. Jagannathan, I. Menache, E. Modiano, G. Zussman, Non-cooperative spectrum access—the dedicated vs. free spectrum choice, IEEE J. Sel. Areas Commun. 30 (11) (2012) 2251–2261.
[8] J. Huang, V. Subramanian, R. Agrawal, R. Berry, Downlink scheduling and resource allocation for OFDM systems, IEEE Trans. Wirel. Commun. 8 (1) (2009) 288–296.
[9] R. Cohen, L. Katzir, Computational analysis and efficient algorithms for micro and macro OFDMA downlink scheduling, IEEE/ACM Trans. Netw. 18 (1) (2010) 15–26.
[10] E. Coffman, T. Leighton, A provably efficient algorithm for dynamic storage allocation, J. Comp. Sys. Sci. 38 (1) (1989) 2–35.
[11] E. Coffman, P. Robert, F. Simatos, S. Tarumi, G. Zussman, A performance analysis of channel fragmentation in dynamic spectrum access systems, Queueing Syst. 71 (3) (2012) 293–320. Special issue of selected papers from ACM SIGMETRICS'10.
[12] C. Kipnis, P. Robert, A dynamic storage process, Stochastic Process. Appl. 34 (1) (1990) 155–169.

Edward G. Coffman, Jr. is a Professor Emeritus in the Computer Science Department and in the Electrical Engineering Department of Columbia University. He retired as a Distinguished Member of Technical Staff of Bell Laboratories in 1999. He is a Fellow of IEEE and an inaugural Fellow of ACM. His work in Computer Engineering started in 1958 as a systems programmer at the System Development Corporation. Highlights of his career include his work as a co-inventor of time-sharing systems and a co-inventor of computer networks. In the performance field, he was a co-founder of the SIGMETRICS and IFIP working group meetings. His current interests are in the mathematical modeling and analysis of processes arising in computer and network engineering.

Robert Margolies is a Ph.D. student in Electrical Engineering at Columbia University. He received the B.S. degree (Summa Cum Laude) from Rensselaer Polytechnic Institute in 2010 and the M.S. degree from Columbia University in 2011, both in Electrical Engineering. His research interests are in cross-layer PHY and MAC layer design for ultra-low-power networks, networks of energy harvesting nodes, and opportunistic cellular scheduling. Robert has received a number of awards including the NSF Graduate Research Fellowship.

Peter Winkler is William Morrill Professor of Mathematics and Computer Science at Dartmouth College, Hanover, NH USA. His research is primarily in combinatorics, probability and the theory of computing, with forays into statistical physics. He is also the author of a dozen patents, two books on mathematical puzzles, another on cryptology in the game of bridge, and a portfolio of pieces for ragtime piano.

Gil Zussman received the Ph.D. degree in electrical engineering from the Technion in 2004 and was a postdoctoral associate at MIT in 2004–2007. He is currently an Associate Professor of Electrical Engineering at Columbia University. He is a co-recipient of 5 paper awards including the ACM SIGMETRICS'06 Best Paper Award and the 2011 IEEE Communications Society Award for Outstanding Paper on New Communication Topics. He received the Fulbright Fellowship, two Marie Curie Fellowships, the DTRA Young Investigator Award, and the NSF CAREER Award, and was a member of a team that won first place in the 2009 Vodafone Foundation Wireless Innovation Project competition.