AofA 2014, Paris, France

Fragmented Random Structures

Ed Coffman[1,2] and Robert Margolies[1] and Peter Winkler[3] and Gil Zussman[1]

[1] *Electrical Engineering, Columbia University, New York, NY 10027*
[2] *Computer Science, Columbia University, New York, NY 10027*
[3] *Mathematics and Computer Science, Dartmouth College, Hanover, NH 03755*

Computers with hard-disk storage and networks with dynamic spectrum access illustrate resources that allow fragmented items – in these examples, items are files and spectra, respectively. In our discrete model of such systems, the resource is a sequence of slots. There is a queue of items awaiting allocations of the resource; the queue is served in FIFO order. Specified for each item are the number of slots needed by the item, and for what period of time. Under the key assumption that an item's allocation can not be changed prior to its departure, fragmentation in the form of alternating gaps and allocated resource builds up as items come and go, regardless of the allocation algorithm adopted. The improvements in resource utilization created by fragmentation are countered by the added cost of manipulating fragmented items, so how fragmentation evolves is an important performance issue. Within a baseline probability model of the system operating at capacity, we prove that, in the stationary limit under any practical algorithm, almost all items are completely fragmented, i.e., items of size i are fragmented into i disjoint slots. In the full paper, this result is balanced by many experimental results which show that, in reality, the times to approach steady-state fragmentation are typically exceptionally long, and hence that even nearly complete fragmentation is often of no concern.

Keywords: Random structures, Fragmentation, Dynamic storage allocation

The model: In our baseline probability model the system operates at capacity; effectively, there is an unbounded queue of waiting items. The resource is a sequence of $M > 1$ slots and item sizes are integers from 1 to K drawn from a given distribution $\mathbf{q} = \{\mathbf{q_1}, \ldots, \mathbf{q_K}\}$. To avoid trivialities, we assume that $1 < K \leq M$ and that at least two item sizes have positive probability. Requests to allocate items are served in FIFO order; once allocated, items remain for times that are independent, each a sample from the same exponential distribution, after which their allocated slots become available, either extending current gaps or creating new gaps.

A state σ of the stochastic fragmentation process is a sequence of M integers, one per slot: a 0 identifies an empty slot, and an integer $i > 0$ means that the slot contains a unit of the i-th item to be allocated. We assume that the resource is initially completely unoccupied, so item i is the i-th item in the queue at time 0. At state transitions an allocation algorithm maps a state σ and the size s of the item at the head of the queue/line, the HOL item, into the same state if it has fewer than s empty slots; otherwise, a new state σ' is obtained from σ by allocating s of σ's empty slots to the units of the HOL item. Algorithms differ by the choices they make for the empty slots.

As a simple example, suppose allocations are made in an increasing scan from slot 1 to slot M of the resource. Assume that $M = 8$ and $K = 4$ and consider the state $\sigma = 0, i_1, i_2, i_1, i_3, 0, i_2, i_2$. Slots 1 and 6 are unused, item i_1 has size 2 and is fragmented, as it occupies slots 2 and 4. Item i_3 has size 1 and occupies slot 5, and item i_2 is fragmented with a size-1 fragment in slot 3 and a size-2 fragment in slots 7 and 8. Suppose the HOL item has size 4 and hence σ is *stable*: no further allocations are possible until 1 or more items depart. If item i_3 were next to depart, nothing else would happen, as there would be only 3 available slots; another departure would have to occur before the HOL item could be allocated. On the other hand, if item i_2 were the next departure, then the HOL item would be allocated to slots 1, 3, 6, and 7, thus producing a new state with 1 unused slot. The new state would be stable if and only if the next (i.e., the new HOL) item in the queue has size greater than 1; if it had size 1 it would be allocated slot 8, creating thereby a fully occupied resource and hence a stable state.

Given the allocation algorithm, the fragmentation process embedded at departure epochs is determined by the parameters M, K, and $\mathbf{q}$, and an initial state. A formal definition of this Markov chain is left to the interested reader. It simplifies the analogous definition in [CRS$^+$10] for the original continuous version of the model. We briefly describe this less realistic version as its shortcomings help showcase our new results. Reviews of the two applications we have mentioned, which apply to both the discrete and continuous versions, can be found in [Knu97, ALC09].

Results past and present: In the continuous model, the resource is represented by the unit interval $[0, 1]$, a convenient normalization. The distribution of item sizes is restricted to $[0, \alpha]$, with maximum item size $0 < \alpha \leq 1$. Note that if we take the resource of the discrete model to be the normalized sequence $1/M, 2/M, \ldots, 1$, then the continuous model is approached as K and M increase with the limiting ratio $K/M \to \alpha$ as $K, M \to \infty$. There are two fundamental questions applicable to both the continuous and discrete models.

In the continuous version, depending on the assumed distribution of item sizes, there need not be a positive lower bound to fragment size, and in this case, the first question to be answered is: For a given α, does a stationary regime exist? In particular, does fragmentation continue indefinitely with fragments becoming smaller and smaller? For distributions of practical interest, the answer to the latter question is essentially NO, so a stationary regime does indeed exist. The proof of this property, however, seems to be very difficult (cf. [CRS$^+$10]). This existence question adapted to our more realistic discrete model is trivial – the answer is affirmative and comes down to an appeal to basic theory of finite Markov chains.

The second, equally significant question in the continuous model is: What happens to the extent of fragmentation as $\alpha \to 0$? Experiments clearly suggest that the fragmentation of items increases without bound, a property that is not particularly intuitive. But this remains a fundamental open question of fragmentation theory. *On the other hand, the analogous question can in fact be answered in the more realistic discrete model of this paper.* The question becomes: In the stationary regime, do states tend to exhibit nearly *total fragmentation,*, i.e., are the units of almost all allocated items in mutually disjoint slots? The answer is provably YES for large M. This result for the discrete model is the centerpiece of this paper and is given more precisely below.

Total Fragmentation in the Discrete Model: At any time in the stationary process for given $\mathbf{q}$ and M, let the random variable $N = N(\mathbf{q}, M)$ count the number of items that are not totally fragmented, i.e., the number of items with one or more fragments of width at least two. The theorem below describes the stationary behavior of the fragmentation process embedded at departure epochs. In the order given, the assertions made in the theorem are progressively stronger but less general. All three hold regardless of the allocation algorithm, even one that can look ahead in the queue of waiting items and exploit the sizes of future items in making current allocations.

Theorem 1 *For given K and $\mathbf{q}$, we have $\mathbb{E}N/M \to 0$ as $M \to \infty$. If in addition, $q_1 > 0$, a stronger result holds: There exists a constant $C = C(\mathbf{q}, K)$ such that, for all M sufficiently large, $\mathbb{E}N < C$. Finally, if we also fix $K = 2$, then $\mathbb{E}N \leq 2(1 - q_1)/q_1$.*

The last inequality can be replaced by an equality for an allocation algorithm that always exploits two adjacent empty slots in allocating a size-2 item, whenever such slots exist.

The more difficult proofs of the first two assertions are based on the notion of *bonds*. A bond exists between any pair of adjacent slots which are both empty or both occupied by units of the same item. The basic idea exploits the fact that, in statistical equilibrium, the rate at which bonds are created by arriving items must equal the rate at which they decrease at departures. For example, with $B(t)$ the number of bonds at time t, the proof of the second assertion shows that bonds are destroyed at a rate proportional to $B(t)$, but created at only a constant rate.

Final Remarks On the surface, our theory makes it hard for the engineer to justify the use of fragmentation to improve resource utilization. But there are potentially critical mitigating factors, that must be considered as well, especially when considered in conjunction with efficient defragmentation techniques. For example, the constant C in the second result can be expected to grow quickly as a function of the maximum item size K, so that for many practical applications "all M sufficiently large" refers to values that are far beyond those expected in practice. In the full paper, our experimental studies with uniform distributions $q_i = 1/K,\ 1 < K \leq M$, quantify this observation.

Acknowledgements: This work was supported in part by NSF grant CNS-10-54856.

References

[ALC09] Ian F. Akyildiz, Won-Yeol Lee, and Kaushik R. Chowdhury. CRAHNs: Cognitive radio ad hoc networks. *Ad Hoc Networks*, 7(5):810–836, 2009.

[CRS$^+$10] E. Coffman, P. Robert, F. Simatos, S. Tarumi, and G. Zussman. Channel fragmentation in dynamic spectrum access systems - a theoretical study. In *Proc. ACM SIGMETRICS'10*, Jun. 2010.

[Knu97] Donald E. Knuth. *The Art of Computer Programming, Vol. 1 - Fundamental Algorithms*. Addison Wesley Longman Publishing Co., Redwood City, CA, USA, 3rd edition, 1997.