institute imdea networks

Power-efficient Assignment of Virtual Machines to Physical Machines

Antonio Fernández Anta

Joint work with Jordi Arjona Aroca, Miguel A. Mosteiro, Christoper Thraves, Lin Wang

[Developing the Science of Networks]

Virtual Machines and Physical Machines

- We consider a data center with *m* physical machines (PM)
- Virtual machines (VM) are executed in the PMs
- All PM are similar and have a capacity *C*
- The power consumed by a PM is a function of its load
- Each VM is permanently assigned to a PM
- Each VM d_i has a load $l(d_i)$
- No PM can be overloaded

Virtual Machine Assignment Problem (VMA)

- Where to assign an incoming VM in a power-efficient way?

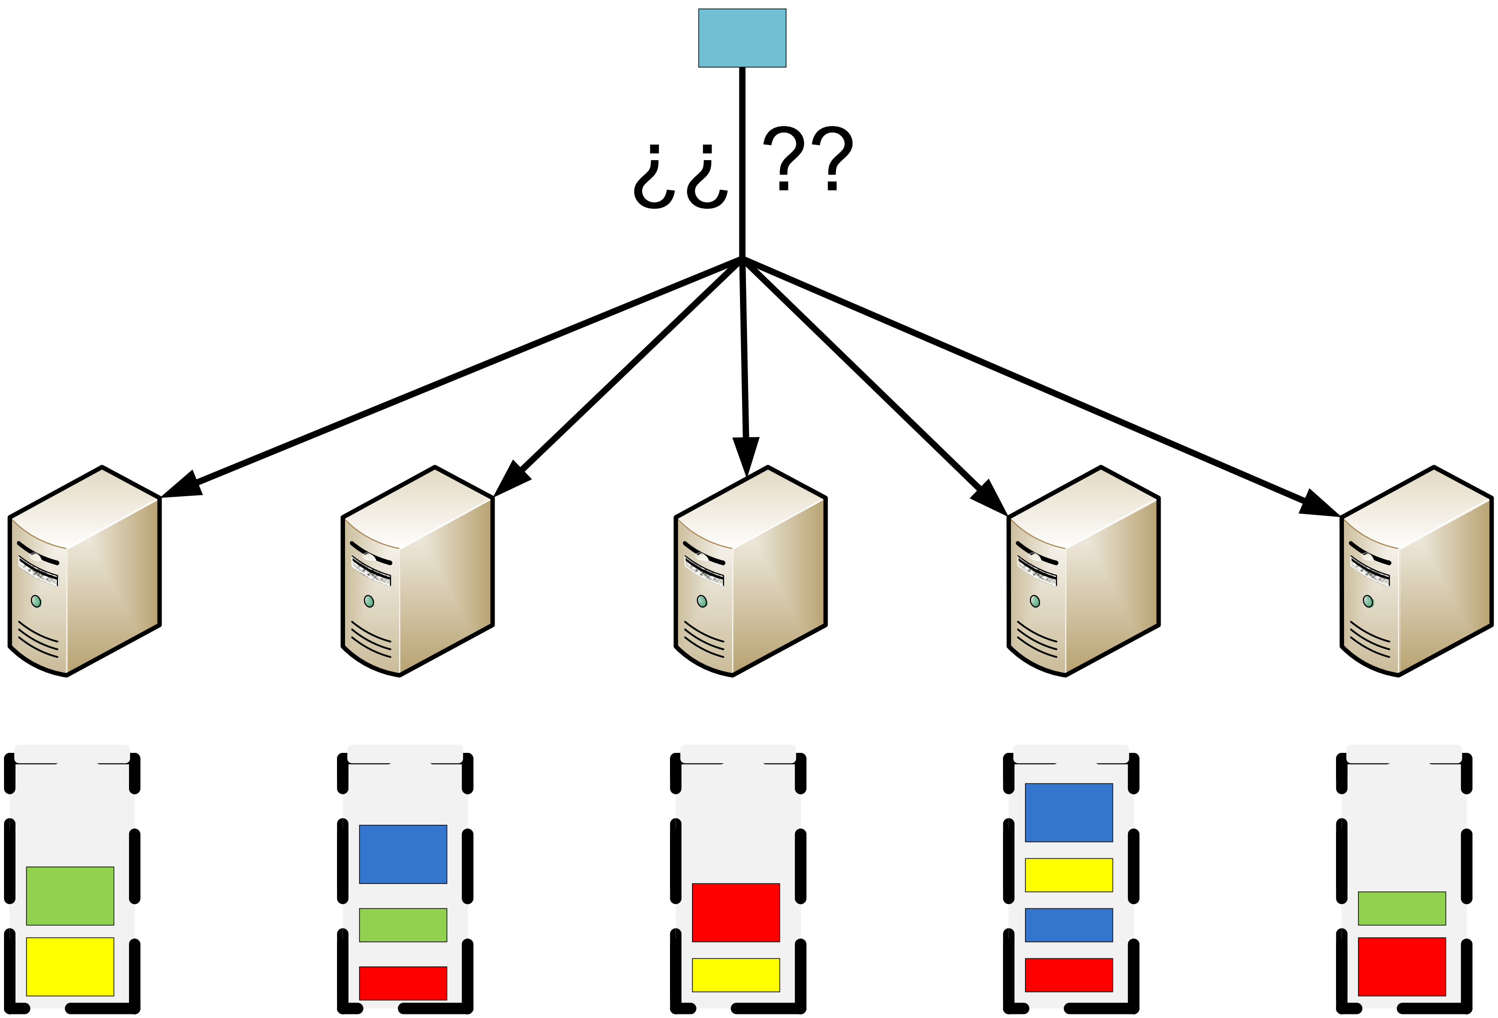

Power Consumption Model

- Power consumption of a PM

$$f(x) = \begin{cases} 0 & x = 0 \\ \mu x^{\alpha} + b & x > 0 \end{cases}$$

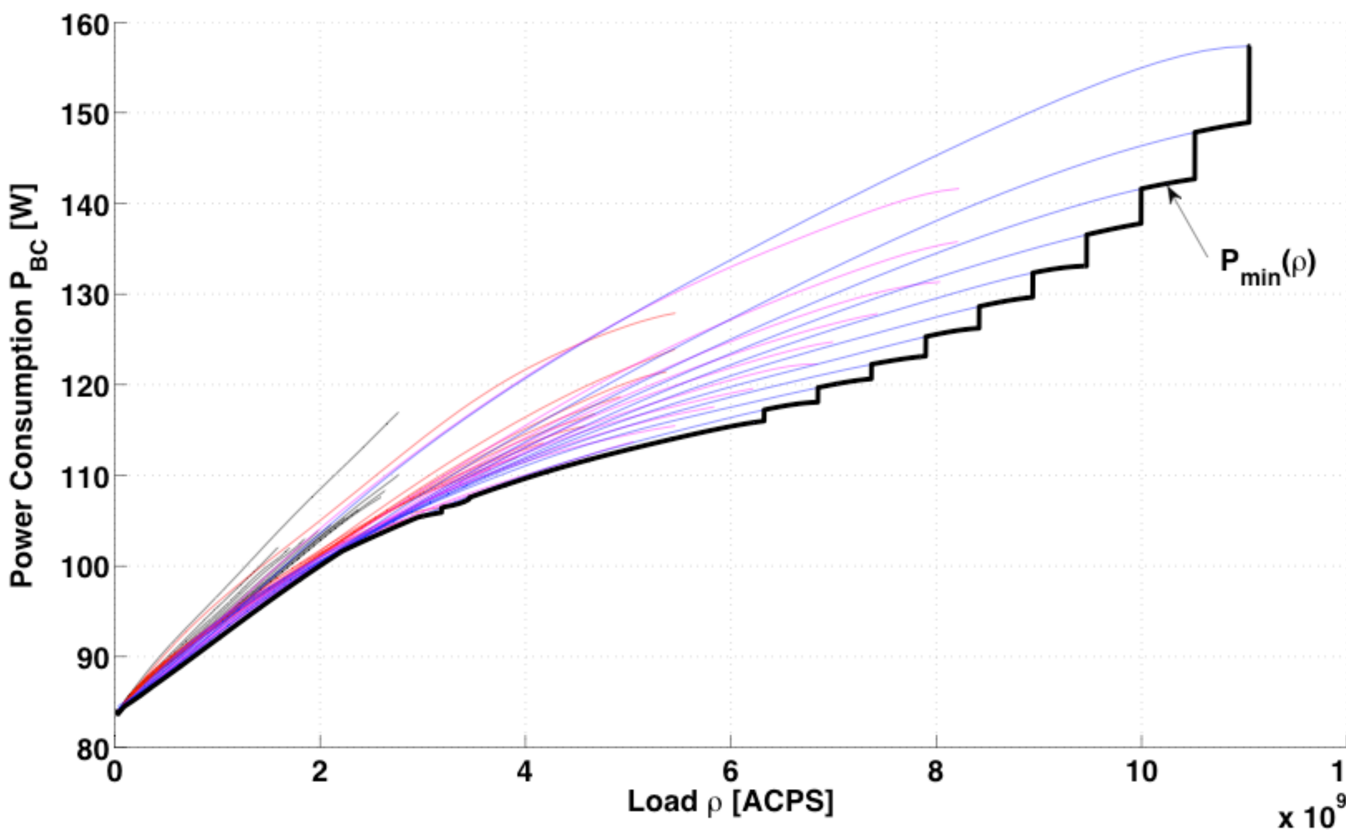

Objective Function

- Given the cost function

$$f(x) = \begin{cases} 0 & x = 0 \\ \mu x^{\alpha} + b & x > 0 \end{cases}$$

- We want to find a partition of the set of VMs that minimizes

$$P(\pi) = \sum_{i \in [1,m]} f(L_i(\pi))$$

  - $L_i(\pi)$ is the aggregated load in the ith PM imposed by π
  - $P(\pi)$ denotes the power consumed by the partition

Virtual Machine Assignment Problem (VMA)

- Variants of VMA:

<u>(C,m)-VMA</u>
Limited Capacity
Finite number of servers

<u>(·,m)-VMA</u>
Unlimited Capacity
Finite number of servers

<u>(C,·)-VMA</u>
Limited Capacity
Unbounded number of servers

<u>(·,·)-VMA</u>
Unlimited Capacity
Unbounded number of servers

- Each variant can be offline or online

Related Work

- Large body of work on consolidation, usually using different power consumption model or not considering energy at all.
- [Alon97, Alon98]: PTAS for the L_p norm problem with $p \geq 1$ and for other more general functions.
- [Epstein04]: Extension of [Alon97,Alon98] for the uniformly related machines case.
- [Srik08]: Energy efficient VMA is not a mere packing problem.

Optimal Load and Power

- There is an optimal load per PM and a corresponding optimal power

$$x^* = (b/(\alpha - 1))^{1/\alpha} \rightarrow \rho^* = f(x^*)/x^*$$

Lemma 1 *Given an instance of the VMA problem with a set of VMs* $D = \{d_1, \ldots, d_n\}$, *any solution* $\pi = \{A_1, \ldots, A_m\}$ *where* $\sum_{d \in A_i} d \neq x^*$ *for some* $i \in [1, m]$, *satisfies*

$$P(\pi) > \rho^* l(D) = \rho^* \sum_{d \in D} l(d).$$

Understanding the Optimal Load

Observation 1 *The optimal load is* $x^* = (b/(\alpha-1))^{1/\alpha}$. *Additionaly, for any* $x \neq x^*$, $f(x)/x > \varphi^*$.

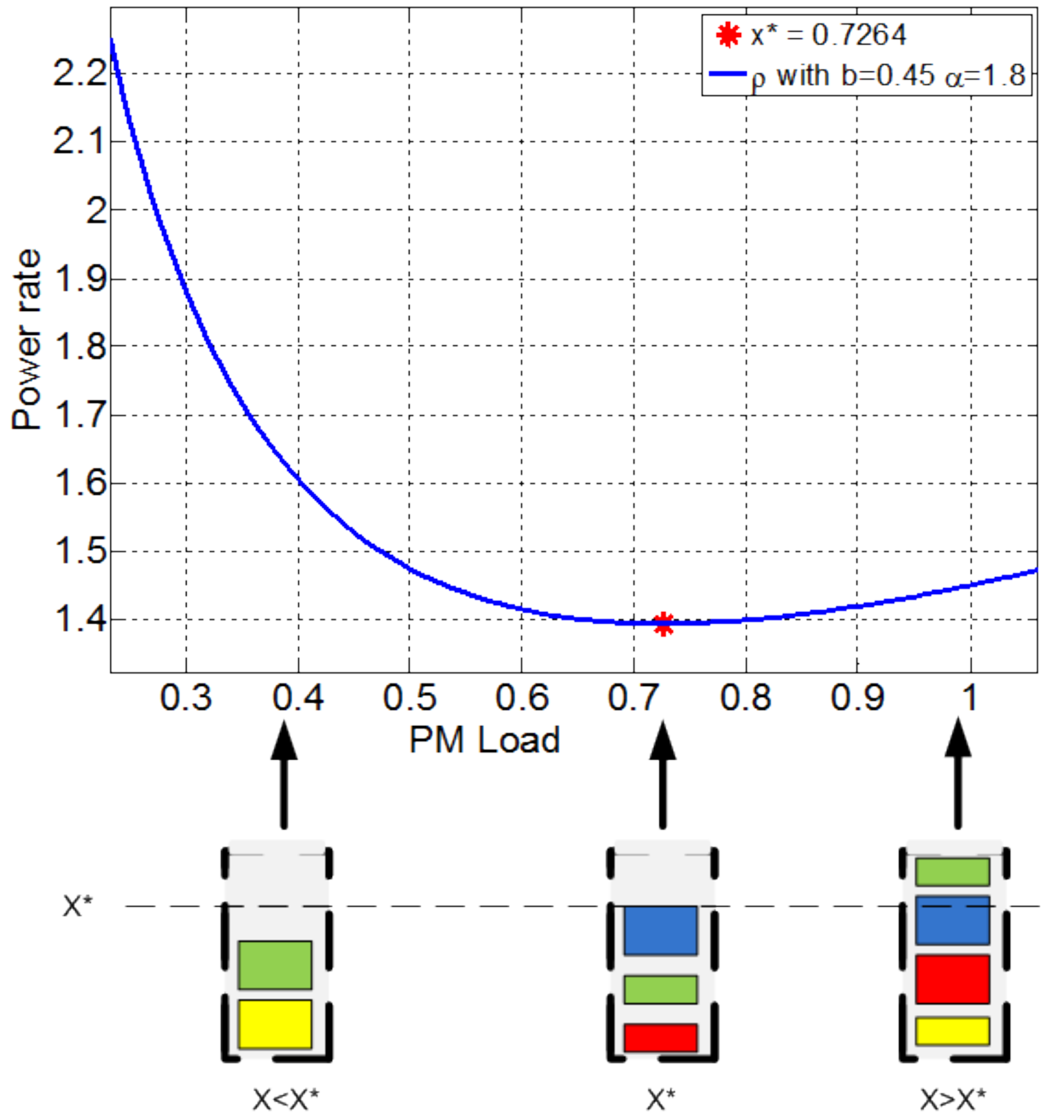

Contributions

<table>
<tr><th colspan="2"></th><th>(C,m)-VMA</th><th>(C,·)-VMA</th><th>(·,m)-VMA</th><th>(·,·)-VMA</th><th>(·,2)-VMA</th></tr>
<tr><td colspan="2">NP-Completeness (decision prob.)</td><td>X</td><td></td><td></td><td></td><td></td></tr>
<tr><td colspan="2">NP-Hardness</td><td></td><td>X</td><td>X</td><td>X</td><td></td></tr>
<tr><td rowspan="2">Offline UB</td><td>x*≥C</td><td></td><td>X</td><td rowspan="4">PTAS</td><td rowspan="4">PTAS</td><td></td></tr>
<tr><td>x*<C</td><td></td><td>X</td><td></td></tr>
<tr><td rowspan="2">Offline LB</td><td>x*≥C</td><td></td><td>X</td><td></td></tr>
<tr><td>x*<C</td><td></td><td>X</td><td></td></tr>
<tr><td rowspan="2">Online UB</td><td>x*≥C</td><td></td><td>X</td><td>N/A</td><td>N/A</td><td>N/A</td></tr>
<tr><td>x*<C</td><td></td><td>X</td><td>X</td><td>X</td><td>X</td></tr>
<tr><td rowspan="2">Online LB</td><td>x*≥C</td><td>X</td><td>X</td><td>N/A</td><td>N/A</td><td>N/A</td></tr>
<tr><td>x*<C</td><td>X</td><td>X</td><td>X</td><td>X</td><td>X</td></tr>
</table>

Contributions

VMA subprob.	$x^* < C$		$x^* \geq C$	
$(C,\cdot)$ offline	$\rho \geq \frac{3}{2}\frac{\alpha-1+(2/3)^\alpha}{\alpha}$	$\rho \geq \frac{11}{9}$	$\rho \geq \frac{3}{2}\frac{\alpha-1+(2/3)^\alpha}{\alpha}$	$\rho \geq \frac{11}{9}$
	$\rho < \frac{\overline{m}}{m^*}\left(1+\epsilon+\frac{1}{\alpha-1}+\frac{1}{m}\right)$	$\rho < \frac{\overline{m}}{m^*}\left(\frac{3}{2}+\epsilon+\frac{1}{m}\right)$	$\rho < 1+\epsilon+\frac{C^\alpha}{b}+\frac{1}{m}$	$\rho < \frac{3}{2}+\epsilon+\frac{1}{m}$
$(C,\cdot)$ online	$\rho \geq \frac{(3/2)2^\alpha-1}{2^\alpha-1}$	$\rho \geq \frac{11}{7}$	$\rho \geq \frac{C^\alpha+2b}{b+\max\{C^\alpha,2(C/2)^\alpha+b\}}$	$\rho \geq \frac{20}{17}$
	$\rho = 1$ if $D_s = \emptyset$, else $\rho \leq \left(1-\frac{1}{\alpha}\left(1-\frac{1}{2^\alpha}\right)\right)\left(2+\frac{x^*}{\ell(D_s)}\right)$	$\rho \leq \frac{17}{12}\left(1+\frac{1}{2\ell(D_s)}\right)$	$\rho \leq \frac{2b}{C}\left(1+\frac{1}{(\alpha-1)2^\alpha}\right)\left(2+\frac{C}{\ell(D)}\right)$	$\rho \leq \frac{17}{2}\left(1+\frac{1}{2\ell(D)}\right)$
(C,m) online	$\rho \geq \frac{(3/2)2^\alpha-1}{2^\alpha-1}$	$\rho \geq \frac{11}{7}$	$\rho \geq \frac{C^\alpha+2b}{b+\max\{C^\alpha,2(C/2)^\alpha+b\}}$	$\rho \geq \frac{20}{17}$
$(\cdot,\cdot)$ online	$\rho \geq \frac{(3/2)2^\alpha-1}{2^\alpha-1}$	$\rho \geq \frac{11}{7}$	not applicable	
	$\rho = 1$ if $D_s = \emptyset$, else $\rho \leq \left(1-\frac{1}{\alpha}\left(1-\frac{1}{2^\alpha}\right)\right)\left(2+\frac{x^*}{\ell(D_s)}\right)$	$\rho \leq \frac{17}{12}\left(1+\frac{1}{2\ell(D_s)}\right)$		
$(\cdot,m)$ online	$\rho \geq \max\{\frac{(3/2)2^\alpha-1}{2^\alpha-1}, \frac{3^\alpha}{2^{\alpha+2}+\epsilon}\}$	$\rho \geq \frac{11}{7}$	not applicable	
	$\rho \leq O(\alpha)^\alpha$ In [?]			
$(\cdot,2)$ online	$\rho \geq \max\{\frac{3^\alpha}{2^{\alpha+1}}, \frac{(3/2)2^\alpha-1}{2^\alpha-1}, \frac{3^\alpha}{2^{\alpha+2}+\epsilon}\}$	$\rho \geq \frac{11}{7}$	not applicable	
	$\rho = 1$ if $\ell(D) \leq \sqrt[\alpha]{b/(2^\alpha-2)}$, else $\rho \leq \max\{2, \left(\frac{3}{2}\right)^{\alpha-1}\}$	$\rho \leq \frac{9}{4}$		

Table 1: Exact values correspond to $\alpha = 3$, $b = 2$, and $C = 2$ on the left and $C = 1$ on the right.

Offline Problems

- The decision version of (C,m)-VMA is NP-complete
  - Directly from 3-partition
- (C,·)-VMA, (·,m)-VMA and (·,·)-VMA are strongly NP-hard
  - Reduction from 3-partition:
    - VMA instance with m PMs, 3m VMs and m·x* total load.
    - Assign 3 VMs to each PM such that the total load in each PM is x*
- (C,·)-VMA, (·,m)-VMA and (·, ·)-VMA have no FPTAS (Fully Polynomial-Time Approximation Scheme)
- There exists a PTAS for (·,·)-VMA and (·,m)-VMA
  - Directly from [Epstein2004]

Online Problem: Lower Bounds

Theorem 6 *There exists an instance of problems* $(\cdot,\cdot)$*-VMA,* $(\cdot,m)$*-VMA,* $(C,\cdot)$*-VMA and* (C,m)*-VMA when* $C > x^*$*, such that no online algorithm can guarantee a competitive ratio smaller than* $\frac{(3/2)2^{\alpha}-1}{2^{\alpha}-1}$.

- Intuition
  - Assume we have VMs of size ϵx^*
  - An adversary starts injecting VMs, which are placed in the same PM, until the algorithm opens a second PM with the kth VM.

ϵx^* ϵx^* ... ϵx^*

  - If $k \leq \frac{1}{\epsilon}\left(\frac{\alpha-1}{1-2^{1-\alpha}}\right)^{1/\alpha}$, then

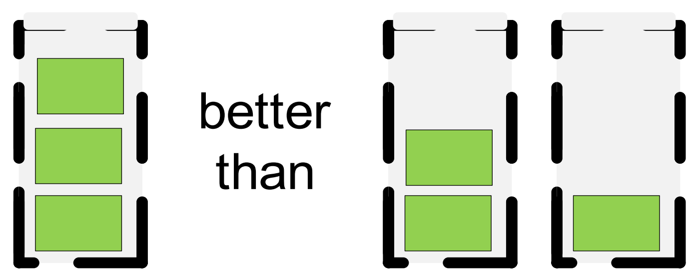

  - If $k > \frac{1}{\epsilon}\left(\frac{\alpha-1}{1-2^{1-\alpha}}\right)^{1/\alpha}$, then

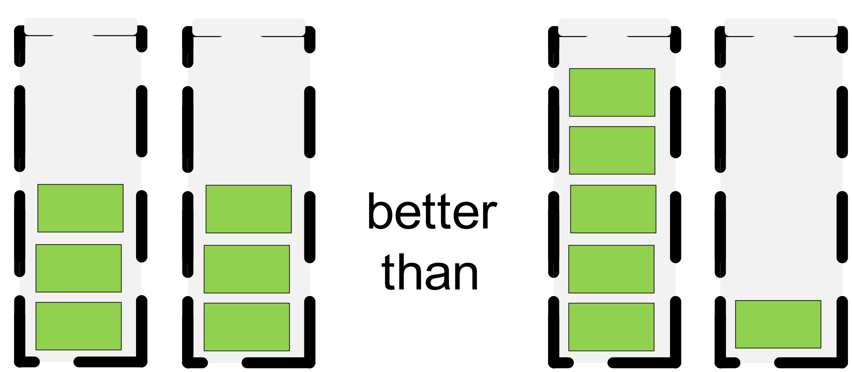

Online Problem: Upper Bounds

Algorithm 1: Online algorithm for $(\cdot,\cdot)$-VMA and $(C,\cdot)$-VMA problems.

for *each VM* d_i **do**
- **if** $\ell(d_i) > \frac{\min\{x^*, C\}}{2}$ **then**
  - d_i is assigned to a new PM
- **else**
  - d_i is assigned to any loaded PM s_j where $\ell(A_j) \leq \frac{\min\{x^*, C\}}{2}$. If such loaded PM does not exist, d_i is assigned to a new PM

Theorem 10 *There exists an online algorithm for $(\cdot,\cdot)$-VMA and $(C,\cdot)$-VMA when $x^* < C$ that achieves the following competitive ratio:*

$$\begin{aligned} \rho &= 1, \text{ if no VM } d_i \text{ has load such that } l(d_i) < x^*, \\ \rho &\leq \left(1 - \frac{1}{\alpha}\left(1 - \frac{1}{2^\alpha}\right)\right)\left(2 + \frac{x^*}{\ell(D_s)}\right), \text{ otherwise.} \end{aligned}$$

Theorem 11 *There exists an online algorithm for $(C,\cdot)$-VMA when $x^* \geq C$ that achieves competitive ratio* $\rho \leq \frac{2b}{C}\left(1 + \frac{1}{(\alpha-1)2^\alpha}\right)\left(2 + \frac{C}{\ell(D)}\right)$.

Online Problem: Upper Bounds

- Intuition
  - VMs of load larger than x*/2 go alone in one PM
  - VMs with load < x*/2 are assigned to the most loaded PM with load < x*/2
  - There can be one PM with load < x*/2
  - The optimal assignent would have been:

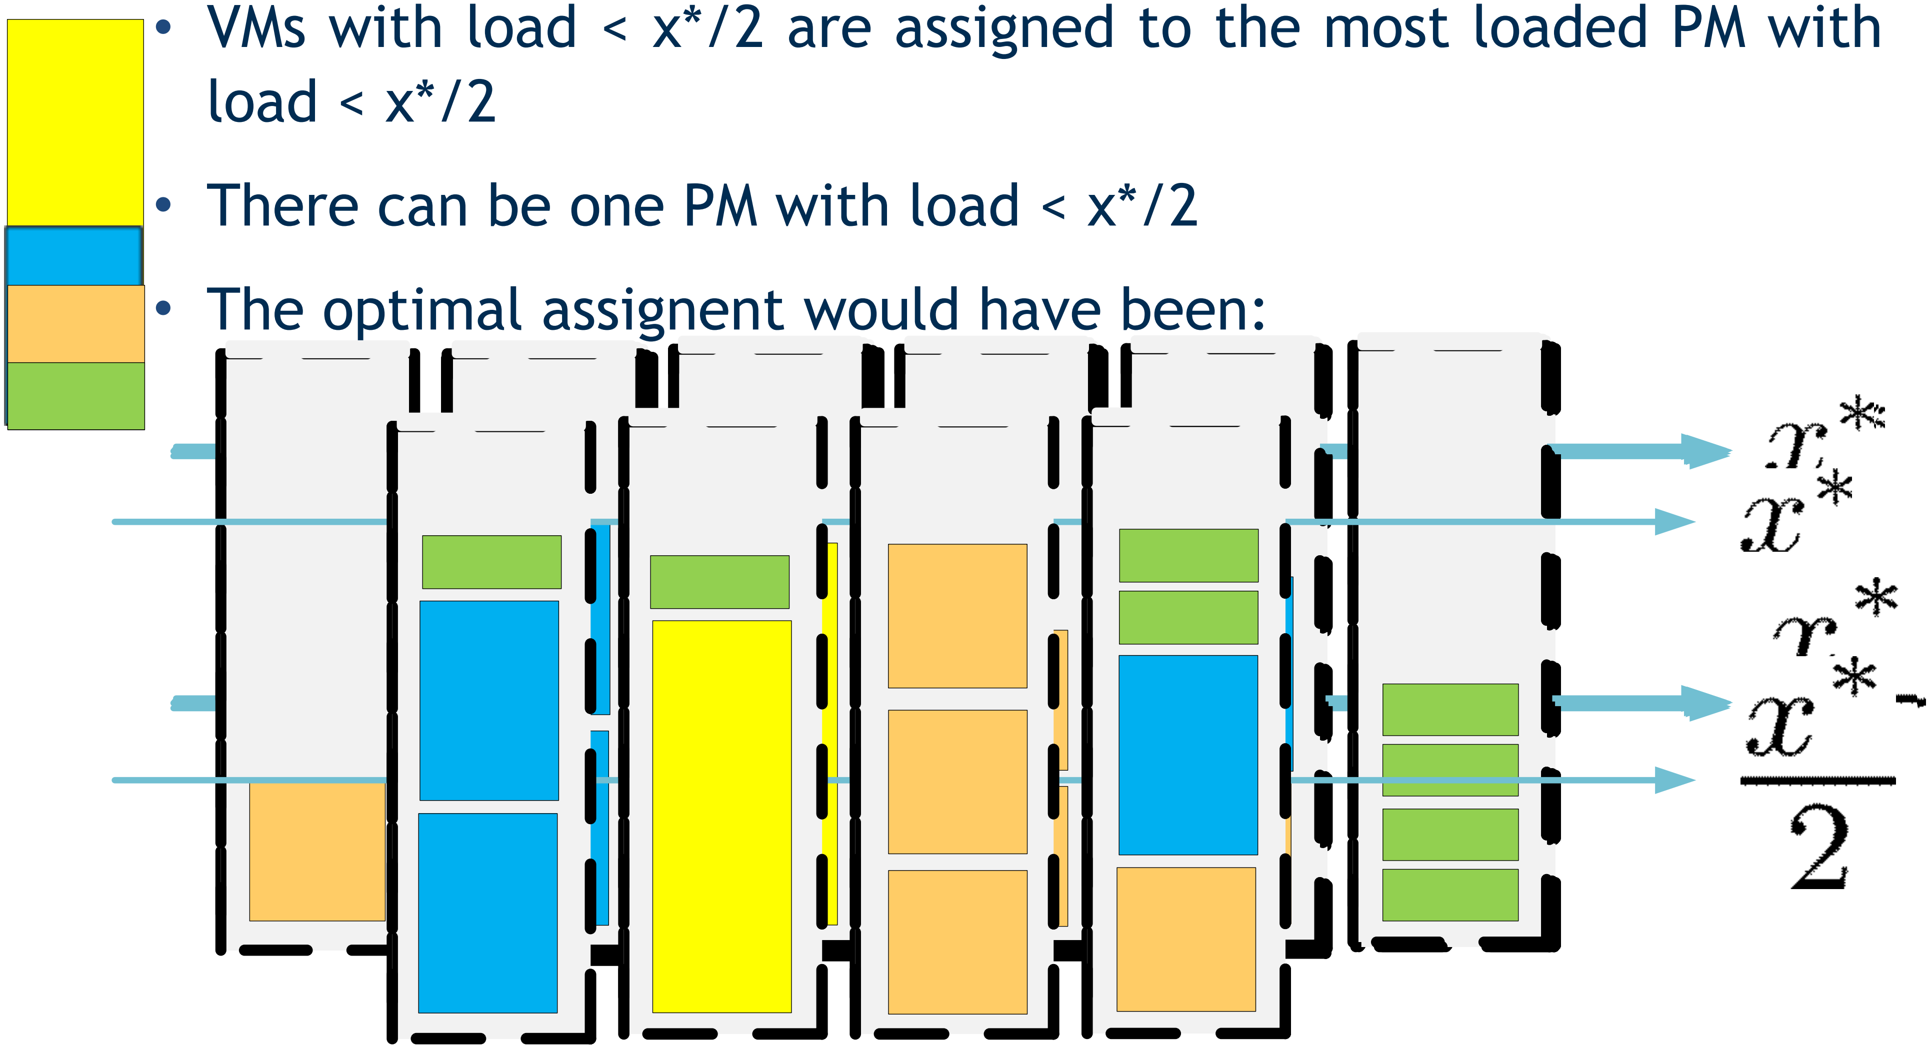

Conclusions

- It is possible to obtain bounds for multiple versions of the VMA problem
- Many variants to be studied
  - Hetereogenity of servers
  - Migration of VM between PMs at a cost
  - VM that arrive and depart
  - VM that chage their load over time
  - Multi-resource scheduling, where the load is not only given by CPU load

Thank you!!

Questions?

Preliminary Claims

Lemma 1 *Consider two solutions* $\pi = \{A_1, \ldots, A_m\}$ *and* $\pi' = \{A'_1, \ldots, A'_m\}$ *of an instance of the VMA problem, such that for some* $x, y \in [1, m]$ *it holds that:* $A_x \neq \emptyset$ *and* $A_y \neq \emptyset$*;* $A'_x = A_x \cup A_y$*,* $A'_y = \emptyset$*, and* $A_i = A'_i$*, for all* $i \neq x$ *and* $i \neq y$*; and* $l(A_x) + l(A_y) \leq \min\{x^*, C\}$*. Then,* $P(\pi') < P(\pi)$*.*

- Intuition:
  - Depending on their load, it might be better to put VMs in the same (different) PM rather than combining (dissociating) them.

Lemma 2 *Consider two solutions* $\pi = \{A_1, \ldots, A_m\}$ *and* $\pi' = \{A'_1, \ldots, A'_m\}$ *of an instance of the VMA problem, such that for some* $x, y \in [1, m]$ *it holds that:* $A_x \cup A_y = A'_x \cup A'_y$*, while* $A_i = A'_i$*, for all* $x \neq i \neq y$*; none of* A_x*,* A_y*,* A'_x*, and* A'_y *is empty; and* $|\ell(A_x) - \ell(A_y)| < |\ell(A'_x) - \ell(A'_y)|$*. Then,* $P(\pi) < P(\pi')$*.*

- Intuition
  - PMs with unbalanced load consume more energy than PMs with the same aggregated load evenly distributed among them.

Corollary 1 *(short)"... power consumption is lower bounded by the power of the (maybe unfeasible) solution that balances the load evenly, i.e.,* $P(\pi) \geq kb + k(Ł/k)^{\alpha} \ldots$*"*

Offline Problem

- Bounds on the approximability of (C,·)-VMA

Theorem 3 *No algorithm achieves an approximation ratio smaller than* $\frac{3}{2} \cdot \frac{\alpha - 1 + (\frac{2}{3})^\alpha}{\alpha}$ *for the* $(C, \cdot)$*-VMA problem unless* $\mathrm{P} = \mathrm{NP}$.

Theorem 4 *For every* $\epsilon > 0$*, there exists an approximation algorithm for the* $(C, \cdot)$*-VMA problem when* $x^* \geq C$ *that achieves an approximation ratio of* $\rho < 1 + \epsilon + \frac{C^\alpha}{b} + \frac{1}{\overline{m}}$*, where* $\overline{m}$ *is the minimum number of PMs required to allocate all the VMs.*

Theorem 5 *For every* $\epsilon > 0$*, there exists an approximation algorithm for the* $(C, \cdot)$*-VMA problem when* $x^* < C$ *that achieves an approximation ratio of* $\rho < \frac{\overline{m}}{m^*}\left((1+\epsilon) + \frac{1}{\alpha - 1}\right) + \frac{1}{m^*}$*, where* m^* *is the number of PMs used by the optimal solution of* $(C, \cdot)$*-VMA, and* $\overline{m}$ *is the minimum number of PMs required to allocate all the VMs without exceeding load* x^* *(i.e., the optimal solution of the bin packing problem).*

- Intuitions
  - LB: Based on the partition problem
  - UB when x*≥C: Optimal solution is lower bounded by a evenly balanced load among the optimal number of PMs (optimal bin packing solution)
  - UB when x*<C: Optimal number of PMs considering Bin Packing with bins of size x*

Online Problem: Lower Bounds

Theorem 7 *There exists an instance of problems $(C,\cdot)$-VMA and (C,m)-VMA when $C \leq x^*$ such that no online algorithm can guarantee a competitive ratio smaller than $(C^\alpha + 2b)/(b + \max(C^\alpha, 2(C/2)^\alpha + b))$.*

- Intuition
  - Consider the same adversarial injection strategy with VMs of size ϵC
  - This time, the thresholds for k and the resulting ratios are:

$$k \leq \frac{1}{\epsilon} \rightarrow \rho(k) \geq \frac{C^\alpha + 2b}{C^\alpha + b} \geq 2 - \frac{1}{\alpha}$$
$$k > \frac{1}{\epsilon} \rightarrow \rho(k) = \frac{C^\alpha + 2b}{2\left(\frac{C}{2}\right)^\alpha + 2b}$$

  - Which, combined, throw the result from Theorem 7.

Online Problem: Lower Bounds

Theorem 8 *There exists an instance of problem $(\cdot, m)$-VMA such that no online algorithm can guarantee a competitive ratio smaller than $3^\alpha/(2^{\alpha+2}+\epsilon)$ for any $\epsilon > 0$.*

- Intuition

Define $\beta > 1$ such that $\epsilon \ge (3^\alpha - 1)/\beta^\alpha$

Input m VMs of load βx^*

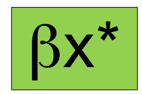

If $\frac{3m}{4}$ or less PMs are used $\rho \ge 2^{-\alpha} + \frac{1}{4}$

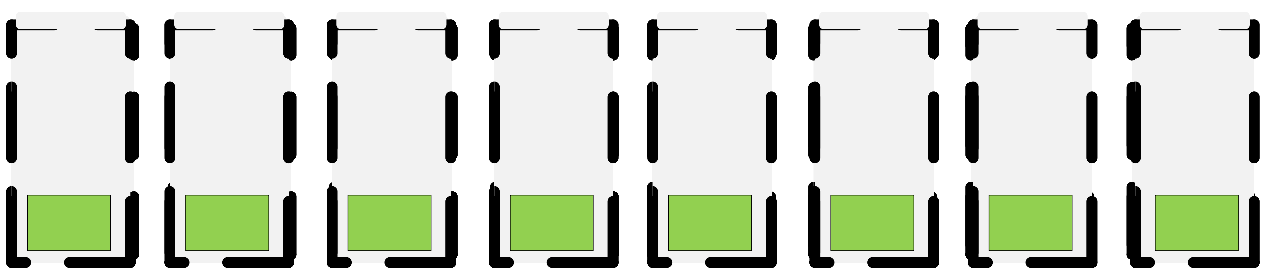

else: we input another set of $\frac{m}{2}$ VMs of size $2\beta x^*$

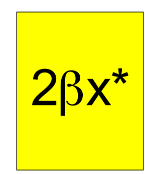

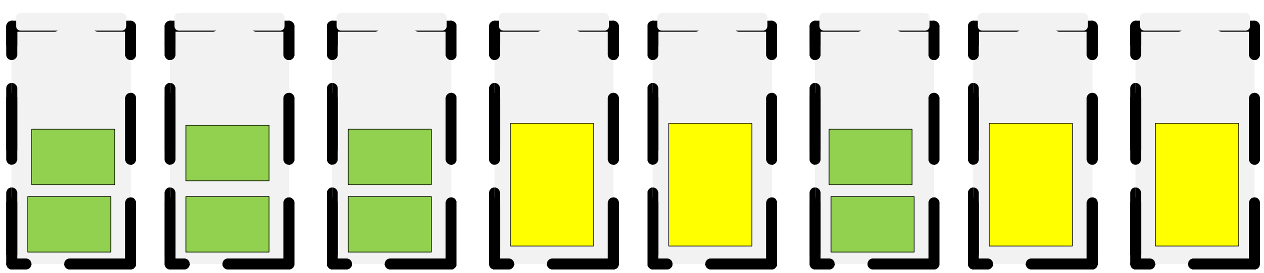

Online Problem: Lower Bounds

Theorem 9 *There exists an instance of problem $(\cdot, 2)$-VMA such that no online algorithm can guarantee a competitive ratio smaller than $3^\alpha/2^{\alpha+1}$.*

- Intuition

Input 2 VMs of size $6x^*$

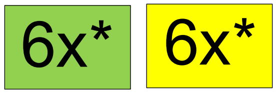

If they are assigned to the same server, we stop

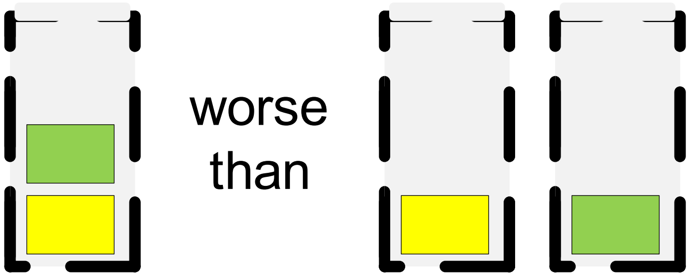

Else, we input a $12x^*$ load VM

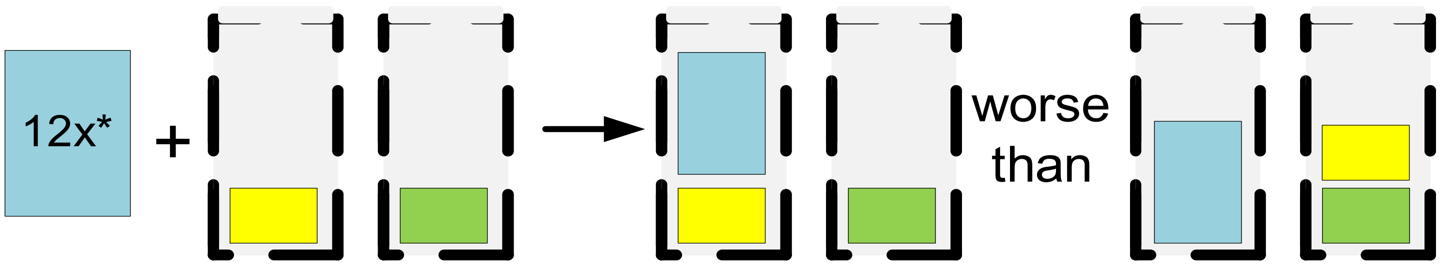

Online Problem: Upper Bounds

Algorithm 2: Online algorithm for $(\cdot, 2)$-VMA.

for *each VM* d_i **do**
  if $l(d_i) + l(A_1) \leq (b/(2^\alpha - 2))^{1/\alpha}$ **or** $l(A_1) \leq l(A_2)$ **then**
    d_i is assigned to s_1
  else
    d_i is assigned to s_2

Theorem 12 *There exists an online algorithm for* $(\cdot, 2)$*-VMA that achieves the following competitive ratios.*

$$\rho = 1, \quad \text{for } l(D) \leq \left(\frac{b}{2^\alpha - 2}\right)^{1/\alpha},$$

$$\rho \leq \max\left\{2, \left(\frac{3}{2}\right)^{\alpha-1}\right\}, \quad \text{for } l(D) > \left(\frac{b}{2^\alpha - 2}\right)^{1/\alpha}.$$

Online Problem: Upper Bounds

- Intuition
  - There are 3 phases

$$1)\quad L \le \left(\frac{b}{2^{\alpha}-2}\right)^{\frac{1}{\alpha}} \qquad\qquad L \to \text{Total Load}$$

$$2)\quad \left(\frac{b}{2^{\alpha}-2}\right)^{\frac{1}{\alpha}} < L < 2\left(\frac{b}{2^{\alpha}-2}\right)^{\frac{1}{\alpha}}$$

$$3)\quad L \ge 2\left(\frac{b}{2^{\alpha}-2}\right)^{\frac{1}{\alpha}}$$