Learning-Augmented Online Algorithms

Stefan Schmid (TU Berlin)

“We cannot direct the wind,
but we can adjust the sails.”

(Folklore)

Acknowledgements:

It‘s a great time to be an algorithm researcher!

Flexibilities everywhere in networked systems!

It‘s a great time to be an algorithm researcher!

Flexibilities everywhere in networked systems!

Enables ***self-adjusting systems***: adapt to demand.

This Talk: Datacenters

Datacenters ("hyper-scale")

Interconnecting networks:
a **critical infrastructure**
of our digital society.

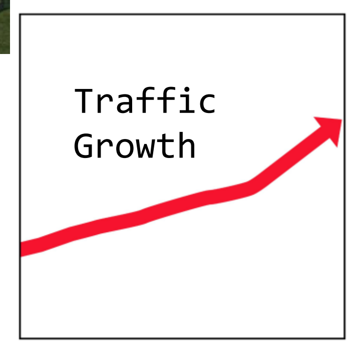

Source: Facebook

This Talk: Datacenters

Datacenters ("hyper-scale")

+network

Interconnecting networks:
a **critical infrastructure**
of our digital society.

Credits: Marco Chiesa

Example 1: Scheduling

Online Re-Partitioning (Sigmetrics‘19, SODA‘21)

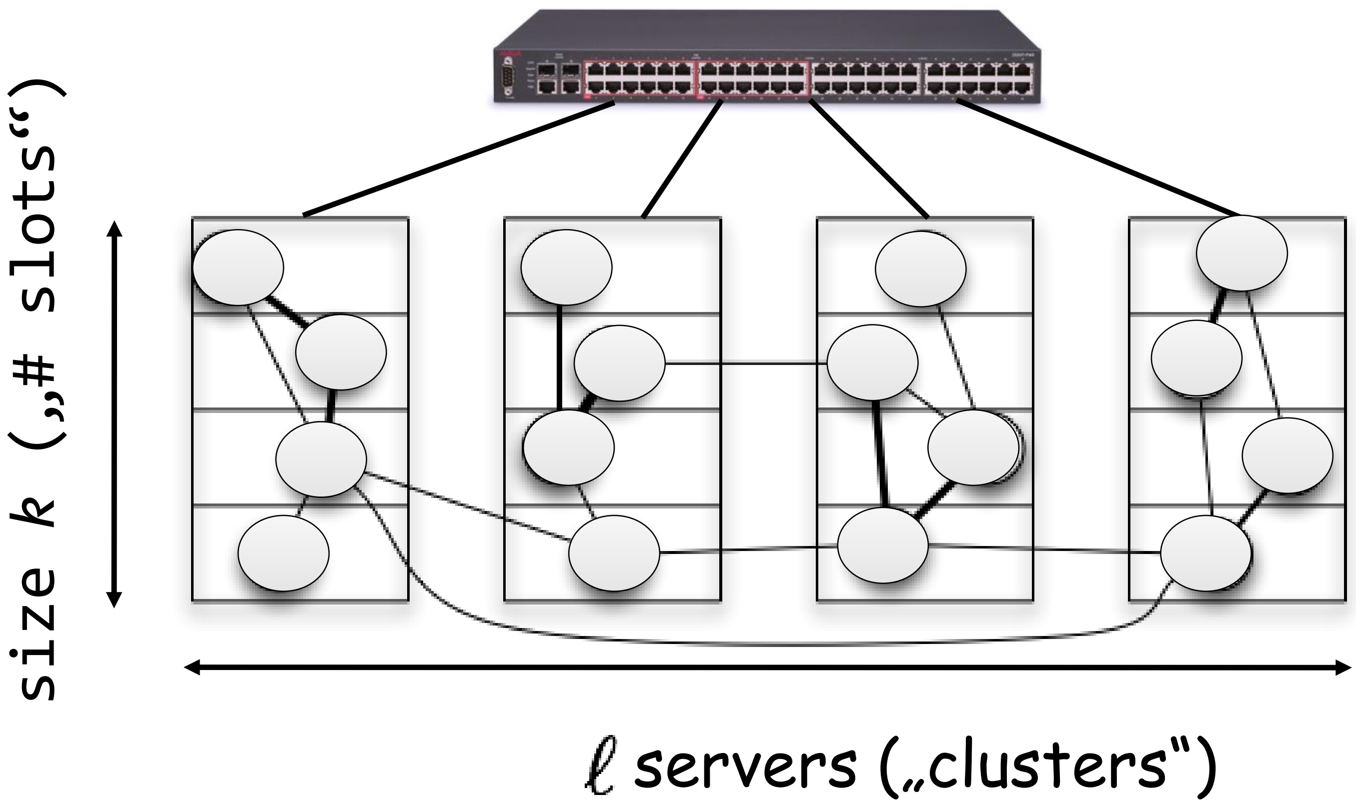

Example 1: Scheduling

Online Re-Partitioning (Sigmetrics‘19, SODA‘21)

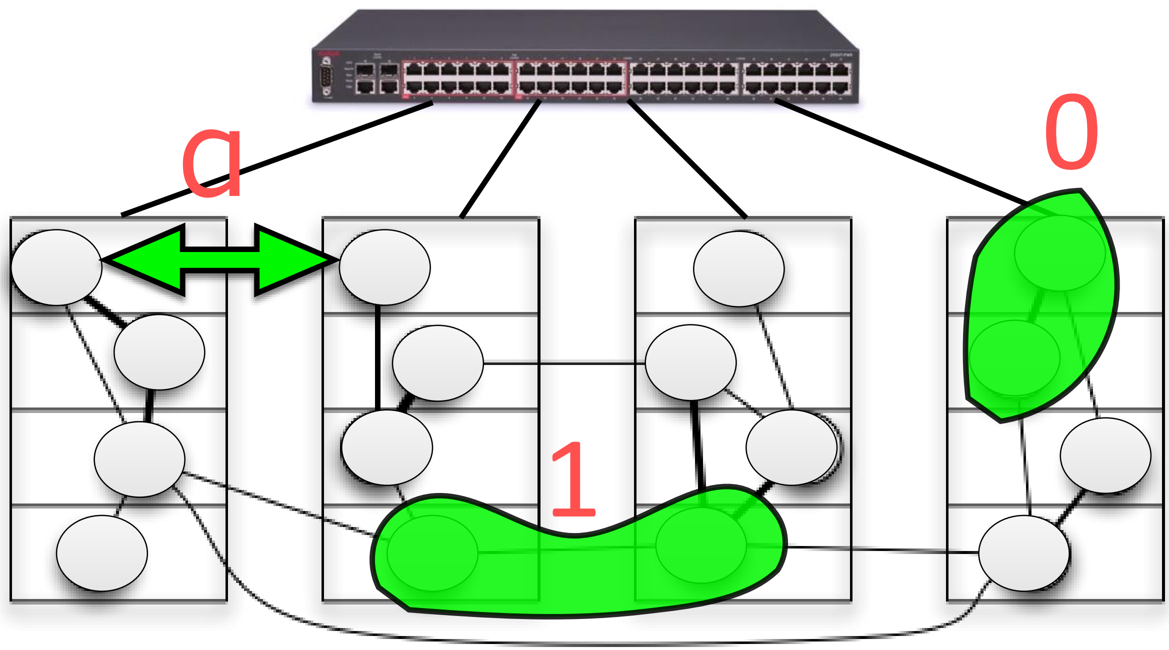

Migrate to reduce communication costs? Tradeoff!

Example 2: Topology

Self-Adjusting Datacenters (Sigmetrics'23)

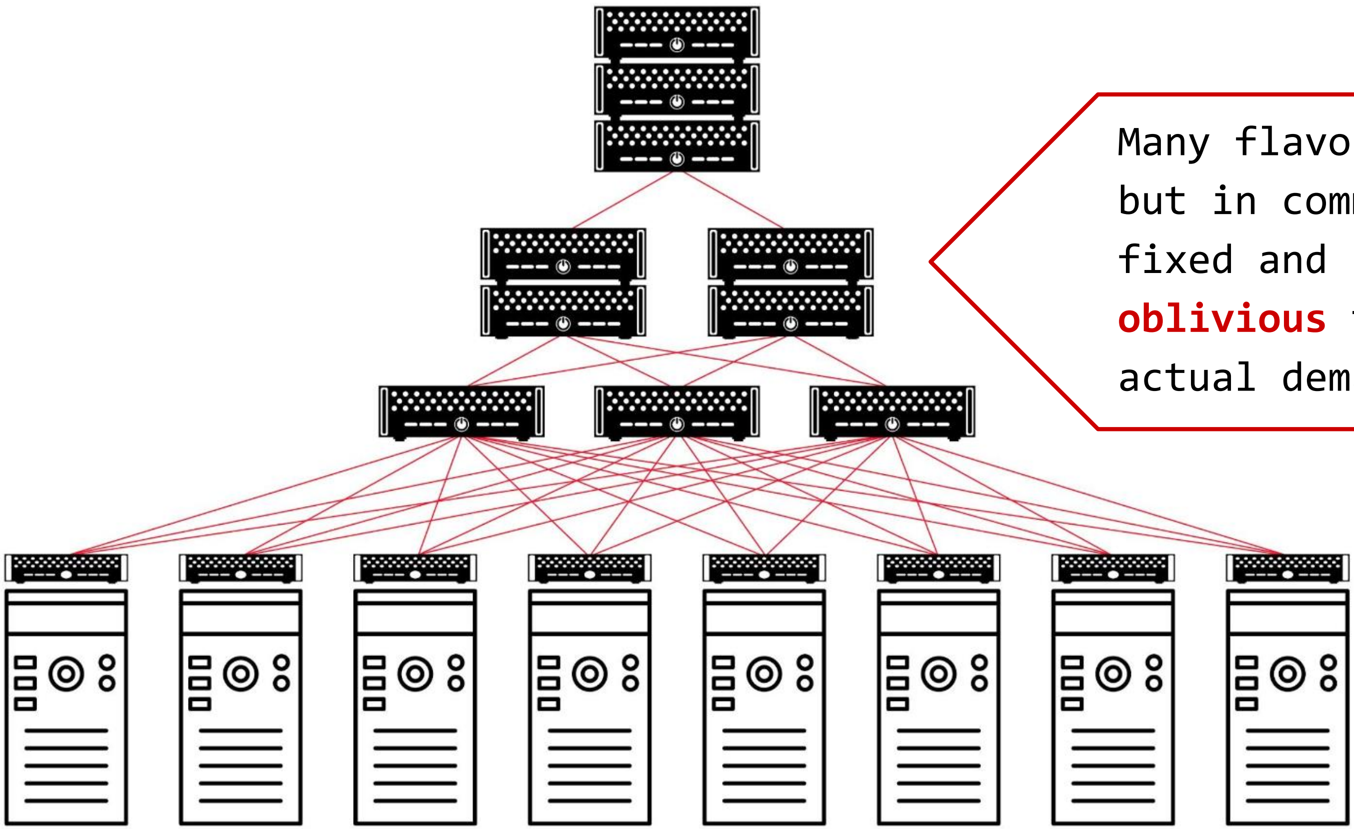

Many flavors, but in common: fixed and **oblivious** to actual demand.

Example 2: Topology

Self-Adjusting Datacenters (Sigmetrics'23)

Highway which ignores actual traffic: **frustrating!**

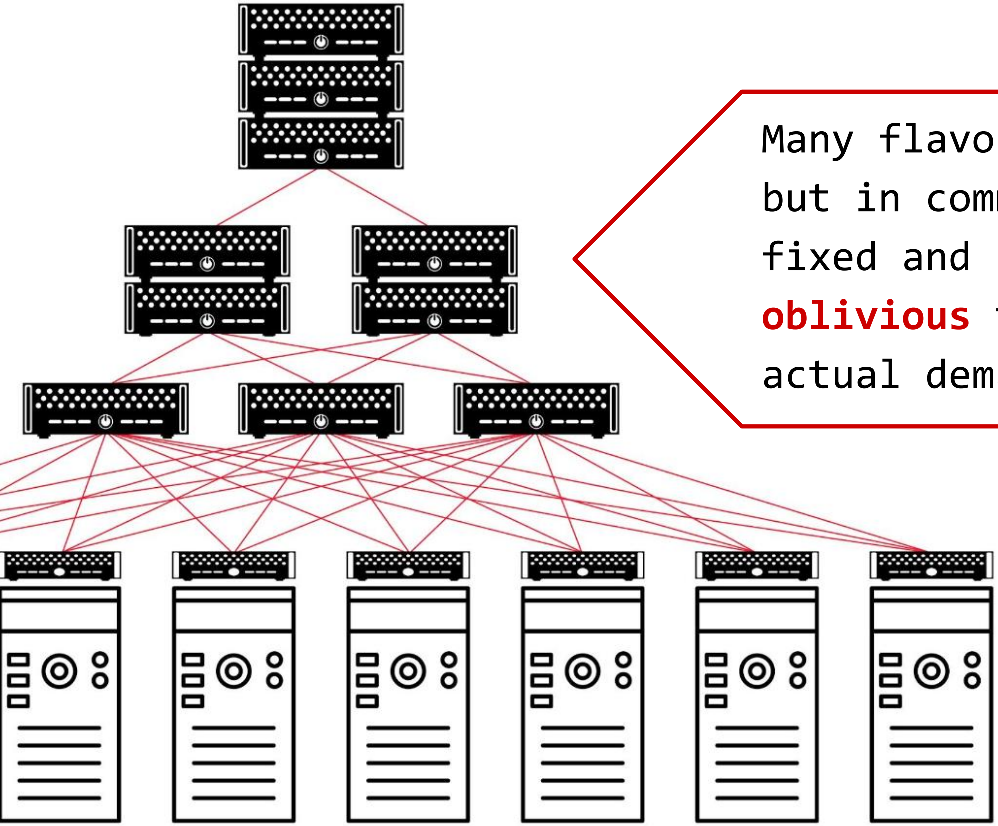

Many flavors, but in common: fixed and **oblivious** to actual demand.

Example 2: Topology

Self-Adjusting Datacenters (Sigmetrics'23)

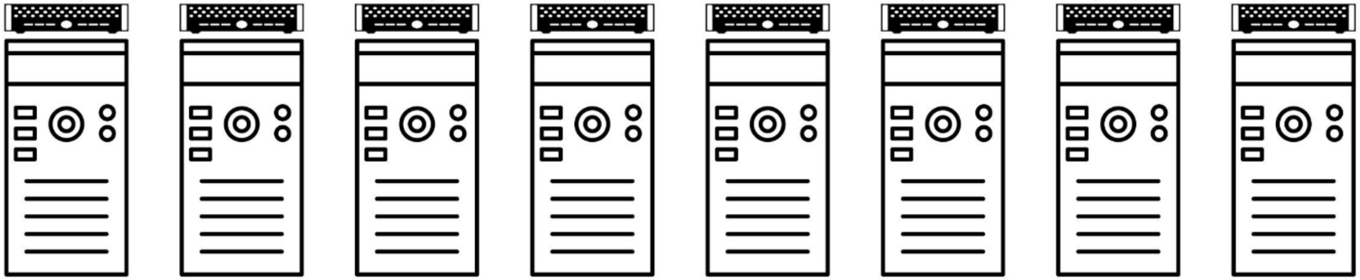

Example 2: Topology

Self-Adjusting Datacenters (Sigmetrics'23)

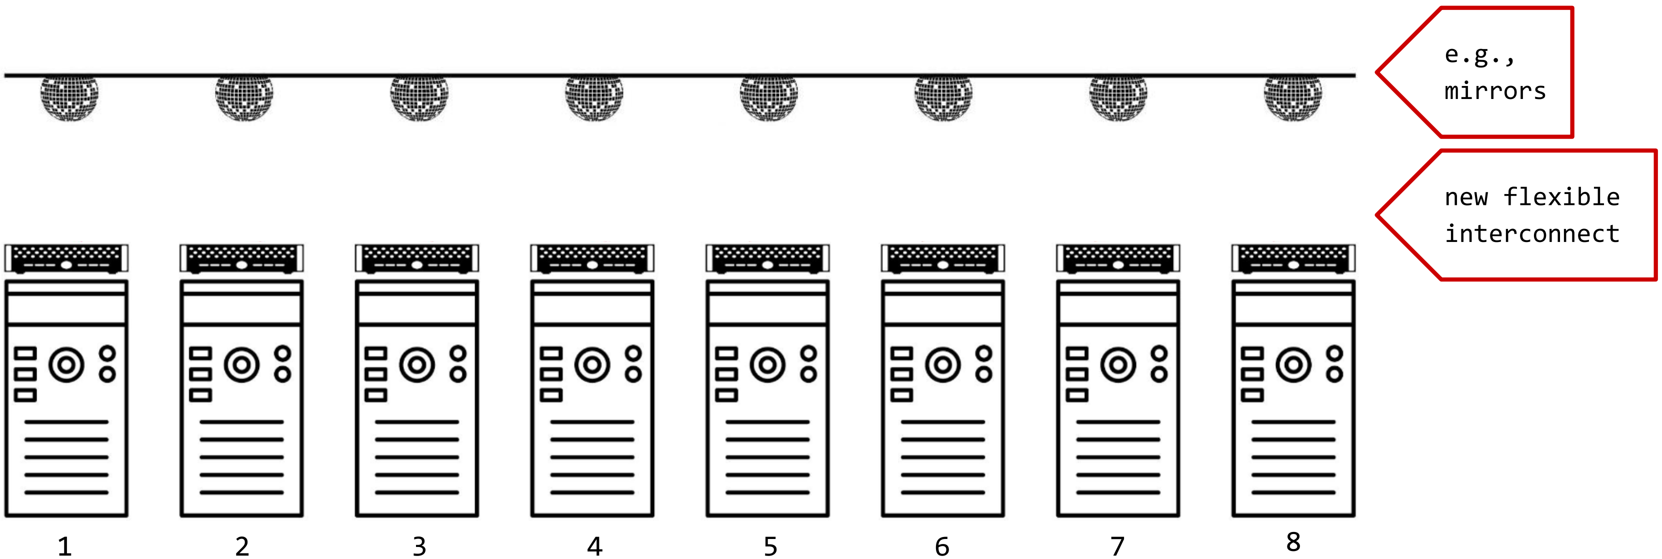

Example 2: Topology

Self-Adjusting Datacenters (Sigmetrics'23)

demand matrix:

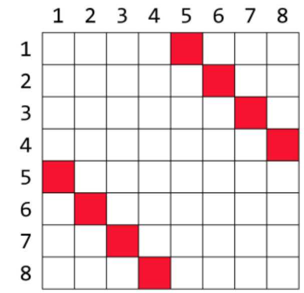

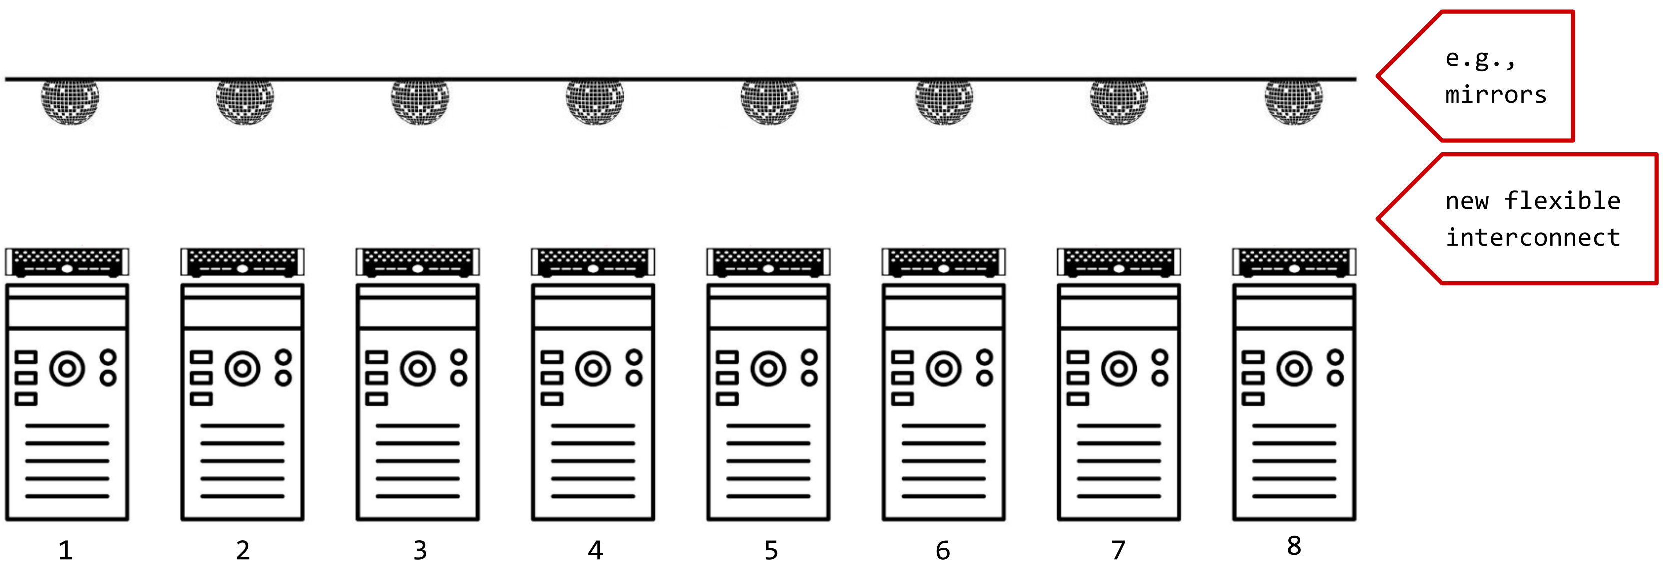

Example 2: Topology

Self-Adjusting Datacenters (Sigmetrics'23)

Matches demand

demand matrix:

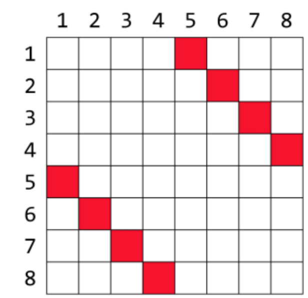

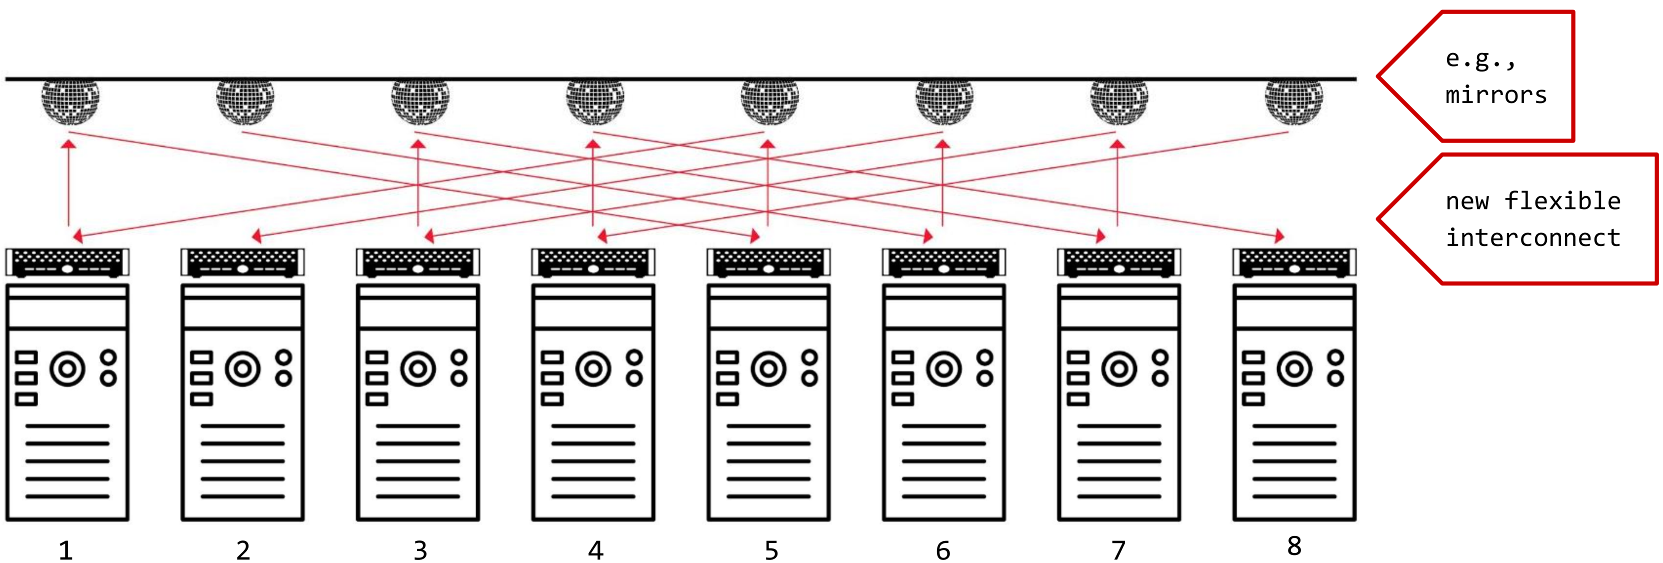

Example 2: Topology

Self-Adjusting Datacenters (Sigmetrics’23)

new demand:

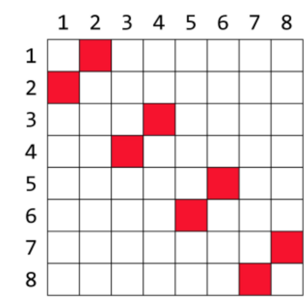

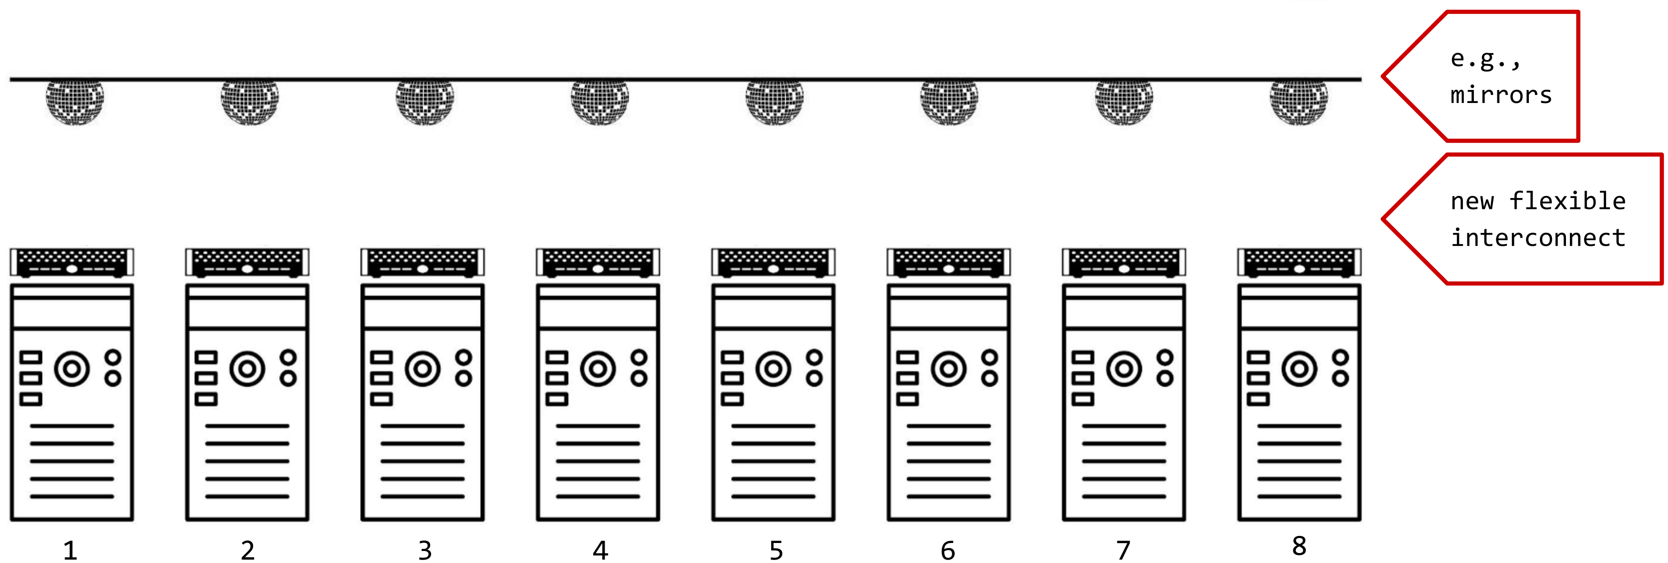

Example 2: Topology

Self-Adjusting Datacenters (Sigmetrics’23)

Matches demand

new demand:

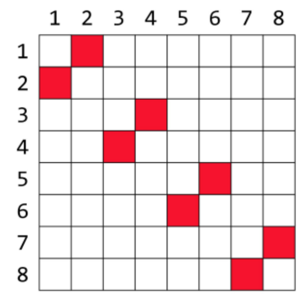

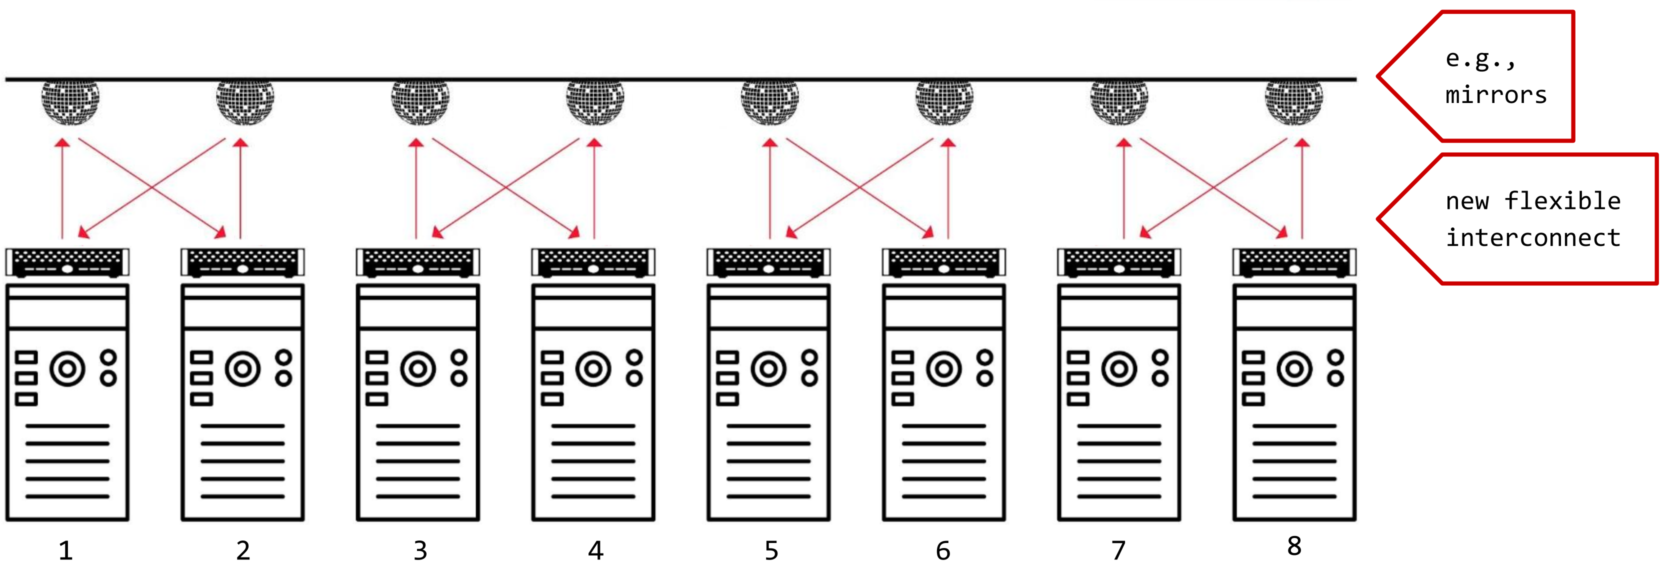

Example 2: Topology

Self-Adjusting Datacenters (Sigmetrics’23)

Self-Adjusting Networks

new demand:

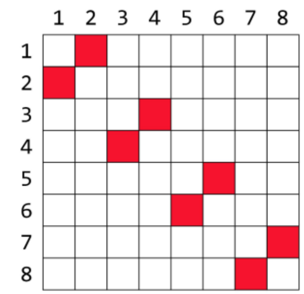

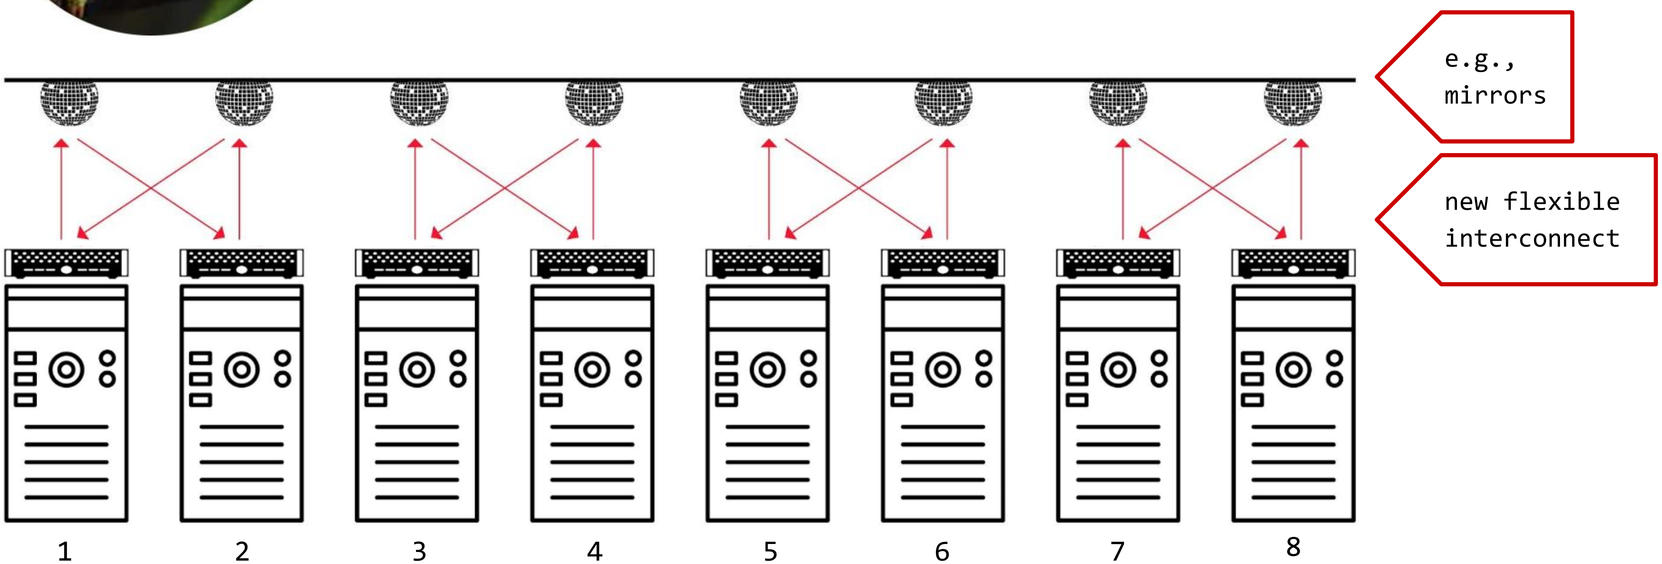

Example 2: Topology

Self-Adjusting Datacenters (Sigmetrics’23)

Self-Adjusting Networks

Tradeoff: benefits vs costs of adjustments

new demand:

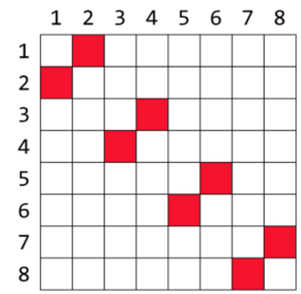

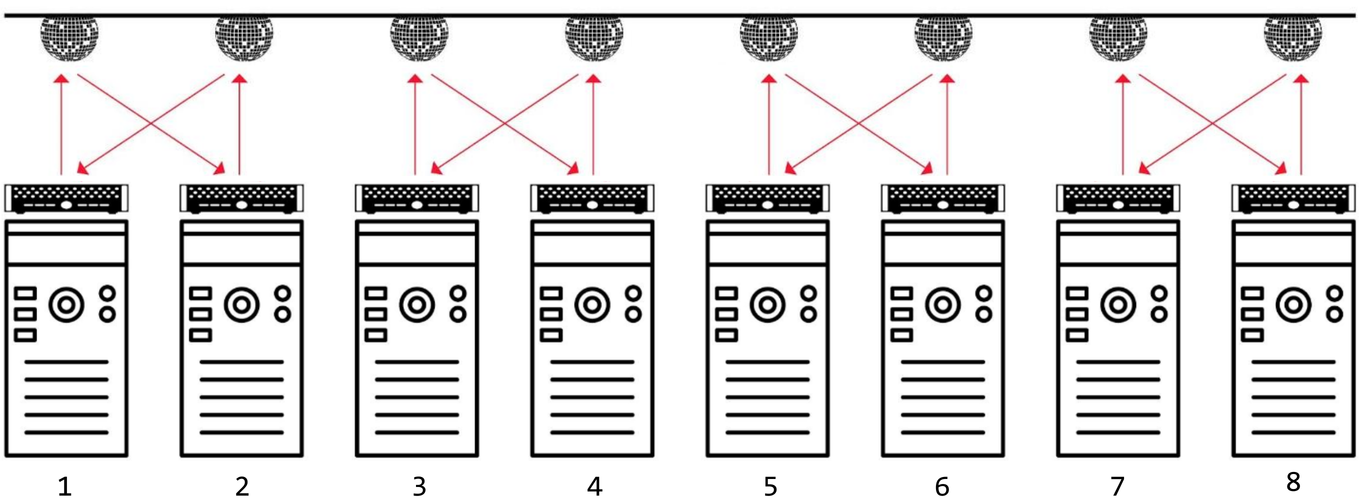

e.g., mirrors

new flexible interconnect

Underlying Technology

Optical Circuit Switch

- ⇢ Optical Circuit Switch rapid adaption of physical layer
  - ⇀ Based on rotating mirrors

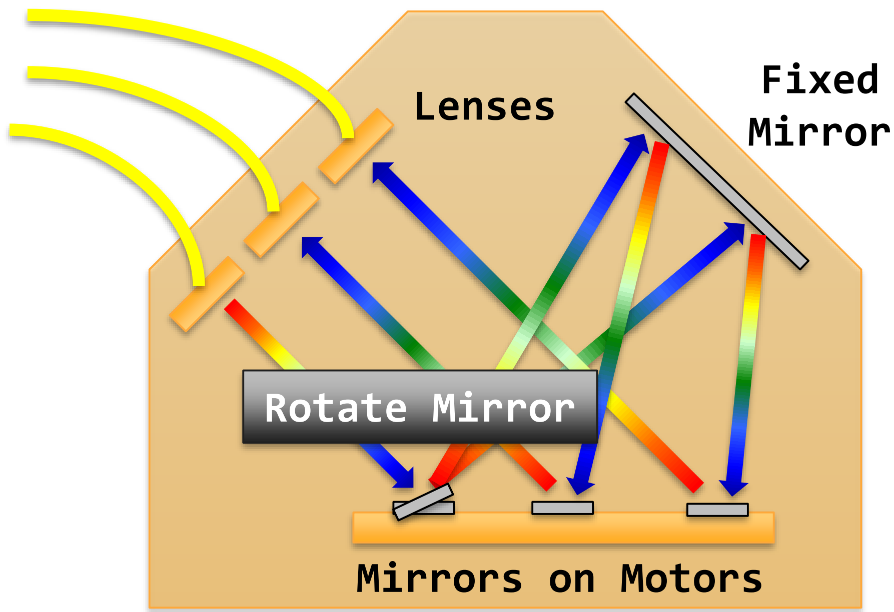

Optical Circuit Switch

By Nathan Farrington, SIGCOMM 2010

Competitive Ratio

Metric for Evaluating Self-Adjusting Systems

$$\rho = \max_{\sigma} \text{Cost\_ON}(\sigma)/\text{Cost\_OFF}(\sigma)$$

Competitive Ratio

Metric for Evaluating Self-Adjusting Systems

$\rho = \max_{\sigma}$ Cost_ON(σ)/Cost_OFF(σ)

Competitive Ratio

Metric for Evaluating Self-Adjusting Systems

$$\rho = \cancel{\max}_\sigma \ \text{Cost\_ON}(\sigma)/\text{Cost\_OFF}(\sigma)$$

Too conservative? Demand often not "worst case".

Too Conservative?

Much Structure in the Demand

Empirical studies:

traffic matrices sparse and skewed

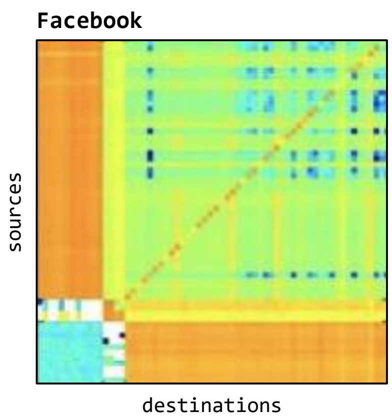

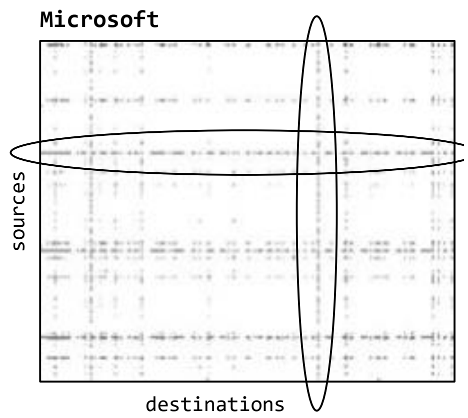

traffic bursty over time

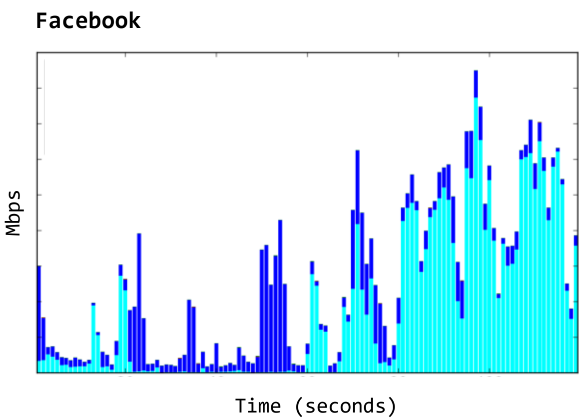

Too Conservative?

Much Structure in the Demand

Empirical studies:

traffic matrices sparse and skewed

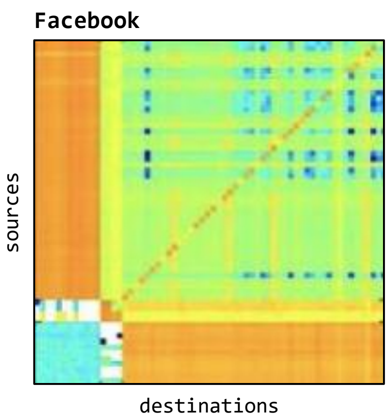

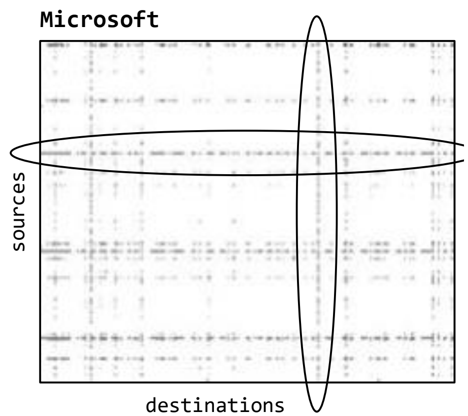

traffic bursty over time

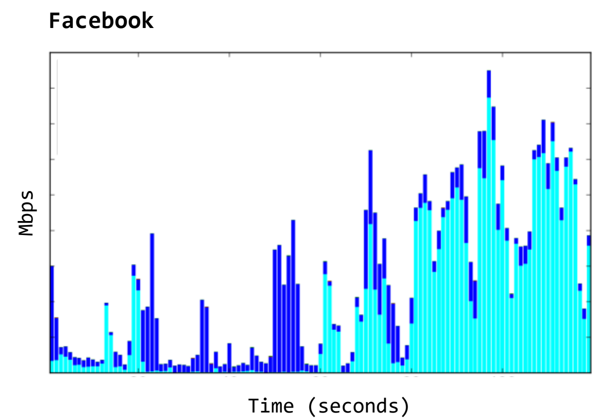

The hypothesis: can be learned/predicted/exploited.

Too Conservative?

Much Structure in the Demand

Empirical studies:

traffic matrices sparse and skewed

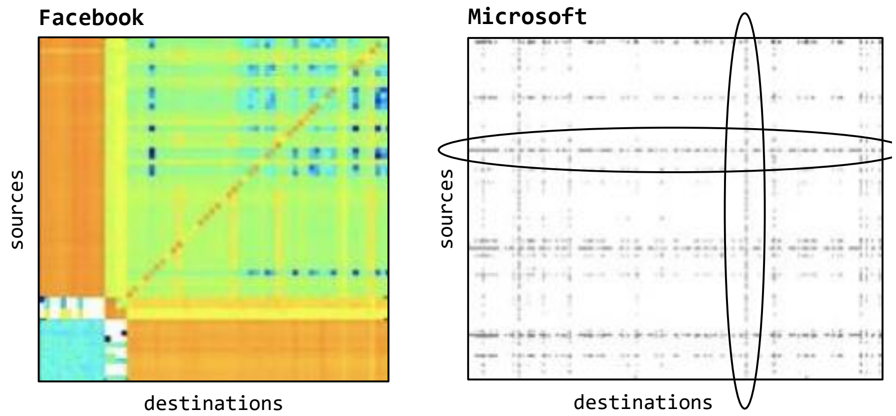

traffic bursty over time

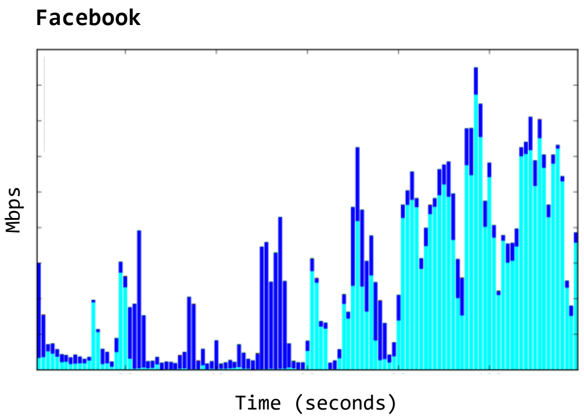

The hypothesis: can be learned/predicted/exploited.

Check out trace complexity website!

Recent Representation of Trace Structure:

Complexity Map (Sigmetrics'20)

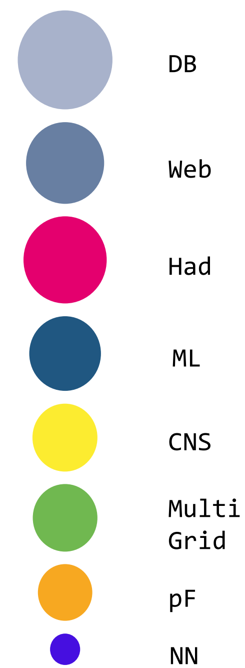

Griner et al., SIGMETRICS 2020

Recent Representation of Trace Structure:

Complexity Map (Sigmetrics’20)

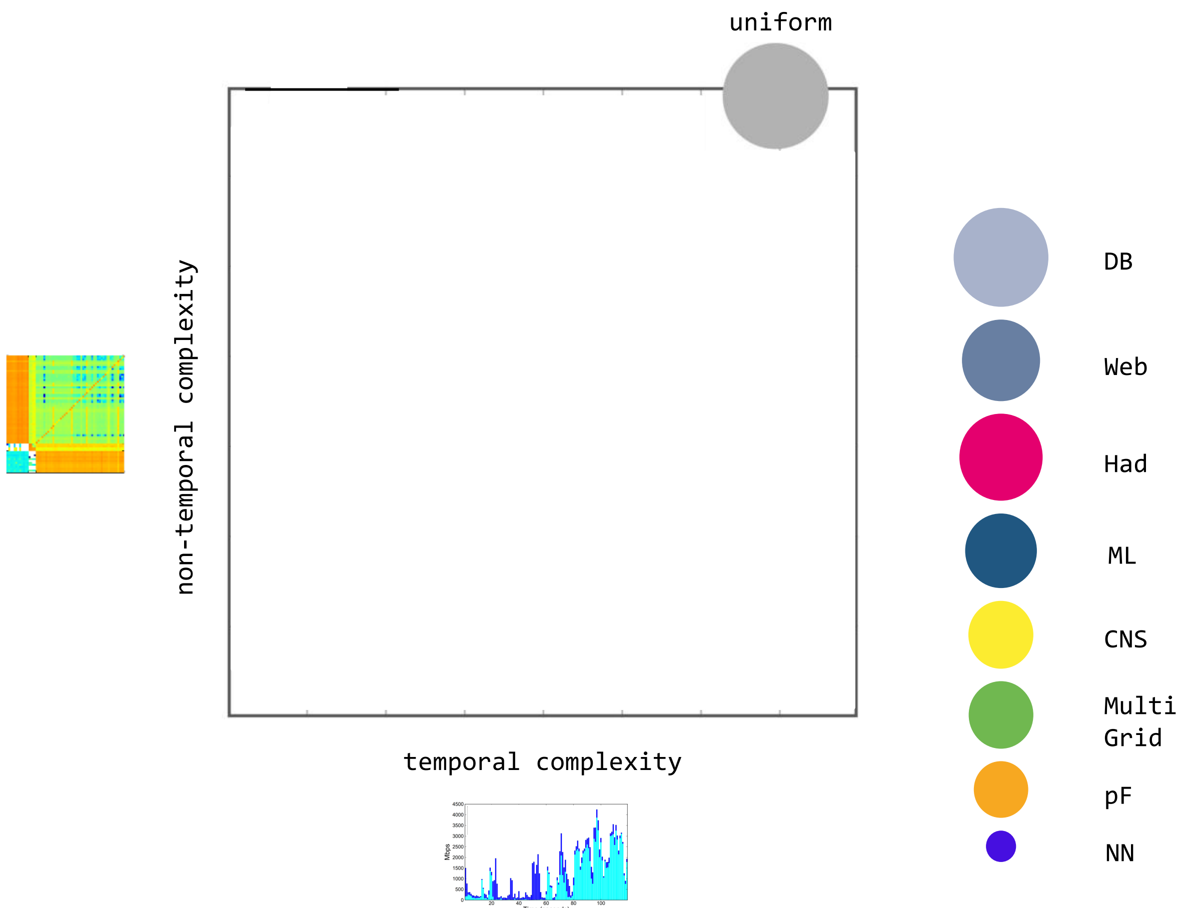

Griner et al., SIGMETRICS 2020

Recent Representation of Trace Structure:

Complexity Map (Sigmetrics’20)

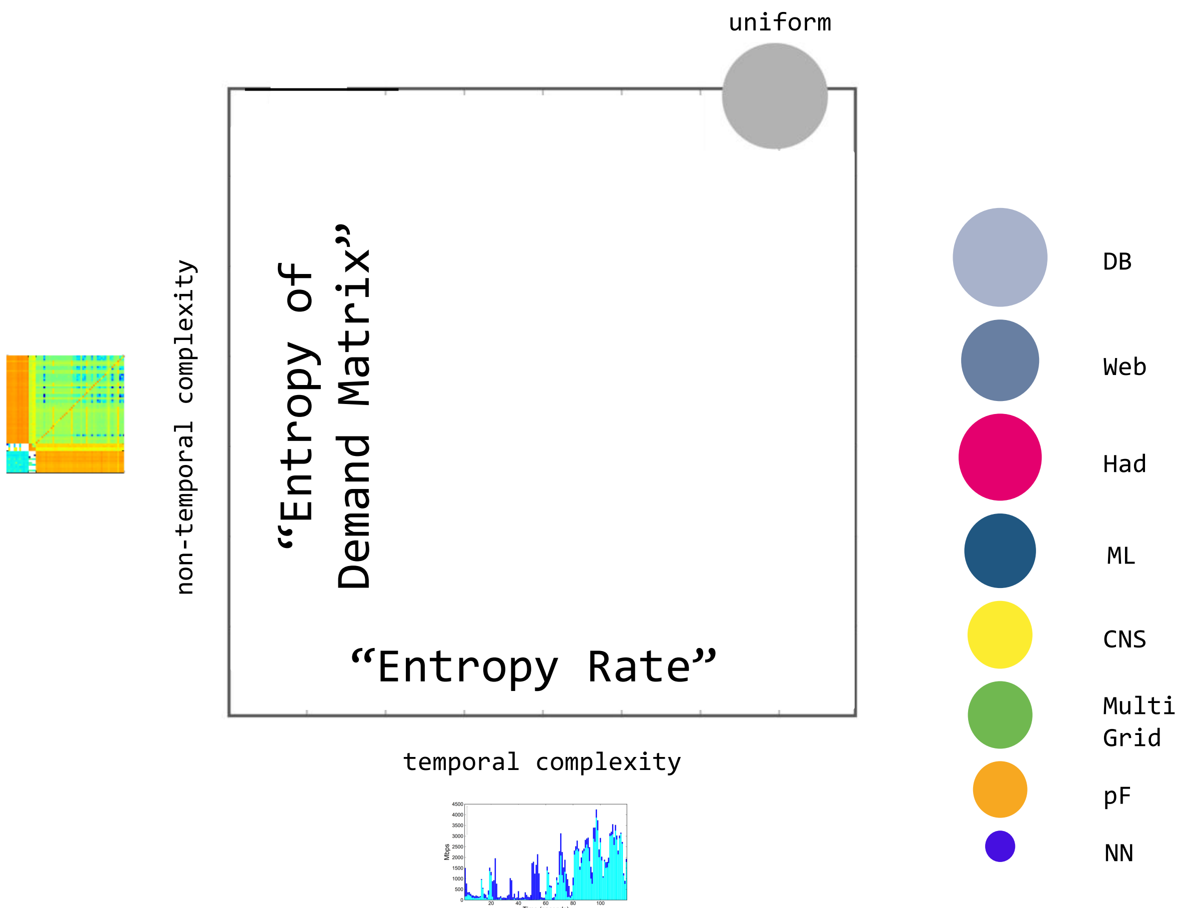

Griner et al., SIGMETRICS 2020

Recent Representation of Trace Structure:

Complexity Map (Sigmetrics'20)

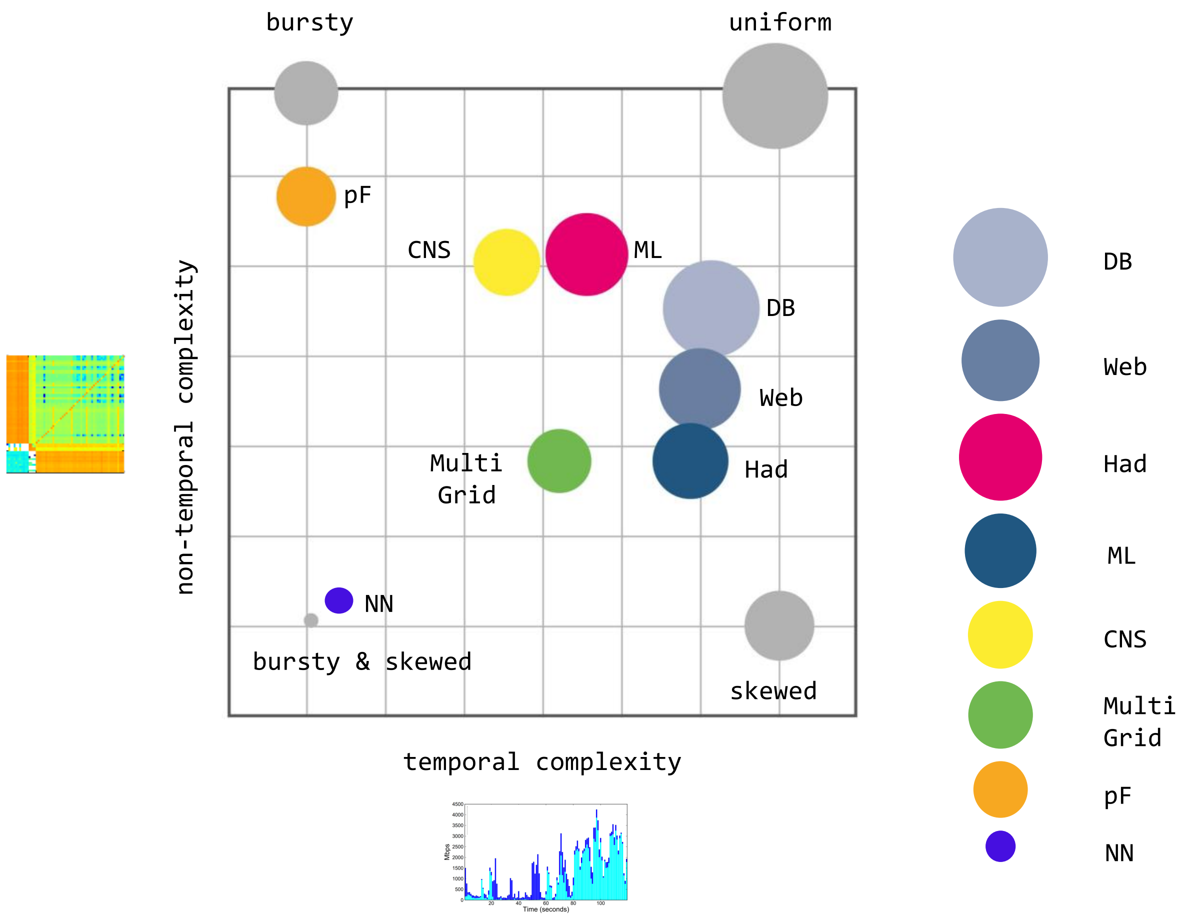

Different structures!

Griner et al., SIGMETRICS 2020

Traffic is also clustered (WWW'23):

Small Stable Clusters

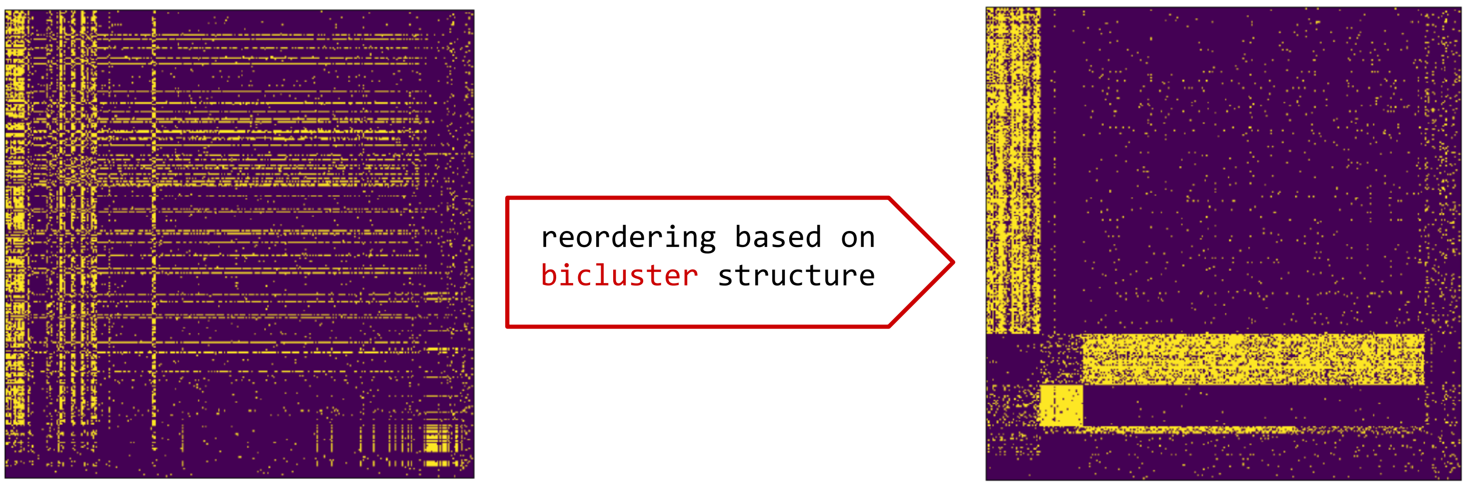

Opportunity: *exploit* with little reconfigurations!

Förster et al., Analyzing the Communication Clusters in Datacenters. WWW 2023

Beyond Worst-Case

Many approaches, also for online algorithms!

- **Restricting inputs:** random arrival order, locality of reference, access graph, smoothed analysis, independent sampling, diffused adversaries, distributional analysis, …
- **Deviating from competitive analysis:** resource augmentation, loose competitiveness, and competitiveness with high probability
- **Advice:** Next slide ☺

Advice Models

Advice from an oracle: ***side-loaded information*** about the future helps online algorithms to make better decisions.

Model 1: perfect advice

- Assumes a powerful, fully ***trustworthy oracle***
- Provides algorithm with any information about the future
- Question ***how many bits*** of the advice an online algorithm needs to achieve a certain competitive ratio c

Model 2: predictions (untrusted advice)

- Introduced by Mitzenmacher and Vassilvitskii: predictor may be faulty, and the competitive ratio depends on its error
- For small error, algorithm should perform close to the offline optimum (***consistency***), for large error, not worse than non-augmented online algorithms (***robustness***)

Model 3: lookahead

- Related to local algorithms of distributed computing*

Limitations

- All these models require new algorithmic features (e.g., a ***designated advice tape***)
- Hence not applicable to ***existing online algorithms***
- Bounds also depend on choice of error function

Infused Advice Model

An alternative

⇢ Idea: for randomized online algorithms, we may feed advice „***non-intrusively***“ via the random bits tape

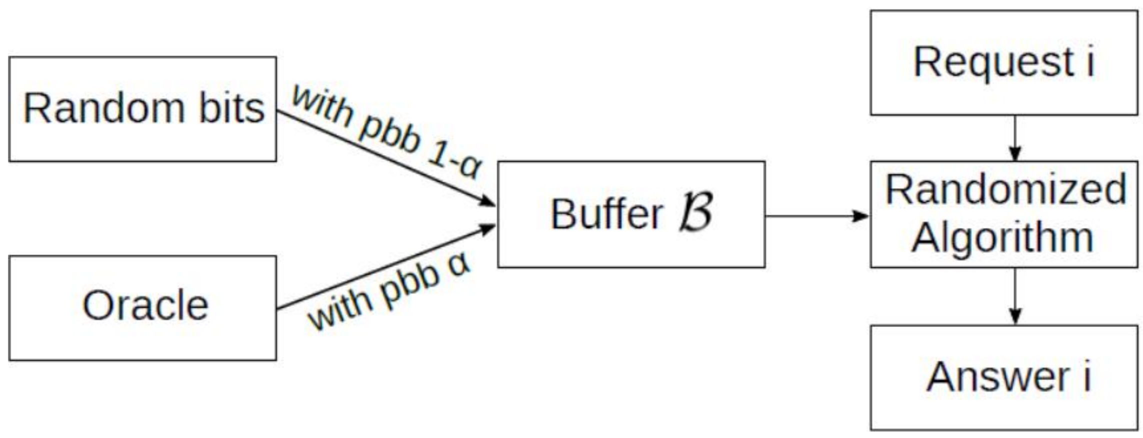

Infused Advice Model

An alternative

⇢ Idea: for randomized online algorithms, we may feed advice „***non-intrusively***“ via the random bits tape

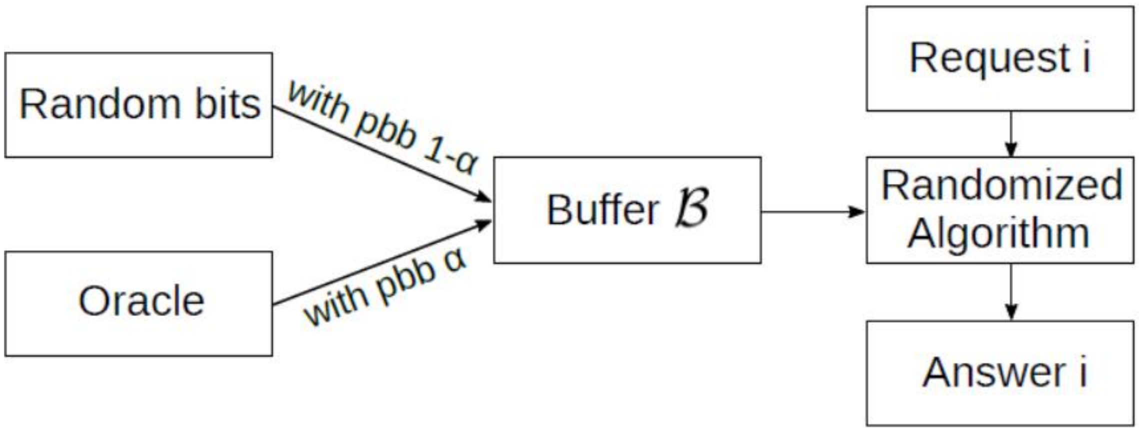

Works also for existing online algorithms

First Results (ESA'23)

- Applies to: ***paging***, ***uniform metrical task systems***, ***online set cover***, etc.
- New ***upper bounds*** which improve when the infusion parameter α increases
- Often tight lower bounds (assuming algorithm cannot access buffer of previous rounds)

Problem	upper bound	lower bound
Caching	$\min\{2H_k, \frac{2}{\alpha}\}$	$\min\{H_k, \frac{1}{\alpha}\}$
Uniform MTS	$\min\{2H_n, \frac{2}{\alpha}+2\}$	$\min\{H_{n-1}, \frac{1}{\alpha}\}$
Online set cover	$O(\min\{\log d \log n, \frac{\log n}{\alpha}\})$	$\min\{\frac{1}{2}\log d, \frac{1}{2\alpha}\}$

First Results (ESA'23)

- Applies to: ***paging***, ***uniform metrical task systems***, ***online set cover***, etc.
- New ***upper bounds*** which improve when the infusion parameter α increases
- Often tight lower bounds (assuming algorithm cannot access buffer of previous rounds)

Problem	upper bound	lower bound
Caching	$\min\{2H_k, \frac{2}{\alpha}\}$	$\min\{H_k, \frac{1}{\alpha}\}$
Uniform MTS	$\min\{2H_n, \frac{2}{\alpha}+2\}$	$\min\{H_{n-1}, \frac{1}{\alpha}\}$
Online set cover	$O(\min\{\log d \log n, \frac{\log n}{\alpha}\})$	$\min\{\frac{1}{2}\log d, \frac{1}{2\alpha}\}$

Application domains for augmentation?

Emerging Applications

Example: Programmable Network Switches

- Classic ***buffer sharing*** algorithm: which packets to accept/drop?
- Typically ***no pushout***: once packet admitted to buffer, cannot be dropped later
- Future arrivals unknown: how to ***maximize throughput***?
- Programmable switches enable new applications: running ML models ***in the data plane***

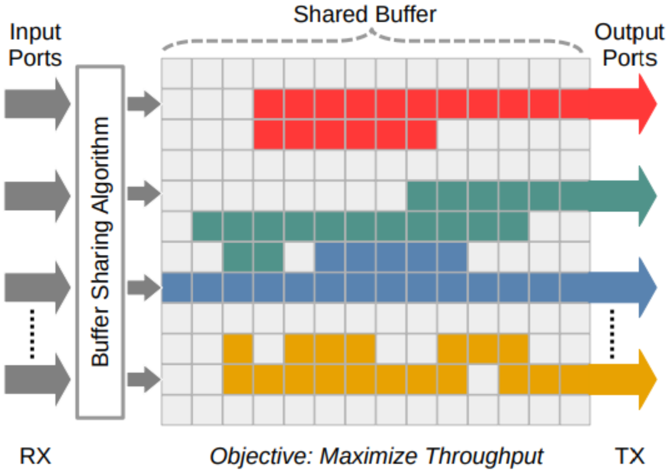

Emerging Applications

Example: Programmable Network Switches

⇢ Challenge: ***drop tail algorithms*** may drop red packets in order to keep buffer for other ports (and transmit in parallel)

⇢ ***Not competitive***: if not differently colored packets arrive

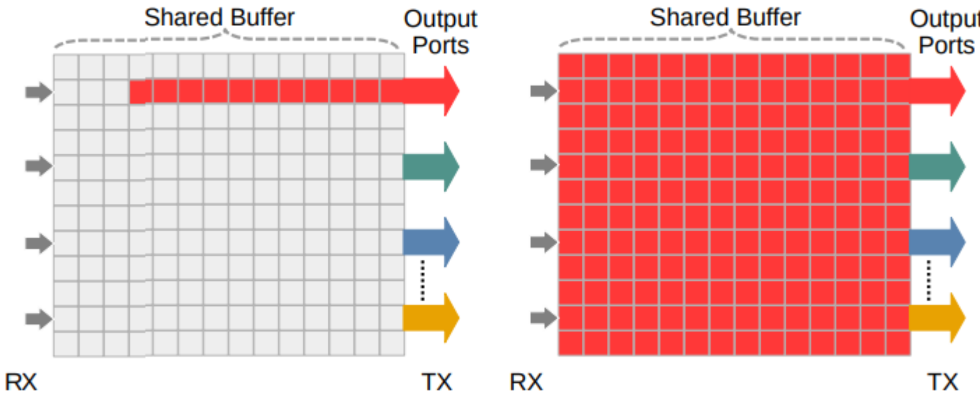

(a) ALG (b) OPT

Emerging Applications

Example: Programmable Network Switches

⇢ Alternatively: if drop tail algorithm ***absorbs burst***, it may lead to reactive drops in future and low throughput

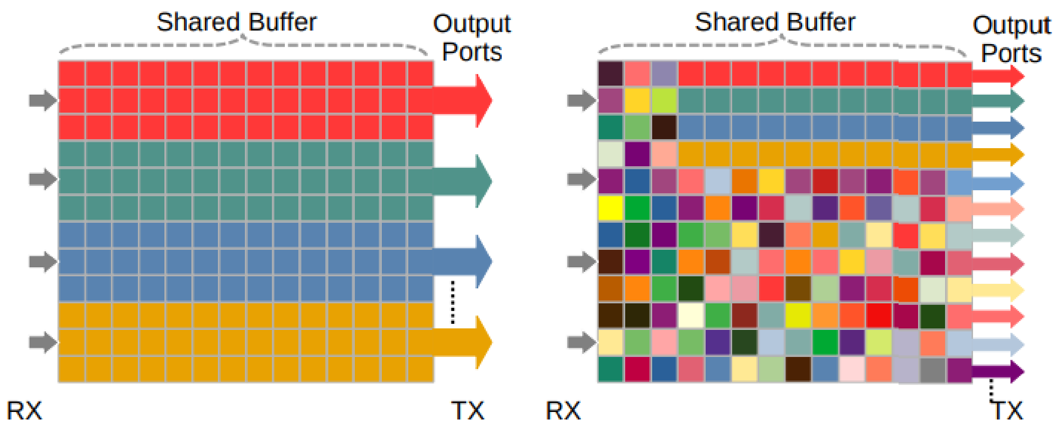

(a) ALG (b) OPT

Emerging Applications

Example: Programmable Network Switches

- Augment switches with predictions
- Simple ***random forest*** approach significantly improves competitive ratio
- Depending on prediction error
- Can be implemented on switch ***hardware***…

Algorithm	Competitive Ratio
Complete Sharing [26]	$N+1$
Dynamic Thresholds [20, 26]	$O(N)$
Harmonic [33]	$\ln(N)+2$
LQD (push-out) [9, 26]	1.707
LateQD (clairvoyant) [14]	1
CREDENCE	$\min(1.707\,\eta, N)$

Emerging Applications

Example: Programmable Network Switches

- ⇢ Further reading: NSDI’24
- ⇢ Video: https://www.youtube.com/watch?v=sAPe78RFsz0

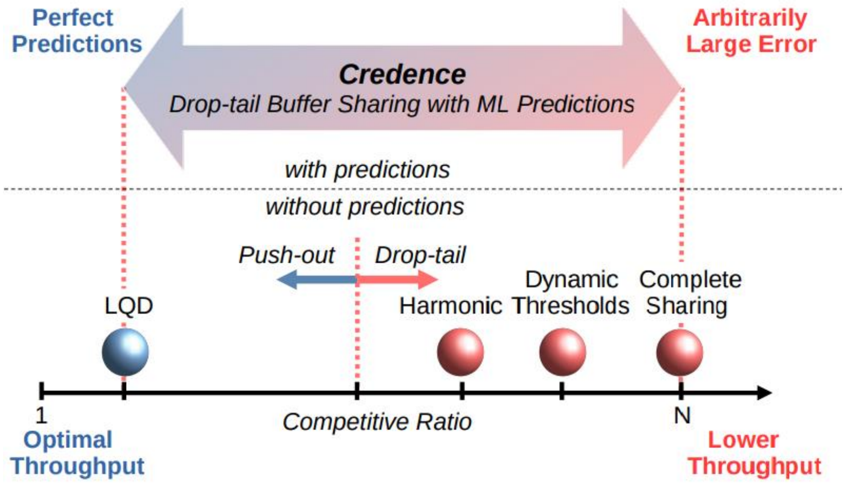

Conclusion

- Augmenting online algorithms with learning can benefit competitive ratio ***in theory but also in practice*** as input is skewed and predictable
- Network equipment supports such augmentations to some extent
- ***Infused advice*** approach allows to augment existing algorithms
- Interesting use cases in networked systems, much to explore

Data Available

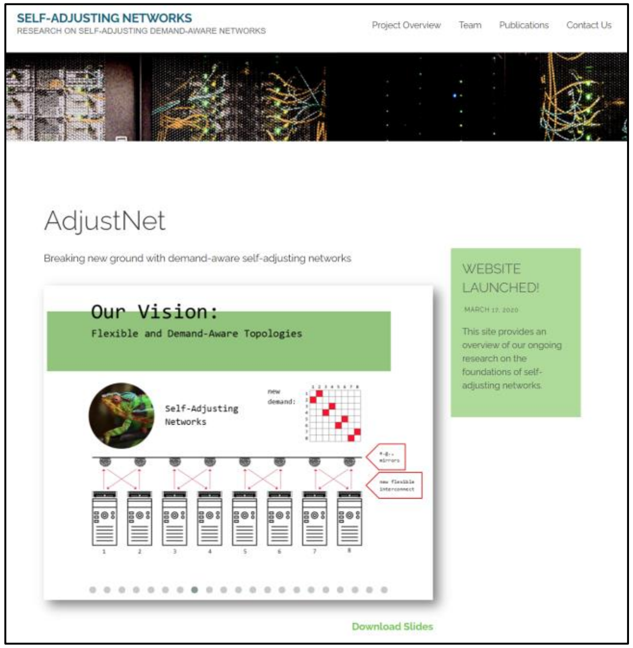

http://self-adjusting.net/
Project website

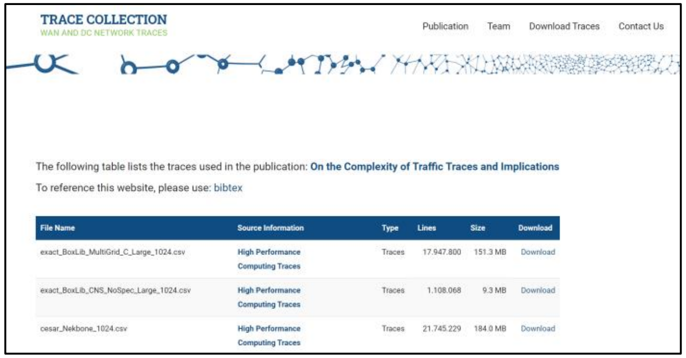

https://trace-collection.net/
Trace collection website

Selected References

On the Complexity of Traffic Traces and Implications
Chen Avin, Manya Ghobadi, Chen Griner, and Stefan Schmid.
ACM SIGMETRICS, Boston, Massachusetts, USA, June 2020.

Survey of Reconfigurable Data Center Networks: Enablers, Algorithms, Complexity
Klaus-Tycho Foerster and Stefan Schmid.
SIGACT News, June 2019.

Analyzing the Communication Clusters in Datacenters
Klaus-Tycho Foerster, Thibault Marette, Stefan Neumann, Claudia Plant, Ylli Sadikaj, Stefan Schmid, and Yllka Velaj.
The Web Conference (WWW), Austin, Texas, USA, April 2023.

Tight Bounds for Online Graph Partitioning
Monika Henzinger, Stefan Neumann, Harald Räcke, and Stefan Schmid.
ACM-SIAM Symposium on Discrete Algorithms (SODA), Alexandria, Virginia, USA, January 2021.

Efficient Distributed Workload (Re-)Embedding
Monika Henzinger, Stefan Neumann, and Stefan Schmid.
ACM/IFIP SIGMETRICS/PERFORMANCE, Phoenix, Arizona, USA, June 2019, and ACM Performance Evaluation Review (PER).

Duo: A High-Throughput Reconfigurable Datacenter Network Using Local Routing and Control
Johannes Zerwas, Csaba Györgyi, Andreas Blenk, Stefan Schmid, and Chen Avin.
ACM SIGMETRICS and ACM Performance Evaluation Review (PER), Orlando, Florida, USA, June 2023.

Online Algorithms with Randomly Infused Advice
Yuval Emek, Yuval Gil, Maciej Pacut, and Stefan Schmid.
European Symposium on Algorithms (ESA), Amsterdam, Netherlands, September 2023.

Credence: Augmenting Datacenter Switch Buffer Sharing with ML Predictions
Vamsi Addanki, Maciej Pacut, and Stefan Schmid.
21st USENIX Symposium on Networked Systems Design and Implementation (NSDI), Santa Clara, California, USA, April 2024.

Thank you! Questions?

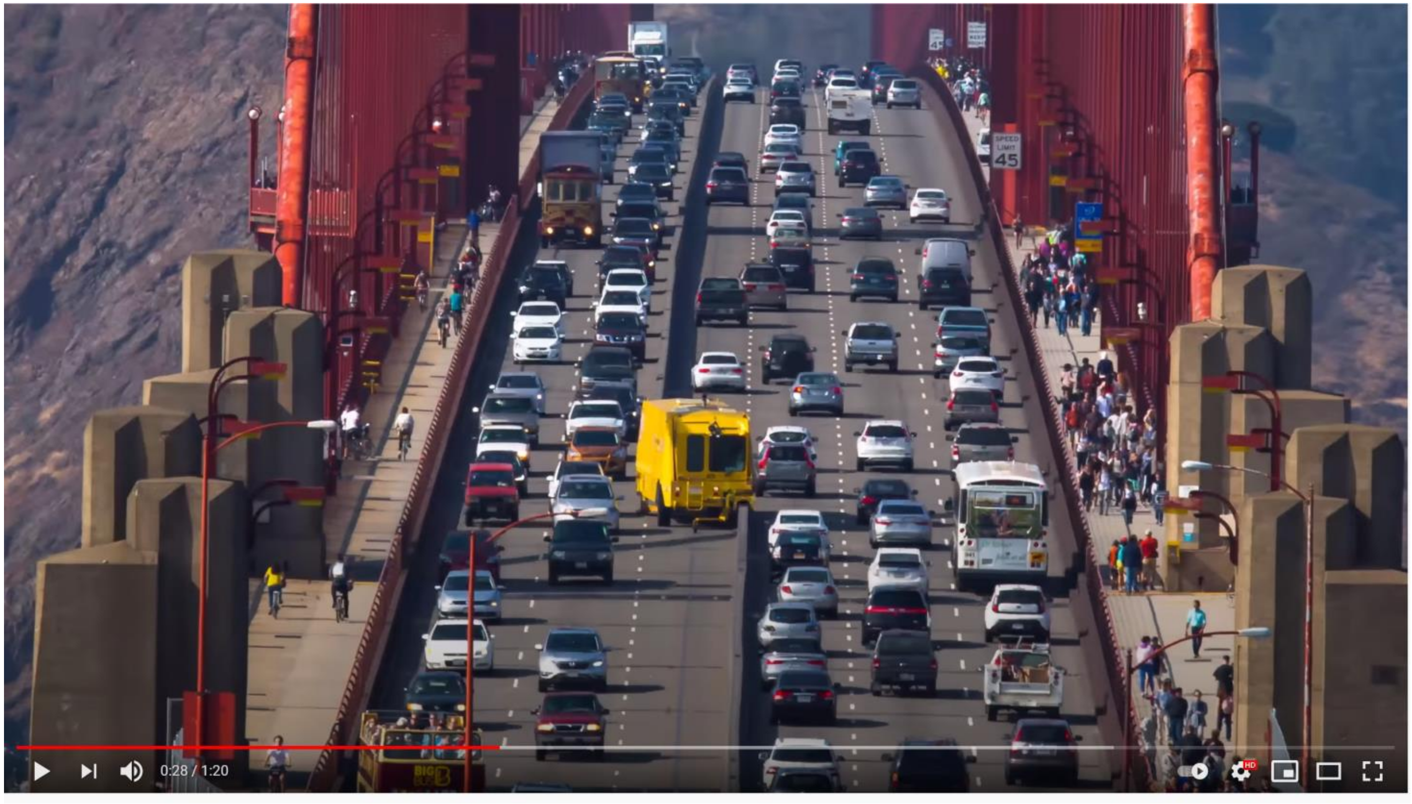

Golden Gate Zipper

Online Video Course

Bonus Material

Hogwarts Stair