

# CSEP 521 - Applied Algorithms

## On-line Algorithms

1

## Introduction

Online Algorithms are algorithms that need to make decisions without full knowledge of the input. They have full knowledge of the past but no (or partial) knowledge of the future.

For this type of problem we will attempt to design algorithms that are competitive with the optimum offline algorithm, the algorithm that has perfect knowledge of the future.

2

## The Ski-Rental Problem

- Assume that you are taking ski lessons. After each lesson you decide (depending on how much you enjoy it, and what is your bones status) whether to continue to ski or to stop totally.
- You have the choice of either renting skis for 1\$ a time or buying skis for  $y$ \$.
- Will you buy or rent?



3

## The Ski-Rental Problem

- If you knew in advance how many times  $t$  you would ski in your life then the choice of whether to rent or buy is simple. If you will ski more than  $y$  times then buy before you start, otherwise always rent.
- The cost of this algorithm is  $\min(t, y)$ .
- This type of strategy, with perfect knowledge of the future, is known as an offline strategy.



4

## The Ski-Rental Problem

- In practice, you don't know how many times you will ski. What should you do?
- An online strategy will be a number  $k$  such that after renting  $k-1$  times you will buy skis (just before your  $k^{\text{th}}$  visit).
- Claim: Setting  $k = y$  guarantees that you never pay more than twice the cost of the offline strategy.
- Example: Assume  $y=7$  Thus, after 6 rents, you buy. Your total payment:  $6+7=13$ .

5

## The Ski-Rental Problem

Theorem: Setting  $k = y$  guarantees that you never pay more than twice the cost of the offline strategy.

Proof: when you buy skis in your  $k^{\text{th}}$  visit, even if you quit right after this time,  $t \geq y$ .

- Your total payment is  $k-1+y = 2y-1$ .
- The offline cost is  $\min(t, y) = y$ .
- The ratio is  $(2y-1)/y = 2-1/y$ . ■

We say that this strategy is  $(2-1/y)$ -competitive.

6

## The Ski-Rental Problem

Is there a better strategy?

- Let  $k$  be any strategy (buy after  $k-1$  rents).
- Suppose you buy the skis at the  $k^{\text{th}}$  time and then break your leg and never ski again.
- Your total ski cost is  $k-1+y$  and the optimum offline cost is  $\min(k,y)$ .
- For every  $k$ , the ratio  $(k-1+y)/\min(k,y)$  is at least  $(2-1/y)$
- Therefore, every strategy is at least  $(2-1/y)$ -competitive. ■

7

## The Ski-Rental Problem

The general rule:

When balancing small incremental costs against a big one-time cost, you want to delay spending the big cost until you have accumulated roughly the same amount in small costs.

8

## The Lost Cow Problem

Old McDonald lost his favorite cow. It was last seen marching towards a junction leading to two infinite roads. None of the witnesses can say if the cow picked the left or the right route.



9

## The Lost Cow Problem

Old McDonald's algorithm:

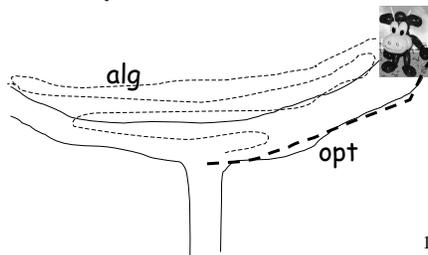
1.  $d=1$ ; current side = right
2. repeat:
  - i. Walk distance  $d$  on current side
  - ii. if find cow then exit
  - iii. else return to starting point
  - iv.  $d = 2d$
  - v. Flip current side

10

## The Lost Cow Problem

Theorem: Old McDonald's algorithm is 9-competitive.

In other words: The distance that Old McDonald might pass before finding the cow is at most 9 times the distance of an optimal offline algorithm (who knows where the cow is).



11

## The Lost Cow Problem

Theorem: Old McDonald's algorithm is 9-competitive.

Proof: The worst case is that he finds the cow a little bit beyond the distance he last searched on this side (why?).

Thus,  $OPT = 2^j + \epsilon$  where  $j = \#$  of iterations and  $\epsilon$  is some small distance. Then,

$$\text{Cost OPT} = 2^j + \epsilon > 2^j$$

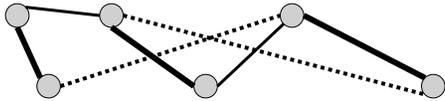
$$\text{Cost ON} = 2(1 + 2 + 4 + \dots + 2^{j+1}) + 2^j + \epsilon$$

$$= 2 \cdot 2^{j+2} + 2^j + \epsilon = 9 \cdot 2^j + \epsilon < 9 \cdot \text{Cost OPT} \quad \blacksquare$$

12

## Edge Coloring

- An Edge-coloring of a graph  $G=(V,E)$  is an assignment,  $c$ , of integers to the edges such that if  $e_1$  and  $e_2$  share an endpoint then  $c(e_1) \neq c(e_2)$ .



- Let  $\Delta$  be the maximal degree of some vertex in  $G$ .
- In the offline case, it is possible to edge-color  $G$  using  $\Delta$  or  $\Delta+1$  colors (which is almost optimal).
- Online edge coloring: The graph is not known in advance. In each step a new edge is added and we need to color it before the next edge is known.

13

## Optimal Online Algorithm for Edge Coloring

- We color the edges with numbers  $1,2,3,\dots$
- Let  $e=(u,v)$  be a new edge.

Color  $e$  with the smallest color which is not used by any edge adjacent to  $u$  or  $v$ .

Claim: The algorithm uses at most  $2\Delta-1$  colors.

Proof outline (was hw6 q.1): assume we need the color  $2\Delta$ . It must be that all the colors  $1,2,\dots,2\Delta-1$  are used by edges adjacent to  $u$  or  $v$ . Therefore, either  $u$  or  $v$  has  $\Delta$  adjacent edges, excluding  $e$ , contradicting the definition of  $\Delta$ .



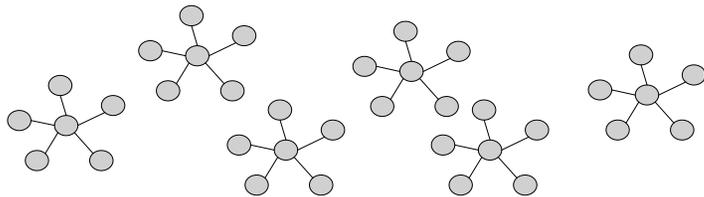
14

## Online Edge Coloring

Claim: Any deterministic algorithm needs at least  $2\Delta-1$  colors.

Proof: Assume  $\exists$  an algorithm that uses only  $2\Delta-2$  colors. Given  $\Delta$  we add to the graph many  $(\Delta-1)$ -stars.

Example:  
 $\Delta=6$

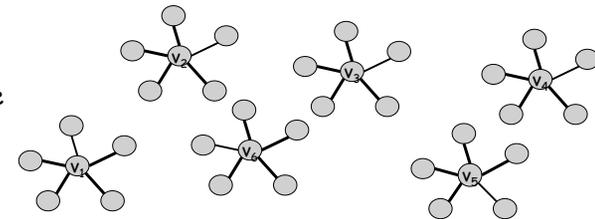


There is a finite number of ways to edge-color a  $(\Delta-1)$ -star with colors from  $\{1,2,\dots,2\Delta-2\}$ , so at some point we must have  $\Delta$  stars, all colored with the the same set of  $\Delta-1$  colors.

15

## Online Edge Coloring

$\Delta$  stars, all colored with the the same set of  $\Delta-1$  colors.



Let  $v_1, v_2, \dots, v_\Delta$  be the centers of these stars.

We are ready to shock the algorithm!

We add a new vertex,  $a$ , and  $\Delta$  edges  $(a-v_1), \dots, (a,v_\Delta)$ .

Each new edge must have a unique color (why?), that is not one of the  $(\Delta-1)$  colors used to color the stars (why?)  $\rightarrow 2\Delta-1$  colors must be used.

Note: the maximal degree is  $\Delta$

16

## Online Scheduling and Load Balancing

Problem Statement:

- A set of  $m$  identical machines,
- A sequence of jobs with processing times  $p_1, p_2, \dots$
- Each job must be assigned to one of the machines.
- When job  $j$  is scheduled, we don't know how many additional jobs we are going to have and what are their processing times.

Goal: schedule the jobs on machines in a way that minimizes the makespan =  $\max_i \sum_{j \text{ on } M_i} p_j$ .  
(the maximal load on one machine)

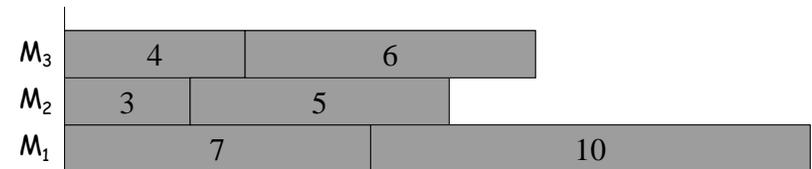
17

## Online Scheduling and Load Balancing

List Scheduling [Graham 1966]:

A greedy algorithm: always schedule a job on the least loaded machine.

Example:  $m=3$   $\sigma = 7 \ 3 \ 4 \ 5 \ 6 \ 10$



Makespan = 17

18

## Online Scheduling and Load Balancing

Theorem: List- Scheduling is  $(2-1/m)$ - competitive.

Proof: Let  $H_j$  denote the last completion time on the  $j^{\text{th}}$  machine. Let  $k$  be the job that finishes last and determines  $C_{LS}$ .

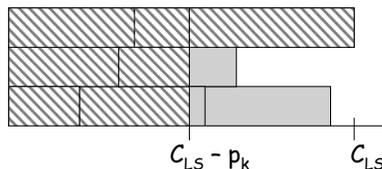
All the machines are busy when  $j$  starts its processing, thus,  $\forall j, H_j \geq C_{LS} - p_k$ .

For at least one machine (that processes  $k$ )  $H_j = C_{LS}$ .

$$\rightarrow \sum_i p_i = \sum_j H_j \geq (m-1)(C_{LS} - p_k) + C_{LS}$$

$$\rightarrow \sum_i p_i + (m-1)p_k \geq mC_{LS}$$

$$\rightarrow C_{LS} \leq 1/m \sum_i p_i + p_k (m-1)/m$$



## Online Scheduling and Load Balancing

$$\rightarrow C_{LS} \leq 1/m \sum_i p_i + p_k (m-1)/m$$

Consider an optimal offline schedule.

$C_{opt} \geq \max_i p_i \geq p_k$  (some machine must process the longest job).

$C_{opt} \geq 1/m \sum_i p_i$  (if the load is perfectly balanced).

Therefore,

$$C_{LS} \leq C_{opt} + C_{opt} (m-1)/m = (2-1/m) C_{opt}$$

20

## Online Scheduling

Are there any better algorithms?

Not significantly. Randomization do help.

| m        | deterministic |             |       | randomized  |             |
|----------|---------------|-------------|-------|-------------|-------------|
|          | lower bound   | upper bound | LS    | lower bound | upper bound |
| 2        | 1.5           | 1.5         | 1.5   | 1.334       | 1.334       |
| 3        | 1.666         | 1.667       | 1.667 | 1.42        | 1.55        |
| 4        | 1.731         | 1.733       | 1.75  | 1.46        | 1.66        |
| $\infty$ | 1.852         | 1.923       | 2     | 1.58        | ---         |

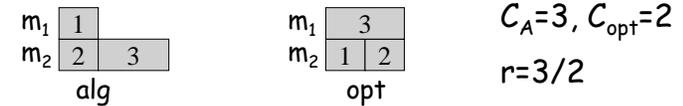
21

## A lower Bound for Online Scheduling

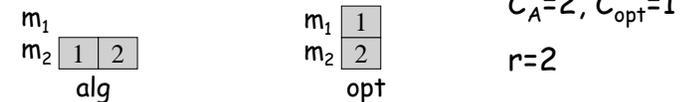
Theorem: For  $m=2$ , no algorithm has  $r < 1.5$

Proof: Consider the sequence  $\sigma = 1,1,2$ .

If the first two jobs are scheduled on different machines, the third job completes at time 3.



If the first two jobs are scheduled on the same machine, the adversary stops.



22

## Paging- Cache Replacement Policies

Problem Statement:

- There are two levels of memory:
  - fast memory  $M_1$  consisting of  $k$  pages (cache)
  - slow memory  $M_2$  consisting of  $n$  pages ( $k < n$ ).
- Pages in  $M_1$  are a strict subset of the pages in  $M_2$ .
- Pages are accessible only through  $M_1$ .
- Accessing a page contained in  $M_1$  has cost 0.
- When accessing a page not in  $M_1$ , it must first be brought in from  $M_2$  at a cost of 1 before it can be accessed. This event is called a page fault.

23

## Paging- Cache Replacement Policies

Problem Statement (cont.):

If  $M_1$  is full when a page fault occurs, some page in  $M_1$  must be evicted in order to make room in  $M_1$ .

How to choose a page to evict each time a page fault occurs in a way that minimizes the total number of page faults over time?

24

# Paging- An Optimal Offline Algorithm

Algorithm LFD (Longest-Forward-Distance)

An optimal off-line page replacement strategy.

On each page fault, evict the page in  $M_1$  that will be requested farthest in the future.

Example:  $M_2 = \{a, b, c, d, e\}$   $n=5$ ,  $k=3$

$\sigma = a, b, c, d, a, b, e, d, e, b, c, c, a, d$

$$M_1 = \begin{cases} a & a & a & a & e & e & e & e & c & c & c & c \\ b & b & b & b & b & b & b & b & b & b & a & a \\ c & d & d & d & d & d & d & d & d & d & d & d \end{cases}$$

\*                    \*                    \*                    \*

4 cache misses in LFD

# Paging- An Optimal Offline Algorithm

A classic result from 1966:

LFD is an optimal page replacement policy.

Proof idea: For any other algorithm  $A$ , the cost of  $A$  is not increased if in the 1<sup>st</sup> time that  $A$  differs from LFD we evict in  $A$  the page that is requested farthest in the future.

However, LFD is not practical.

It is not an *online* algorithm!

## Online Paging Algorithms

FIFO: first in first out: evict the page that was entered first to the cache.

Example:  $M_2 = \{a, b, c, d, e\}$   $n=5$ ,  $k=3$

$\sigma = a, b, c, d, a, b, e, d, e, b, c, c, a, d$

$$M_1 = \begin{cases} a & d & d & d & e & e & e & e & e & e & a & a \\ b & b & a & a & a & d & d & d & d & d & d & d \\ c & c & c & b & b & b & b & b & c & c & c & c \end{cases}$$

\*   \*   \*   \*   \*                    \*                    \*

7 cache misses in FIFO

Theorem: FIFO is  $k$ -competitive: for any sequence,  $\#misses(FIFO) \leq k \#misses(LFD)$

## Online Paging Algorithms

LIFO: last in first out: evict the page that was entered last to the cache.

Example:  $M_2 = \{a, b, c, d, e\}$   $n=5$ ,  $k=3$

$\sigma = a, b, c, d, a, b, e, d, e, b, c, c, a, d$

$$M_1 = \begin{cases} a & a & a & a & a & a & a & a & a & a & a & a \\ b & b & b & b & b & b & b & b & b & b & b & b \\ c & d & d & d & e & d & e & e & c & c & c & d \end{cases}$$

\*                    \*   \*   \*                    \*                    \*

6 cache misses in LIFO

Theorem: For all  $n > k$ , LIFO is not competitive: For any  $c$ , there exists a sequence of requests such that  $\#misses(FIFO) \geq c \#misses(LFD)$

Proof idea: Consider  $\sigma = 1, 2, \dots, k, k+1, k, k+1, k, k+1, \dots$

## Online Paging Algorithms

LRU: least recently used: evict the page with the earliest last reference.

Example:  $M_2 = \{a, b, c, d, e\}$   $n=5$ ,  $k=3$

$\sigma = a, b, c, d, a, b, d, e, d, e, b, c$   
 $\begin{array}{cccccccccc} a & d & d & d & d & d & d & d & d & c \\ b & b & a & a & a & e & e & e & e & e \\ c & c & c & b & b & b & b & b & b & b \\ * & * & * & & * & & & & & * \end{array}$

Theorem: LRU is  $k$ -competitive

Proof: Not here

29

## Paging- a bound for any deterministic online algorithm

Theorem: For any  $k$  and any deterministic on-line algorithm  $A$ , the competitive ratio of  $A \geq k$ .

Proof: Assume  $n = k+1$  (there are  $k+1$  distinct pages).

What will the adversary do?

Always request the page that is not currently in  $M_1$

This causes a page fault in every access. The total cost of  $A$  is  $|\sigma|$ .

30

## Paging- a bound for any deterministic online algorithm

What is the price of LFD in this sequence?

- At most a single page fault in any  $k$  accesses (LFD evicts the page that will be needed in the  $k+1^{\text{th}}$  request or later)
- The total cost of LFD is at most  $|\sigma|/k$ .

Therefore: Worst-case analysis is not so important in analyzing paging algorithm

• Can randomization help? Yes!! There is a randomized  $2H_k$ -competitive algorithm. ( $H_k = 1 + 1/2 + 1/3 + \dots + 1/k$ )

31

## Online Bin Packing

The input: A sequence of items (numbers),  $a_1, a_2, \dots, a_n$ , such that for all  $i$ ,  $0 < a_i < 1$

The goal: 'pack' the items in bins of size 1. Use as few bins as possible.

Example: The input:  $1/2, 1/3, 2/5, 1/6, 1/5, 2/5$ .

Optimal packing in two bins:

$(1/2, 1/3, 1/6), (2/5, 2/5, 1/5)$ .

Legal packing in three bins:

$(1/2, 1/3), (2/5, 1/6, 1/5), (2/5)$

Online BP:  $a_i$  must be packed before we know  $a_{i+1}, \dots, a_n$

32

## Online Bin Packing

Next-fit Algorithm:

1. Open an *active* bin.
2. For all  $i=1,2,\dots,n$  :
  - If possible, place  $a_i$  in the current active bin;
  - Otherwise, open a new active bin and place  $a_i$  in it.

Example: The input: {0.3, 0.9, 0.2}.

Next-fit packing (three bins): (0.3), (0.9), (0.2).

Theorem: Next-fit is 2-competitive.

Proof: Identical to 2-approximation (see lecture 6)

33

## Online Bin Packing

**First fit algorithm:** place the next item in the first open bin that can accommodate it. Open a new bin only if no open bin has enough room.

Theorem:  $h_{ff} \leq 1.7opt + 2$  (proof not here)

~~**First fit Decreasing:** sort the items from largest to smallest. Run FF according to the resulting order.~~

~~Theorem:  $h_{ffd} \leq 1.22opt + 3$  (proof not here)~~

~~Was fine for approximation, not online!!~~

34

## Online Bin Packing

**Current lower bound:** There is no online bin-packing algorithm with  $r < 1.54$

**Current upper bound:**  $r = 1.5889$

A very important problem, with many applications and variants:

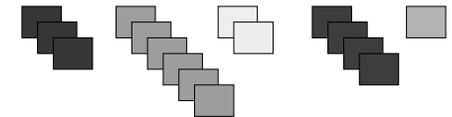
- Multidimensional items and bins
- Online with look-ahead (know the next  $m$  items, or the total length, or other information)
- Limited active bins (you can 'load' into  $k$  trucks)
- Temporary items (with known or unknown duration).
- Many other ad-hoc variants.

35

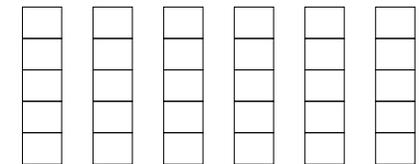
## Online Class-constrained Packing

We need to pack items into bins.

• All items have the same (unit) size. Each item has a color (type).



• All the bins have the same capacity.



• Each bin can accommodate items from a bounded number of colors.

36

## Notations

$n$  - number of items in the instance.

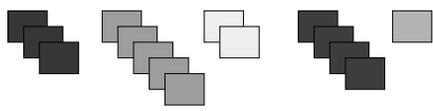
$M$  - number of distinct colors in the instance.

$v$  - bin's capacity.

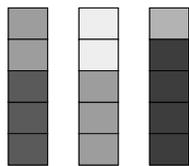
$c$  - number of compartments in a bin.

(The bin can accommodate  $v$  items of  $c$  distinct colors)

### Example of class-constrained packing



$n=15, M=5$



$v = 5, c = 2$

37

## Optimization Goal:

Class-constrained bin-packing (CCBP):

Pack all the items in a minimal number of bins.

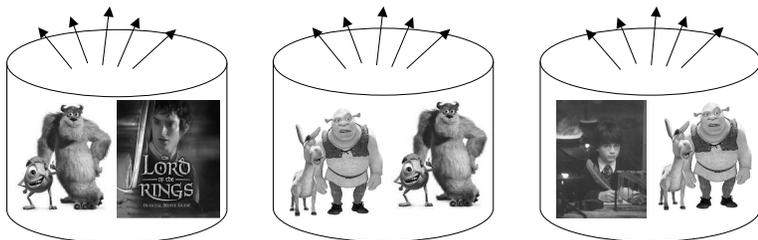
Claim:  $\max(\lceil n/v \rceil, \lceil M/c \rceil)$  is a lower bound for the number of required bins.

Performance measure: For an instance  $I$ , let  $N_{\text{opt}}$  denote the number of bins used by an optimal algorithm to pack all the items of  $I$ . An algorithm is  $r$ -competitive if it packs all the items of  $I$  in at most  $r \cdot N_{\text{opt}}$  bins.

38

## Applications

Multimedia on Demand Systems: The system receives requests for broadcasts of  $M$  movies. The requested movie should be transmitted by a shared disk. Each disk has limited load capacity,  $v$ , and limited storage capacity,  $c$ .



$c=2$

$v=5$

Production Planning: Each device possesses some amount,  $v$ , of a shared resource and can be set to  $c$  distinct configurations. There are  $M$  distinct products.

39

## Online Class-constrained Packing :

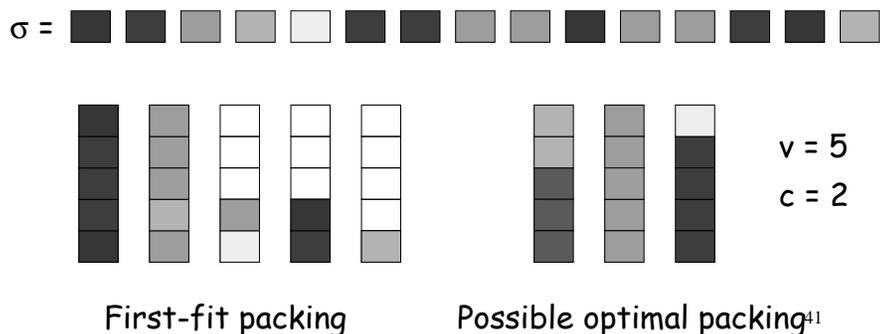
- The items arrive one at a time. In each step we get one unit size item of some color.
- We need to pack this item without any knowledge of the subsequent items.
- Formally, the instance is given as a sequence,  $\sigma = a_1, a_2, \dots$  of length  $n$ , such that  $\forall k, a_k \in \{1, \dots, M\}$ .

• Online CCBP: use a minimal number of bins to pack all the items in  $\sigma$ .

40

# Online Class-constrained Packing, An Example of CCBP.

First-fit Algorithm: Put an arriving item in the leftmost bin that can accommodate it.



# Upper Bound for First-fit

Theorem:  $r_{ff} \leq 2$ .

Proof: Let  $S_1$  be the set of full bins ( $v$  items);  $S_2$  is the set of occupied bins ( $c$  distinct colors).

Claim 1: Each bin (except maybe for the last one) is either full or occupied.

Claim 2: Each bin in  $S_2$  (except maybe for the last one) contains the last item of each of its  $c$  colors.  
(any additional appearance of a color can fit into this bin).

From Claim 1,  $N_{ff} = |S_1| + |S_2|$ .

From Claim 2,  $|S_2| \leq \lceil M/c \rceil$ . Also,  $|S_1| \leq \lceil n/v \rceil$ .

Since  $N_{opt} \geq \max(\lceil n/v \rceil, \lceil M/c \rceil)$  we get a 2-approximation.

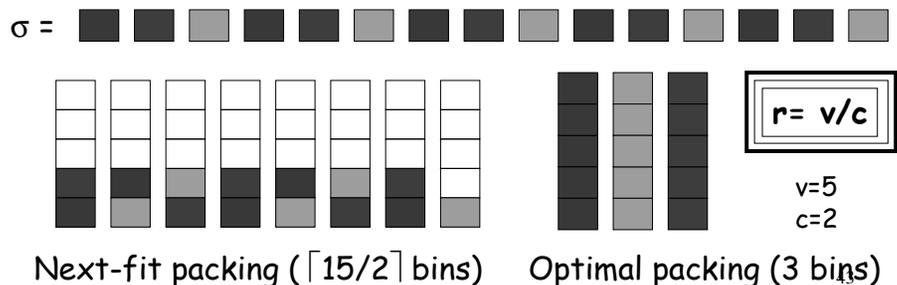
<sup>42</sup>

# Next-fit Algorithm

Next-fit algorithm: Put an arriving item in the currently active bin. Open a new bin if the active bin cannot accommodate this item.

For traditional online bin-packing, it is known that Next-fit uses at most  $2 \cdot N_{OPT}$  bins.

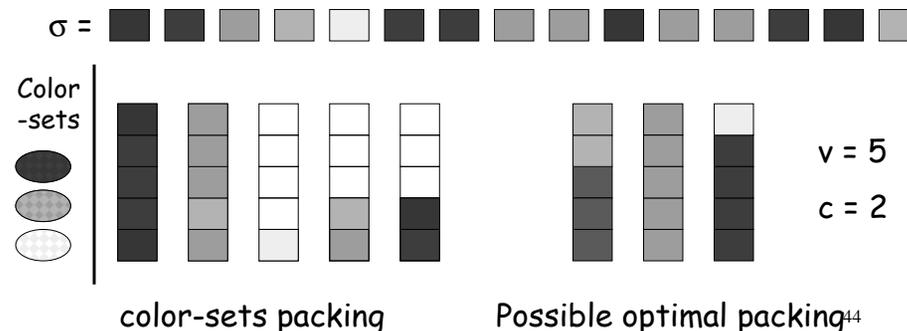
This is not the case for class-constrained packing:



# The Color-sets Algorithm

The algorithm: Partition (online) the  $M$  colors in  $\sigma$  into  $\lceil M/c \rceil$  color-sets. Pack the items of each color-set greedily.

Theorem:  $r_{CS} < 2$



## The Color-sets Algorithm

Theorem:  $r_{CS} < 2$

Proof: Assume that when  $A_{CS}$  terminates there are  $d$  active bins, containing  $x_1, \dots, x_d$  items. Since we open a new active bin for some color-set only when the current active bin of that color set is full, we have

$$N_{CS}(\sigma) = (n - (x_1 + x_2 + \dots + x_d)) / v + d \leq (n/v) + d$$

Since  $d \leq \lceil M/c \rceil$ , we have

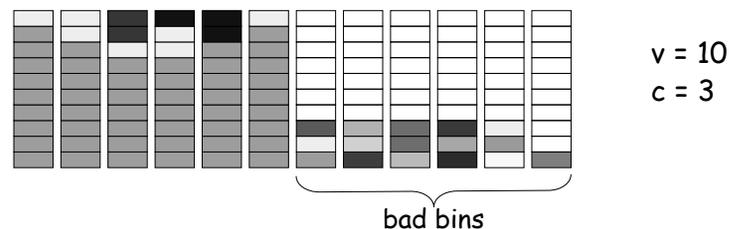
$$N_{CS}(\sigma) \leq \lceil n/v \rceil + \lceil M/c \rceil \leq 2N_{opt}(\sigma)$$

45

## Lower Bound for Deterministic Algorithms

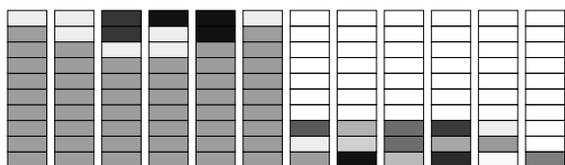
Theorem: for any deterministic algorithm  $A$ ,  $r_A \geq 2$ .

Proof: The adversary constructs the sequence,  $\sigma$ , online such that some bins include  $v$  items from few colors and some bins include  $c$  items from  $c$  distinct colors.

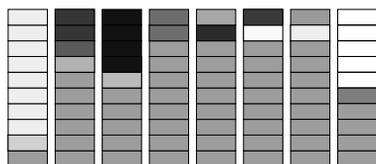


The idea: repeat in  $\sigma$  items of the same color. Switch to the next color whenever an item is packed in one of the 'bad bins'. This color will not be repeated anymore in  $\sigma$ .

46



The packing of the algorithm



An optimal packing:  
The rear colors spread among the bins

A closer analysis of the adversary's strategy yields a lower bound that depends on the ratios  $n/v$  and  $M/c$ .

47