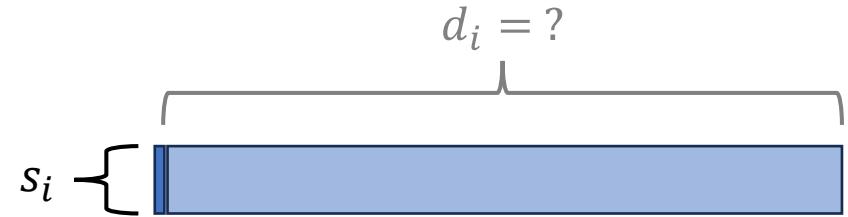


# The Power of Migrations in Dynamic Bin Packing

Konstantina Mellou   Marco Molinaro   **Rudy Zhou**  
Microsoft Research   Carnegie Mellon ⇒ Microsoft

# Dynamic Bin Packing



- Items arrive online at their **arrival time** with **sizes**
- Items depart after their **(unknown) duration**
- Must pack into **unit-size bins**

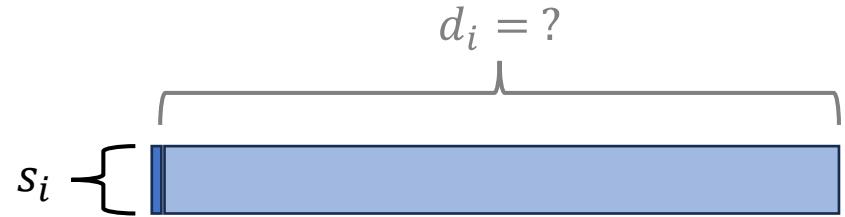
[

[

:

[

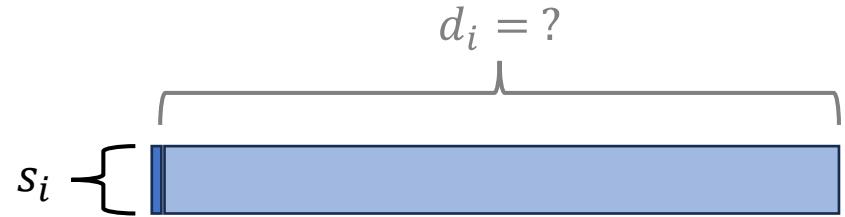
# Dynamic Bin Packing



- Items arrive online at their **arrival time** with **sizes**
- Items depart after their **(unknown) duration**
- Must pack into **unit-size bins**



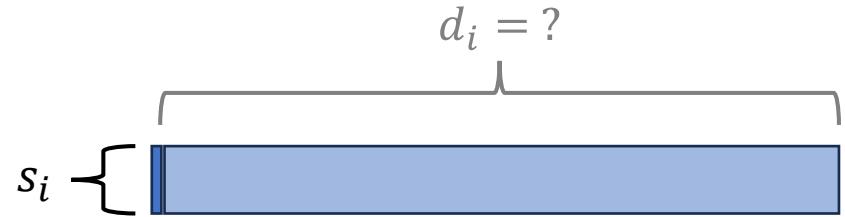
# Dynamic Bin Packing



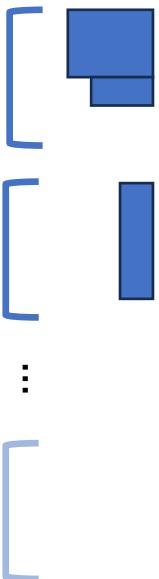
- Items arrive online at their **arrival time** with **sizes**
- Items depart after their **(unknown) duration**
- Must pack into **unit-size bins**



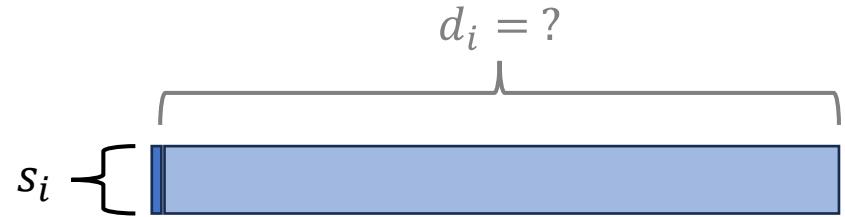
# Dynamic Bin Packing



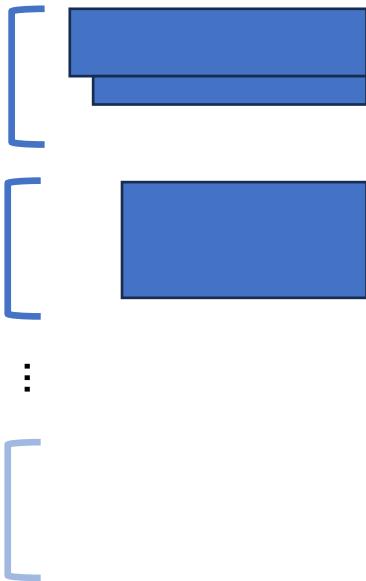
- Items arrive online at their **arrival time** with **sizes**
- Items depart after their **(unknown)** duration
- Must pack into **unit-size bins**



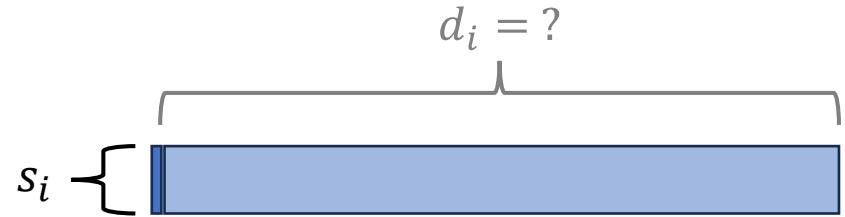
# Dynamic Bin Packing



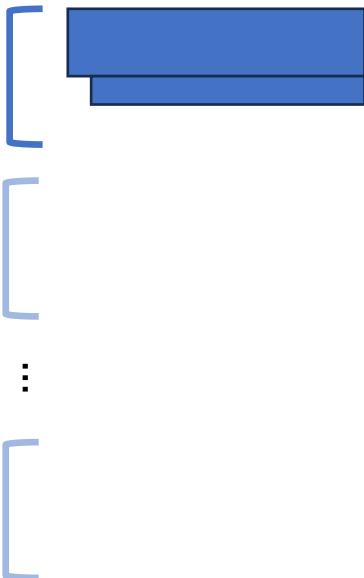
- Items arrive online at their **arrival time** with **sizes**
- Items depart after their **(unknown) duration**
- Must pack into **unit-size bins**



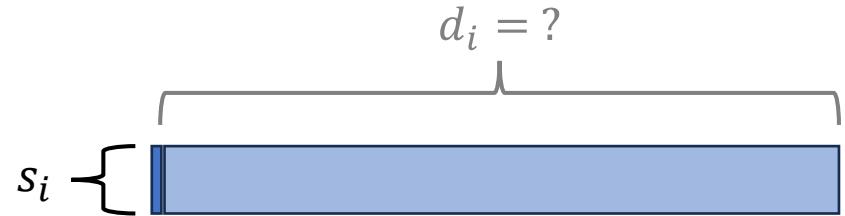
# Dynamic Bin Packing



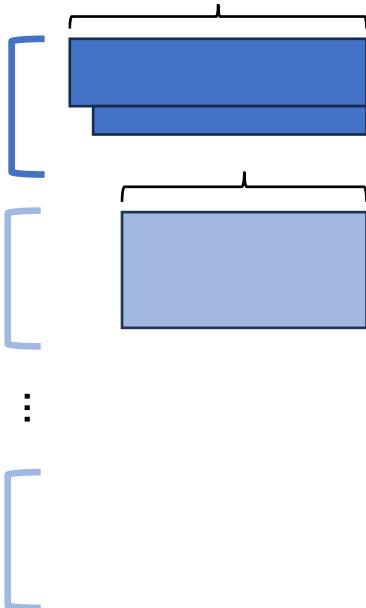
- Items arrive online at their **arrival time** with **sizes**
- Items depart after their **(unknown) duration**
- Must pack into **unit-size bins**



# Dynamic Bin Packing

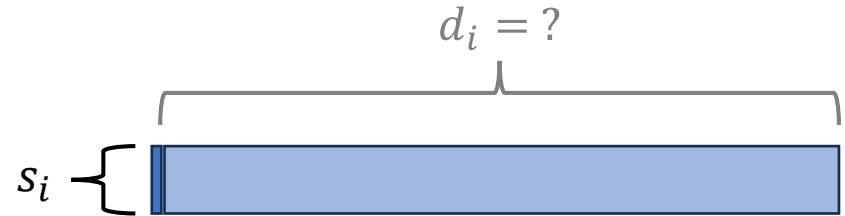


- Items arrive online at their **arrival time** with **sizes**
- Items depart after their **(unknown)** duration
- Must pack into **unit-size bins**

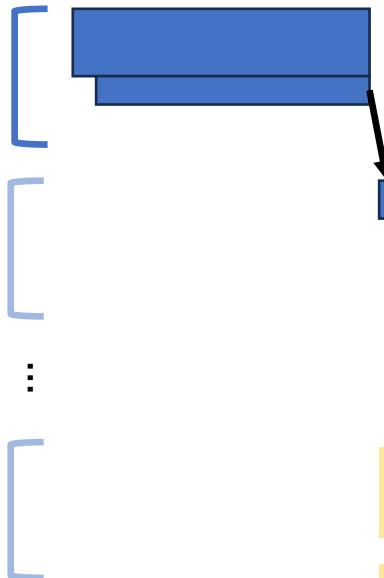


Minimize total active time over all bins

# Dynamic Bin Packing



- Items arrive online at their **arrival time** with **sizes**
- Items depart after their **(unknown) duration**
- Must pack into **unit-size bins**



Minimize total active time over all bins

Can also **migrate** items

# Related Work

$$\mu = \frac{\text{max duration}}{\text{min duration}}$$
$$n = \# \text{ items}$$

- **No migrations:**  $\Theta(\mu)$ -approximation (first fit)

Yusen Li, Xueyan Tang, Wentong Cai:

On dynamic bin packing for resource allocation in the cloud. SPAA 2014

- **>  $n$  migrations:**  $\approx 1.387 + \epsilon$ -approximation with  $O(\frac{n}{\epsilon^2})$  migrations

Björn Feldkord, Matthias Feldotto, Anupam Gupta, Guru Guruganesh, Amit Kumar, Sören Riechers, David Wajc:

Fully-Dynamic Bin Packing with Little Repacking. ICALP 2018

- Can get better guarantees with no migrations if know item durations exactly or approximately (predictions)

Yossi Azar, Danny Vainstein:

Tight Bounds for Clairvoyant Dynamic Bin Packing. SPAA 2017

Mozhengfu Liu, Xueyan Tang:

Dynamic Bin Packing with Predictions. SIGMETRICS 2023

Niv Buchbinder, Yaron Fairstein, Konstantina Mellou, Ishai Menache, Joseph (Seffi) Naor:

Online Virtual Machine Allocation with Lifetime and Load Predictions. SIGMETRICS 2021

# Related Work

$$\mu = \frac{\text{max duration}}{\text{min duration}}$$
$$n = \# \text{ items}$$

- **No migrations:**  $\Theta(\mu)$ -approximation (first fit)

Yusen Li, Xueyan Tang, Wentong Cai:

On dynamic bin packing for resource allocation in the cloud. SPAA 2014

- **>  $n$  migrations:**  $\approx 1.387 + \epsilon$ -approximation with  $O(\frac{n}{\epsilon^2})$  migrations

Björn Feldkord, Matthias Feldotto, Anupam Gupta, Guru Guruganesh, Amit Kumar, Sören Riechers, David Wajc:

Fully-Dynamic Bin Packing with Little Repacking. ICALP 2018

- Can get better guarantees with no migrations if know item durations exactly or approximately (predictions)

Yossi Azar, Danny Vainstein:

Tight Bounds for Clairvoyant Dynamic Bin Packing. SPAA 2017

Mozhengfu Liu, Xueyan Tang:

Dynamic Bin Packing with Predictions. SIGMETRICS 2023

Niv Buchbinder, Yair Even, Kfir Levy, Michael Mitzenmacher, and David Zuckerman (2023). N

Online Virtual Machine Migration with Predictions. FOCS 2023

What can we do with  $< n$  migrations?  $\epsilon n$ ?  $\sqrt{n}$ ?

# Bad Example

- $\mu^2$  items; all arriving at time 0 with size  $\frac{1}{\mu}$ 
  - $\mu$  of them are **long** with duration  $\mu$
  - Rest are **short** with duration 1

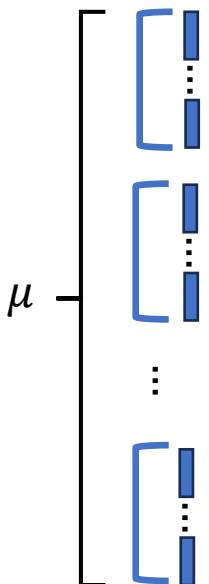


⋮



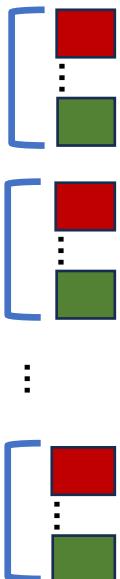
# Bad Example

- $\mu^2$  items; all arriving at time 0 with size  $\frac{1}{\mu}$ 
  - $\mu$  of them are **long** with duration  $\mu$
  - Rest are **short** with duration 1



# Bad Example

- $\mu^2$  items; all arriving at time 0 with size  $\frac{1}{\mu}$ 
  - $\mu$  of them are **long** with duration  $\mu$
  - Rest are **short** with duration 1



# Bad Example

- $\mu^2$  items; all arriving at time 0 with size  $\frac{1}{\mu}$ 
  - $\mu$  of them are **long** with duration  $\mu$
  - Rest are **short** with duration 1



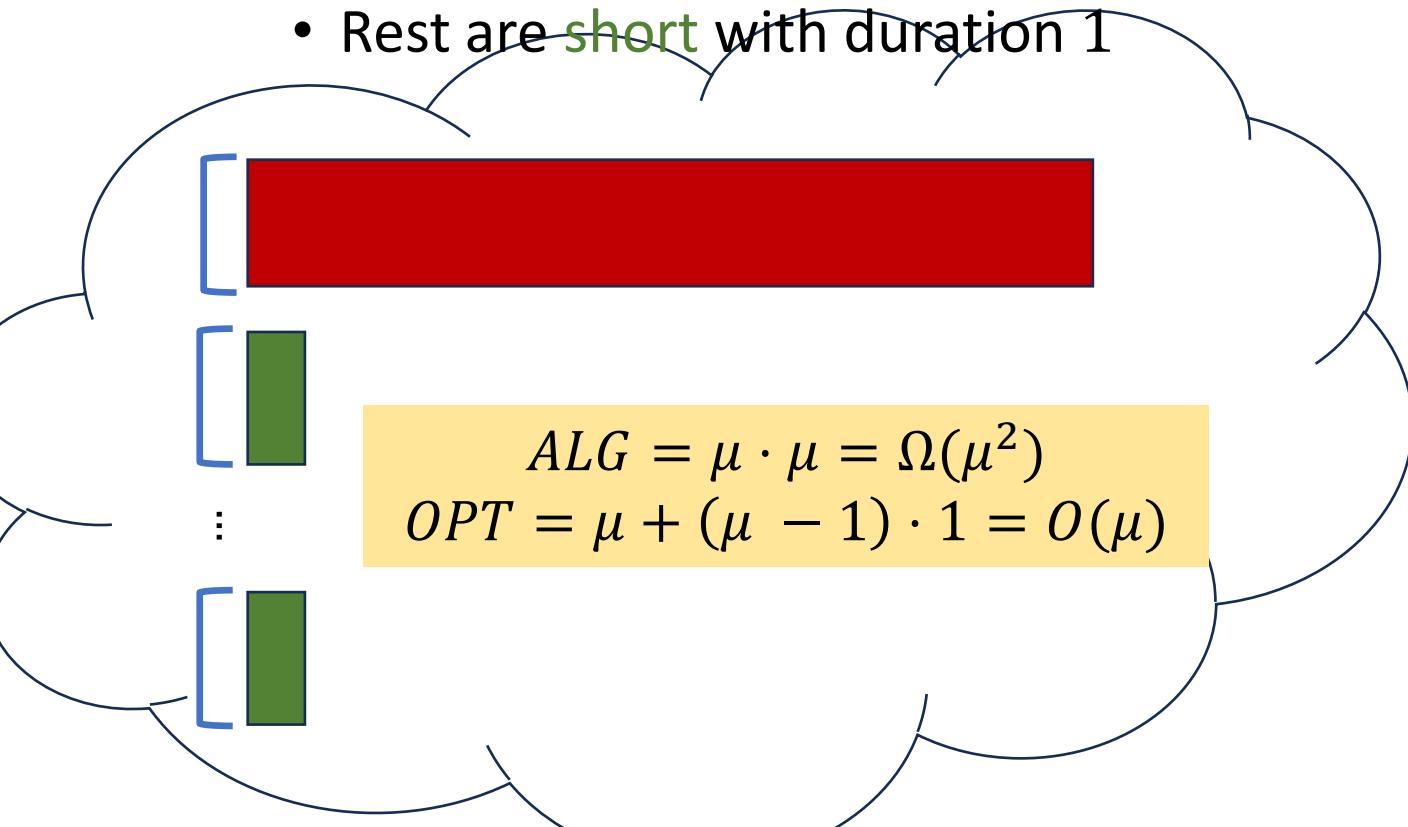
# Bad Example

- $\mu^2$  items; all arriving at time 0 with size  $\frac{1}{\mu}$ 
  - $\mu$  of them are **long** with duration  $\mu$
  - Rest are **short** with duration 1



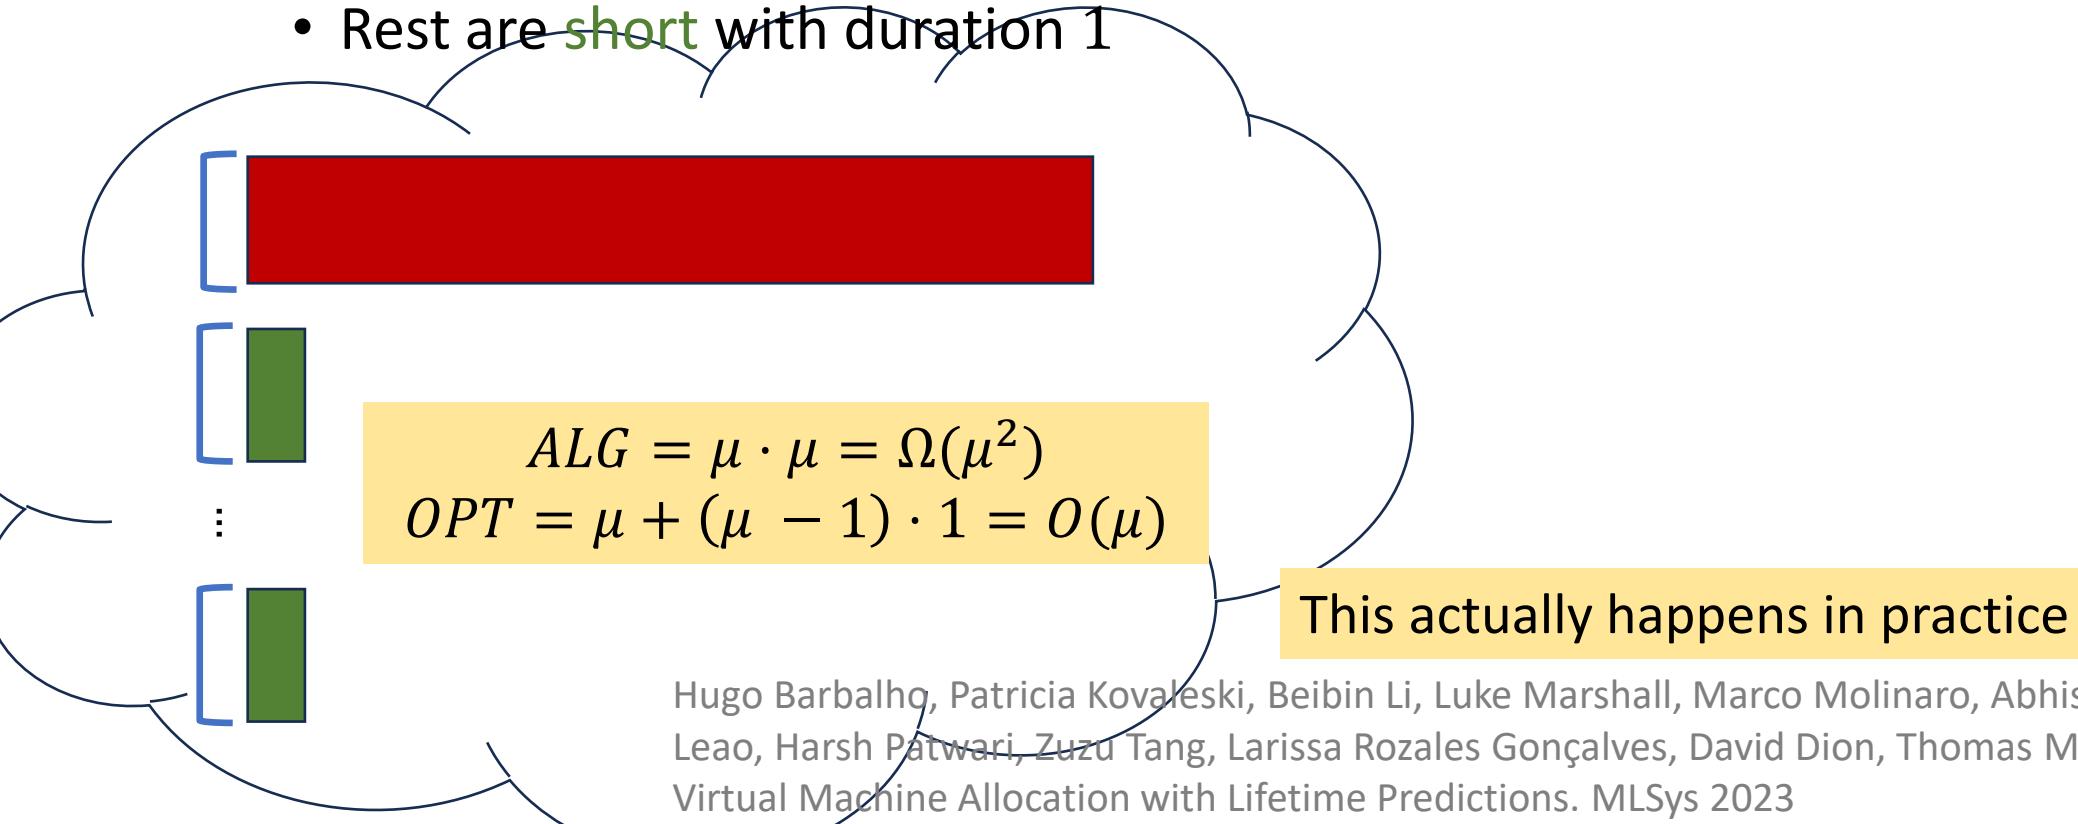
# Bad Example

- $\mu^2$  items; all arriving at time 0 with size  $\frac{1}{\mu}$ 
  - $\mu$  of them are **long** with duration  $\mu$
  - Rest are **short** with duration 1



# Bad Example

- $\mu^2$  items; all arriving at time 0 with size  $\frac{1}{\mu}$ 
  - $\mu$  of them are **long** with duration  $\mu$
  - Rest are **short** with duration 1



# Our Results (Part 1)

- **Sublinear migrations:** Any algorithm that does  $o(n)$  migrations must be  $\Omega(\mu)$ -approximate
- **$< n$  migrations:** For any  $\epsilon \in (0, 1)$ , can get  $\approx \frac{1}{\epsilon}$ -approximation using  $\epsilon n$  migrations, and this is best possible\*

\* up to a  $\log n$  additive term in approximation

# Our Results (Part 1)

- **Sublinear migrations:** Any algorithm that does  $o(n)$  migrations must be  $\Omega(\mu)$ -approximate
- **$< n$  migrations:** For any  $\epsilon \in (0, 1)$ , can get  $\approx \frac{1}{\epsilon}$ -approximation using  $\epsilon n$  migrations, and this is best possible\*
- **$> n$  migrations:**  $\approx 1.387 + \epsilon$ -approximation with  $O(\frac{n}{\epsilon^2})$  migrations

\* up to a  $\log n$  additive term in approximation

# $< n$ migrations

- Want  $\approx \frac{1}{\epsilon}$ -approximation  $\Rightarrow$  suffices to ensure most bins are  $\geq \epsilon$ -full
- Assume all items have same size\*

\* standard discretization argument; lose additive  $\log n$  here

# $< n$ migrations

- Want  $\approx \frac{1}{\epsilon}$ -approximation  $\Rightarrow$  suffices to ensure most bins are  $\geq \epsilon$ -full
- Assume all items have same size\*
- Classify bins as **bad** and **good**

\* standard discretization argument; lose additive  $\log n$  here

# $< n$ migrations

- Want  $\approx \frac{1}{\epsilon}$ -approximation  $\Rightarrow$  suffices to ensure most bins are  $\geq \epsilon$ -full
- Assume all items have same size\*
- Classify bins as **bad** and **good**

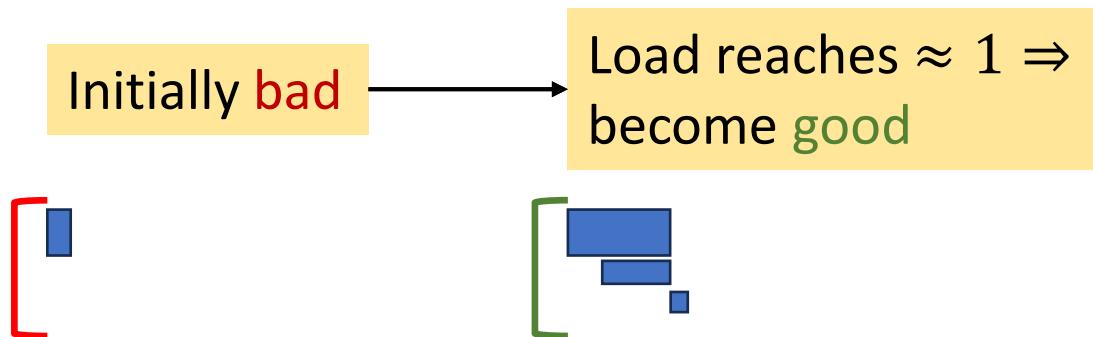
Initially **bad**



\* standard discretization argument; lose additive  $\log n$  here

# $< n$ migrations

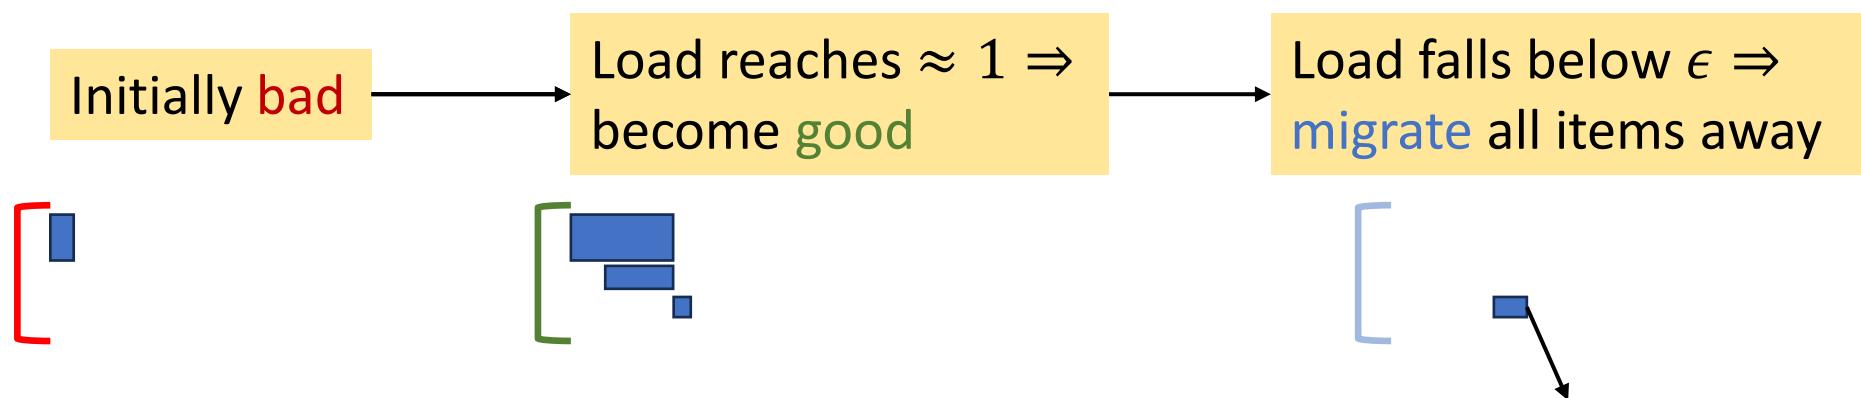
- Want  $\approx \frac{1}{\epsilon}$ -approximation  $\Rightarrow$  suffices to ensure most bins are  $\geq \epsilon$ -full
- Assume all items have same size\*
- Classify bins as **bad** and **good**



\* standard discretization argument; lose additive  $\log n$  here

# $< n$ migrations

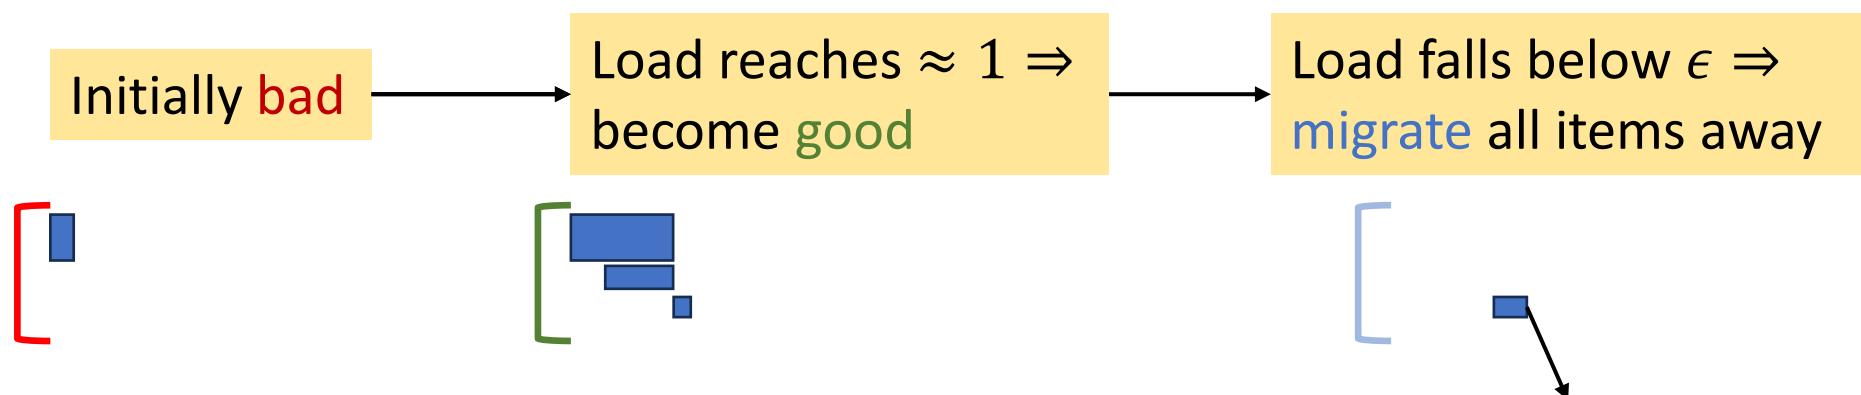
- Want  $\approx \frac{1}{\epsilon}$ -approximation  $\Rightarrow$  suffices to ensure most bins are  $\geq \epsilon$ -full
- Assume all items have same size\*
- Classify bins as **bad** and **good**



\* standard discretization argument; lose additive  $\log n$  here

# $< n$ migrations

- Want  $\approx \frac{1}{\epsilon}$ -approximation  $\Rightarrow$  suffices to ensure most bins are  $\geq \epsilon$ -full
- Assume all items have same size\*
- Classify bins as **bad** and **good**



- Can ensure  $\leq 1$  bad bins at any time
- Migrate  $\epsilon$ -fraction of bin load when  $1 - \epsilon$ -fraction departs

\* standard discretization argument; lose additive  $\log n$  here

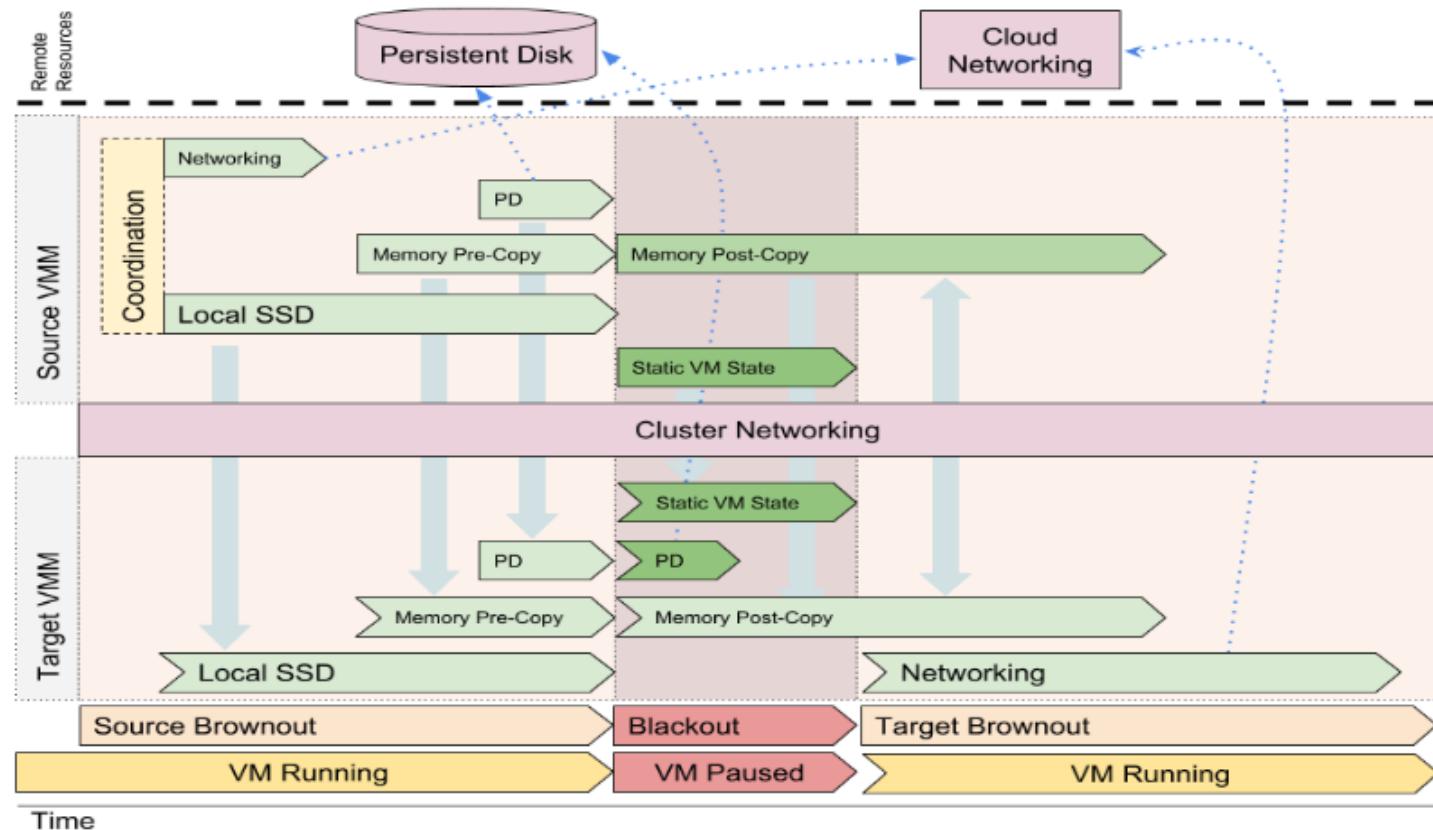
# Our Results (Part 1)

- **Sublinear migrations:** Any algorithm that does  $o(n)$  migrations must be  $\Omega(\mu)$ -approximate
- **$< n$  migrations:** For any  $\epsilon \in (0, 1)$ , can get  $\approx \frac{1}{\epsilon}$ -approximation using  $\epsilon n$  migrations, and this is best possible\*
- **$> n$  migrations:**  $\approx 1.387 + \epsilon$ -approximation with  $O(\frac{n}{\epsilon^2})$  migrations

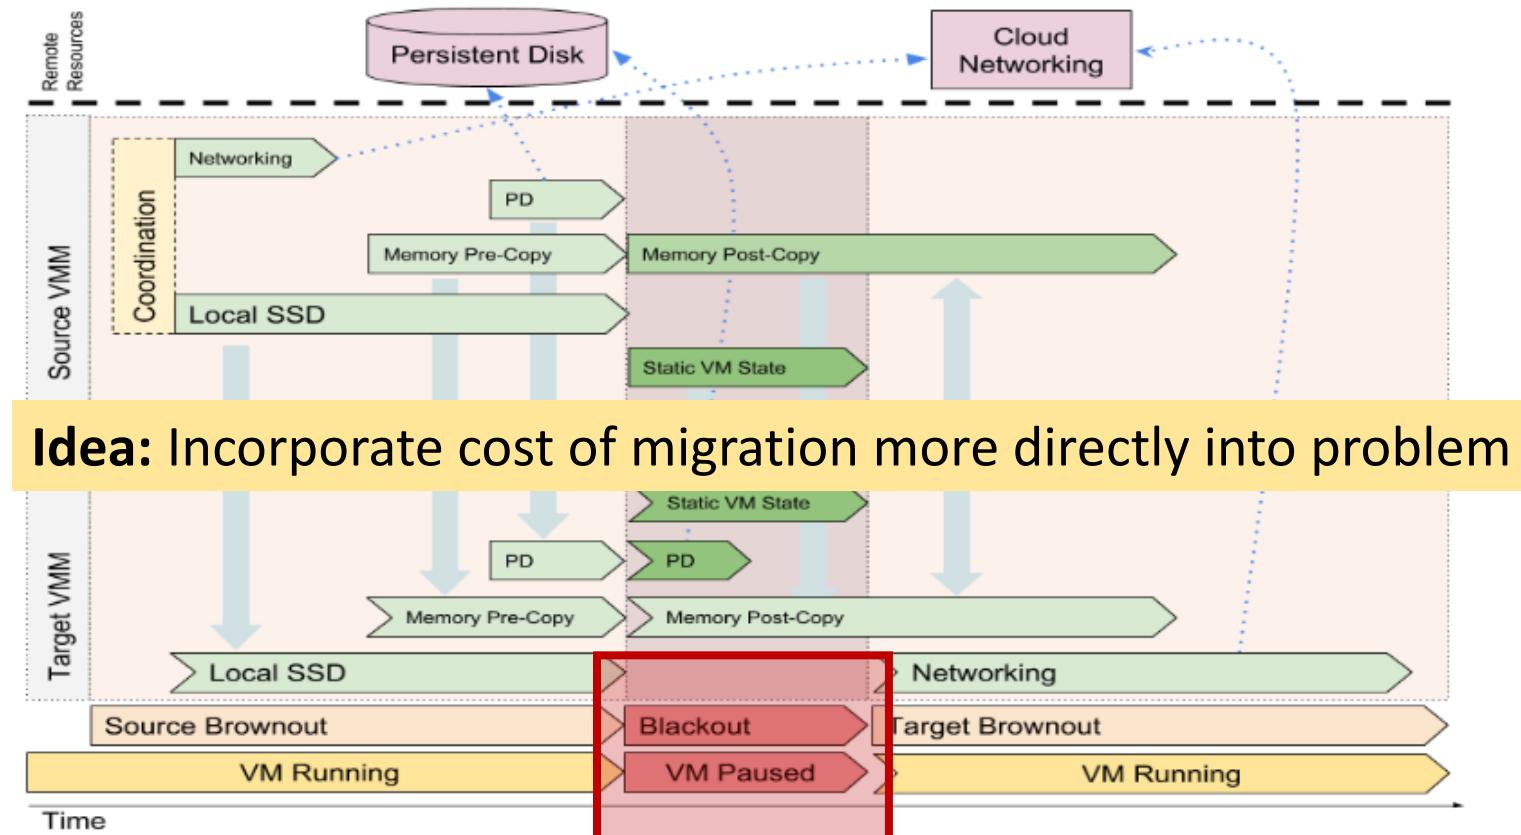
Can we do better?

\* up to a  $\log n$  additive term in approximation

# Real cost of migration

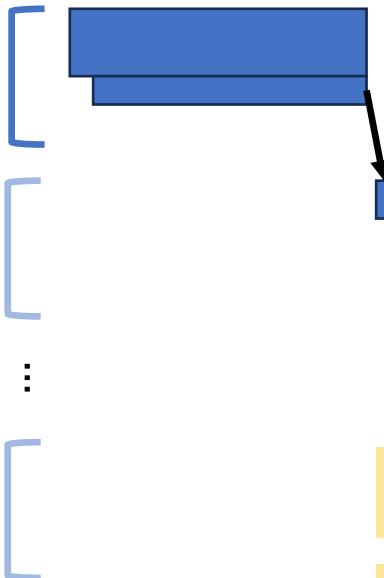


# Real cost of migration



# Dynamic Bin Packing with Delays

- Items arrive online at their **arrival time** with sizes
- Items depart after their **(unknown)** duration
- Must pack into **unit-size bins**

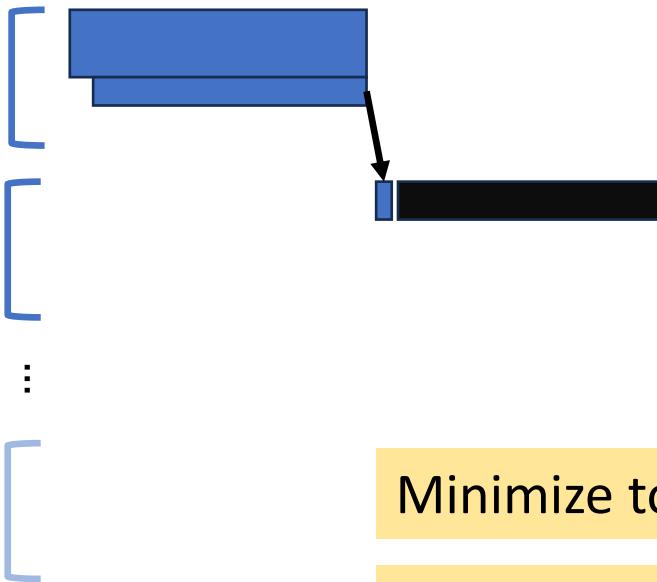


Minimize total active time over all bins

Can also **migrate** items

# Dynamic Bin Packing with Delays

- Items arrive online at their **arrival time** with sizes
- Items depart after their **(unknown)** duration
- Must pack into **unit-size bins**

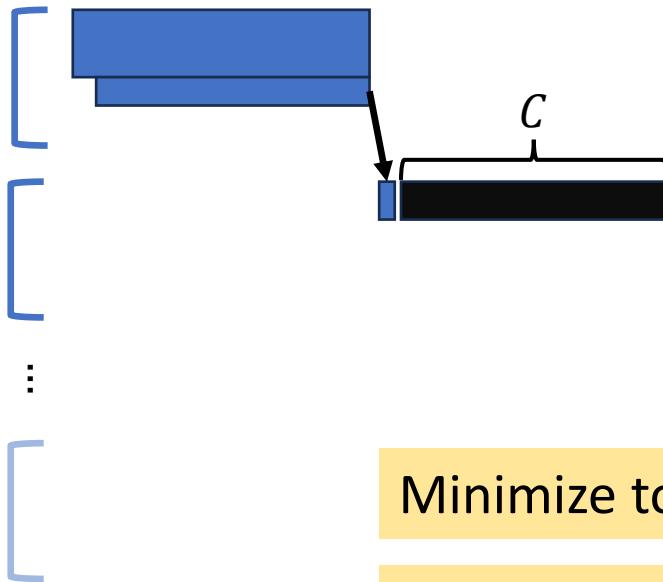


Minimize total active time over all bins

Can also migrate items, but migrating increases item duration by  $+C$

# Dynamic Bin Packing with Delays

- Items arrive online at their **arrival time** with sizes
- Items depart after their **(unknown)** duration
- Must pack into **unit-size bins**

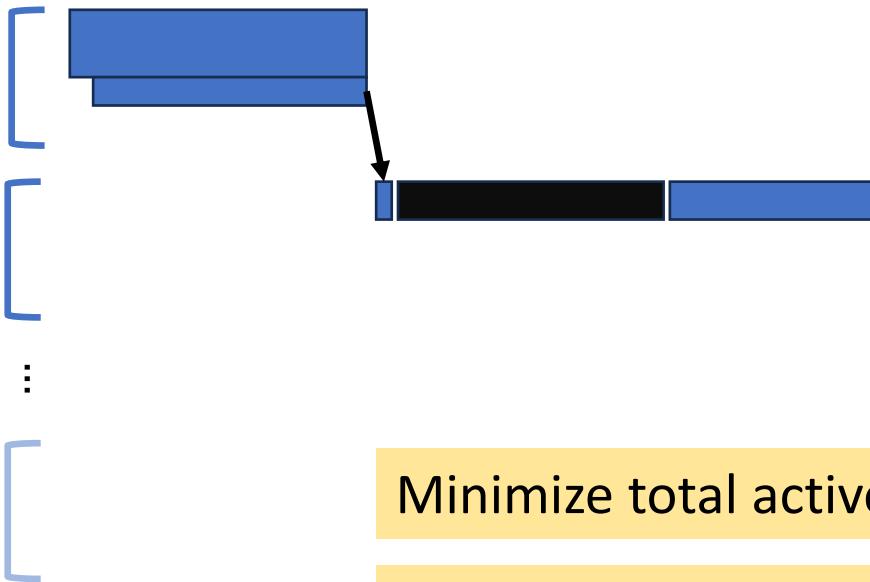


Minimize total active time over all bins

Can also migrate items, but migrating increases item duration by  $+C$

# Dynamic Bin Packing with Delays

- Items arrive online at their **arrival time** with sizes
- Items depart after their **(unknown)** duration
- Must pack into **unit-size bins**



Minimize total active time over all bins

Can also migrate items, but migrating increases item duration by  $+C$

# Our Results (Part 2)

- $O(\min(\sqrt{C}, \mu))$ -approximation for Dynamic bin packing **with delays**, and this is best possible\*

\* some normalization assumptions apply

# Our Results (Part 2)

- $O(\min(\sqrt{C}, \mu))$ -approximation for Dynamic bin packing [with delays](#), and this is best possible\*
- Best of both worlds: never worse than doing no migrations, but can be much better

\* some normalization assumptions apply

# Our Results (Part 2)

- $O(\min(\sqrt{C}, \mu))$ -approximation for Dynamic bin packing [with delays](#), and this is best possible\*
- Best of both worlds: never worse than doing no migrations, but can be much better
- In practice, the minimum duration can be  $\approx 1$  millisecond, and the maximum  $\approx 1$  year  $\Rightarrow \mu \approx 10^{10}$
- However, migrating an item incurs a delay of  $\approx 1$  second  $\Rightarrow C \approx 10^3$

\* some normalization assumptions apply

# Dynamic Bin Packing with Delays

- When to migrate?
- Recall bad example: all items have same size, some short, some **long**



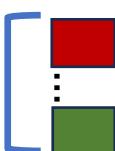
⋮

# Dynamic Bin Packing with Delays

- When to migrate?
- Recall bad example: all items have same size, some short, some **long**



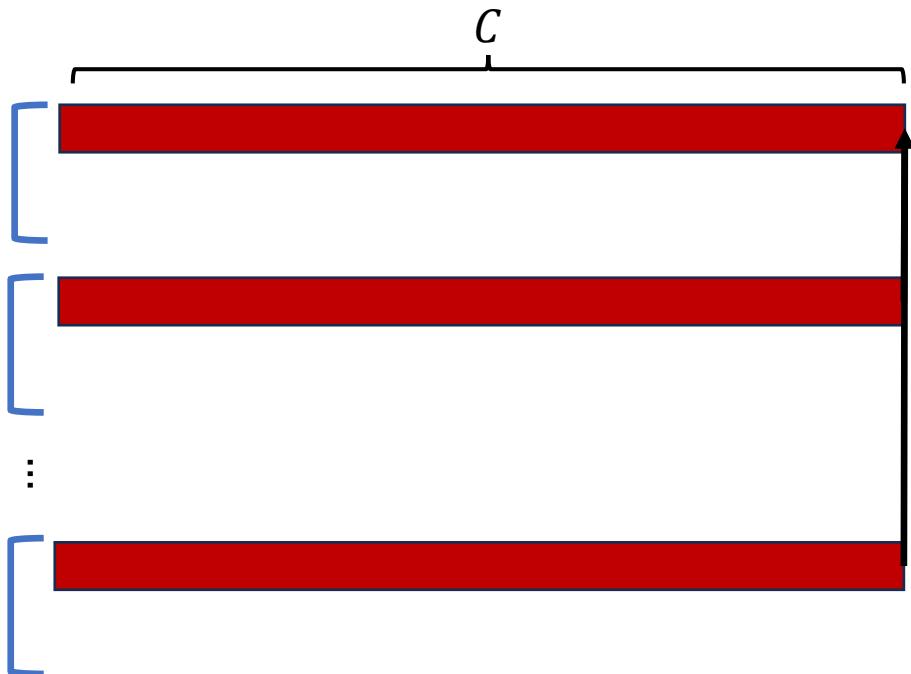
⋮



# Dynamic Bin Packing with Delays

First Try: Migrate at time  $C$

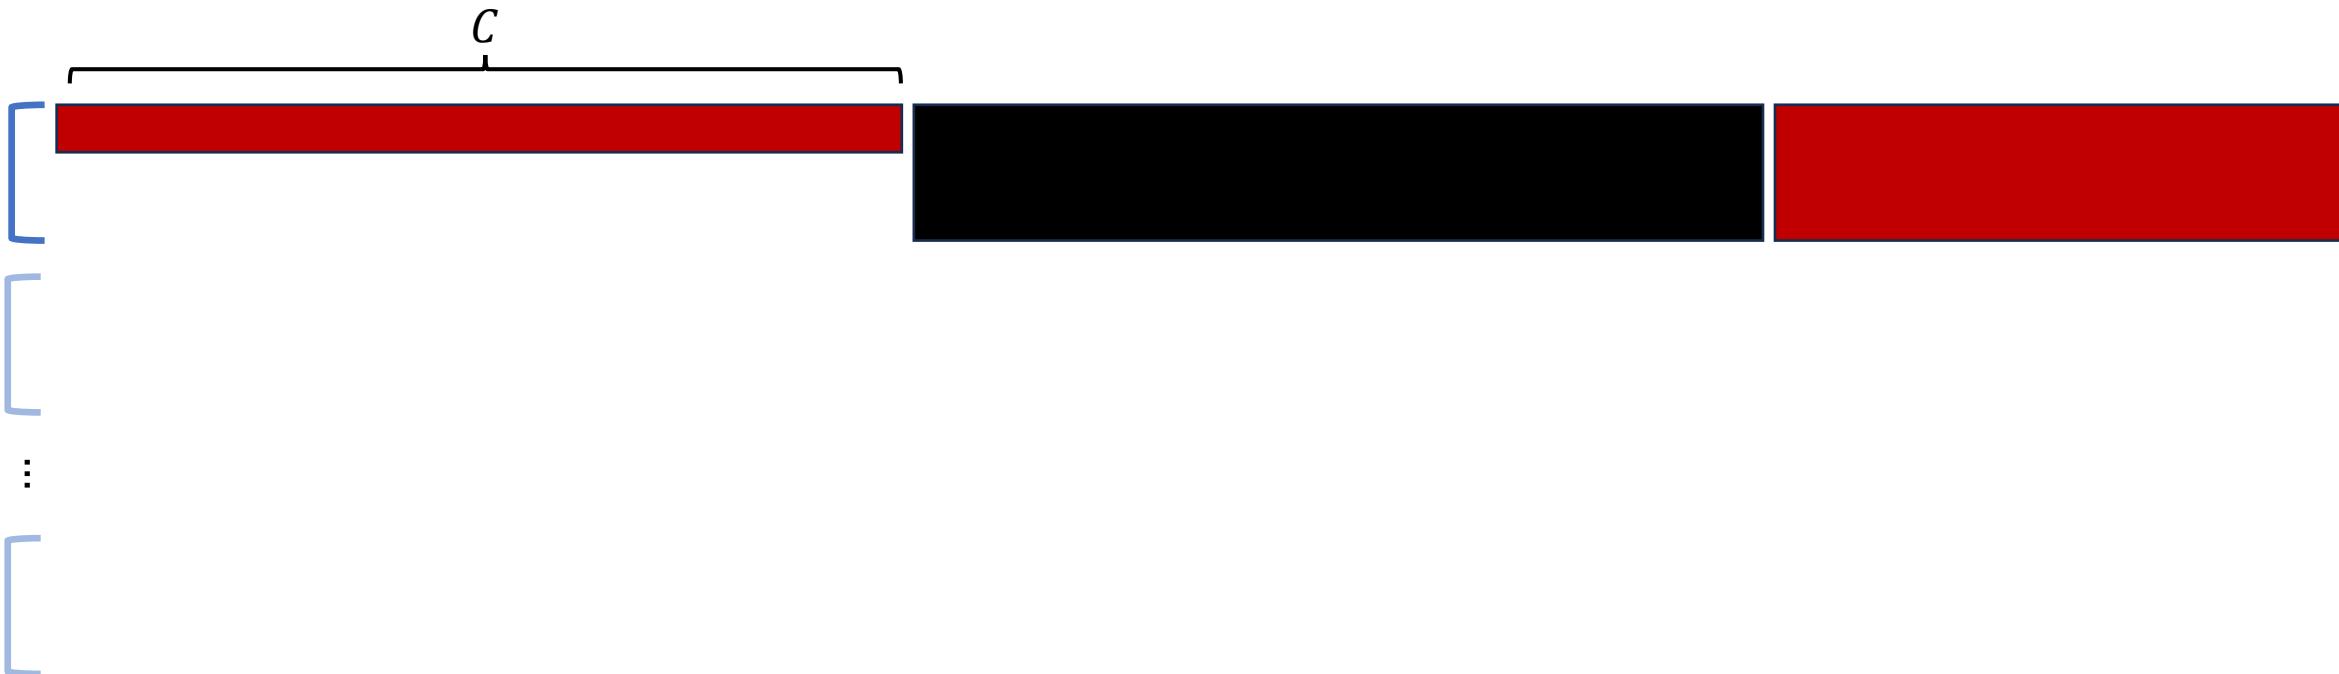
- When to migrate?
- Recall bad example: all items have same size, some short, some **long**



# Dynamic Bin Packing with Delays

First Try: Migrate at time  $C$

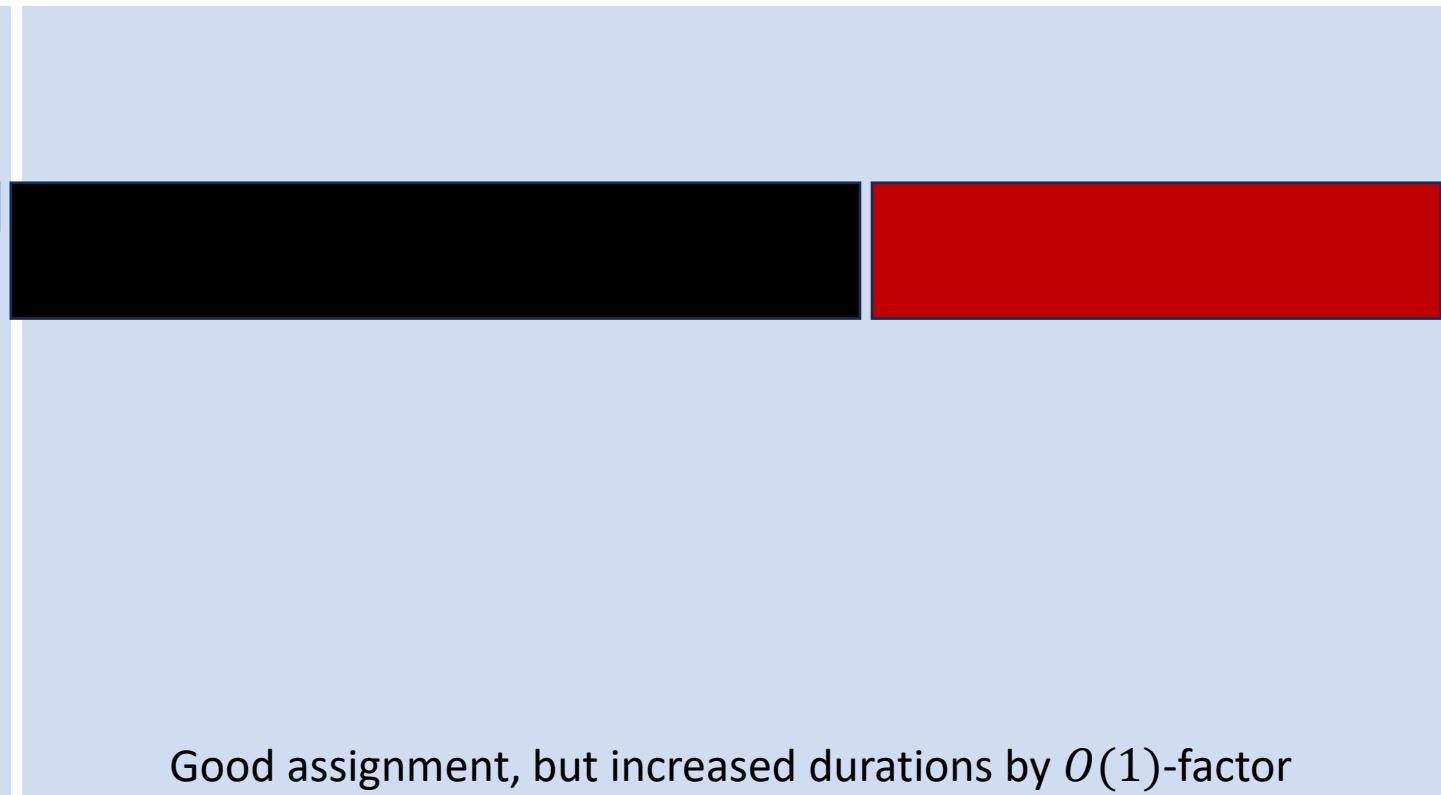
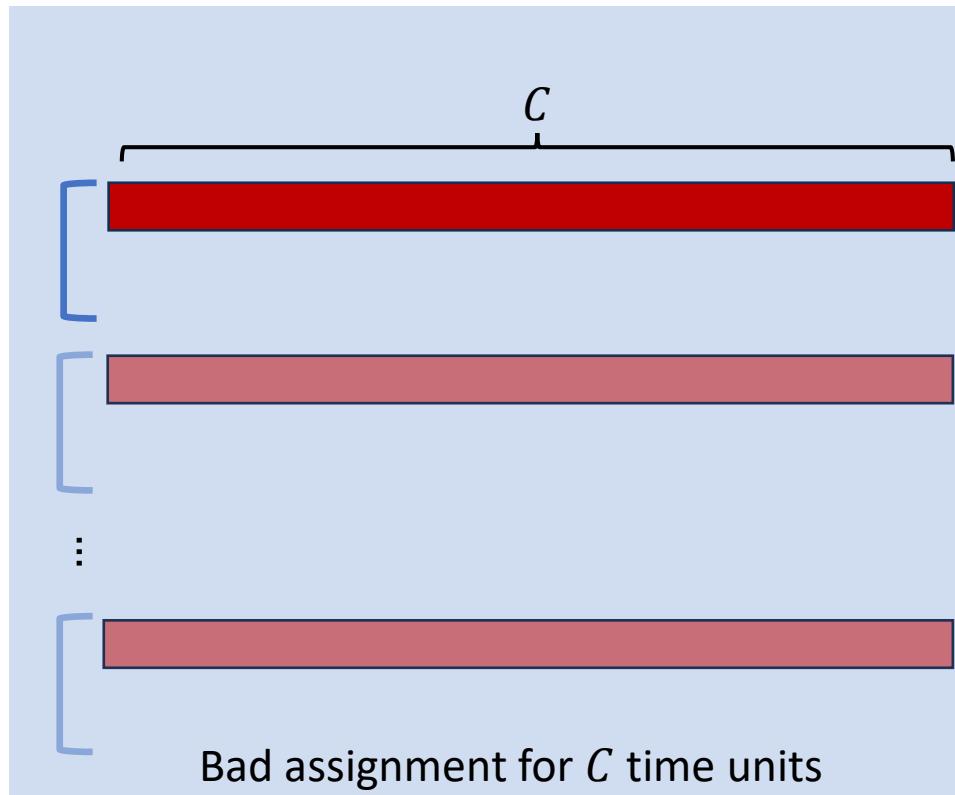
- When to migrate?
- Recall bad example: all items have same size, some short, some **long**



# Dynamic Bin Packing with Delays

First Try: Migrate at time  $C$

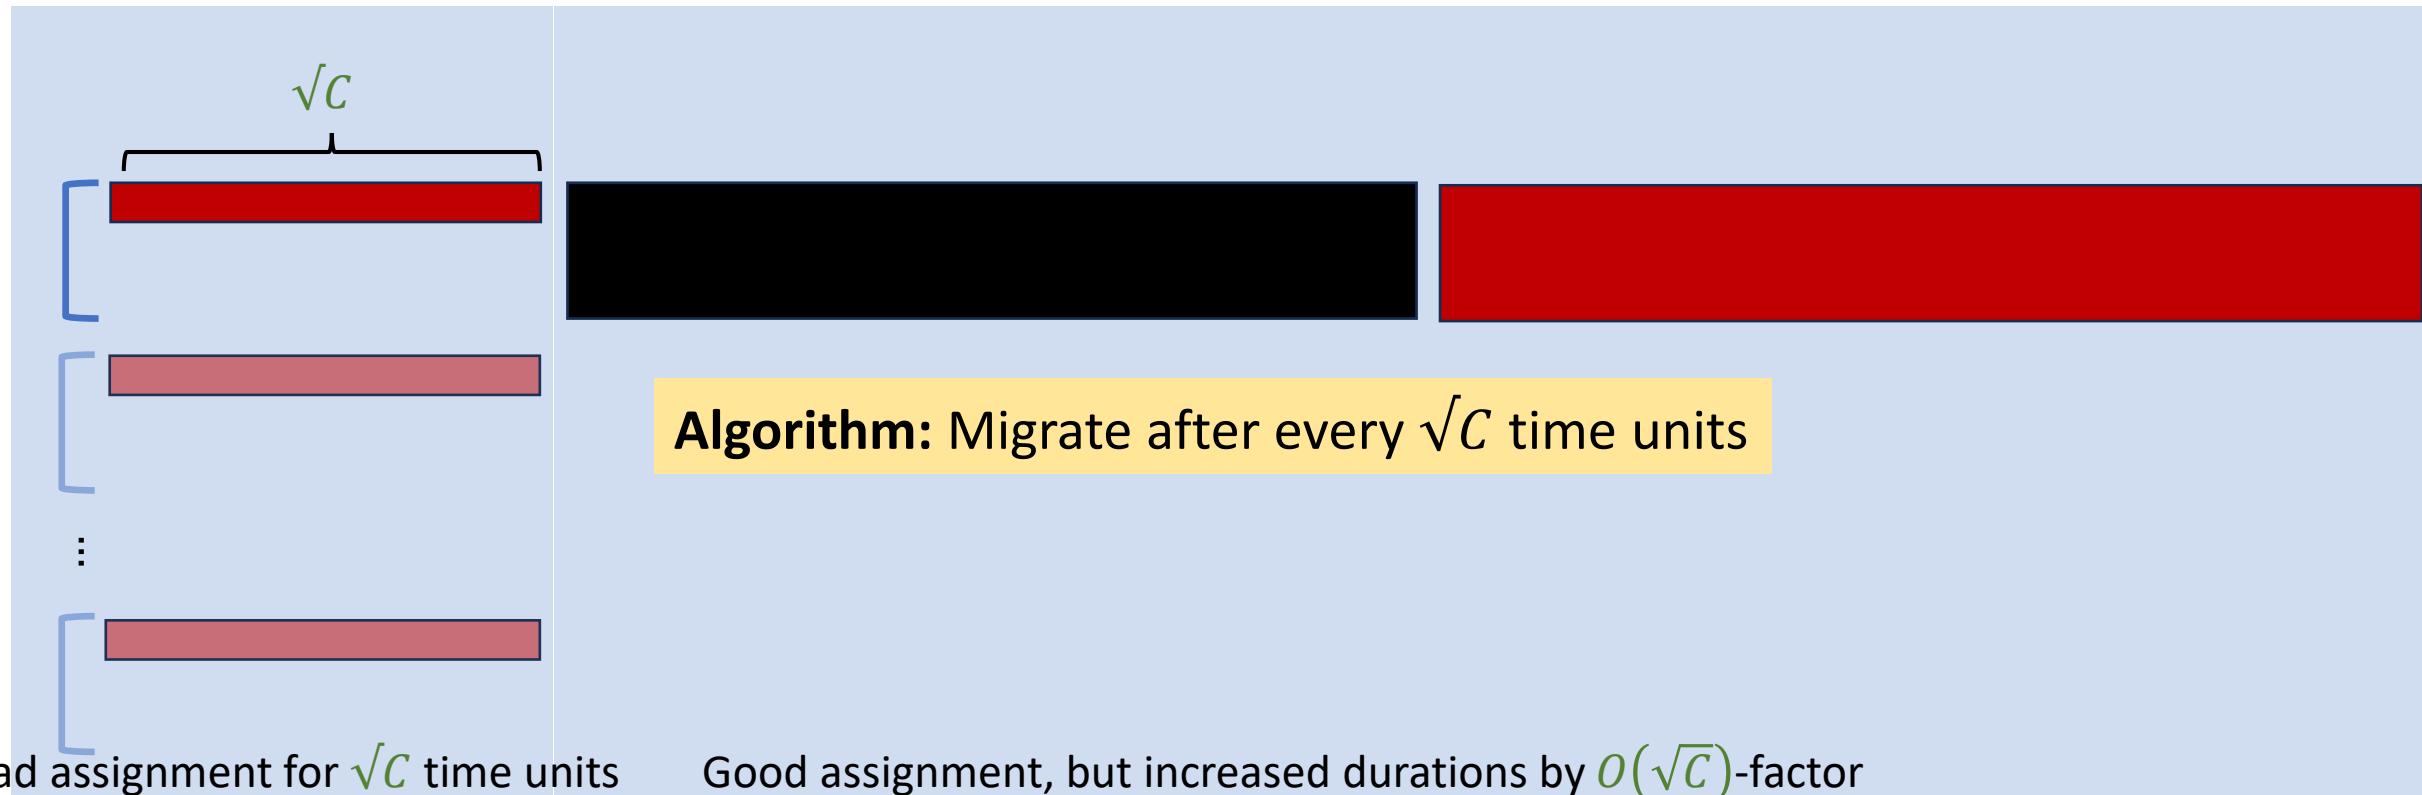
- When to migrate?
- Recall bad example: all items have same size, some short, some **long**



# Dynamic Bin Packing with Delays

First Try: Migrate at time  $C$

- When to migrate?
- Recall bad example: all items have same size, some short, some long



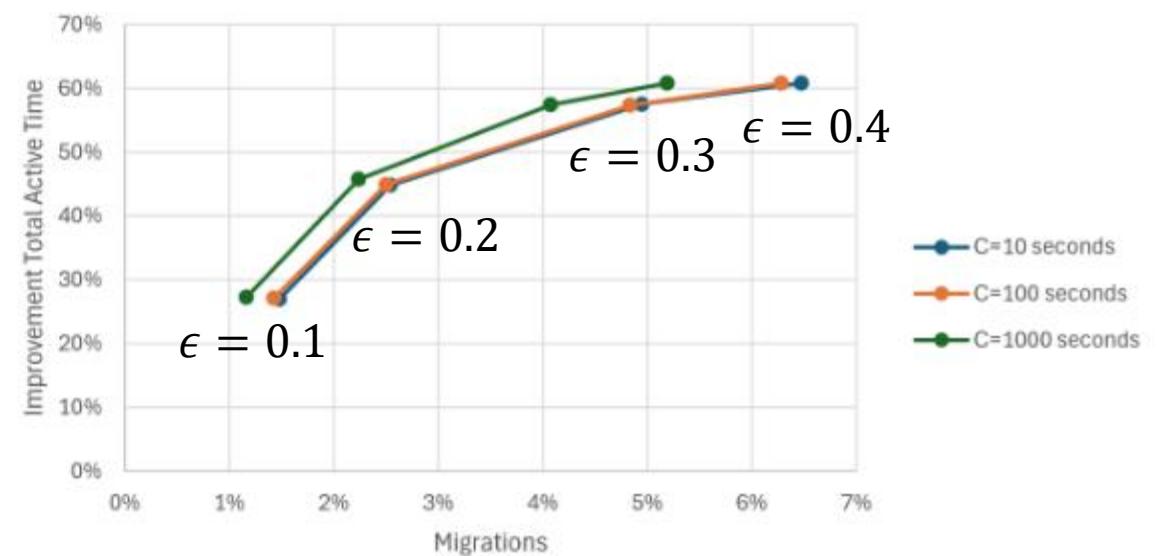
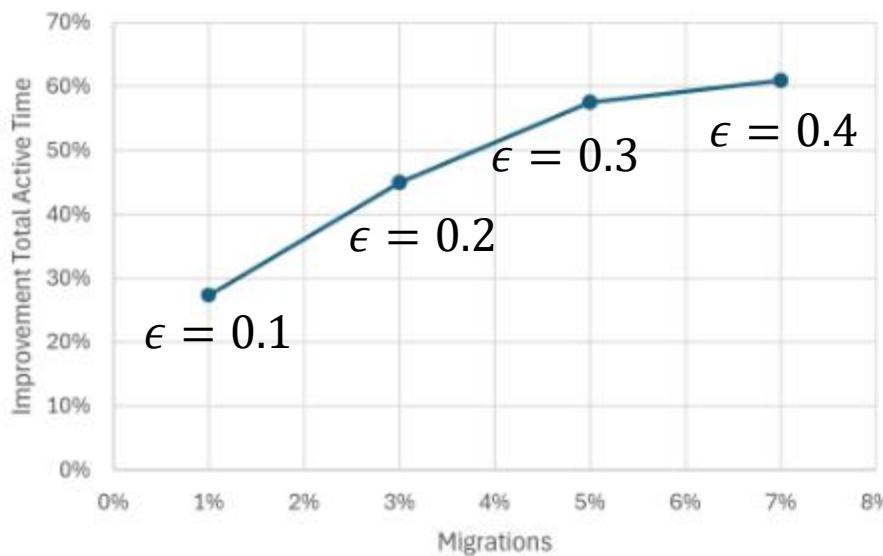
# Our Results (Part 2)

- $O(\min(\sqrt{C}, \mu))$ -approximation for Dynamic bin packing [with delays](#), and this is best possible\*
- Best of both worlds: never worse than doing no migrations, but can be much better
- In practice, the minimum duration can be  $\approx 1$  millisecond, and the maximum  $\approx 1$  year  $\Rightarrow \mu \approx 10^{10}$
- However, migrating an item incurs a delay of  $\approx 1$  second  $\Rightarrow C \approx 10^3$

\* some normalization assumptions apply

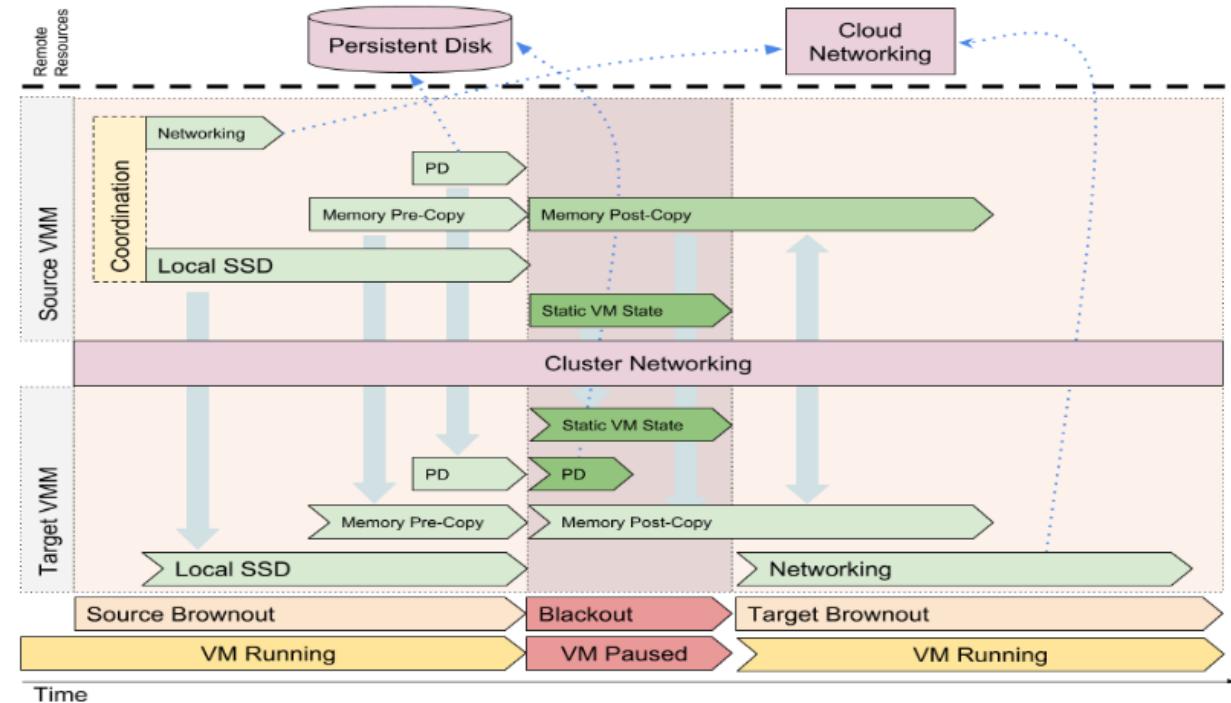
# Experiments

- Data from Microsoft Azure (VM requests)
- Combine both  $< n$  migration and delay algorithms (classify bins as bad and good; only migrate items after every  $C$  time units)



# Conclusion

- Fill gaps in our understanding for  $o(n)$  and  $\epsilon n$  migrations
- Introduce **delays** to dynamic bin packing



# Conclusion

- Fill gaps in our understanding for  $o(n)$  and  $\epsilon n$  migrations
- Introduce **delays** to dynamic bin packing

## Open Questions:

- Remove additive  $\log n$  in  $\epsilon n$  migration case
- Beyond worst case model
- Consider networking limits

