Performance Evaluation of Fragmented Structures: A Theoretical Study

Ed Coffman[1], **Robert Margolies[1]**, Peter Winkler[2], Gil Zussman[1]
[1]Columbia University, New York, NY
[2]Dartmouth College, Hanover, NH

Motivation

Hard and Solid-state Disk Drives

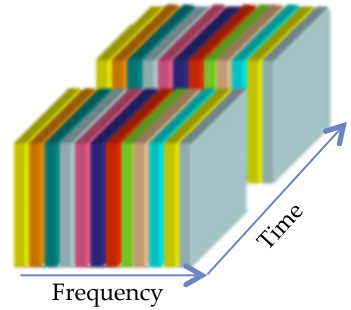

OFDMA , LTE, and Cognitive Radio

Motivation

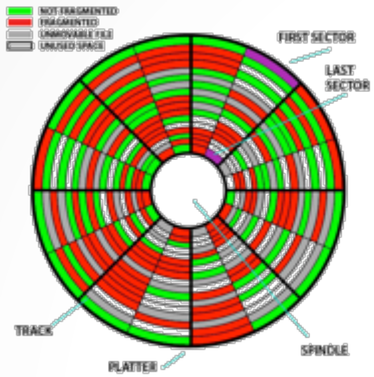

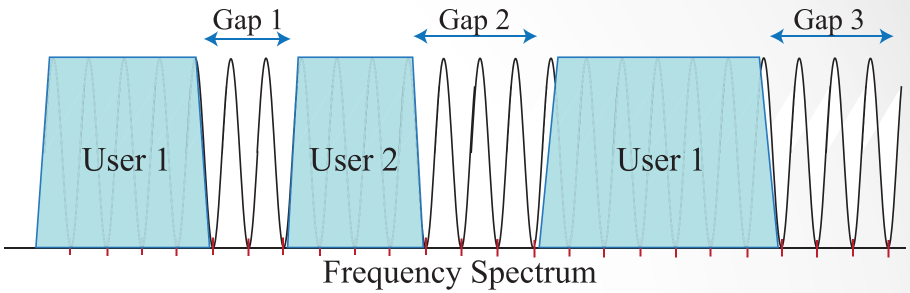

Hard and Solid-state Disk Drives

- Provides interleaving of data allowing for faster read times in some cases
- Disk head motion significantly increases write times

OFDMA , LTE, and Cognitive Radio

- Provides frequency diversity
- Requires complex hardware solutions

Related Work

- Solid-state Disk Drives and OFDMA
  - F. Chen, D.A. Koufaty, X. Zhang, "Understanding intrinsic characteristics and system implications of flash memory based solid state drives", SIGMETRICS Perform. Eval. Rev., vol. 37, no. 1, pp. 81–192, 2009.
  - H. Mahmoud, T. Yucek, and H. Arslan, "OFDM for cognitive radio: Merits and challenges," IEEE Wireless Commun., vol. 16, no. 2, pp. 6–15, Apr. 2009.
  - B. Van Houdt, "A mean field model for a class of garbage collection algorithms in flash-based solid state drives", SIGMETICS Perform. Eval. Rev., vol., no. 1, pp. 191–202, 2013.
  - Y. Li, P.P. Lee, J.C. Lui, "Stochastic modeling of large-scale solid-state storage systems: analysis, design tradeoffs and optimization", SIGMETRICS Perform. Eval. Rev., vol. 41, no. 1, pp. 179–190, 2013.
- Classical Fragmentation Problems
  - E. G. Coffman and F. T. Leighton. "A provably efficient algorithm for dynamic storage allocation." Proceedings of the eighteenth annual ACM symposium on Theory of computing. ACM, 1986.
  - D. E. Knuth, "The Art of Computer Programming, Vol.1 – Fundamental Algorithms", 3rd Edition, 1997.
  - E. Coffman, P. Robert, F. Simatos, S. Tarumi, and G. Zussman, "A performance analysis of channel fragmentation in dynamic spectrum access systems," Queuing Systems, special issue of selected papers from ACM SIGMETRICS'10, vol. 71, no. 3, pp. 293–320, Jul. 2012.

Outline

- ✓ Motivation
- Model + Example
- Motivating Numerical Results
- Asymptotic Theory of Complete Fragmentation
  - Case 1: Items of size 1 or 2
  - Case 2: Items up to size K (size 1 items have positive probability)
  - Case 3: Items up to size K
- Convergence to Complete Fragmentation
- Conclusions

Model

- Resource is modeled as a sequence of $M > 1$ slots
- FIFO queue under full load: there are always waiting items.
- Item sizes are i.i.d. with distribution $q = \{q_1, \ldots, q_K\}$ and have independent i.i.d. exponential residence times.
- An allocation algorithm allocates available gaps in the resource to waiting items

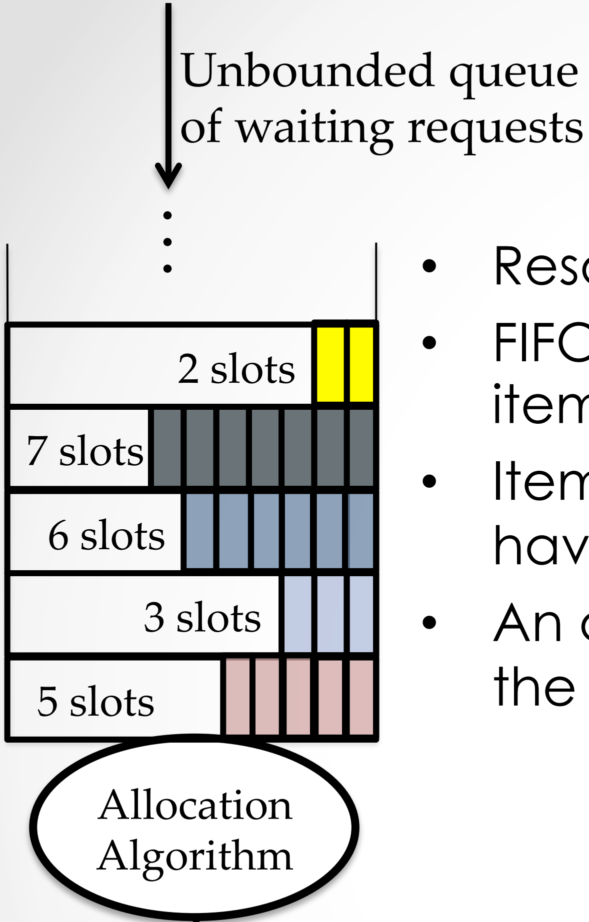

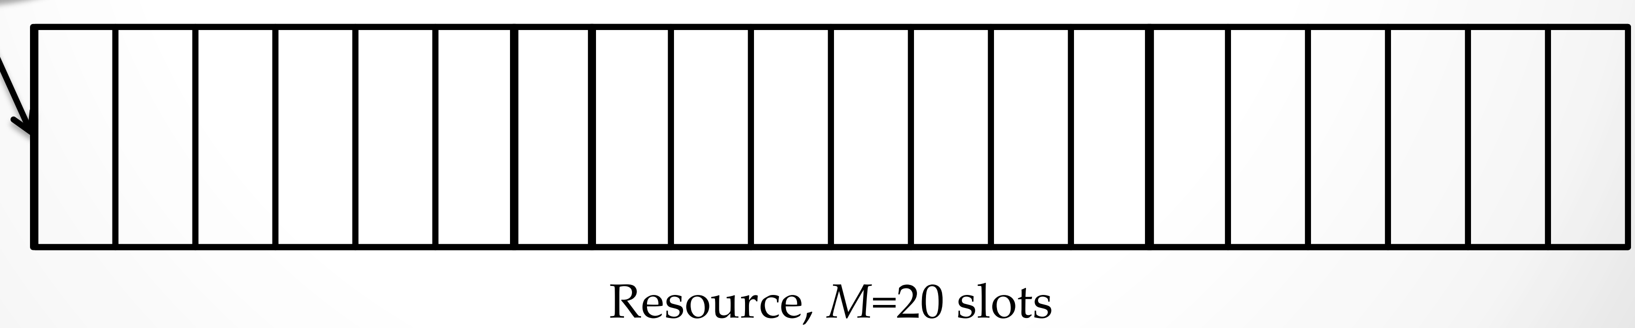

Model

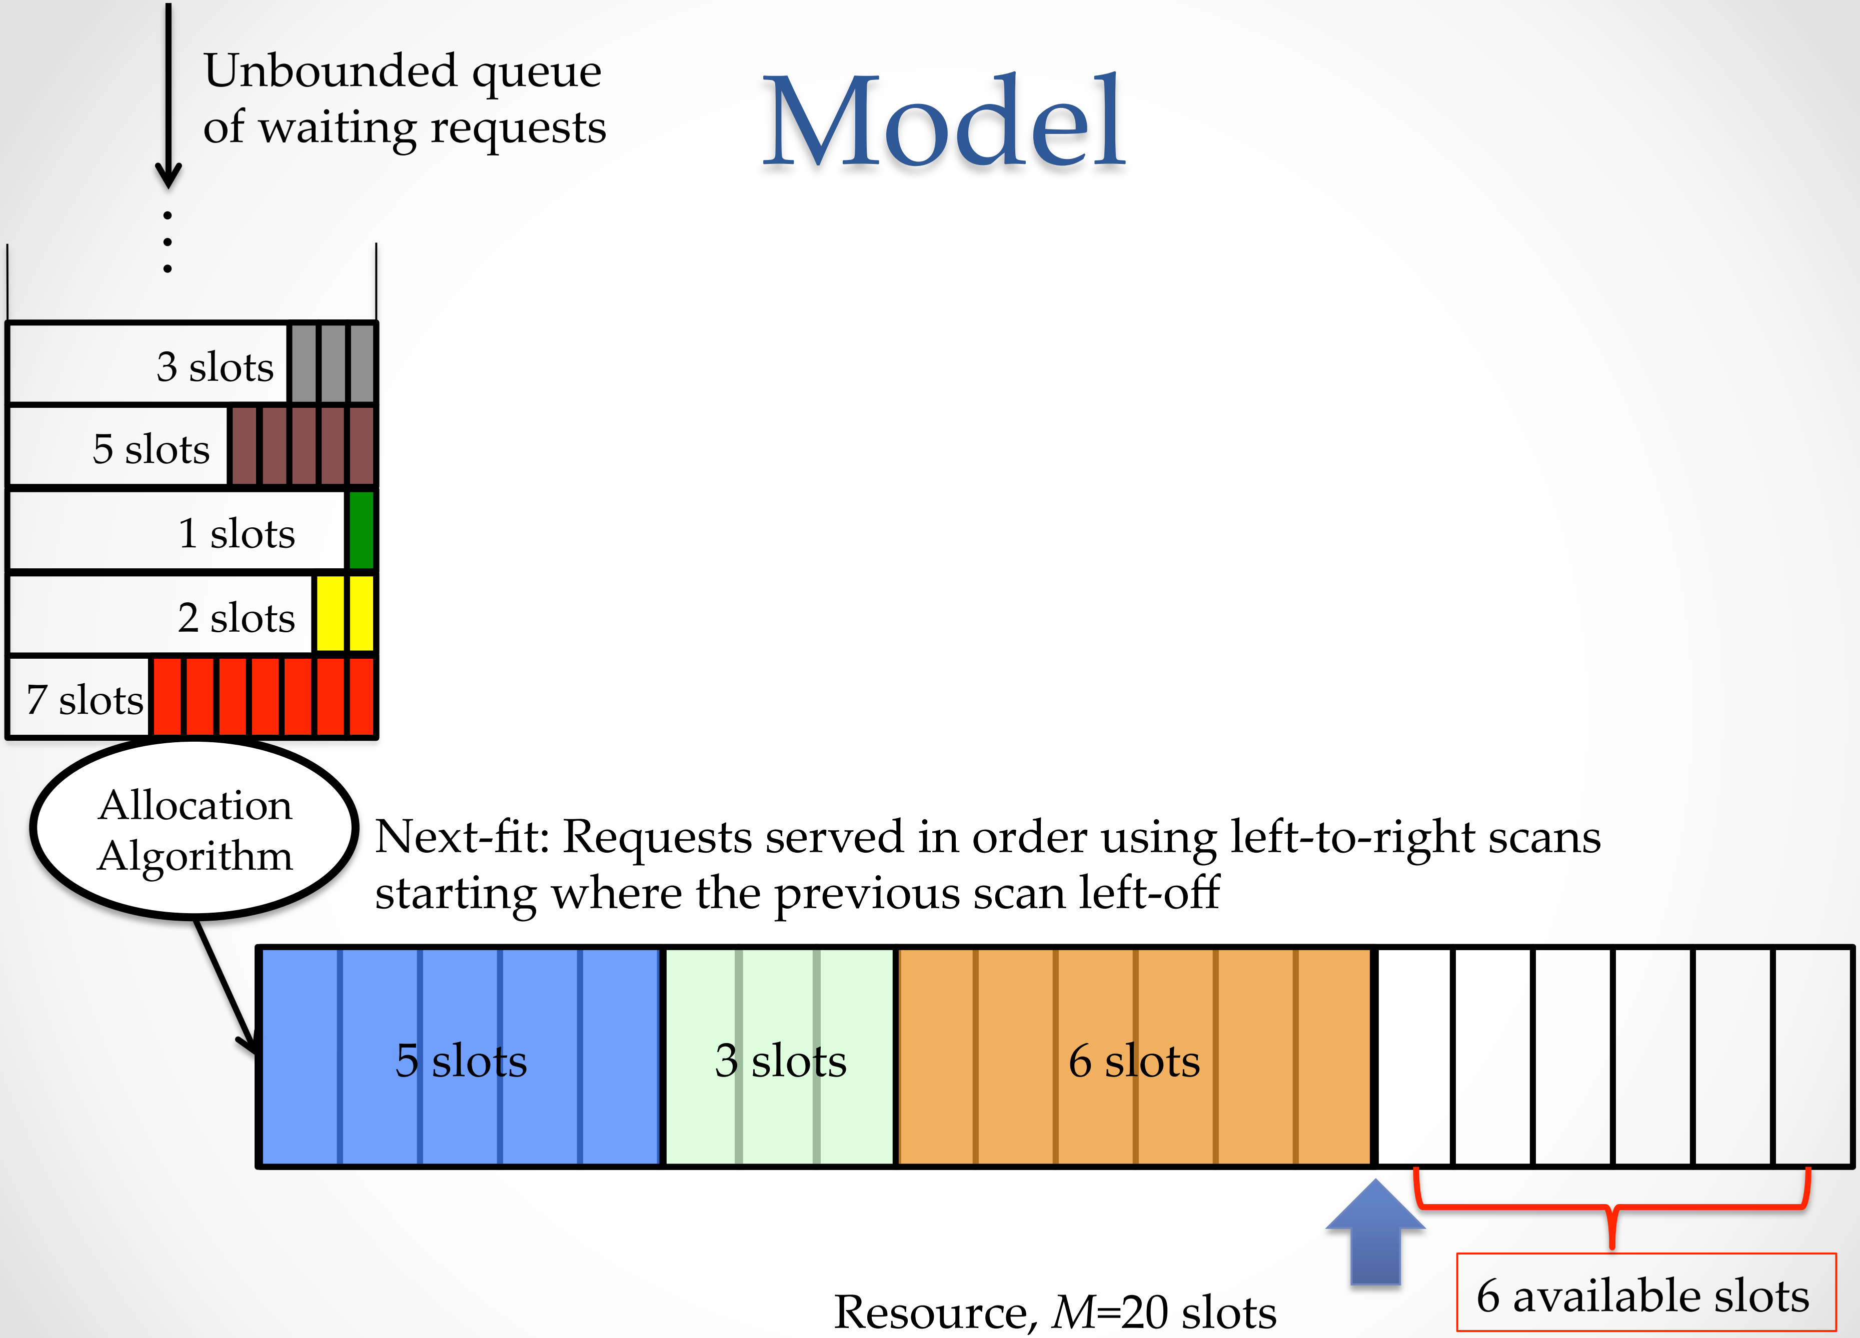
Unbounded queue
of waiting requests
3 slots
5 slots
1 slots
2 slots
7 slots
Allocation
Algorithm
Next-fit: Requests served in order using left-to-right scans
starting where the previous scan left-off
5 slots
3 slots
6 slots
6 available slots
Resource, M=20 slots

Model

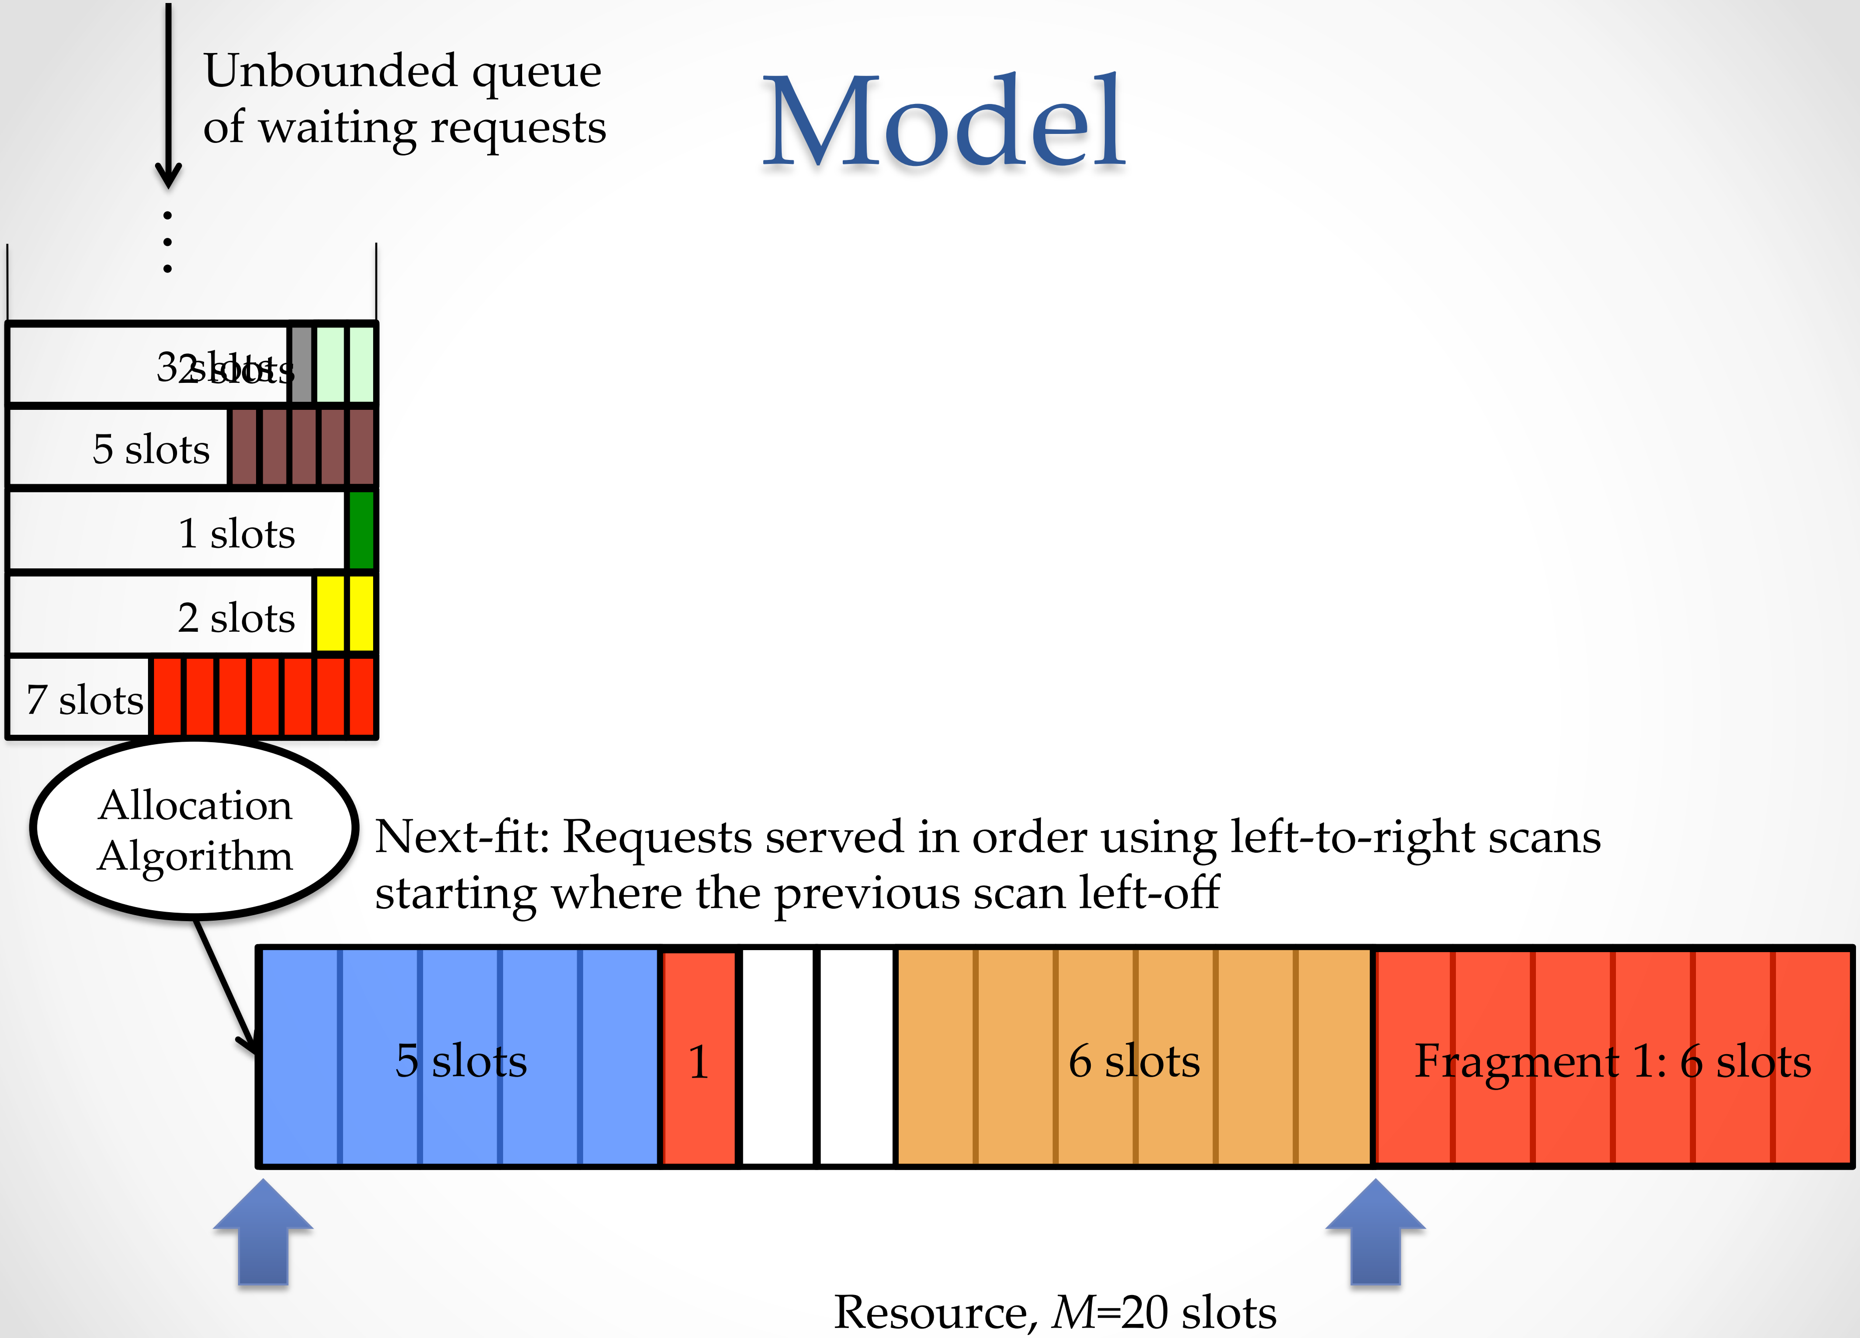
Unbounded queue
of waiting requests
5 slots
1 slots
2 slots
7 slots
Allocation
Algorithm
Next-fit: Requests served in order using left-to-right scans
starting where the previous scan left-off
5 slots
1
6 slots
Fragment 1: 6 slots
Resource, M=20 slots

Complete Fragmentation

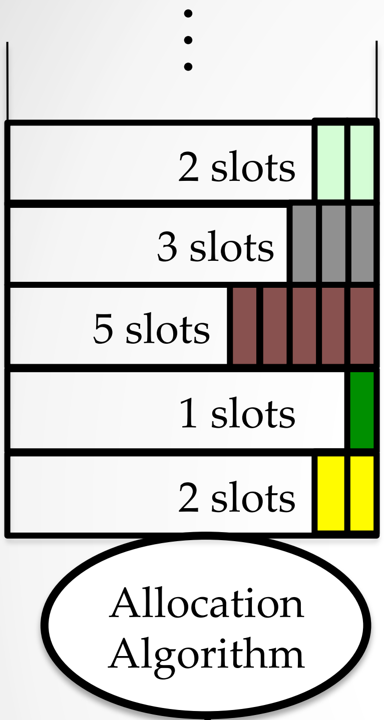

After a *long* time, does fragmentation progress to a point where nearly all items are *completely fragmented?*

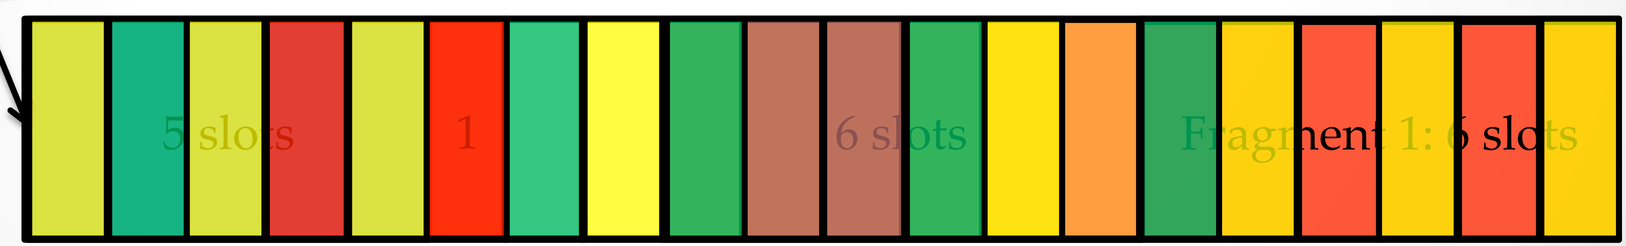

Complete fragmentation: When no two allocated slots of an item are adjacent.

Numerical Examples of Fragmentation

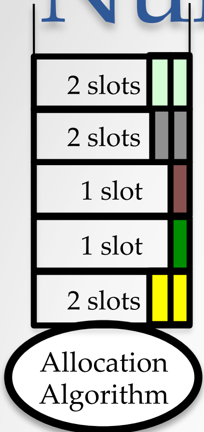

- Unbounded queue restricted to waiting size-1 or size-2 items
- P(size-1 item) = q_1=1/2
- P(size-2 item) = q_2=1/2

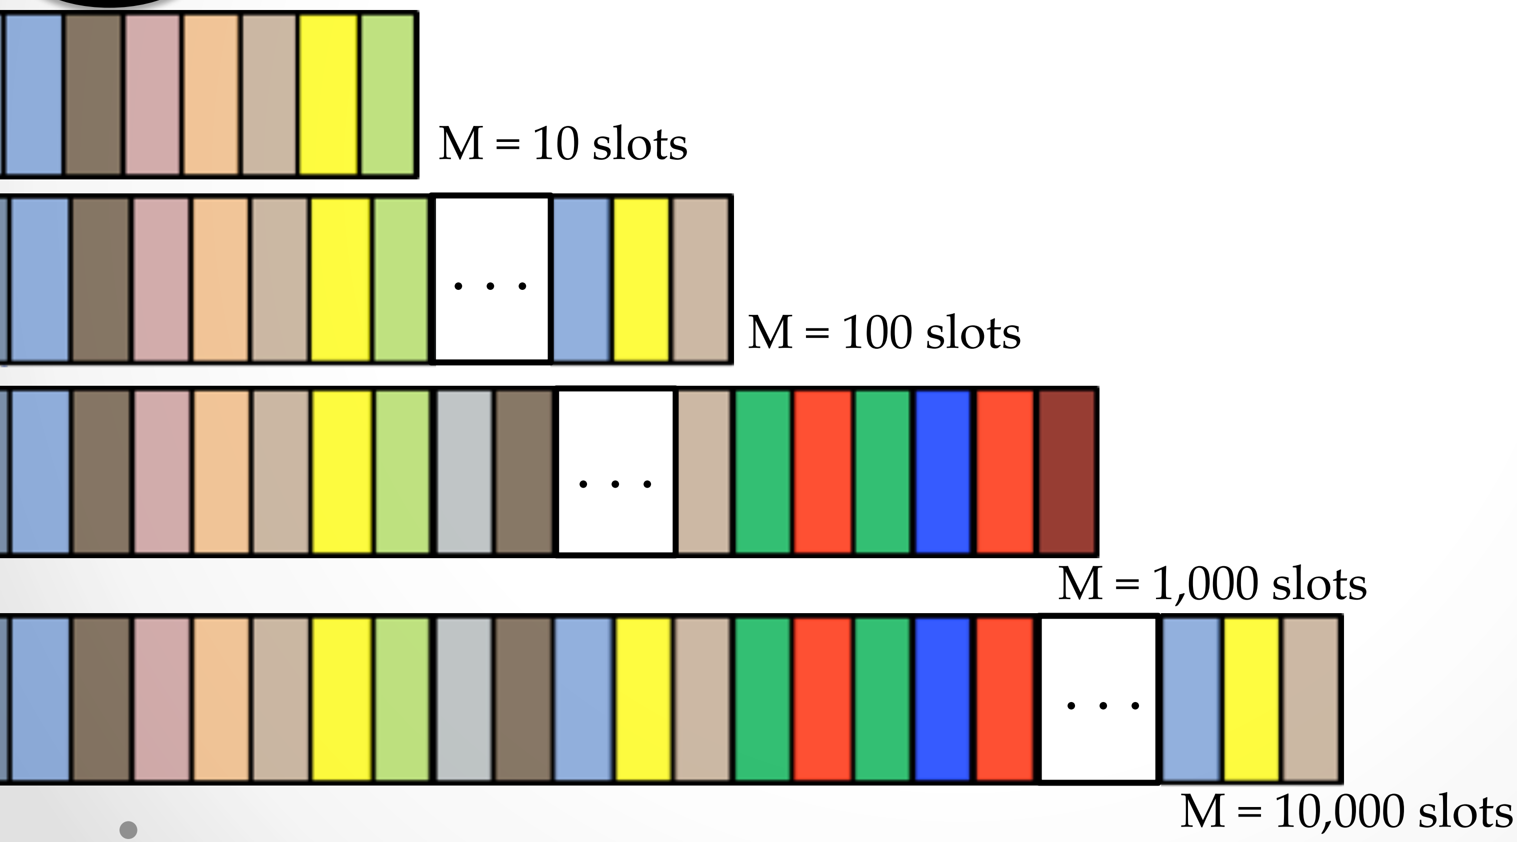

Average # of Unfragmented Items
0.79
0.89
0.91
0.92

Numerical Examples of Fragmentation

Probability of a size-*j* item allocated *j* fragmented slots.

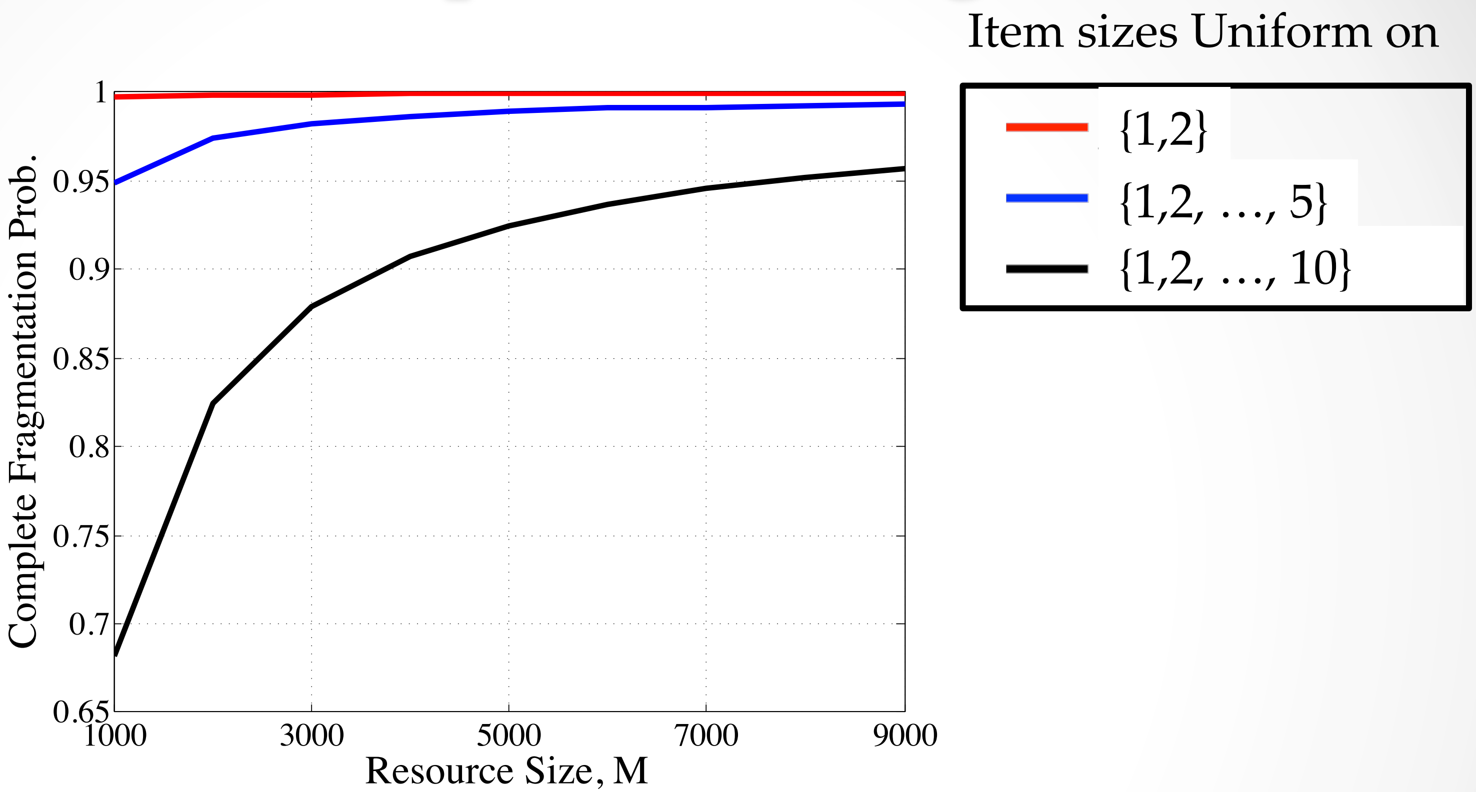

Nearly all items are completely fragmented as $M \rightarrow \infty$

Outline

- ✓ Motivation
- ✓ Model + Example
- ✓ Motivating Numerical Results
- Asymptotic Theory of Complete Fragmentation
  - Case 1: Items of size 1 or 2
  - Case 2: Items up to size K (size 1 items have positive probability)
  - Case 3: Items up to size K
- Convergence to Complete Fragmentation
- Conclusions

Complete Fragmentation: Case 1

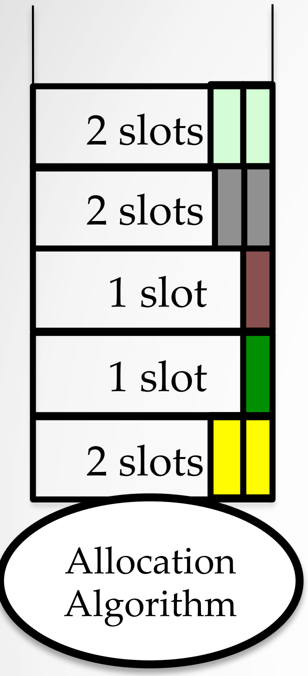

- Infinite queue restricted to waiting size-1 or size-2 items
- P(size-1 item) = $q_1 > 0$
- P(size-2 item) = $q_2 > 0$
- Stable state represented as (G,H) with $G < H$
    - G is size of available gaps
    - H is size of Head-of-line item

Goal: Show that as the size of the resource (*M*) grows, *nearly* all of the size-2 items are *fragmented*.

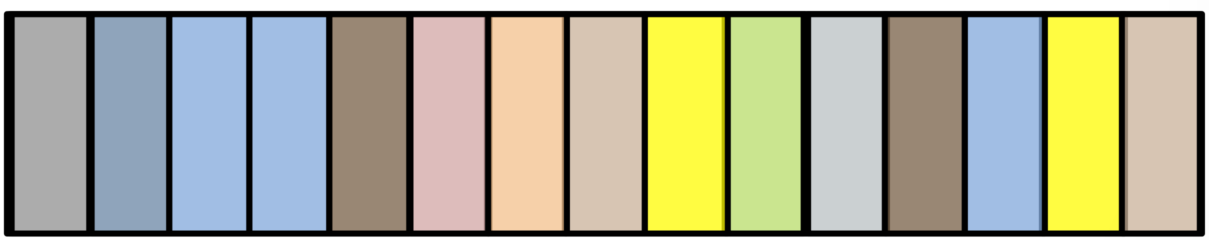

Resource, *M* slots

Complete Fragmentation: Case 1

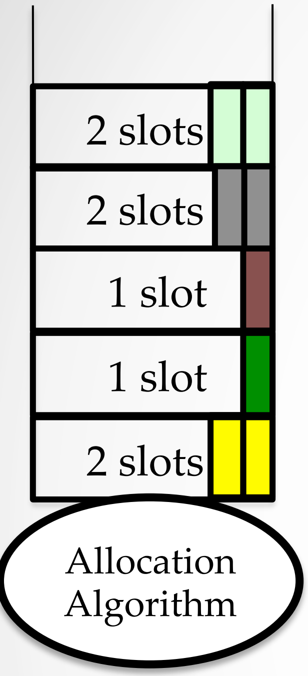

- Infinite queue restricted to waiting size-1 or size-2 items
- P(size-1 item) = $q_1 > 0$
- P(size-2 item) = $q_2 > 0$
- Stable state represented as (G,H) with G<H
  - G is size of available gaps
  - H is size of Head-of-line item

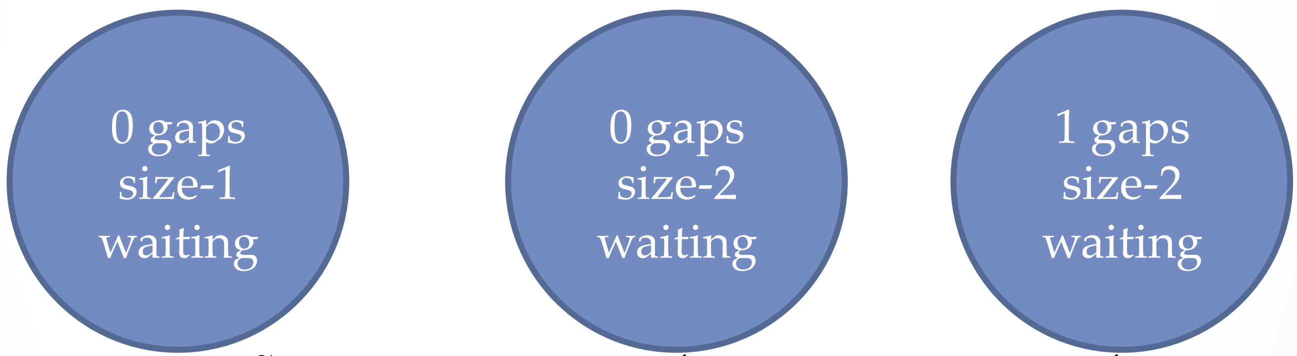

$$\pi_{0,1} = \frac{q_1}{2-q_1} \qquad \pi_{0,2} = \frac{1-q_1}{2-q_1} \qquad \pi_{1,2} = \frac{1-q_1}{2-q_1}$$

Balance the rate of increase/decrease of the number of unfragmented size 2 items, denoted $N_{2,1}$

C. Kipnis, P. Robert, A dynamic storage process, Stochastic processes and their applications 34 (1) (1990) 155–169.

Complete Fragmentation: Case 1

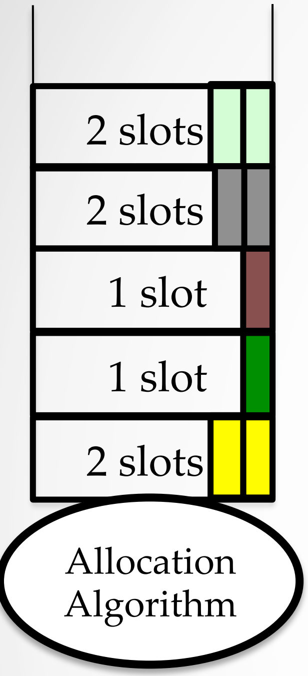

0 gaps size-2 waiting

Event – size 1 item departs with probability q_1

Result: No change in $N_{2,1}$ and move to state (1,2)

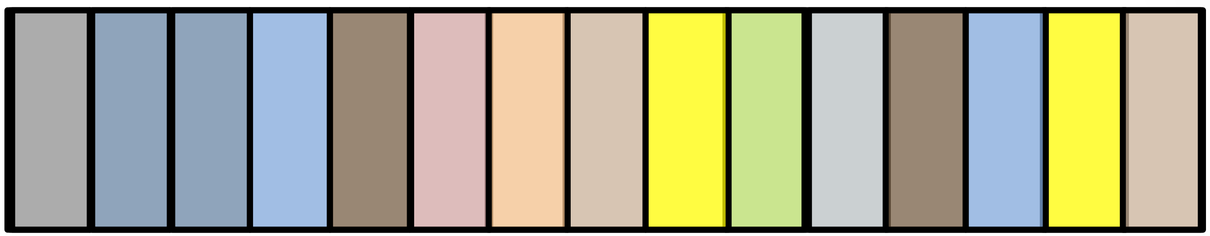

Event – size 2 item departs with probability q_2

Result: Waiting size-2 item will occupy departed size 2 item and no change in $N_{2,1}$

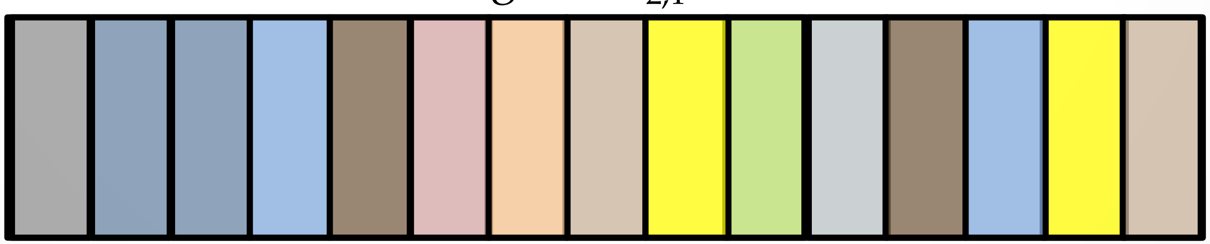

Conclusion: $N_{2,1}$ does not change in state (0,2)

Complete Fragmentation: Case 1

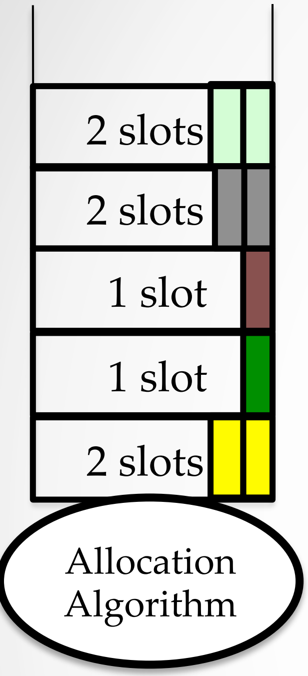

1 gap size-2 waiting

Event – size 1 item departs with probability q_1

Result: $N_{2,1}$ increases at rate 1 if a size-1 item departs adjacent to the present gap

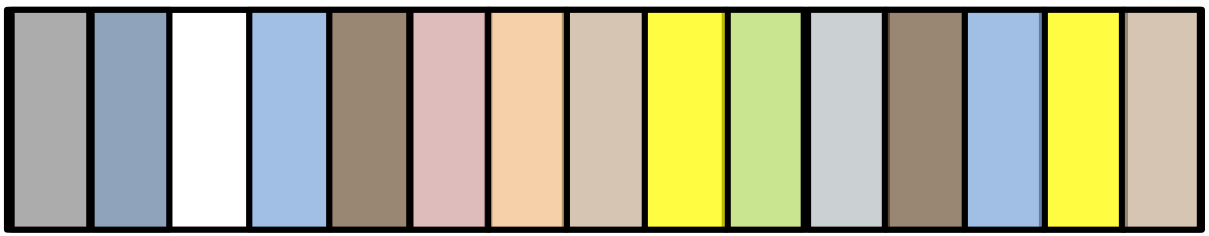

Event – size 2 item departs with probability q_2

Result: $N_{2,1}$ increases at rate at most 2 if one of the 2 slots in the departing size-2 item is adjacent to the present gap.

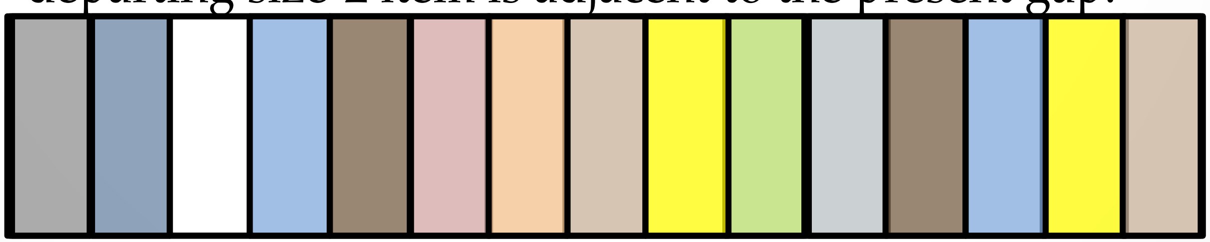

Conclusion: $N_{2,1}$ increases at rate at most 2 when in state (1,2)

Complete Fragmentation: Case 1

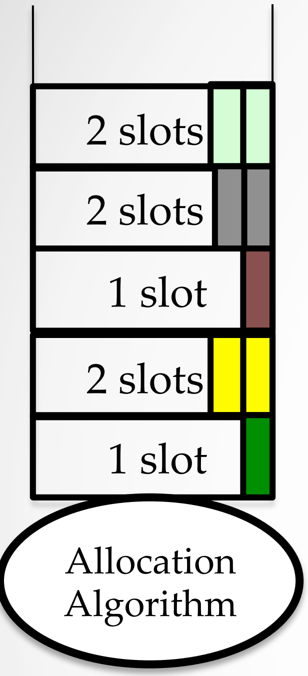

0 gaps size-1 waiting

Event – size 1 item departs with probability q_1

Result: No change in N_{21} and move to state (0,1) or (0,2)

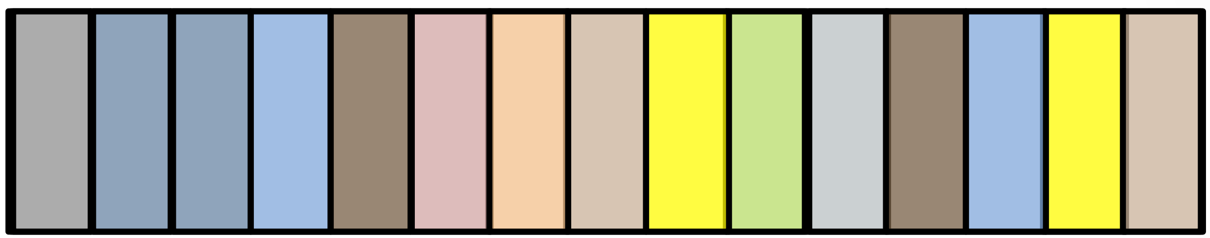

Event – size 2 item departs with probability q_2

Result: N_{21} decreases by 1 if an unfragmented size-2 item departs which occurs with rate $N_{21}(t)$

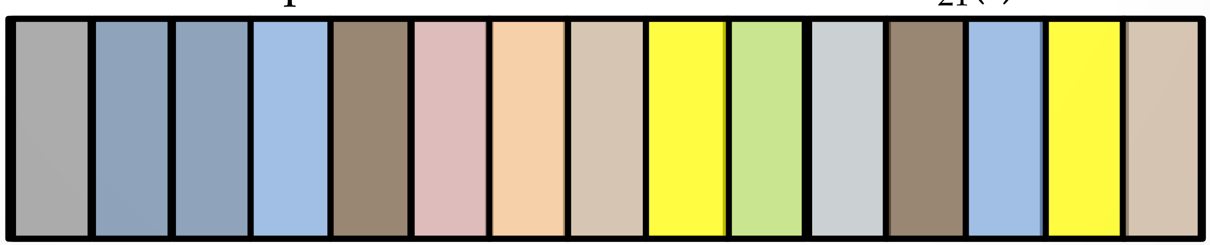

Conclusion: N_{21} decreases at rate $N_{21}(t)$ when in state (0,1)

Complete Fragmentation: Case 1

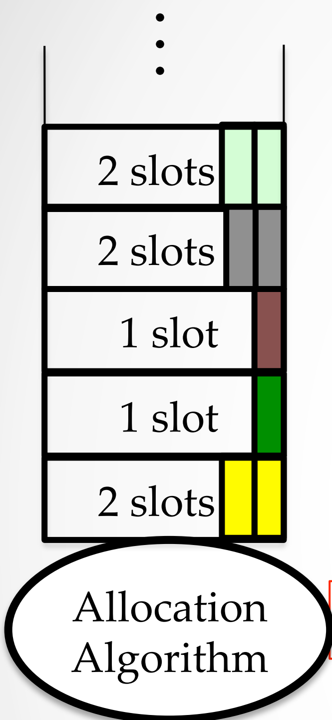

0 gaps size-1 waiting	0 gaps size-2 waiting	1 gaps size-2 waiting
$\pi_{0,1} = \frac{q_1}{2 - q_1}$	$\pi_{0,2} = \frac{1 - q_1}{2 - q_1}$	$\pi_{1,2} = \frac{1 - q_1}{2 - q_1}$
Decreases at rate $N_{21}(t)$	No Change in N_{21}	N_{21} Increases at rate at most 2

$$\mathbb{E}[N_{21}] \leq 2\frac{1 - q_1}{q_1}$$

All but a constant number of size-2 items become completely fragmented

Resource-size grows to infinity

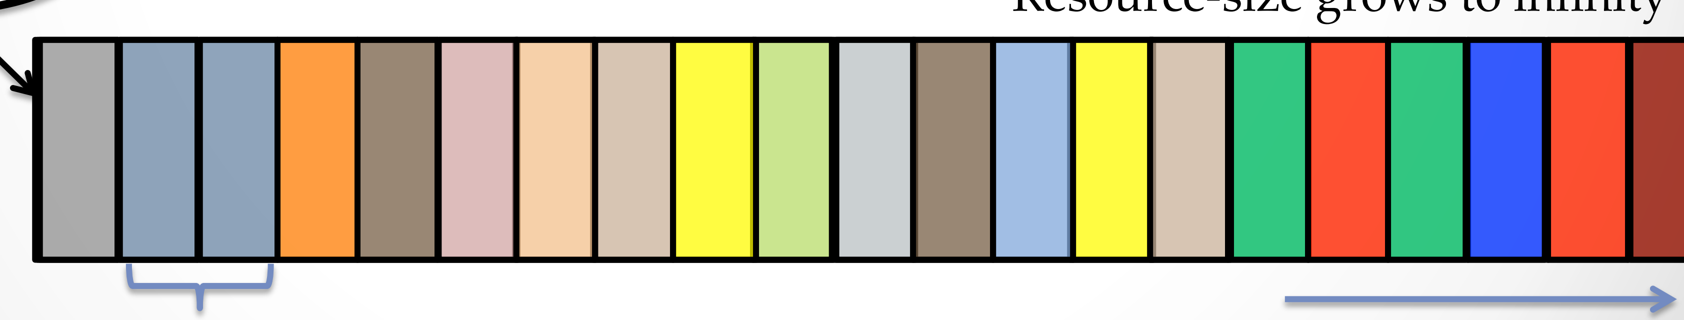

$\leq 2(1 - q_1)/q_1$

Complete Fragmentation: Case 2

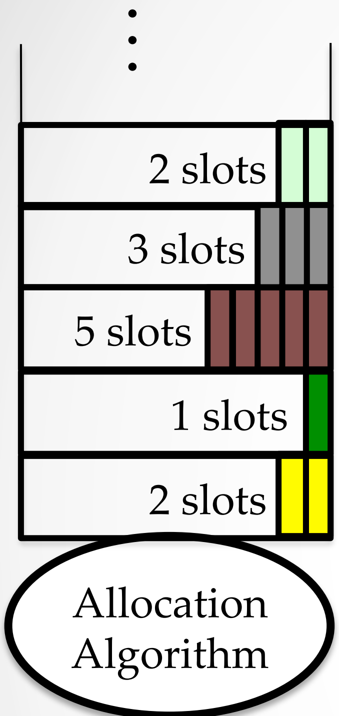

- Infinite queue of waiting items up to size-K
- P(item of size j) = q_j
- Positive probability of size-1 items: $q_1>0$

All but a constant number of items become completely fragmented

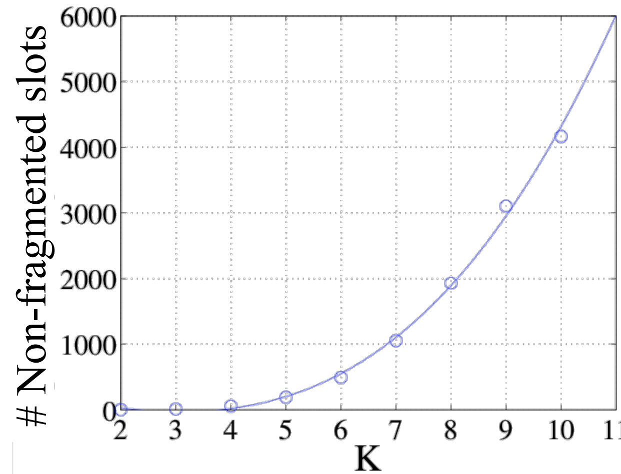

Resource-size grows to infinity

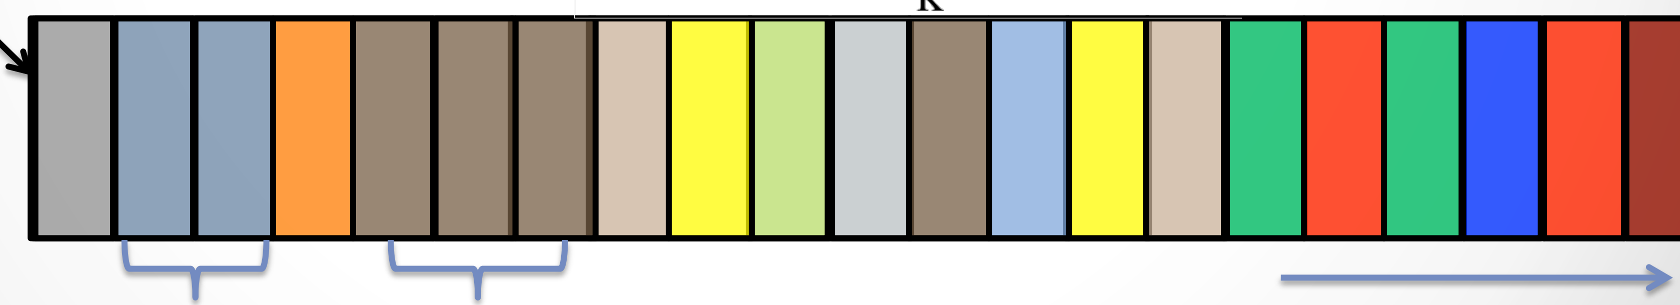

non-completely fragmented items

Complete Fragmentation: Case 3

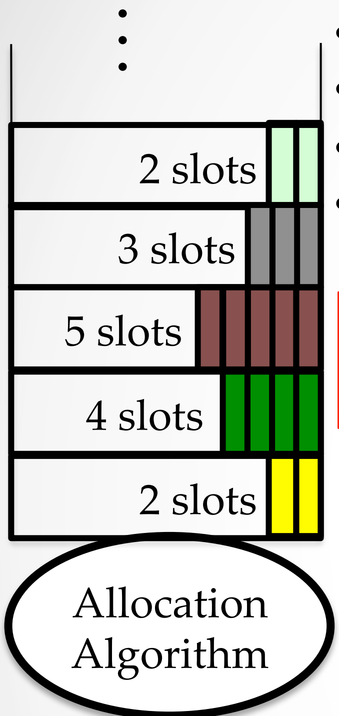

- Infinite queue of waiting items up to size-K
- P(item of size j) = q_j
- ~~Positive probability of size-1 items: q_1>0~~
- Item sizes with positive probability have a non-trivial common divisor

The fraction of completely fragmented items tends to 1 as resource size grows to infinity.

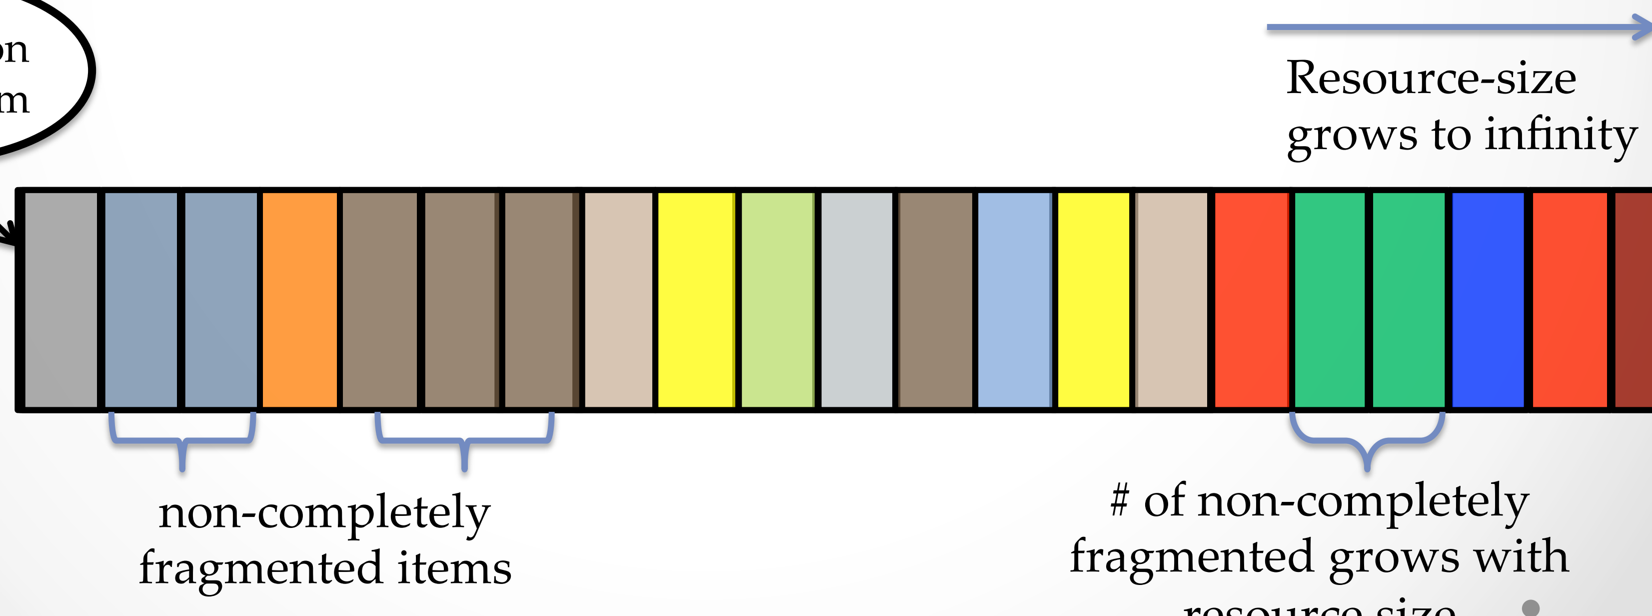

Complete Fragmentation

	Number of unfragmented items:
Case 1:Items of Size 1 and 2	Bounded by $\mathbb{E}[N_{21}] \leq 2\frac{1-q_1}{q_1}$
Case 2: Items up to size K (size 1 items have positive probability)	Bounded by a constant C
Case 3: Items up to size K	Grow at a rate o(M)

Implication: For all cases, complete fragmentation is approached as

$$\lim_{M\to\infty} \frac{\#\text{ of Unfragmented Items}}{M} = 0$$

where M is the size of the resource.

Outline

- ✓ Motivation
- ✓ Model + Example
- ✓ Initial Experimental Results
- ✓ Asymptotic Theory of Complete Fragmentation
  - ✓ Case 1: Items of size 1 or 2
  - ✓ Case 2: Items up to size K (size 1 items w.p.p)
  - ✓ Case 3: Items up to size K
- Convergence to Complete Fragmentation
- Conclusions

Convergence to Nearly Complete Fragmentation: Time

Fragmentation growth is most logarithmic in the number of departures.

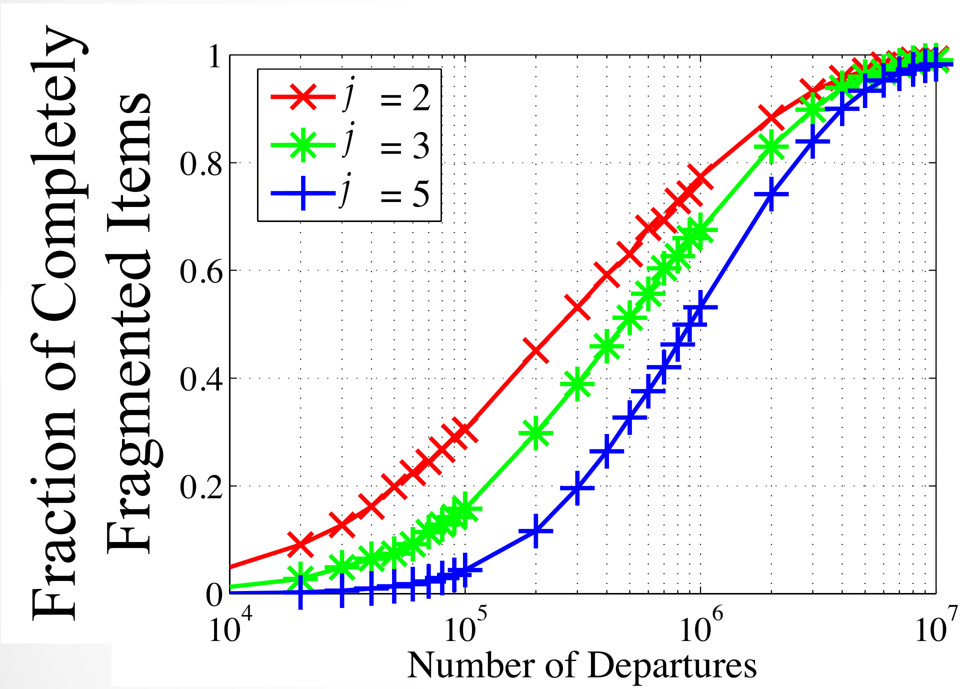

When items are of size 1 or 2, the number of departures until the fraction of fragmented requests increase by ϒ is

$$\approx M \frac{1-q_1}{2-q_1} \ln \frac{1}{1-\gamma}$$

M = 100,000 slots

Item size distribution ~ uniform on {1,…,5}

Convergence to Nearly Complete Fragmentation: Space

Recall that: $\lim_{M\to\infty} \frac{\text{\# of Unfragmented Items}}{M} = 0$

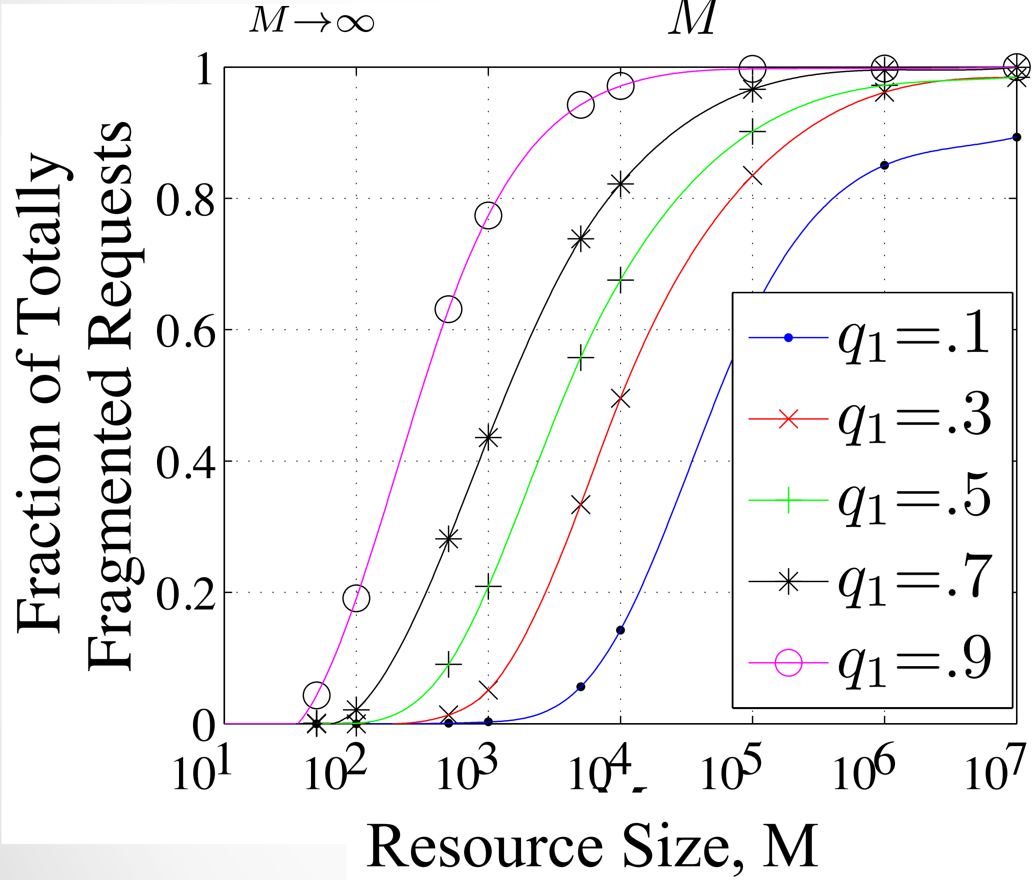

The size of the resource (M) required to approach complete fragmentation can be enormous even for simple item size distributions

Item size distribution ~ {1,9} with probabilities q_1 and q_9=(1-q_1)

Conclusions

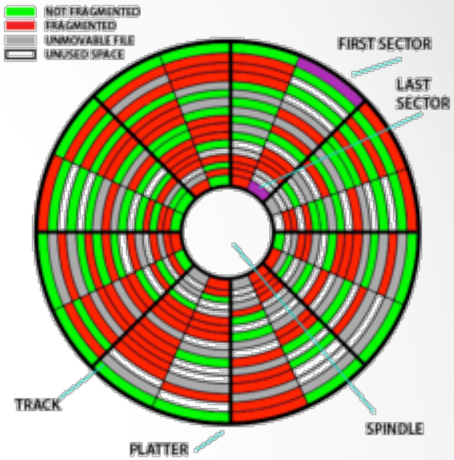

- Nearly all items become *completely fragmented* in statistical equilibrium when the resource size grows to infinity
  - Proofs for cases 1 and 2 balance the rates at which the number of non-fragmented items increases and decreases in equilibrium
  - Good news: frequency diversity in OFDMA, defragmentation software
  - Bad news: requires complex hardware solutions
- Convergence rates can be surprisingly slow in the limits of time and resource size

- Robert Margolies
- robm@ee.columbia.edu