

The Controller Placement Problem

Brandon Heller

Rob Sherwood, Nick McKeown

HotSDN 2012

Helsinki, Finland

contact: brandonh@stanford.edu

from HotSDN 2012 CFP:

SDN, MPLS PCE, BGP RR, 4D, RCP, ...

~~“Software Defined Networking (SDN)–~~

refactors the relationship between network devices and the software that controls them.”

SDN is not alone.


(decoupled from)

Control Software ········· Network Devices

Each architecture faces a design choice in the WAN that fully distributed networks don't have...

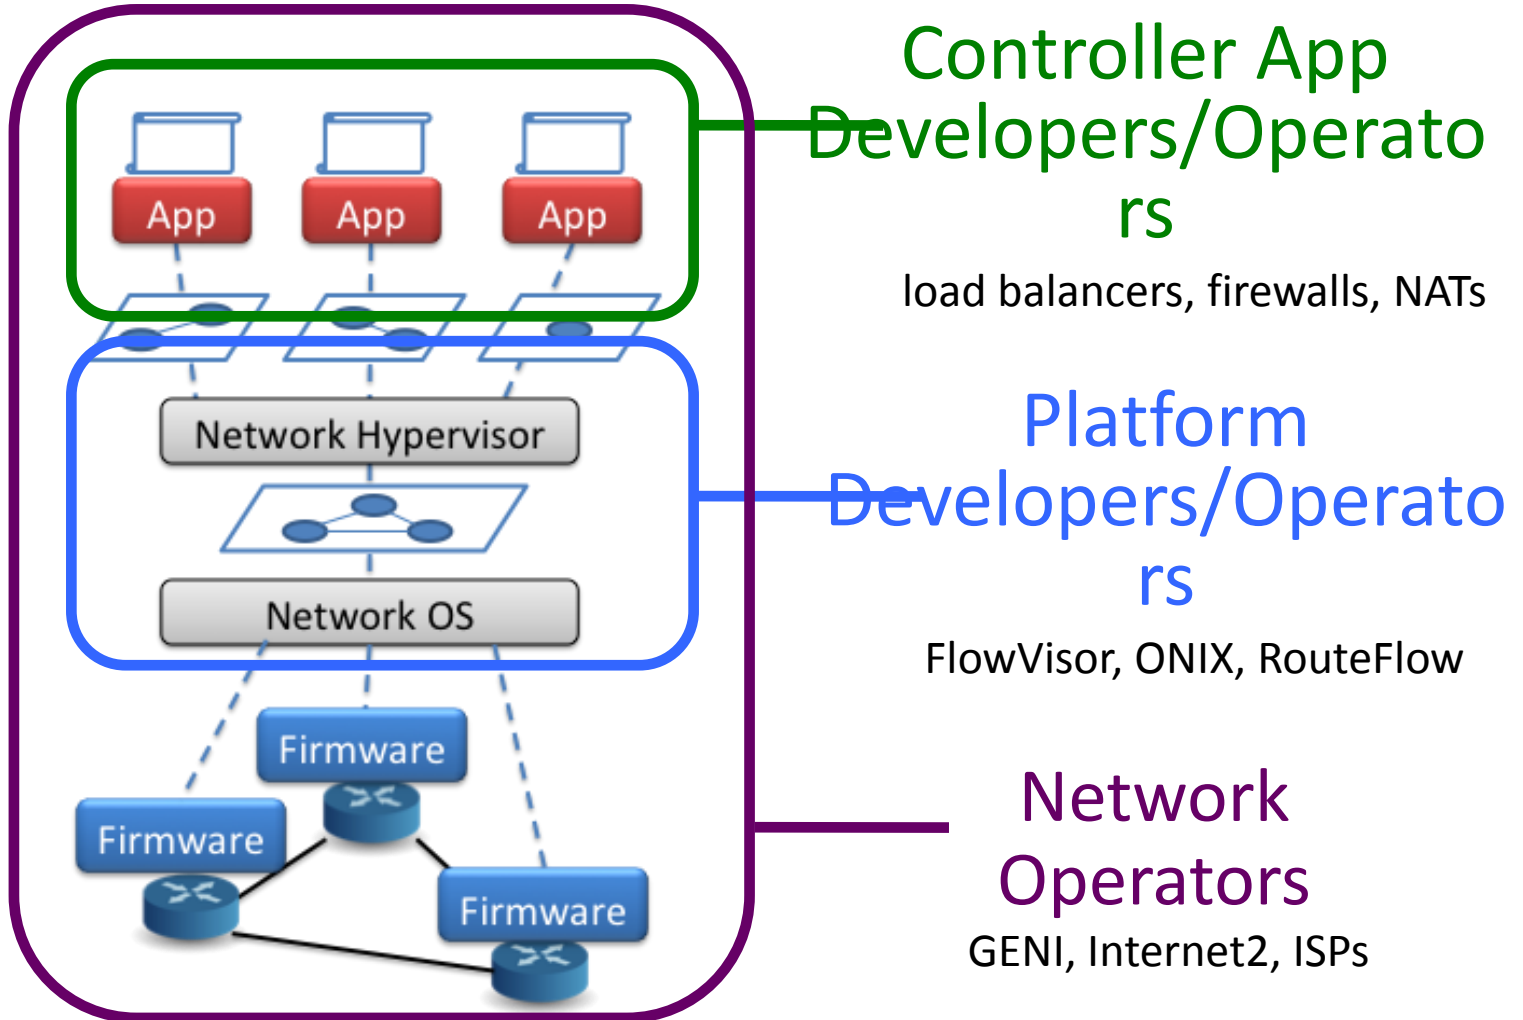
The Controller Placement Problem

- (1) How many controllers?
- (2) Where should they go?



[Internet2 OS³E]

Who Cares?




Optimal placement depends on goals

fault tolerance, load balancing, latency, ...

Latency “Methodology”

WAN topologies




[www.topologyzoo.org]

brute-force metric
computation

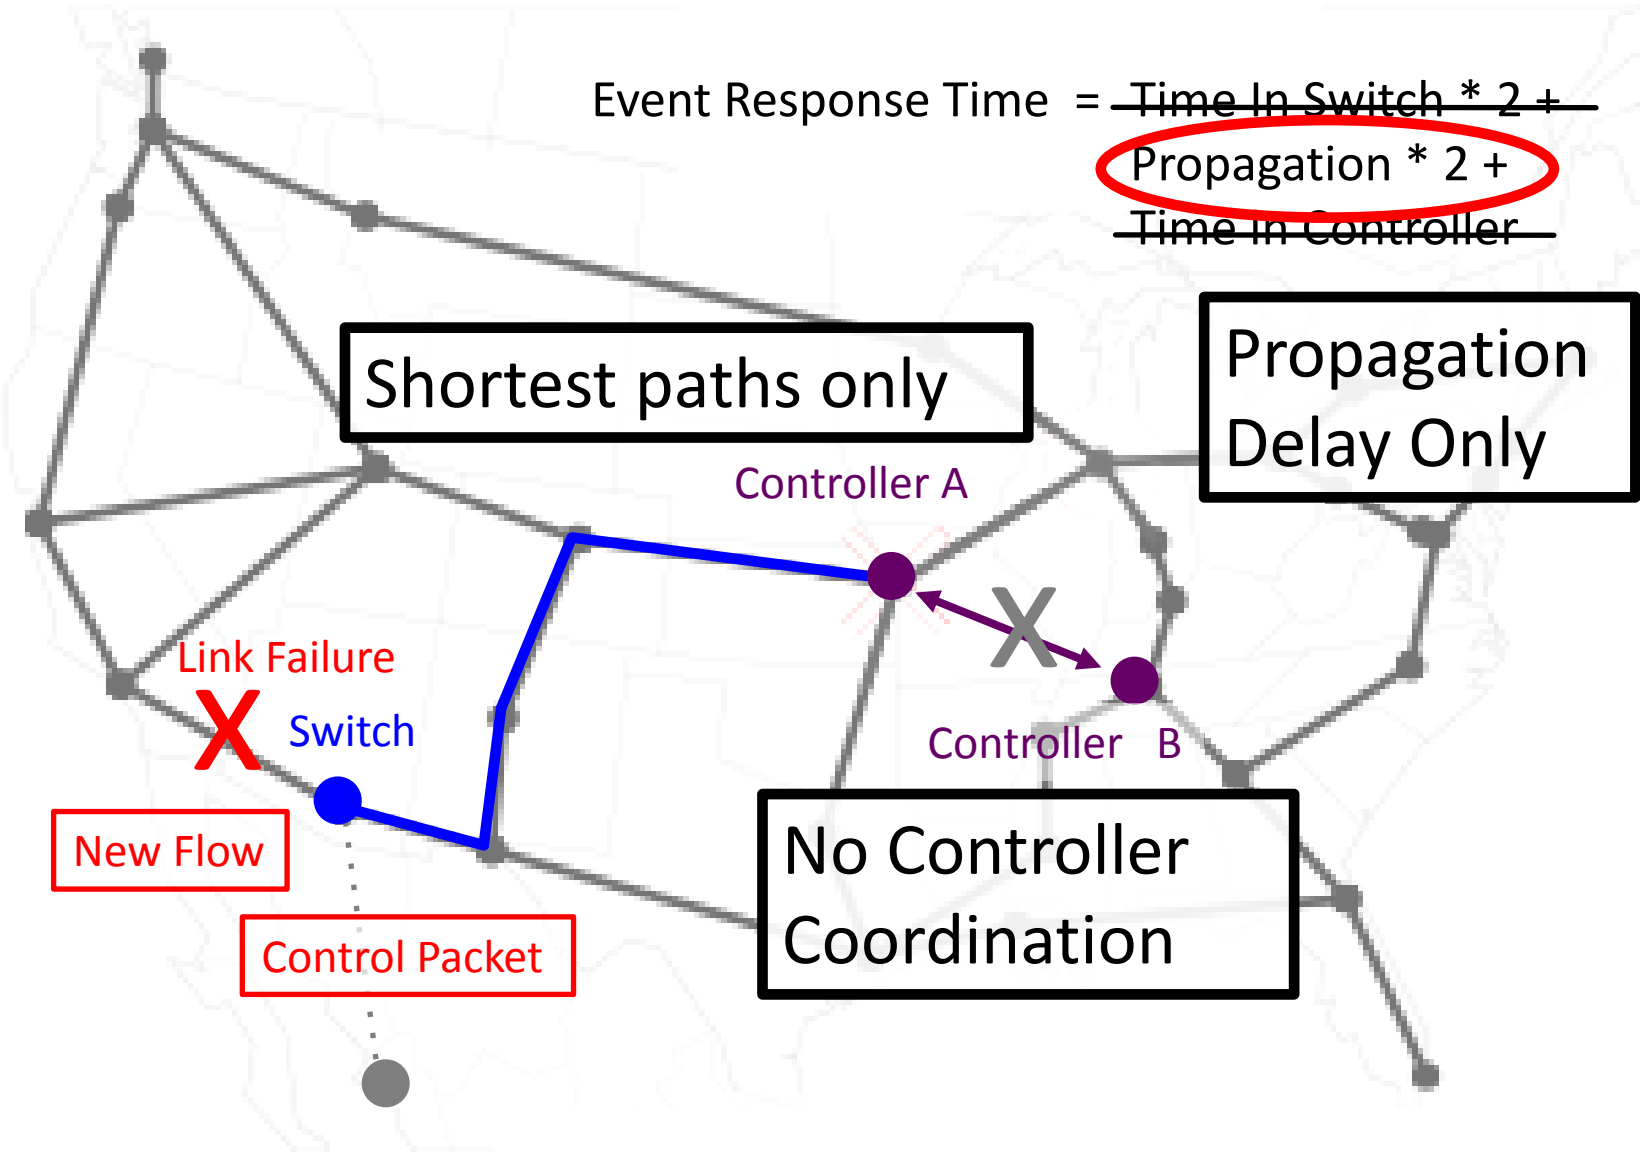


optimal
controller
placements



Response Time

$$\text{Event Response Time} = \cancel{\text{Time In Switch} * 2} + \text{Propagation} * 2 + \cancel{\text{Time in Controller}}$$



Shortest paths only

Propagation Delay Only

Controller A

Controller B

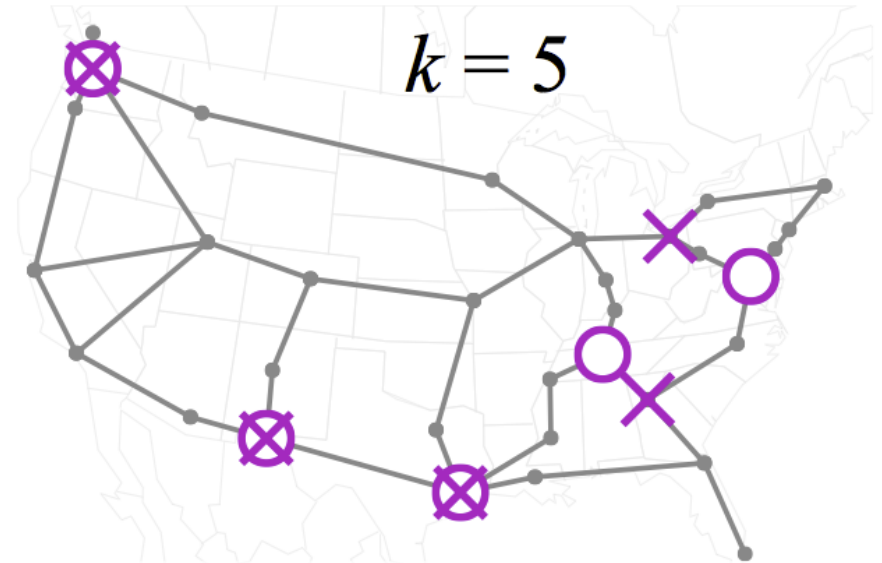
No Controller Coordination

Link Failure
X
Switch

New Flow

Control Packet

Placement Examples



○ location in average-latency-optimized placement

✕ location in worst-case-latency-optimized placement

Topics

- Is one location enough?
- What is a good placement?
- How do extra controllers reduce latency?
- Are random placements good?
- How to efficiently choose placements?
- Tradeoffs between metrics?
- How to combine metrics?

Rest of this talk:

- Is one location enough?
- What is a good placement?
- How do extra controllers reduce latency?
- Are random placements good?
- How to efficiently choose placements?
- Tradeoffs between metrics?
- How to combine metrics?

See paper for grayed-out topics.

Is one location
enough?
[for latency!]

Delay Comparison

Round-trip Latency Target

Name	Delay
Through the switch (measured last year)	10 ms
Optical ring protection in SONET, Wireless handover in WiMAX	50 ms
IP-level restoration	200 ms

Delay Comparison

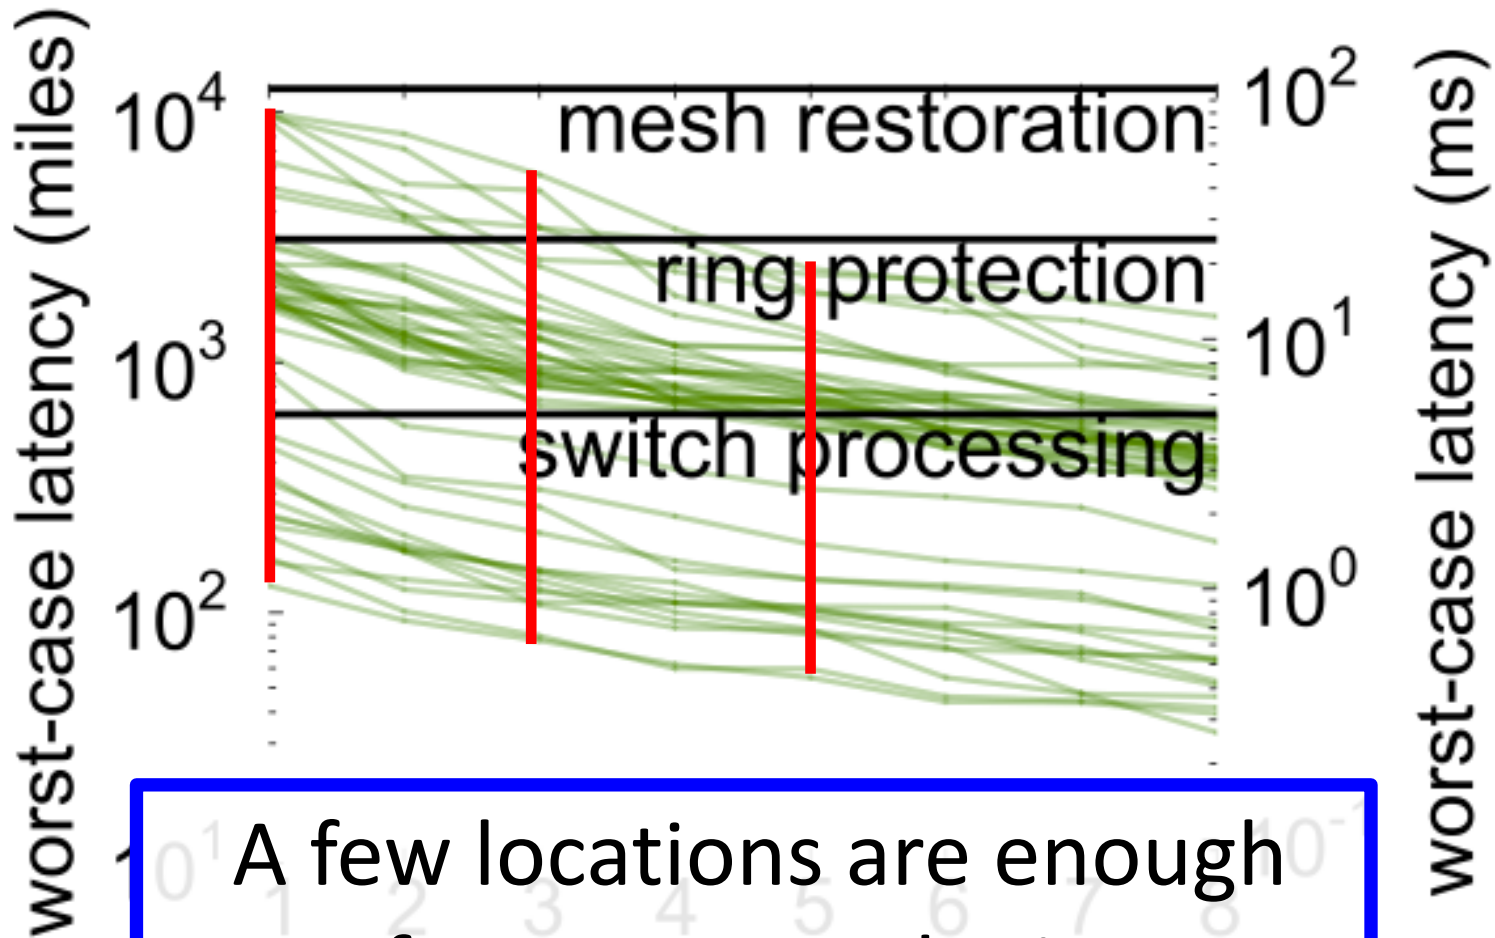
Round-trip Latency Target

Safety Margin

Name	Delay	1.0x	1.5x	2.0x
Through the switch (measured last year)	10 ms	27%	22%	18%
Optical ring protection in SONET, Wireless handover in WiMAX	50 ms	82%	60%	33%
IP-level restoration	200 ms	100%	91%	89%

One controller location may be enough,
for many (smaller) topologies

Opt worst-case latencies



A few locations are enough for most topologies

Wrapping Up

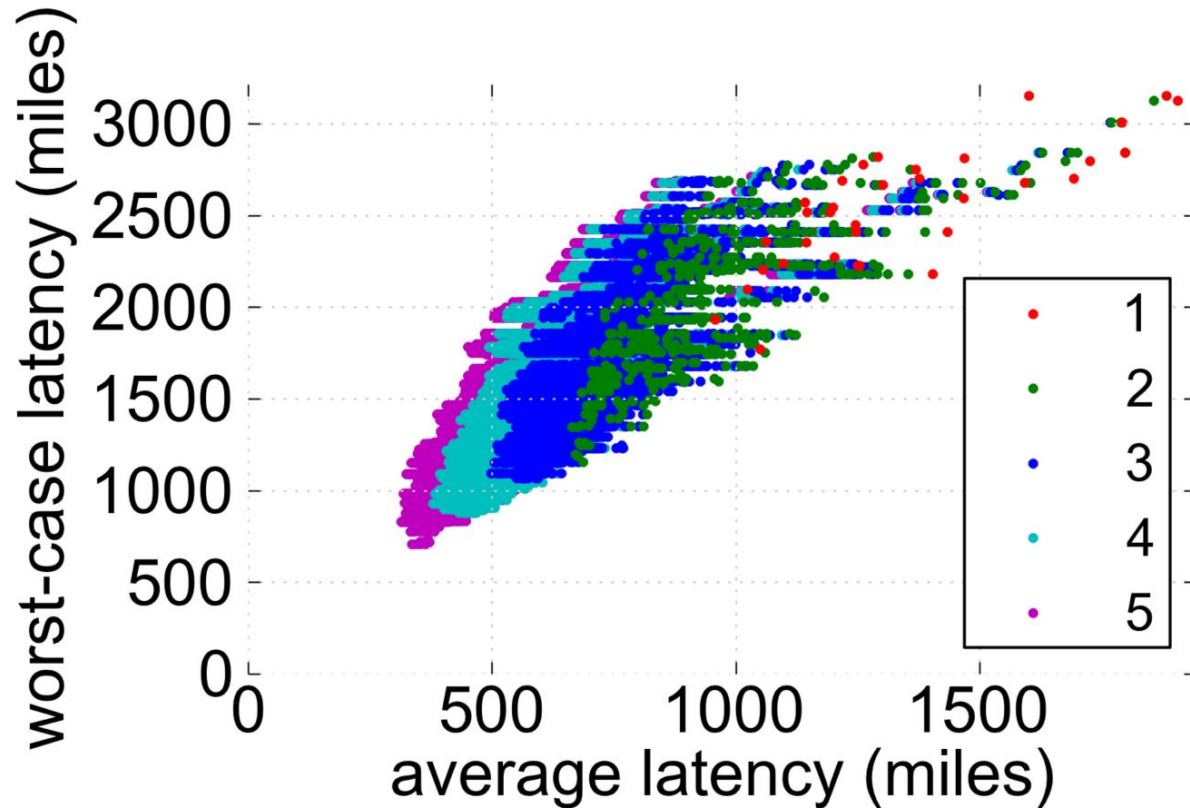
My takeaways

- One: probably good enough for many networks
- Few: probably good enough for most
- Placement likely to be dictated by other concerns
 - e.g. load balancing, fault tolerance

What's Next

- Fault Tolerance
- State Distribution
- Controller Selection

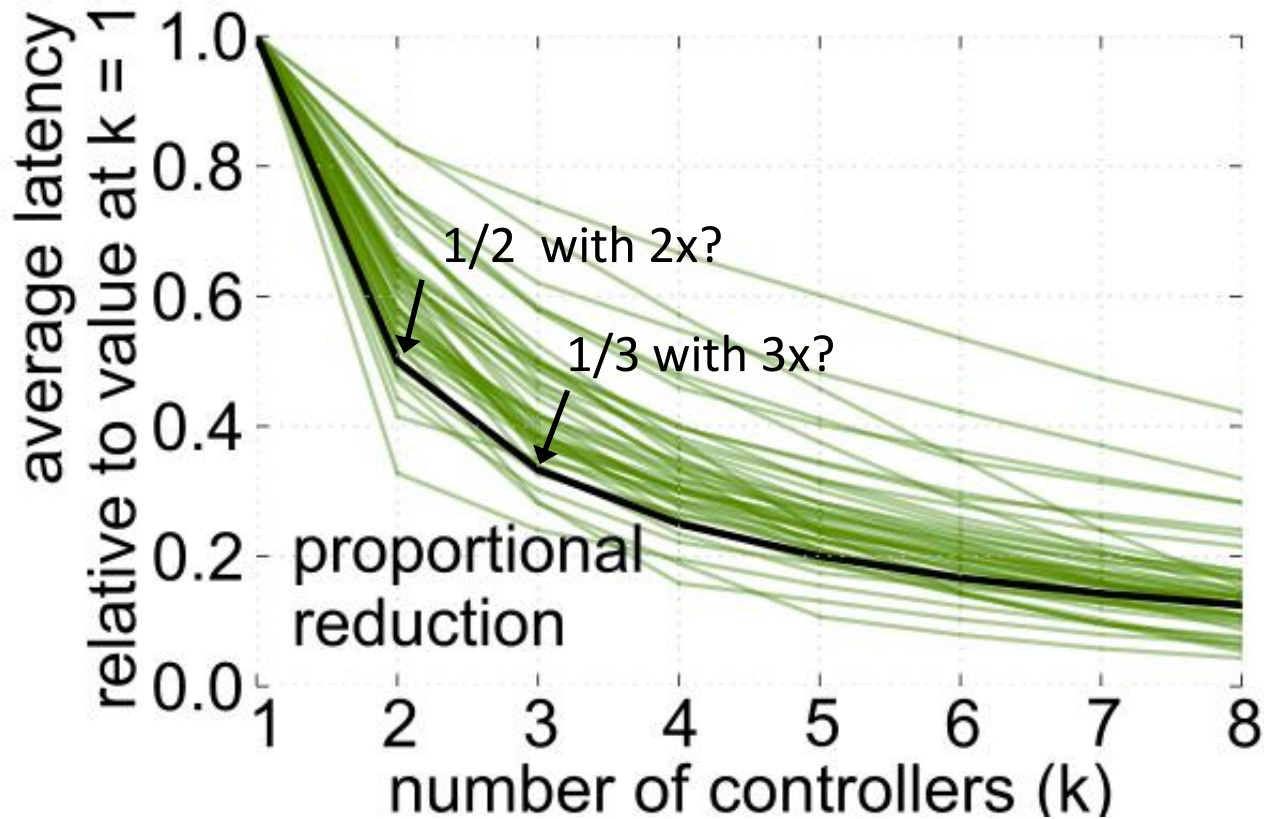
Thanks!



- Code at github.com/brandonheller/cpp.git

Backup Slides

Latency reduction w/more controllers



→ Mostly proportional reduction

Optimal-latency placements

metric: corresponding **Facility Location** problem

- Average-case latency: minimum k-median
- Worst-case latency: minimum k-center
- Nodes within a latency bound: maximum cover

Networks are not alone

Other domains have graphs to
optimize for latency:

Fire/police/hospital locations

Amazon warehouses

Netflix/Content servers

→ Facility location problems

Software Defined Network (SDN)

