

Self-Adjusting Networks

Stefan Schmid (TU Berlin)

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

(Folklore)


Datacenter Traffic

Empirical Observation: Structure


Empirical studies:

traffic matrices **sparse** and **skewed**

Facebook




Microsoft




traffic **bursty** over time

Facebook



Datacenter Traffic

Observation: Depends on Application



Different structures!

Datacenter Traffic


Challenge: Diverse Requirements

Diverse patterns:

- Shuffling/Hadoop:
all-to-all
- All-reduce/ML: **ring** or **tree** traffic patterns
 - **Elephant** flows
- Query traffic: skewed
 - **Mice** flows
- Control traffic: does not evolve but has non-temporal structure

Diverse requirements:

- ML is **bandwidth** hungry, small flows are **latency**-sensitive




The SOTA Datacenter

Fixed and Demand-Oblivious Topology



Highway which ignores
actual traffic:
frustrating!




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

Alternative?

Emerging Reconfigurable Optical Datacenters

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



Optical Circuit Switch

By Nathan Farrington, SIGCOMM 2010

First Deployments

E.g., Google

Systems

Jupiter evolving: Reflecting on Google's data center network transformation

August 24, 2022

[Twitter](#) [LinkedIn](#) [Facebook](#) [Email](#)

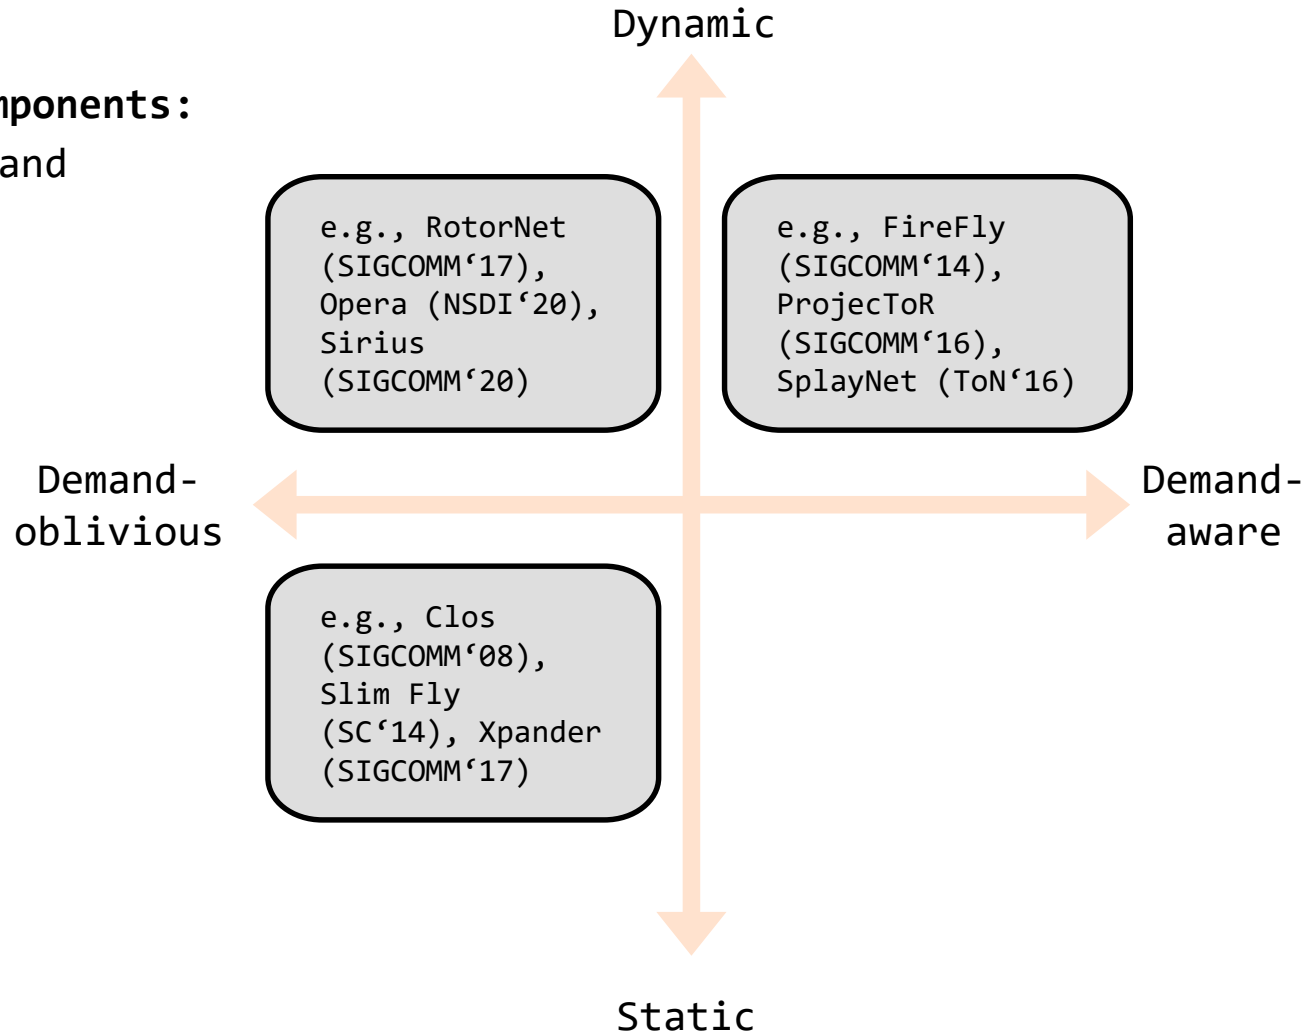


Amin Vahdat
VP & GM, Systems and Services Infrastructure

Opportunity: Tech Diversity

Diverse topology components:


- demand-**oblivious** and demand-**aware**
- static vs dynamic



Opportunity: Tech Diversity

Diverse topology components:

- demand-**oblivious** and demand-**aware**
- static vs dynamic



Which approach
is best?

Challenges

- Many challenges of reconfigurable networks
 - Shock wave through *Layers*:
 - impact on *routing* and congestion control?
 - *Scalability* of control in dynamic graphs:
 - Local algorithms*? Greedy routing?
- Complexity of demand-aware graphs
 - (*pure vs hybrid*, e.g., SplayNet)
 - *Application-specific* self-adjusting networks:
 - e.g., for AI, or similar to *active dynamic networks* (independent sets, consensus, ...)
 - etc.

Thank you!