# Implementation of Algorithms for Right-Sizing Data Centers

Author: Jonas Hübotter

Supervisor: Prof. Dr. Susanne Albers
Advisor: Jens Quedenfeld

Department of Informatics
Technical University of Munich

August 13, 2021

# Outline

# Motivation

- data centers use between 1% and 3% of global energy[1], which is estimated to increase[2]
- most data centers are statically provisioned, leading to average utilization levels between 12% and 18%[3]
- typically servers operate at energy efficiency levels between 20% and 30%[4]
- when idling, servers consume half of their peak power[4]

---

[1] Arman Shehabi et al. *United States Data Center Energy Usage Report*. Tech. rep. Lawrence Berkeley National Laboratory, June 2016.

[2] Nicola Jones. "How to stop data centres from gobbling up the world's electricity". In: *Nature* 561.7722 (2018), pp. 163–167.

[3] Josh Whitney and Pierre Delforge. *Data Center Efficiency Assessment*. Natural Resources Defense Council, Aug. 2014.

[4] Luiz André Barroso and Urs Hölzle. "The case for energy-proportional computing". In: *Computer* 40.12 (2007), pp. 33–37.

# Problem

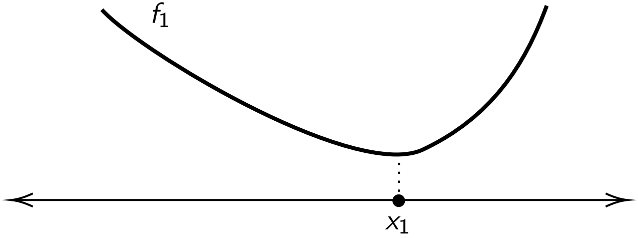

# Problem

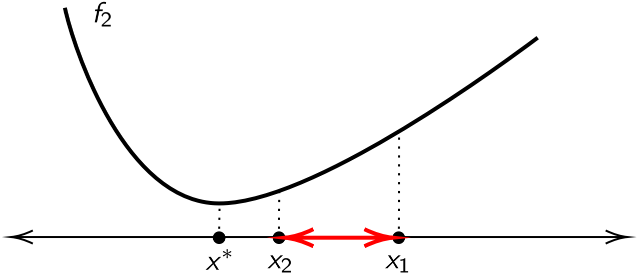

# Model

What is the cost of operating a data center with $x_t \in \mathbb{N}_0$ active servers and under load $\lambda_t \in \mathbb{N}_0$?

- How to distribute jobs across the active servers?
  Distribute evenly across all servers of the same type[5].
- What is the cost associated with such an assignment?
  Consisting of energy costs and the revenue loss incurred by a delayed processing of jobs.
  Algorithms need to *balance* energy costs and revenue loss.

Movement costs are on the order of operating an idling server for 1-4 hours[6].

---

[5] Susanne Albers and Jens Quedenfeld. "Algorithms for Right-Sizing Heterogeneous Data Centers". In: *Proceedings of the 33rd ACM Symposium on Parallelism in Algorithms and Architectures*. 2021, pp. 48–58.

[6] Minghong Lin et al. "Dynamic right-sizing for power-proportional data centers". In: *IEEE/ACM Transactions on Networking* 21.5 (2012), pp. 1378–1391.

# Algorithms for one dimension

| problem | algorithm | results |
|---|---|---|
| fractional | Lazy Capacity Provisioning[7] | 3-competitive |
| | Memoryless[8] | 3-competitive |
| | Probabilistic[8] | 2-competitive |
| | Randomly Biased Greedy[9], $\theta \geq 1$ | $(1+\theta)$-competitive, $\mathcal{O}(\max\{T/\theta, \theta\})$-regret |
| integral | Lazy Capacity Provisioning[10] | 3-competitive |
| | Randomized[10] | 2-competitive |

[7] Minghong Lin et al. "Dynamic right-sizing for power-proportional data centers". In: *IEEE/ACM Transactions on Networking* 21.5 (2012), pp. 1378–1391.

[8] Nikhil Bansal et al. "A 2-competitive algorithm for online convex optimization with switching costs". In: *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2015)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik. 2015.

[9] Lachlan Andrew et al. "A tale of two metrics: Simultaneous bounds on competitiveness and regret". In: *Conference on Learning Theory*. PMLR. 2013, pp. 741–763.

[10] Susanne Albers and Jens Quedenfeld. "Optimal algorithms for right-sizing data centers". In: *Proceedings of the 30th on Symposium on Parallelism in Algorithms and Architectures*. 2018, pp. 363–372.

# Algorithms for multiple dimensions

| problem | algorithm | results |
|---|---|---|
| integral; linear, time-indep. cost | Lazy Budgeting[11] (deterministic) | $2d$-competitive |
| | Lazy Budgeting[11] (randomized) | $\approx 1.582d$-competitive |
| integral; hom. load | Lazy Budgeting[12] | $(2d + 1 + \epsilon)$-competitive |
| fractional; $\alpha$-loc. polyhedral costs; $\ell_2$ movement | Primal OBD[13] | $3 + \mathcal{O}(1/\alpha)$-competitive |
| | Dual OBD[13] | $\mathcal{O}(\sqrt{T})$-regret |
| fractional; prediction window | RHC[14] | $(1 + \mathcal{O}(1/w))$ -competitive in 1d |
| | AFHC[14] | $(1 + \mathcal{O}(1/w))$-competitive |

---

[11] Susanne Albers and Jens Quedenfeld. "Algorithms for Energy Conservation in Heterogeneous Data Centers.". In: *CIAC*. 2021, pp. 75–89.

[12] Susanne Albers and Jens Quedenfeld. "Algorithms for Right-Sizing Heterogeneous Data Centers". In: *Proceedings of the 33rd ACM Symposium on Parallelism in Algorithms and Architectures*. 2021, pp. 48–58.

[13] Niangjun Chen, Gautam Goel, and Adam Wierman. "Smoothed online convex optimization in high dimensions via online balanced descent". In: *Conference On Learning Theory*. PMLR. 2018, pp. 1574–1594.

[14] Minghong Lin et al. "Online algorithms for geographical load balancing". In: *2012 international green computing conference (IGCC)*. IEEE. 2012, pp. 1–10.

# Traces

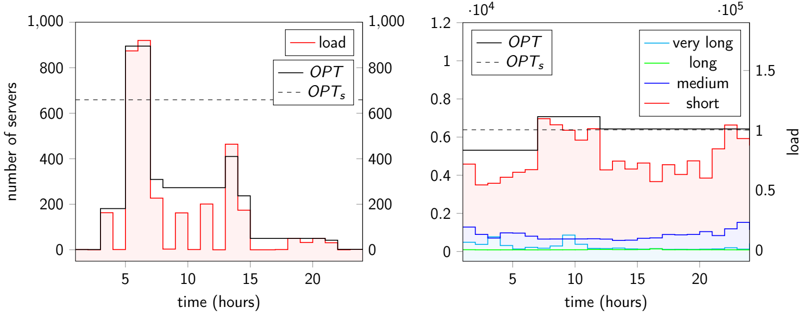


(a) **LANL Mustang** (b) **Alibaba**

# Performance metrics

- normalized cost: $c(ALG)/c(OPT)$
- cost reduction:

$$\frac{c(OPT_s) - c(ALG)}{c(OPT_s)}$$

- static/dynamic ratio: $c(OPT_s)/c(OPT)$

# Results in one dimension

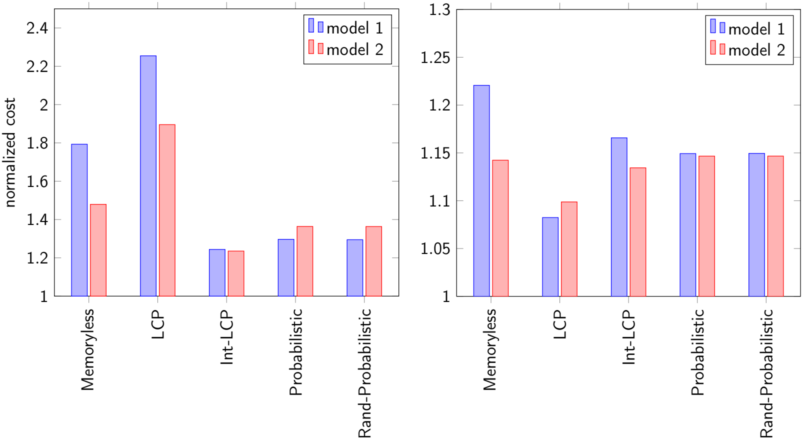


**(a) LANL Mustang** **(b) Alibaba**

# Results in one dimension

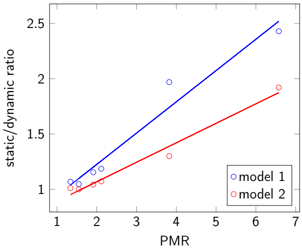


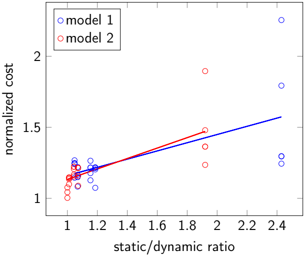


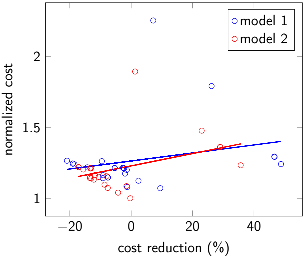

# Other results

## Multiple dimensions

- lazy budgeting algorithms perform nearly optimally (normalized cost $\in [1.05, 1.25]$), without consideration of revenue loss
- descent methods achieve normalized costs of $\approx 2.5$

## With predictions

- even a short prediction window of several hours can significantly improve the results (by $\approx 5\%$)
- robust to imperfect (realistic) predictions

# Future work

- compare performance to algorithms for convex body chasing
- performance of algorithms in other applications
- better algorithms to make use of predictions

Thanks for your attention! Questions?

# Problem

Smoothed online convex optimization (or *convex function chasing*)[15]: Given a convex decision space $\mathcal{X} \subset \mathbb{R}^d$, a norm $\|\cdot\|$ on $\mathbb{R}^d$, and a sequence $F$ of non-negative convex functions $f_t : \mathcal{X} \to \mathbb{R}_{\geq 0}$, find $x \in \mathcal{X}^T$ such that

$$\sum_{t=1}^{T} f_t(x_t) + \|x_t - x_{t-1}\|$$

is minimized where $T$ is the time horizon and $x_0 = 0$.

[15] Minghong Lin et al. "Dynamic right-sizing for power-proportional data centers". In: *IEEE/ACM Transactions on Networking* 21.5 (2012), pp. 1378–1391.

# Problem

- similar to *online convex optimization* with movement costs and lookahead 1
- equivalent to *convex body chasing* in $d + 1$
- fundamental incompatibility between competitive ratio and regret even for linear hitting costs in one dimension