

clISP: A Speed-of-Light Internet Service Provider

Debopam Bhattacherjee*, Waqar Aqeel*¹, Sangeetha Abdu Jyothi, Ilker Nadi Bozkurt, William Sentosa, Muhammad Tirmazi, Anthony Aguirre, Balakrishnan Chandrasekaran, P. Brighten Godfrey, Gregory Laughlin, Bruce Maggs, Ankit Singla



* Equal contribution

¹ Currently at Google

The Internet is slow?!

A toy experiment

The Internet is slow?!

A toy experiment

- \$ wget theoatmeal.com 490ms (fetch size 20KB)

The Internet is slow?!

A toy experiment

- \$ wget theoatmeal.com 490ms (fetch size 20KB)
- The lower bound is the speed of light! (c-latency)

The Internet is slow?!

A toy experiment

- \$ wget theoatmeal.com 490ms (fetch size 20KB)
- The lower bound is the speed of light! (c-latency)
- Washington DC (client) to Portland (server) roundtrip

The Internet is slow?!

A toy experiment

- \$ wget theoatmeal.com 490ms (fetch size 20KB)
- The lower bound is the speed of light! (c-latency)
- Washington DC (client) to Portland (server) roundtrip
 - 7,560 kilometers (c-latency = 25ms)

The Internet is slow?!

A toy experiment

- \$ wget theoatmeal.com 490ms (fetch size 20KB)
- The lower bound is the speed of light! (c-latency)
- Washington DC (client) to Portland (server) roundtrip
 - 7,560 kilometers (c-latency = 25ms)
 - 20x inflation!

The Internet is slow?!

A toy experiment

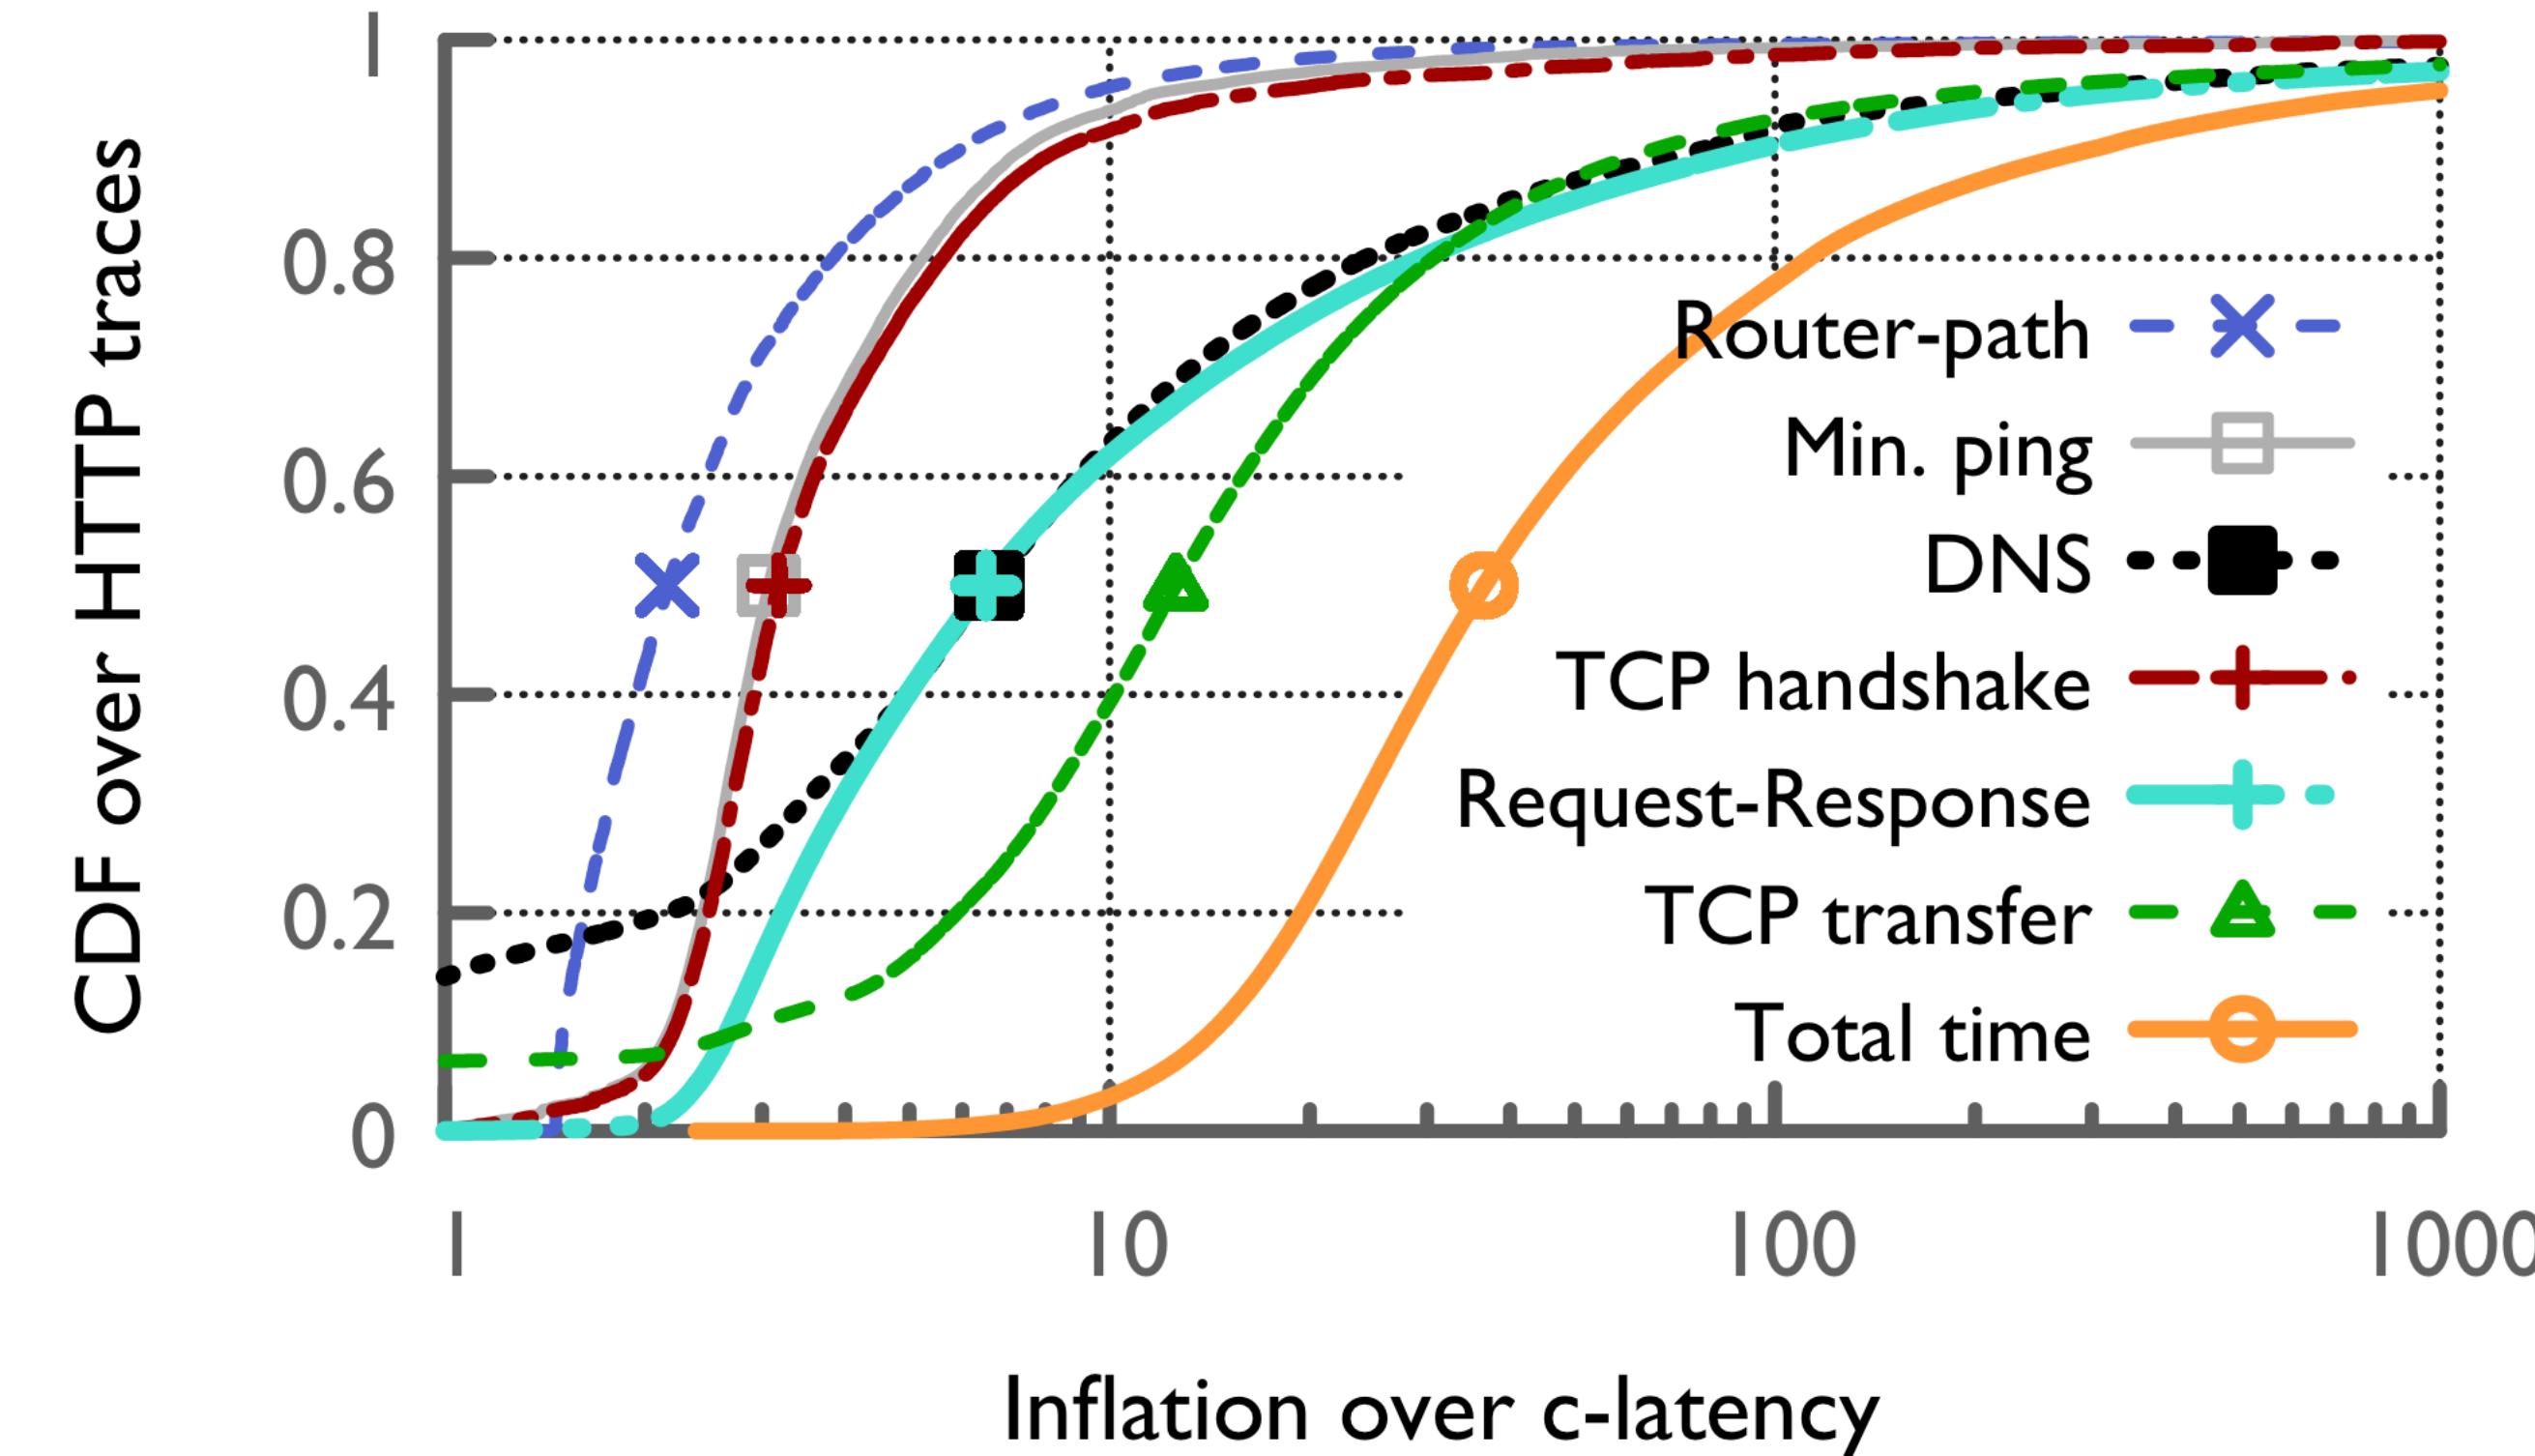
- \$ wget theoatmeal.com 490ms (fetch size 20KB)
- The lower bound is the speed of light! (c-latency)
- Washington DC (client) to Portland (server) roundtrip
 - 7,560 kilometers (c-latency = 25ms)
 - 20x inflation!

37x inflation in larger experiments!^[1]

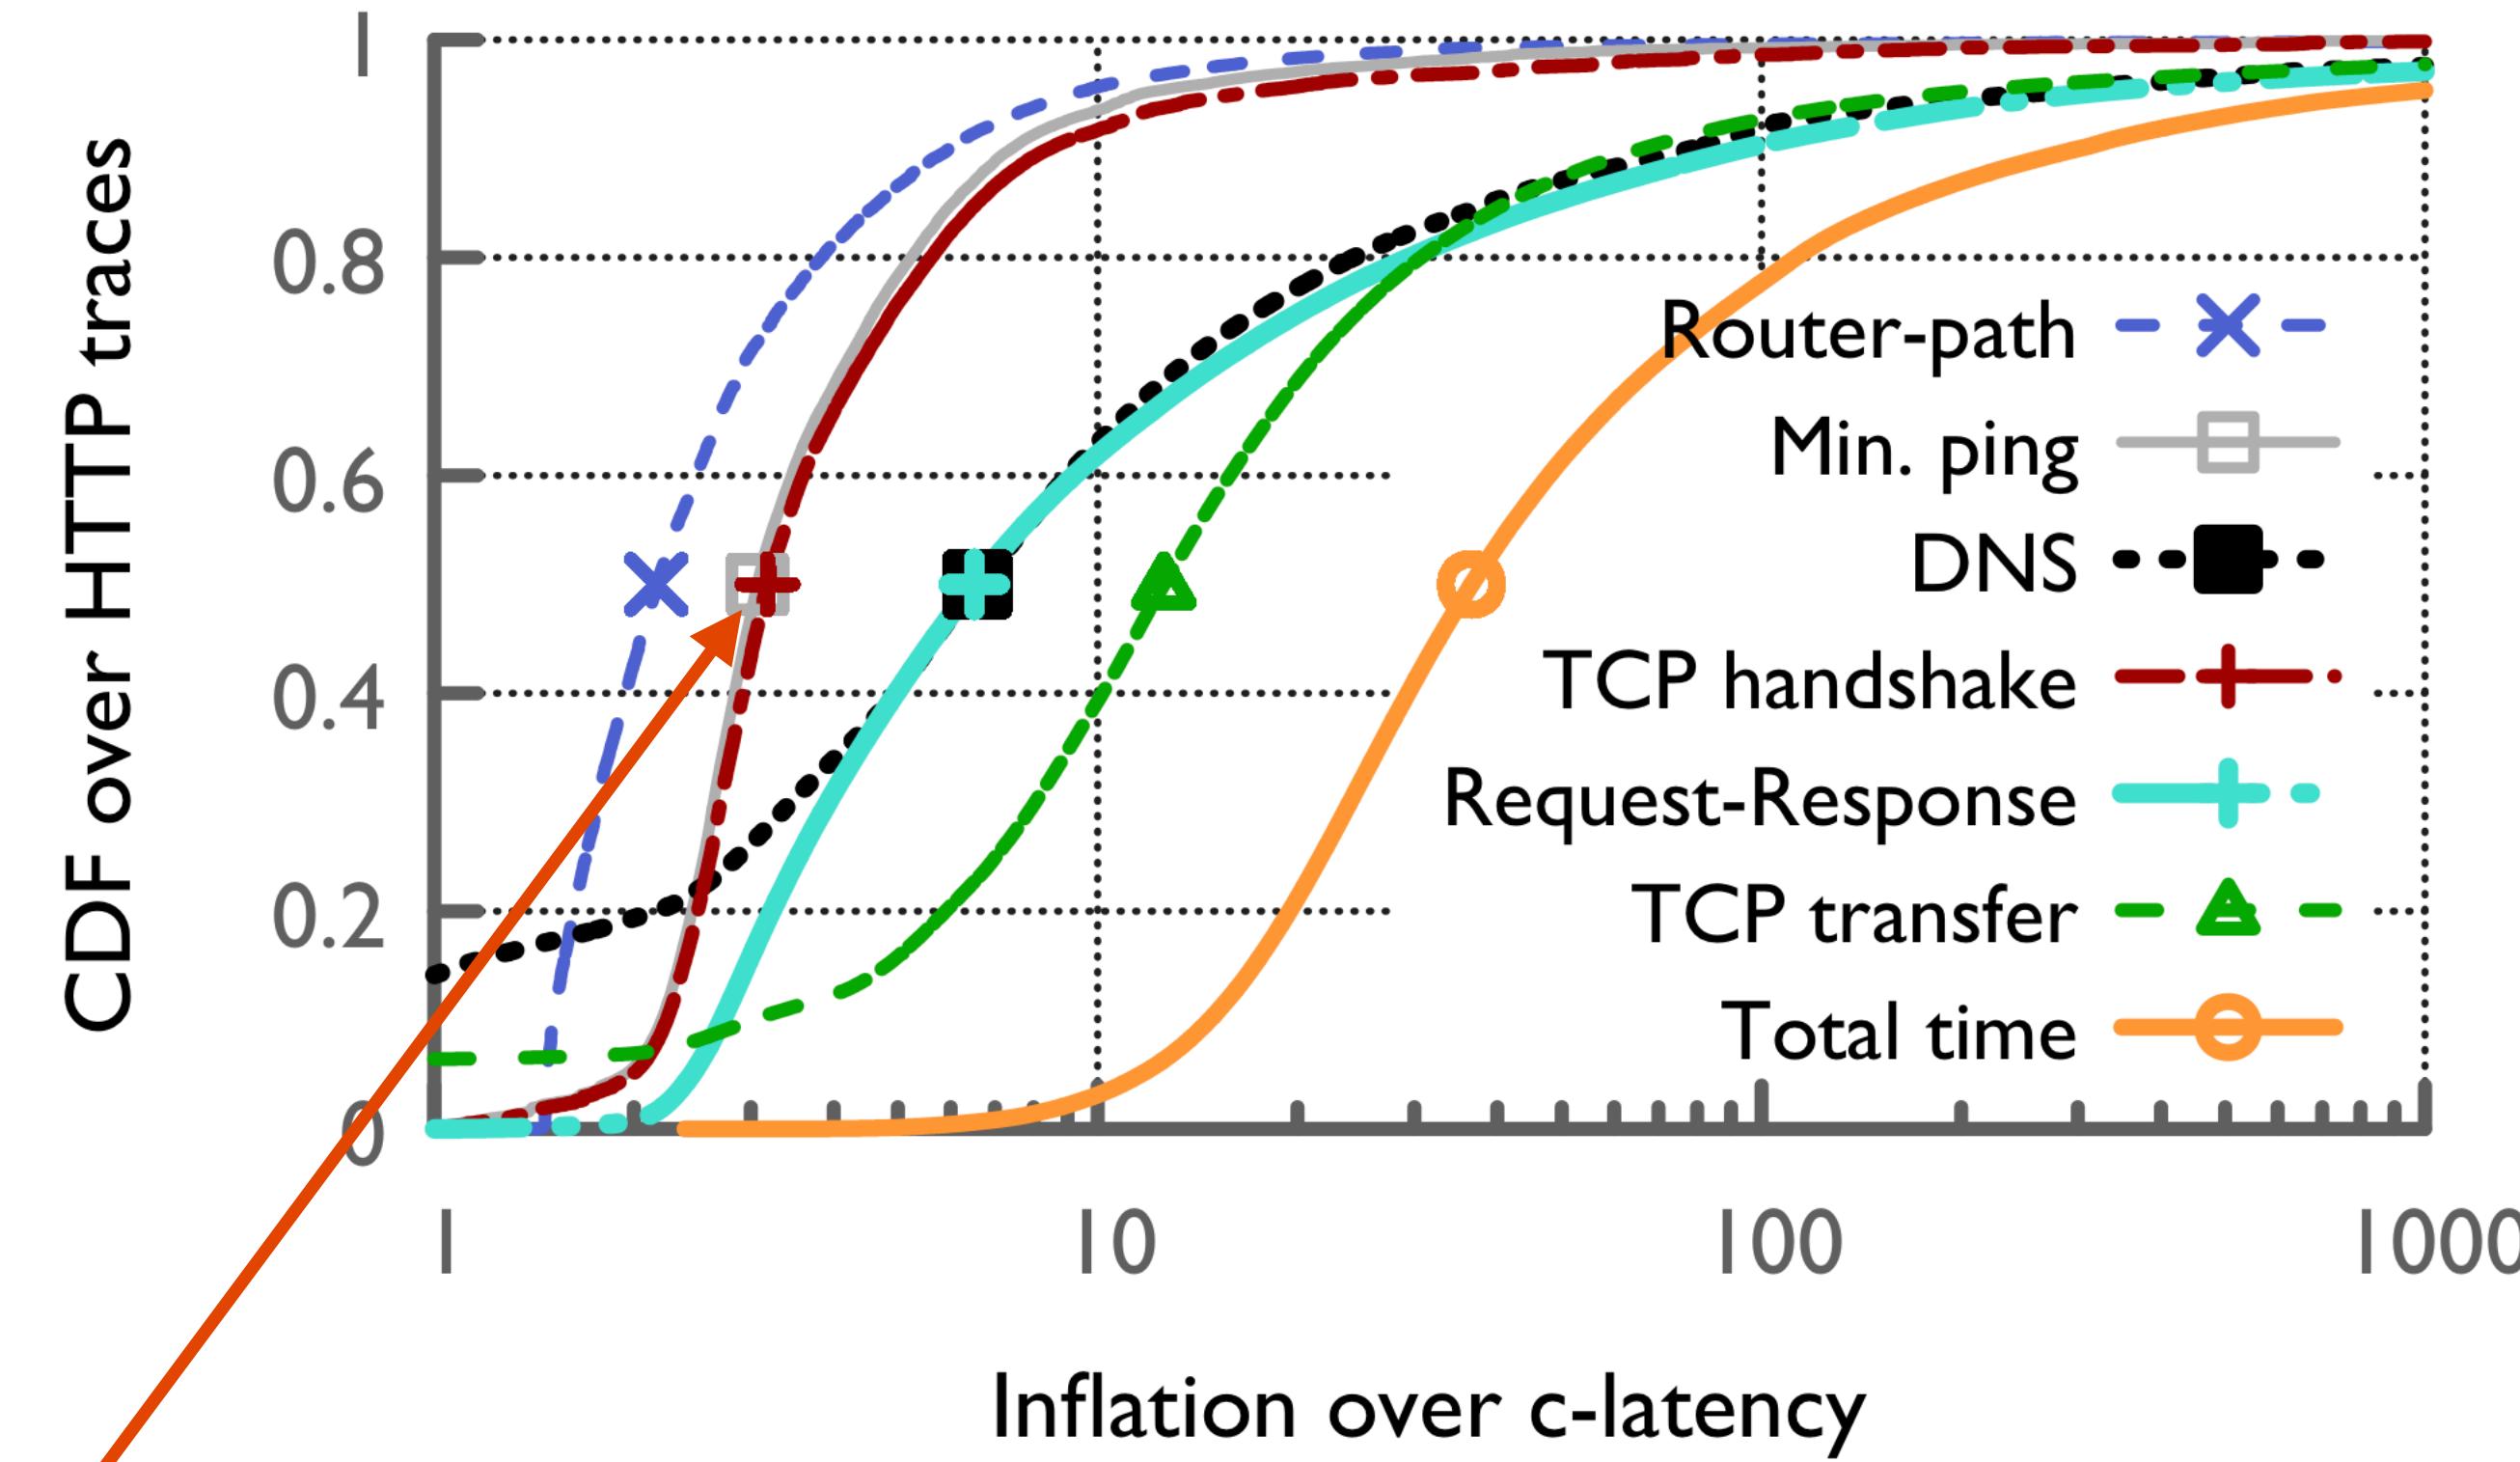
[1] Bozkurt et al. "Why is the internet so slow?!. International Conference on Passive and Active Network Measurement. 2017.

Let's break down the 37x inflation

Let's break down the 37x inflation



Let's break down the 37x inflation



3-4x median inflation for delivering a single packet!

Deliver packets on the Internet at the speed of light!

Deliver packets on the Internet at the speed of light!

cISP = ISP operating at the speed of light

Why is packet delivery slow?



Why is packet delivery slow?



Circuitous routing

Why is packet delivery slow?



Circuitous routing

Fiber is not line-of-sight straight

Why is packet delivery slow?



Circuitous routing

Fiber is not line-of-sight straight

Glass has higher propagation delay

World's fastest wide-area networks

Microwave networks used by the High Frequency Trading industry

World's fastest wide-area networks

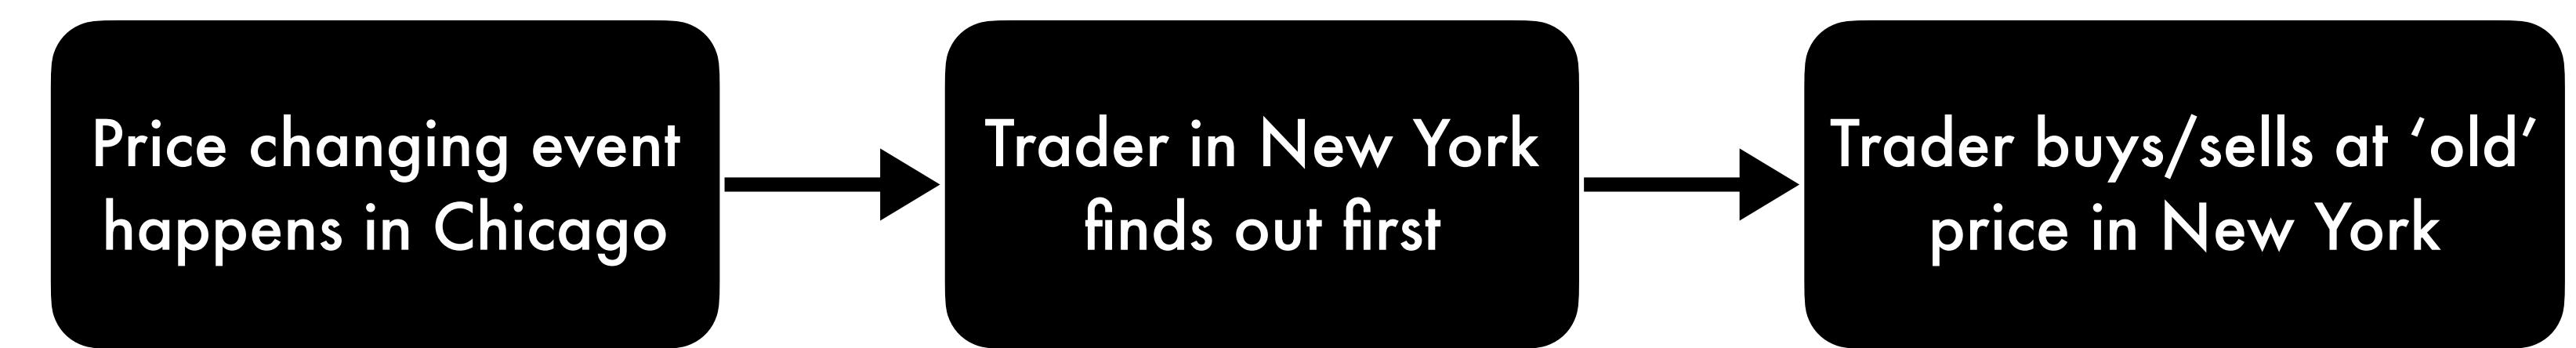
Microwave networks used by the High Frequency Trading industry

- Latency arbitrage

World's fastest wide-area networks

Microwave networks used by the High Frequency Trading industry

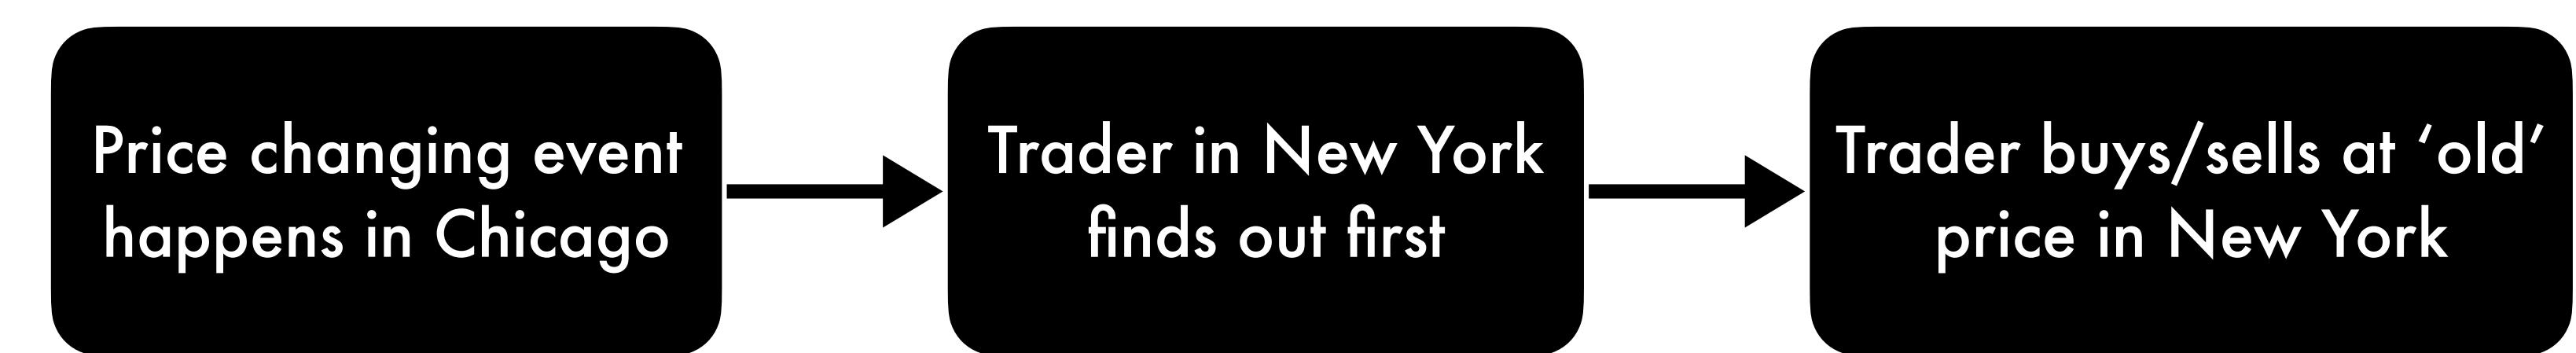
- Latency arbitrage



World's fastest wide-area networks

Microwave networks used by the High Frequency Trading industry

- Latency arbitrage

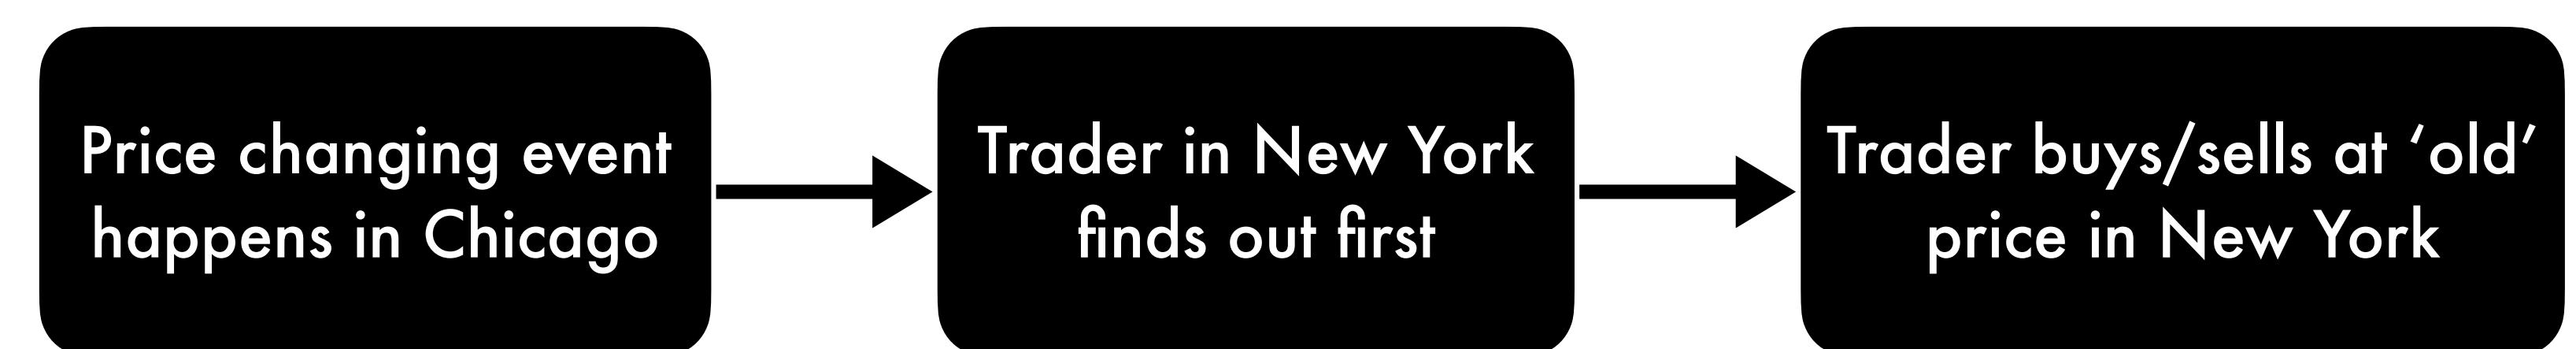


- Millions of dollars for a millisecond or less

World's fastest wide-area networks

Microwave networks used by the High Frequency Trading industry

- Latency arbitrage



- Millions of dollars for a millisecond or less

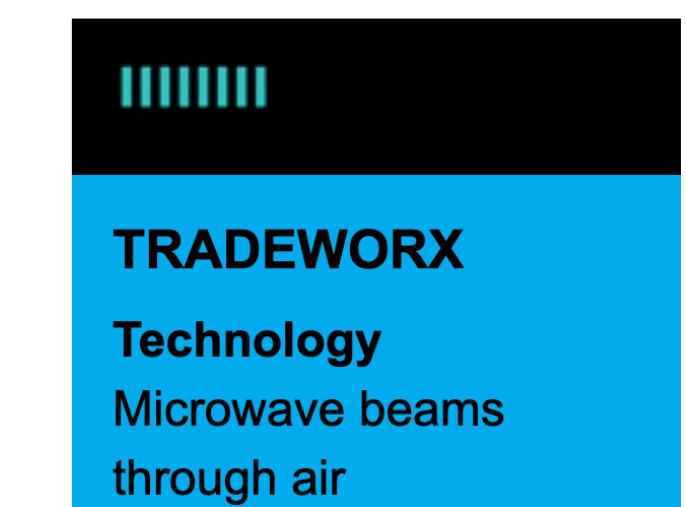
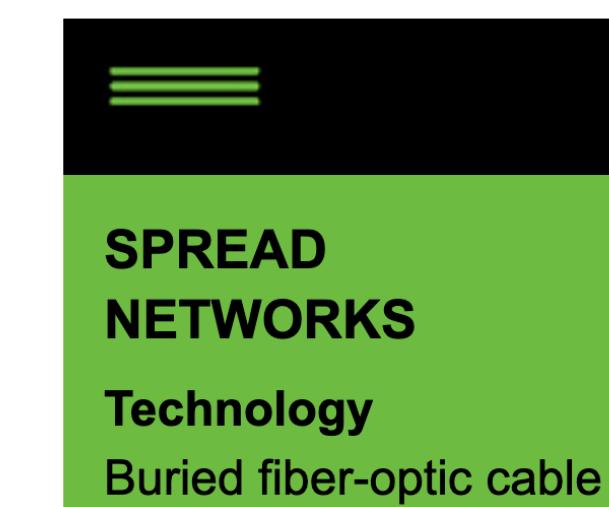
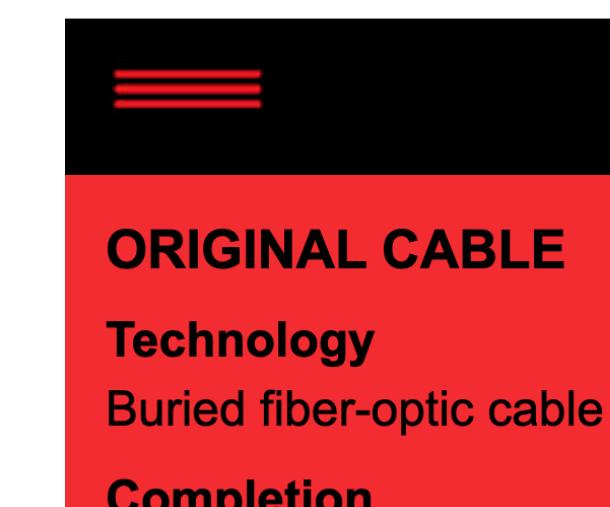
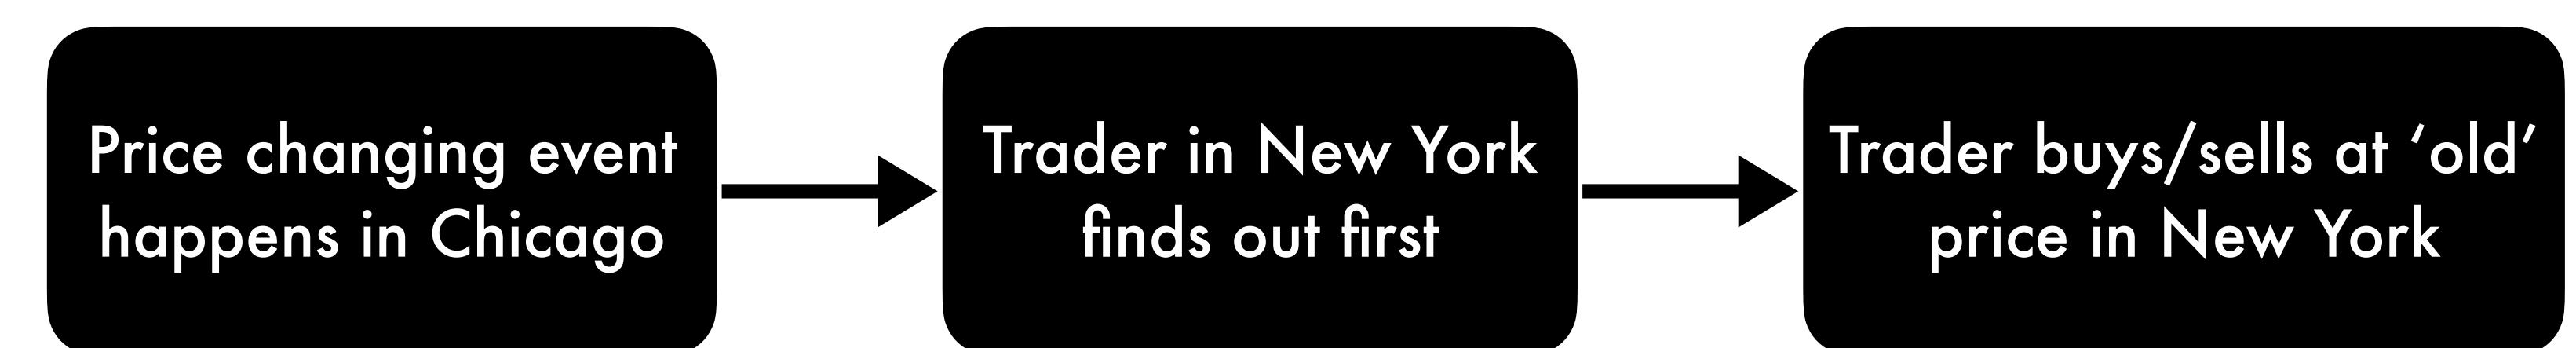


Photo from wired.com

World's fastest wide-area networks

Microwave networks used by the High Frequency Trading industry

- Latency arbitrage



- Millions of dollars for a millisecond or less



- Some operate at within 1% of c-latency

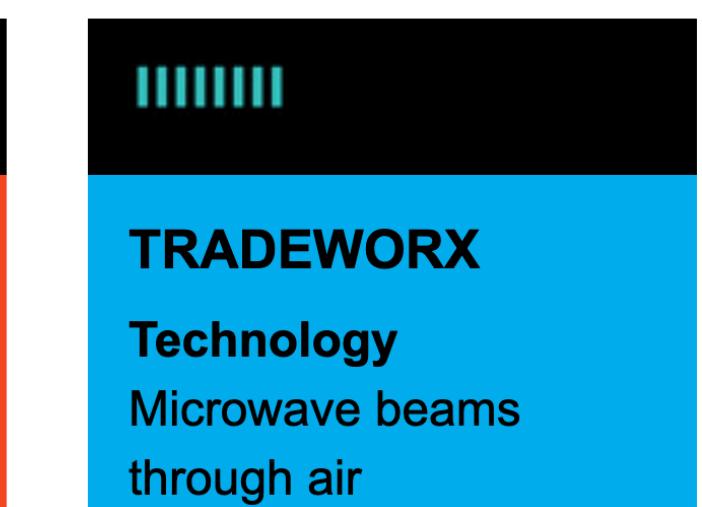
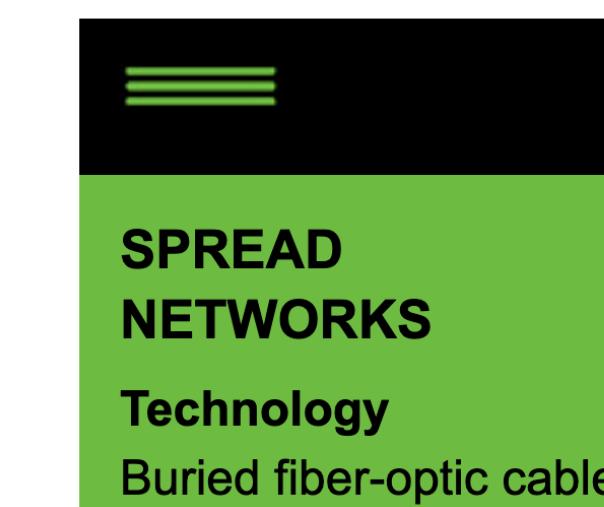
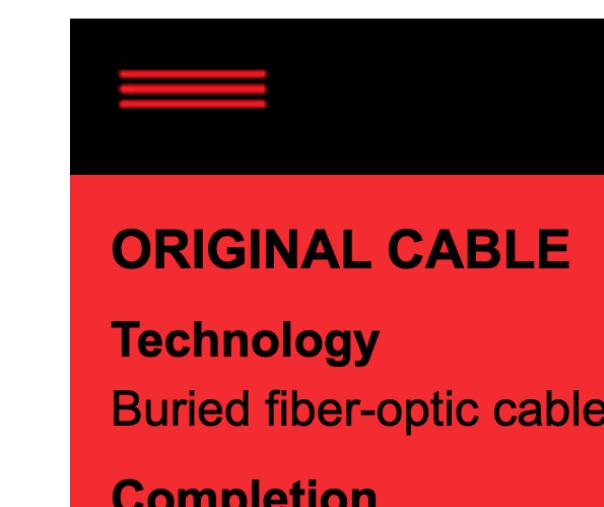


Photo from [wired.com](#)

Microwave Transmission

Quick background



Microwave Transmission

Quick background

- Switching latency within a few microseconds



Microwave Transmission

Quick background

- Switching latency within a few microseconds
- 1 Gbps bandwidth achievable



Microwave Transmission

Quick background

- Switching latency within a few microseconds
- 1 Gbps bandwidth achievable
- Average cost:



Microwave Transmission

Quick background

- Switching latency within a few microseconds
- 1 Gbps bandwidth achievable
- Average cost:
 - \$75K (\$150K) for a 500Mbps (1Gbps) link



Microwave Transmission

Quick background

- Switching latency within a few microseconds
- 1 Gbps bandwidth achievable
- Average cost:
 - \$75K (\$150K) for a 500Mbps (1Gbps) link
 - \$100K for a new tower



Microwave Transmission

Quick background

- Switching latency within a few microseconds
- 1Gbps bandwidth achievable
- Average cost:
 - \$75K (\$150K) for a 500Mbps (1Gbps) link
 - ~~\$100K for a new tower~~
- Lots of usable towers are already there!



Sure, microwave is fast. But how do we make clSP?

Sure, microwave is fast. But how do we make clSP?

Tower placement?

Sure, microwave is fast. But how do we make clSP?

Tower placement?

Bandwidth?

Sure, microwave is fast. But how do we make clSP?

Tower placement?

Bandwidth?

Resilience?

Sure, microwave is fast. But how do we make clSP?

Tower placement?

Bandwidth?

Resilience?

Cost?

Designing clSP

Which towers should be connected?

Designing clSP

Which towers should be connected?

1. Select cities to connect

Designing clSP

Which towers should be connected?

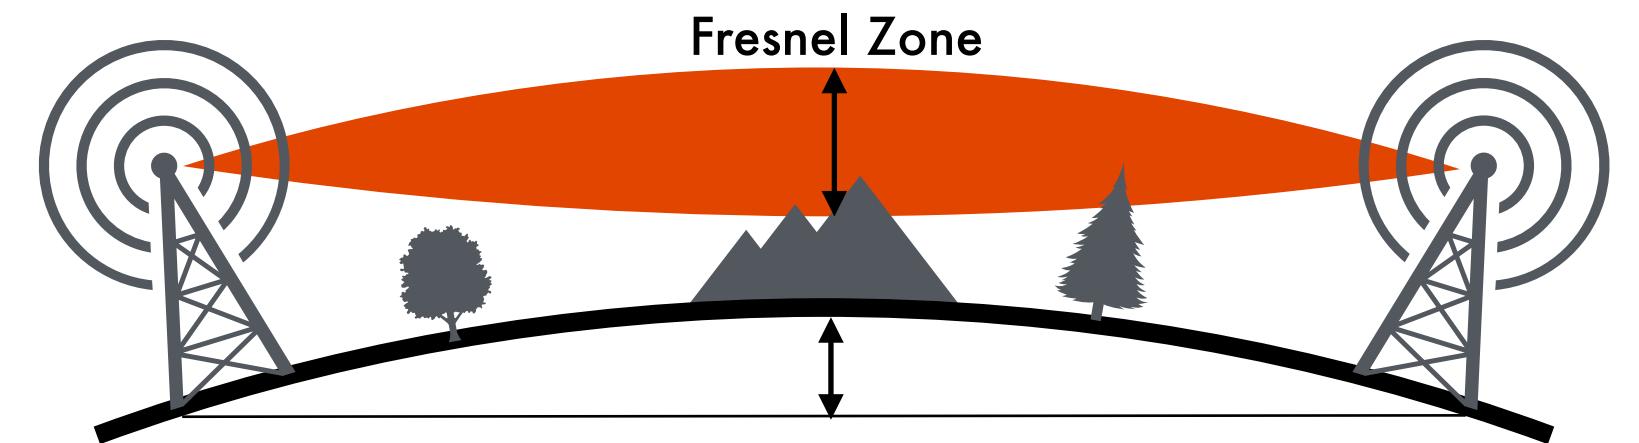
1. Select cities to connect
2. Identify feasible hops: a pair of connected towers must be

Designing clSP

Which towers should be connected?

1. Select cities to connect
2. Identify feasible hops: a pair of connected towers must be

$$h_{\text{Fres}} \simeq 8.7 \text{ m} \left(\frac{D}{1 \text{ km}} \right)^{1/2} \left(\frac{f}{1 \text{ GHz}} \right)^{-1/2}$$



$$h_{\text{Earth}} \simeq \frac{1 \text{ m}}{50 \text{ K}} \left(\frac{D}{1 \text{ km}} \right)^2$$

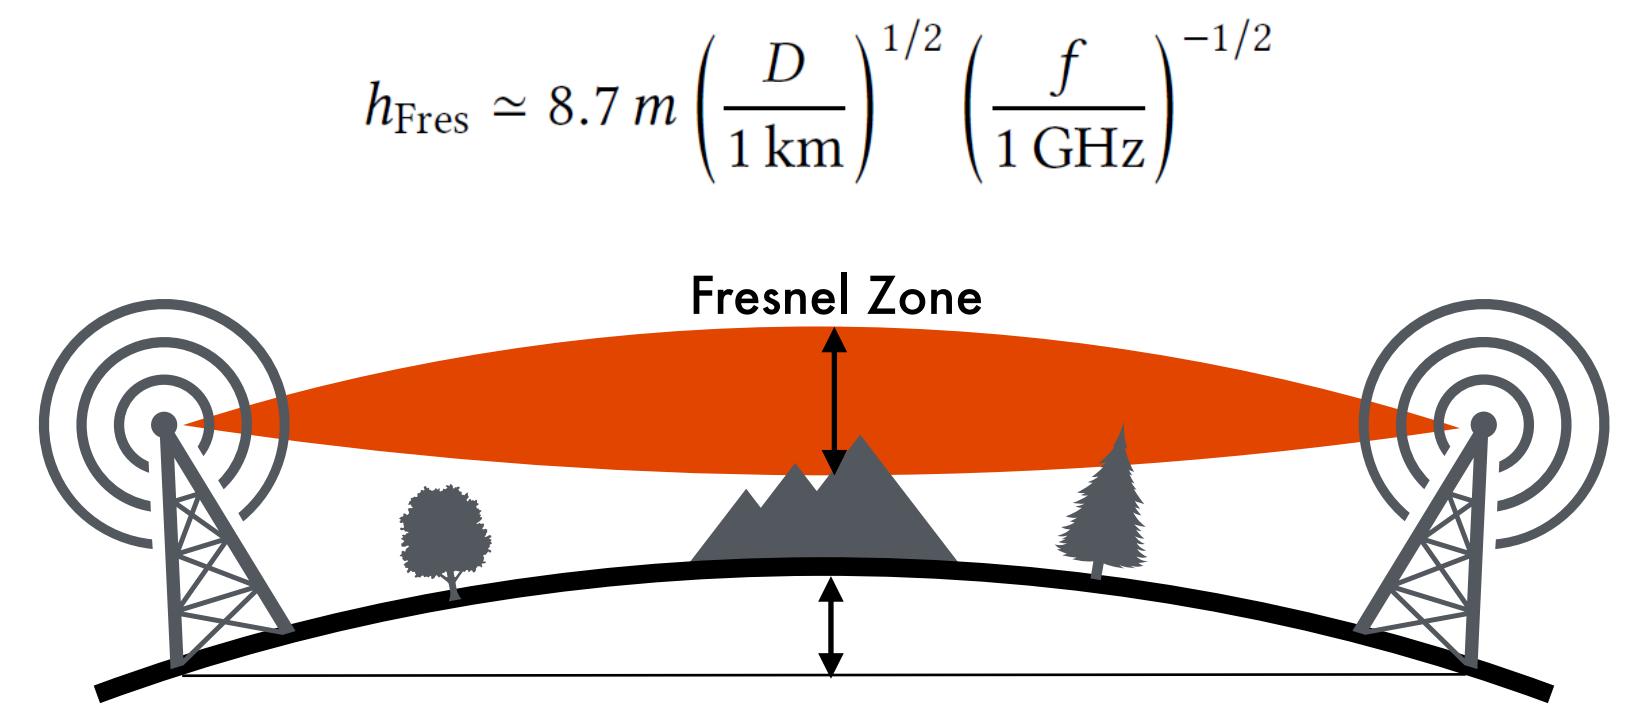
Designing clSP

Which towers should be connected?

1. Select cities to connect

2. Identify feasible hops: a pair of connected towers must be

- high enough to clear Earth's curvature and obstructions in the Fresnel Zone



$$h_{\text{Earth}} \simeq \frac{1 \text{ m}}{50 \text{ K}} \left(\frac{D}{1 \text{ km}} \right)^2$$

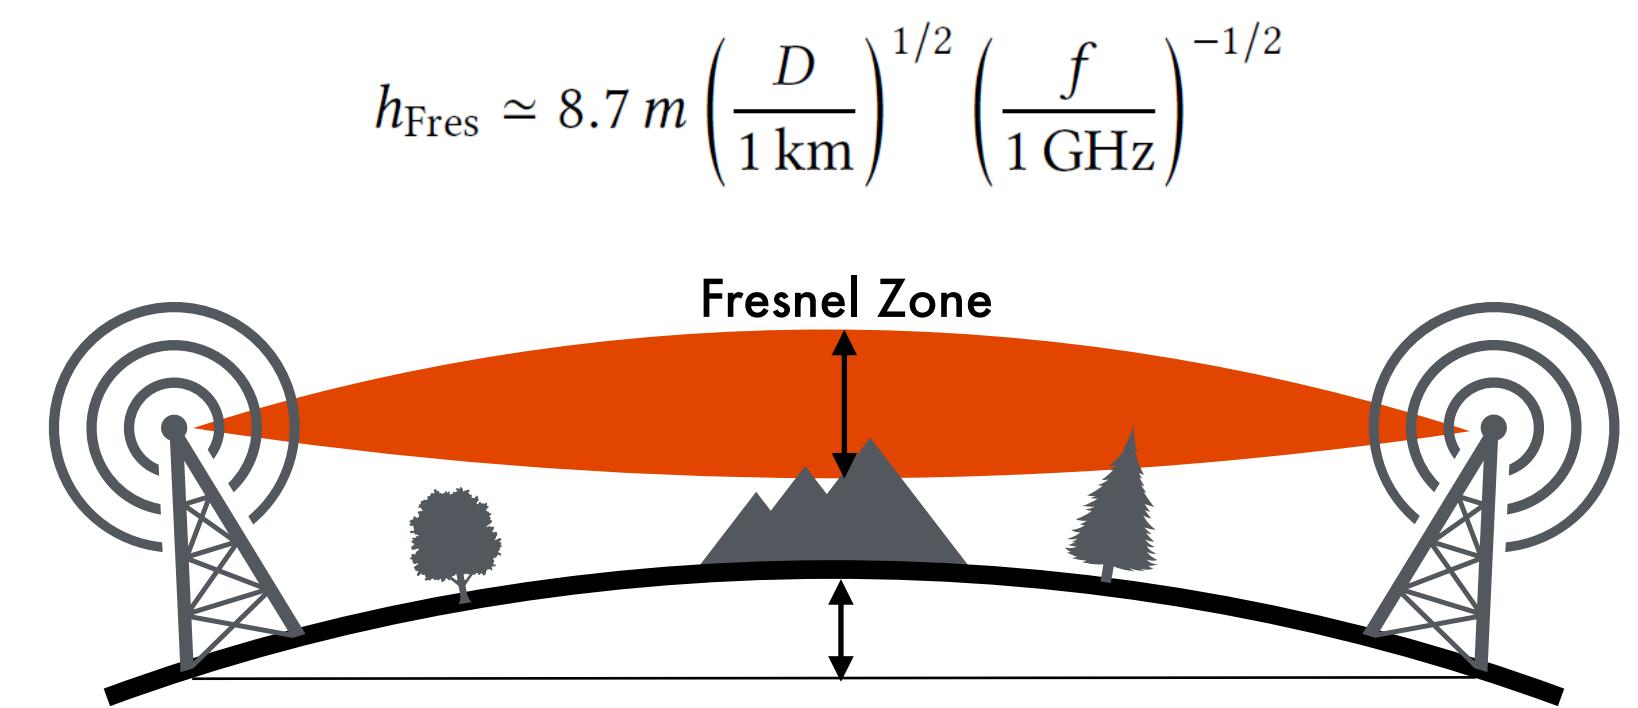
Designing clSP

Which towers should be connected?

1. Select cities to connect

2. Identify feasible hops: a pair of connected towers must be

- high enough to clear Earth's curvature and obstructions in the Fresnel Zone
- at most 100km apart

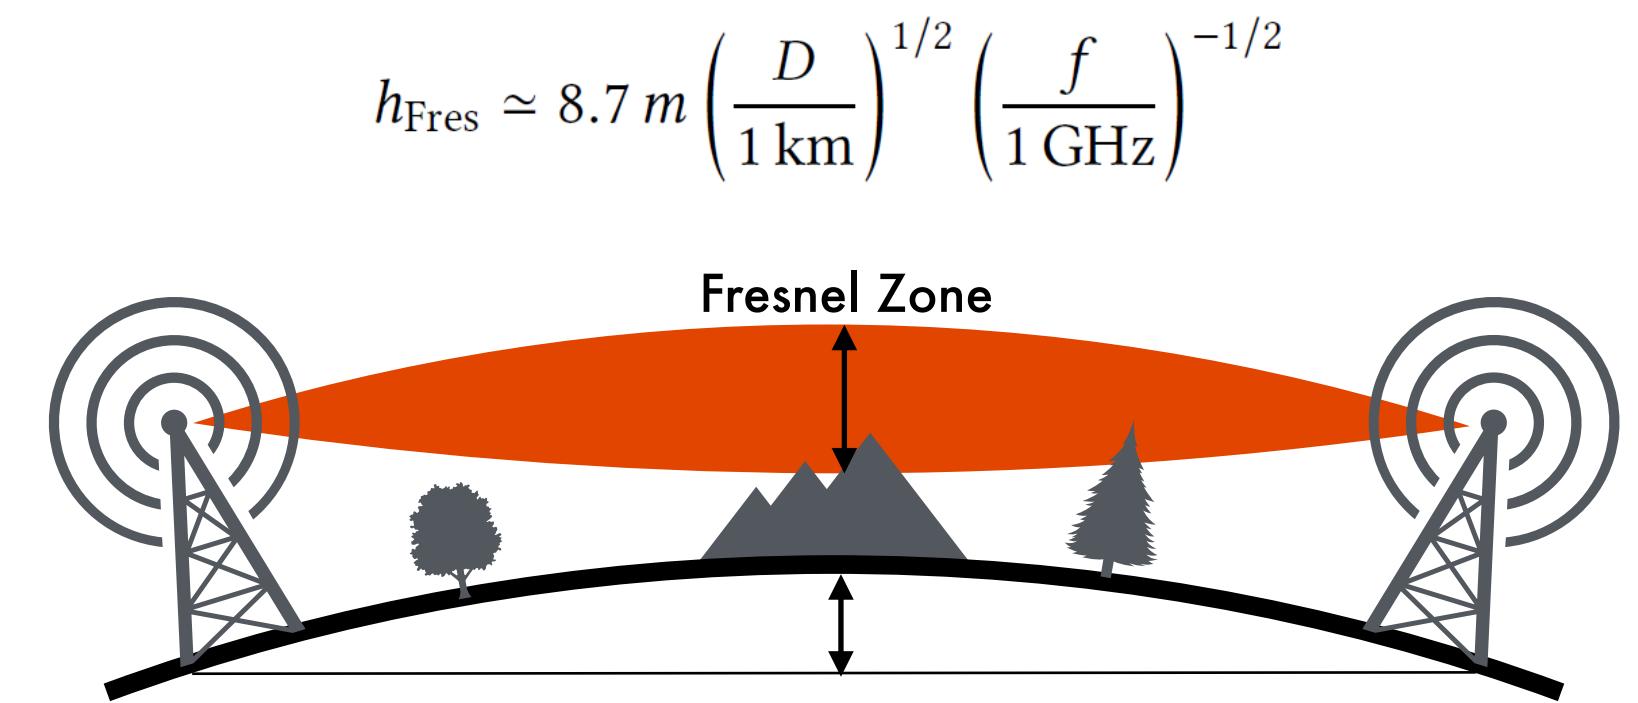


$$h_{\text{Earth}} \simeq \frac{1 \text{ m}}{50 \text{ K}} \left(\frac{D}{1 \text{ km}} \right)^2$$

Designing clSP

Which towers should be connected?

1. Select cities to connect
2. Identify feasible hops: a pair of connected towers must be
 - high enough to clear Earth's curvature and obstructions in the Fresnel Zone
 - at most 100km apart
3. Pick shortest path of towers to link a pair of cities



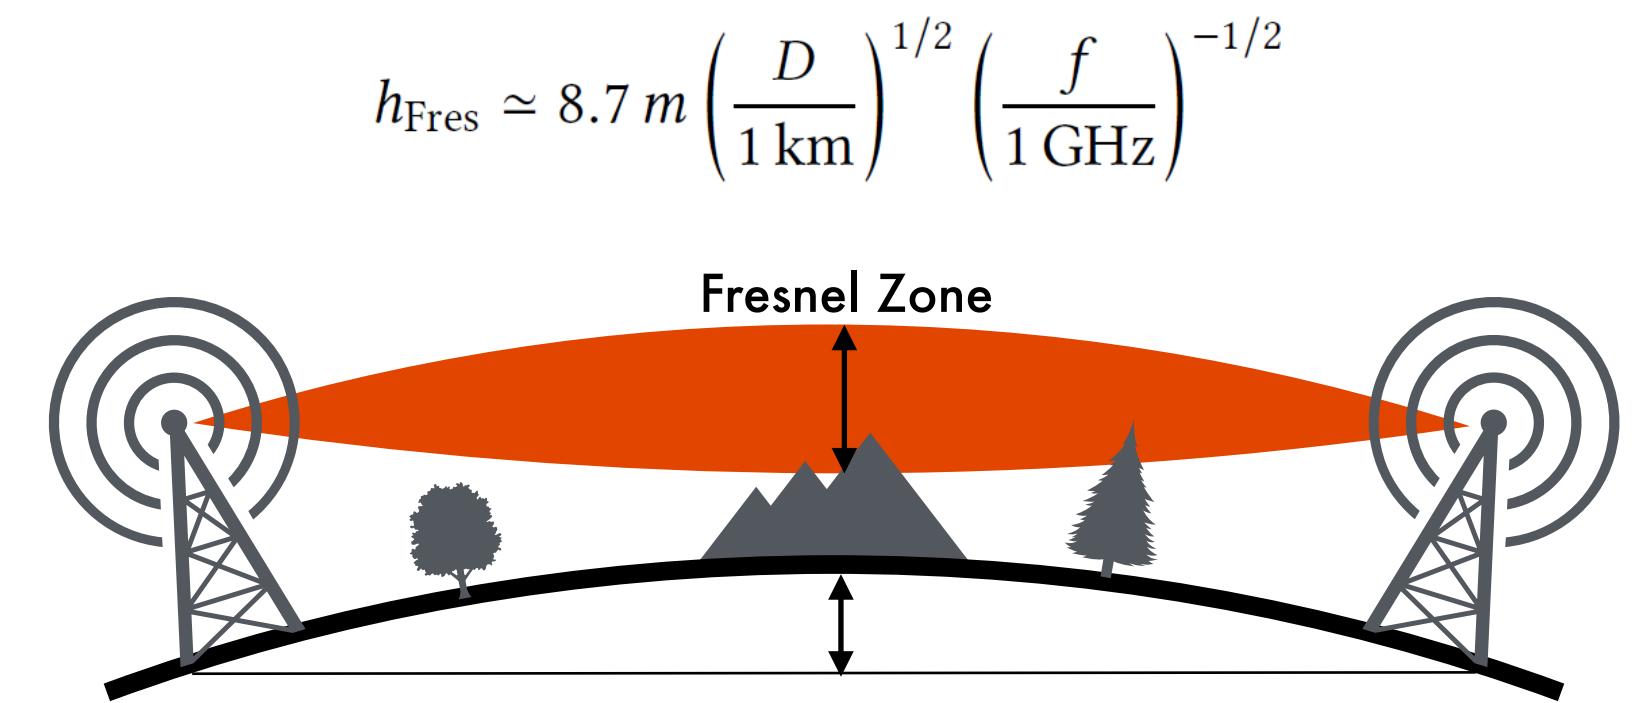
$$h_{\text{Fres}} \simeq 8.7 \text{ m} \left(\frac{D}{1 \text{ km}} \right)^{1/2} \left(\frac{f}{1 \text{ GHz}} \right)^{-1/2}$$

$$h_{\text{Earth}} \simeq \frac{1 \text{ m}}{50 K} \left(\frac{D}{1 \text{ km}} \right)^2$$

Designing clSP

Which towers should be connected?

1. Select cities to connect
2. Identify feasible hops: a pair of connected towers must be
 - high enough to clear Earth's curvature and obstructions in the Fresnel Zone
 - at most 100km apart
3. Pick shortest path of towers to link a pair of cities
4. Select an **optimal** subset of links



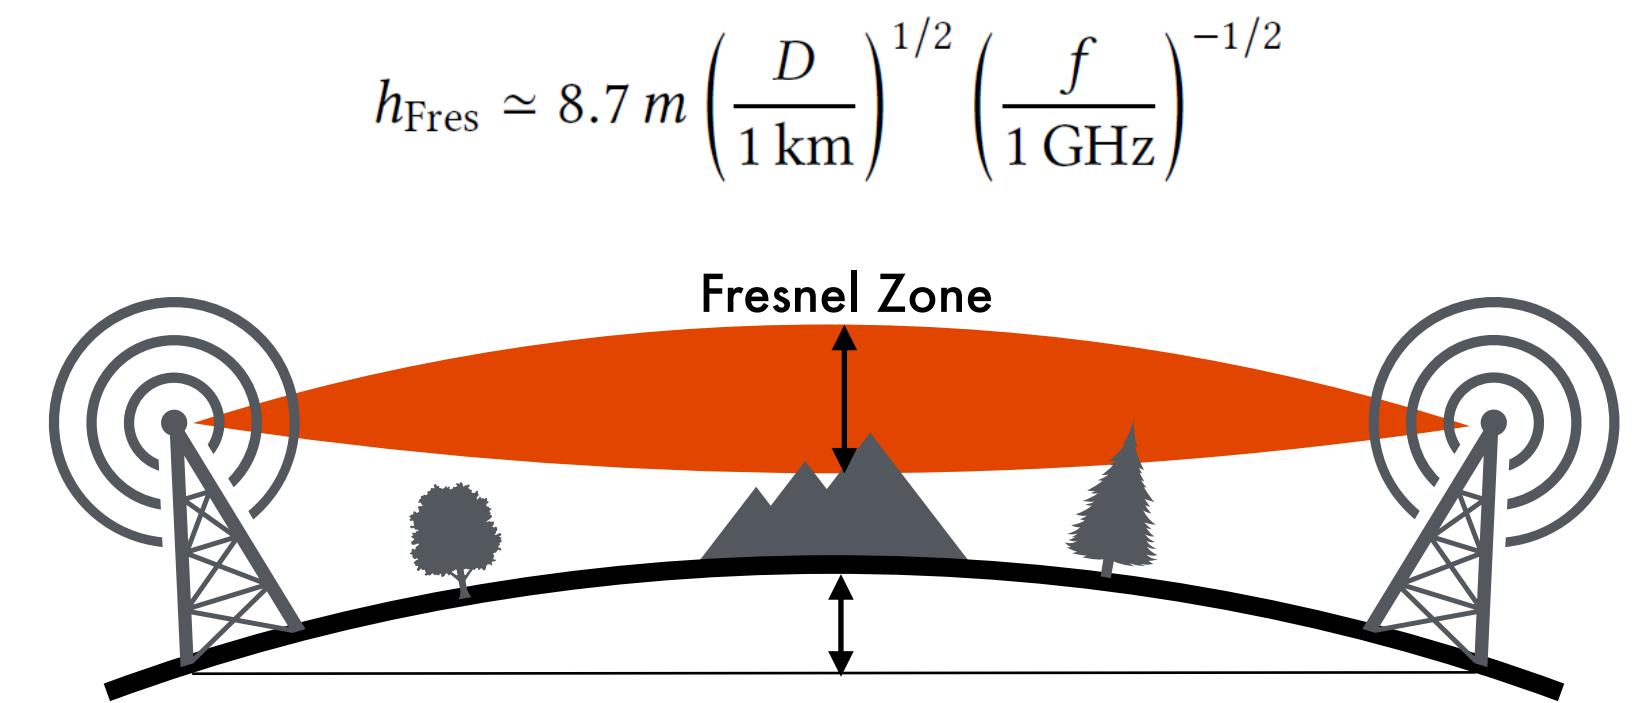
$$h_{\text{Fres}} \simeq 8.7 \text{ m} \left(\frac{D}{1 \text{ km}} \right)^{1/2} \left(\frac{f}{1 \text{ GHz}} \right)^{-1/2}$$

$$h_{\text{Earth}} \simeq \frac{1 \text{ m}}{50 K} \left(\frac{D}{1 \text{ km}} \right)^2$$

Designing clSP

Which towers should be connected?

1. Select cities to connect
2. Identify feasible hops: a pair of connected towers must be
 - high enough to clear Earth's curvature and obstructions in the Fresnel Zone
 - at most 100km apart
3. Pick shortest path of towers to link a pair of cities
4. Select an **optimal** subset of links



$$h_{\text{Fres}} \simeq 8.7 \text{ m} \left(\frac{D}{1 \text{ km}} \right)^{1/2} \left(\frac{f}{1 \text{ GHz}} \right)^{-1/2}$$

$$h_{\text{Earth}} \simeq \frac{1 \text{ m}}{50 K} \left(\frac{D}{1 \text{ km}} \right)^2$$

$$\min \sum_{s,t} \frac{h_{st}}{d_{st}} \sum_{i,j} (o_{i,j} f_{stij,o} + m_{i,j} f_{stij,m})$$

cISP for the United States

Budget = 3000 towers

cISP for the United States

Budget = 3000 towers

Input

120 largest population centers

cISP for the United States

Budget = 3000 towers

Input

120 largest population centers

Gravity model traffic matrix

cISP for the United States

Budget = 3000 towers

Input

120 largest population centers

Gravity model traffic matrix

Tower data from FCC and others

cISP for the United States

Budget = 3000 towers

Input

120 largest population centers

Gravity model traffic matrix

Tower data from FCC and others

Terrain data from NASA

cISP for the United States

Budget = 3000 towers

Input

120 largest population centers

Gravity model traffic matrix

Tower data from FCC and others

Terrain data from NASA

InterTubes fiber map^[3]

[3] Durairajan et al. "InterTubes: A study of the US long-haul fiber-optic infrastructure." ACM SIGCOMM. 2015.

cISP for the United States

Budget = 3000 towers

Input

120 largest population centers

Gravity model traffic matrix

Tower data from FCC and others

Terrain data from NASA

InterTubes fiber map^[3]

100Gbps target aggregate bandwidth

[3] Durairajan et al. "InterTubes: A study of the US long-haul fiber-optic infrastructure." ACM SIGCOMM. 2015.

cISP for the United States

Budget = 3000 towers

Input

120 largest population centers

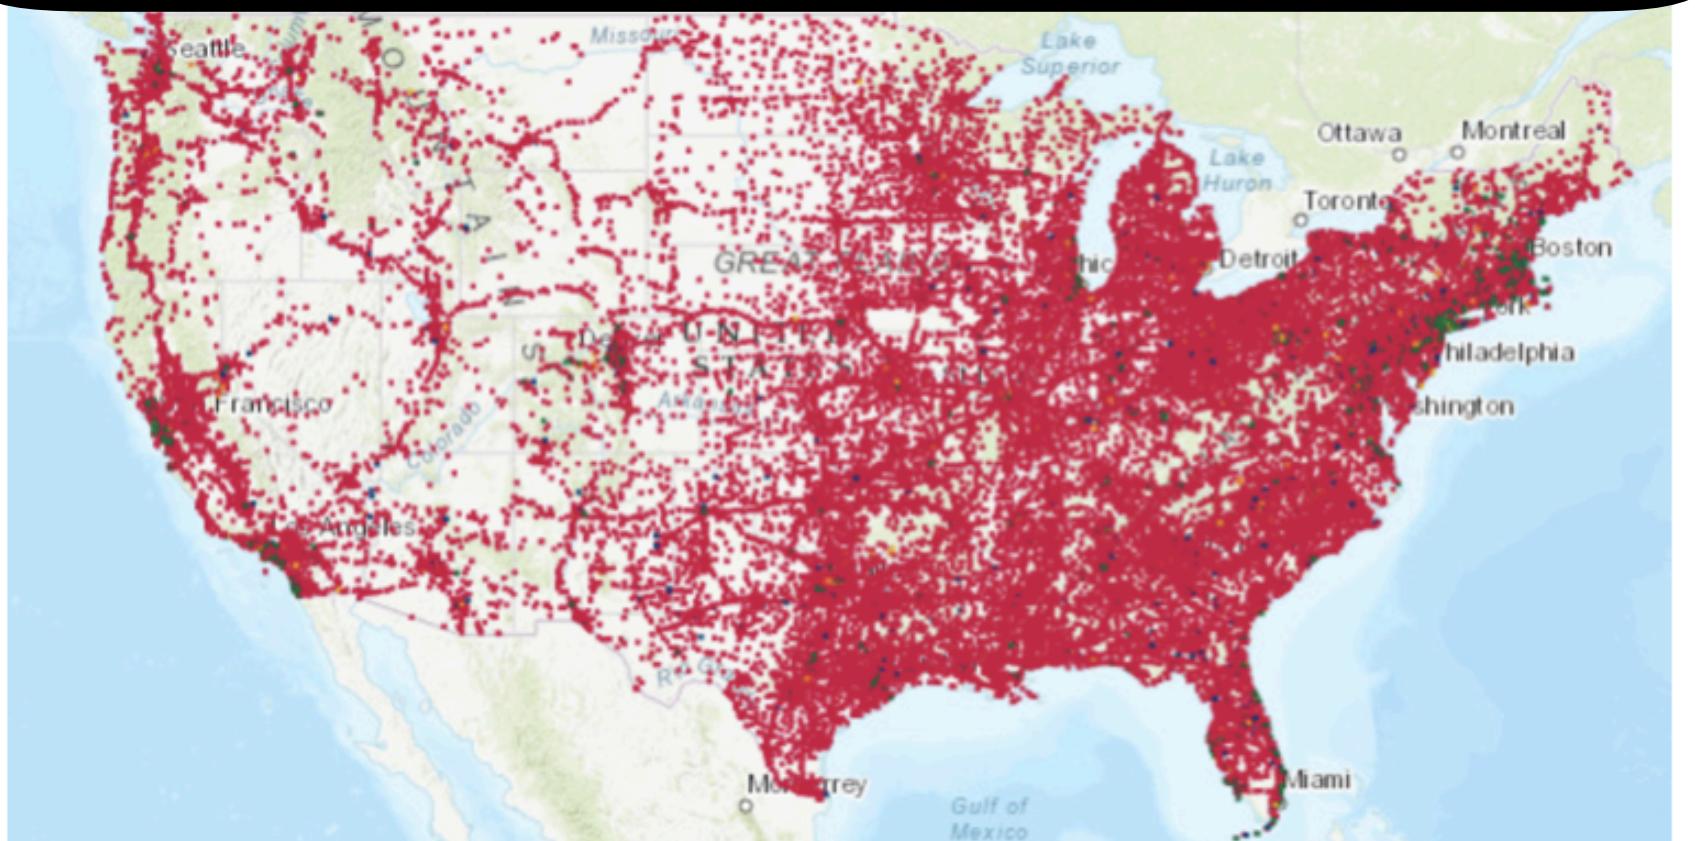
Gravity model traffic matrix

Tower data from FCC and others

Terrain data from NASA

InterTubes fiber map^[3]

100Gbps target aggregate bandwidth



[3] Durairajan et al. "InterTubes: A study of the US long-haul fiber-optic infrastructure." ACM SIGCOMM. 2015.

cISP for the United States

Budget = 3000 towers

Input

120 largest population centers

Gravity model traffic matrix

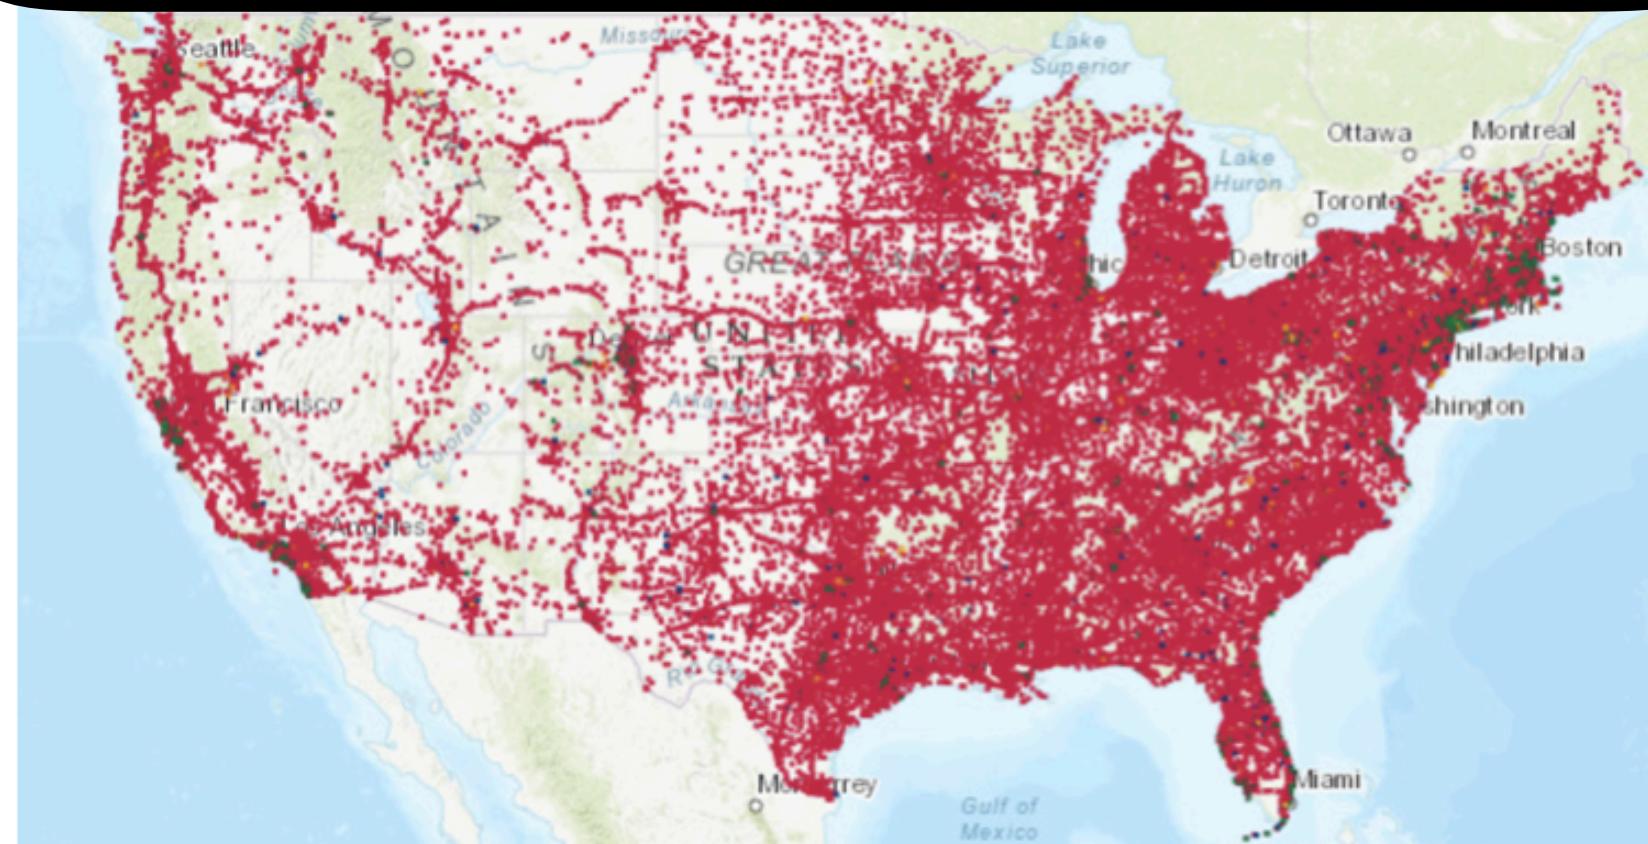
Tower data from FCC and others

Terrain data from NASA

InterTubes fiber map^[3]

100Gbps target aggregate bandwidth

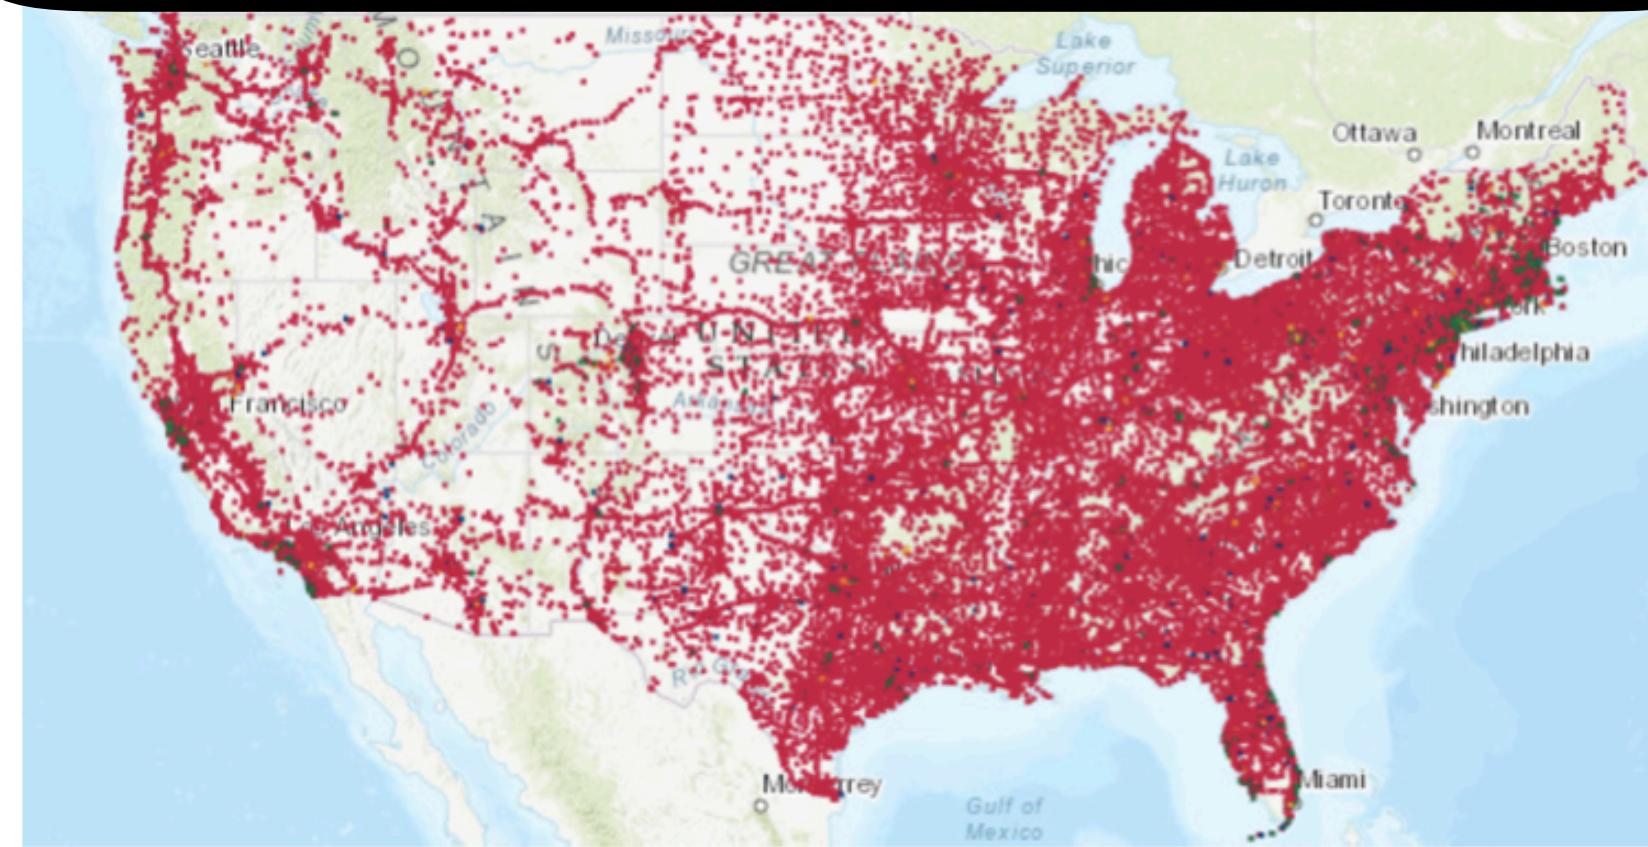
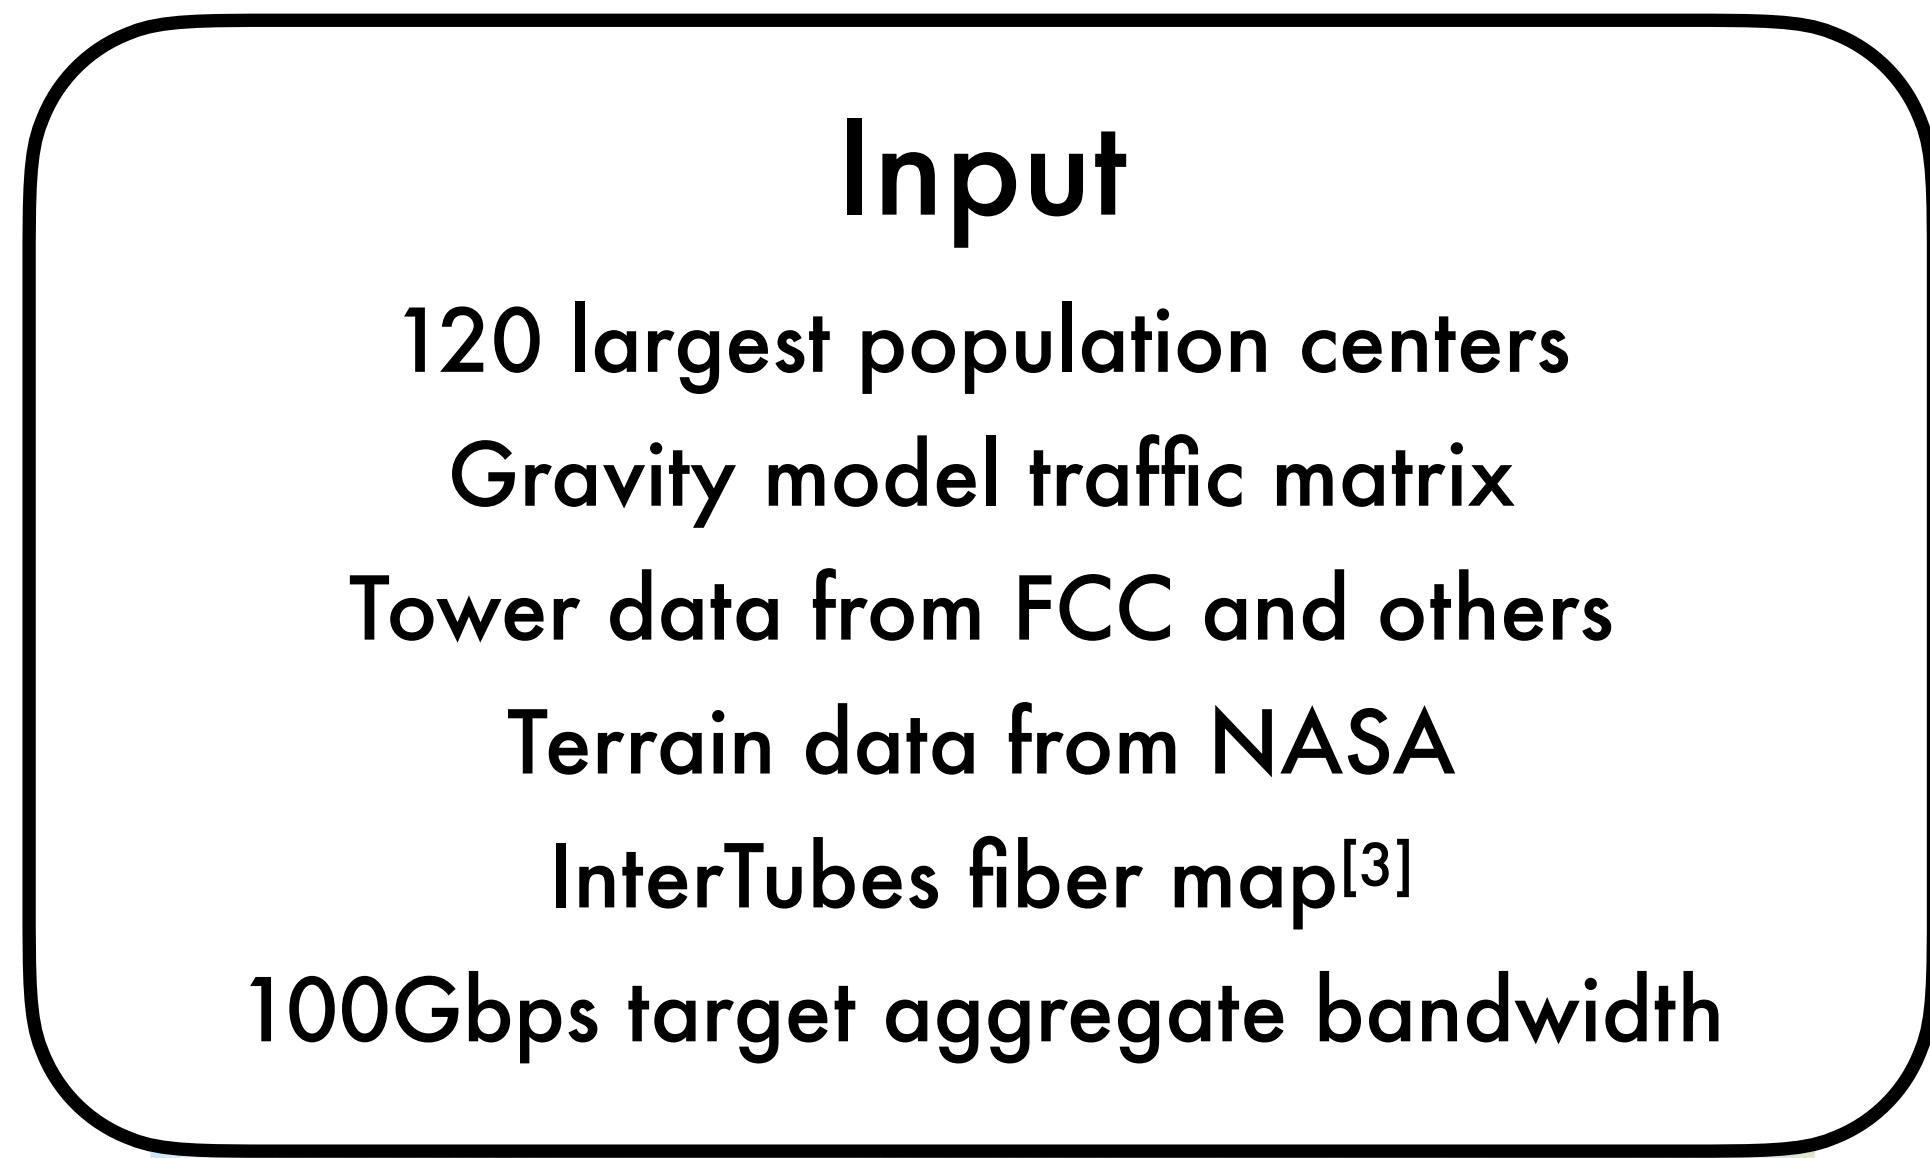
Identify feasible hops
Select shortest links



[3] Durairajan et al. "InterTubes: A study of the US long-haul fiber-optic infrastructure." ACM SIGCOMM. 2015.

cISP for the United States

Budget = 3000 towers



Identify feasible hops
Select shortest links

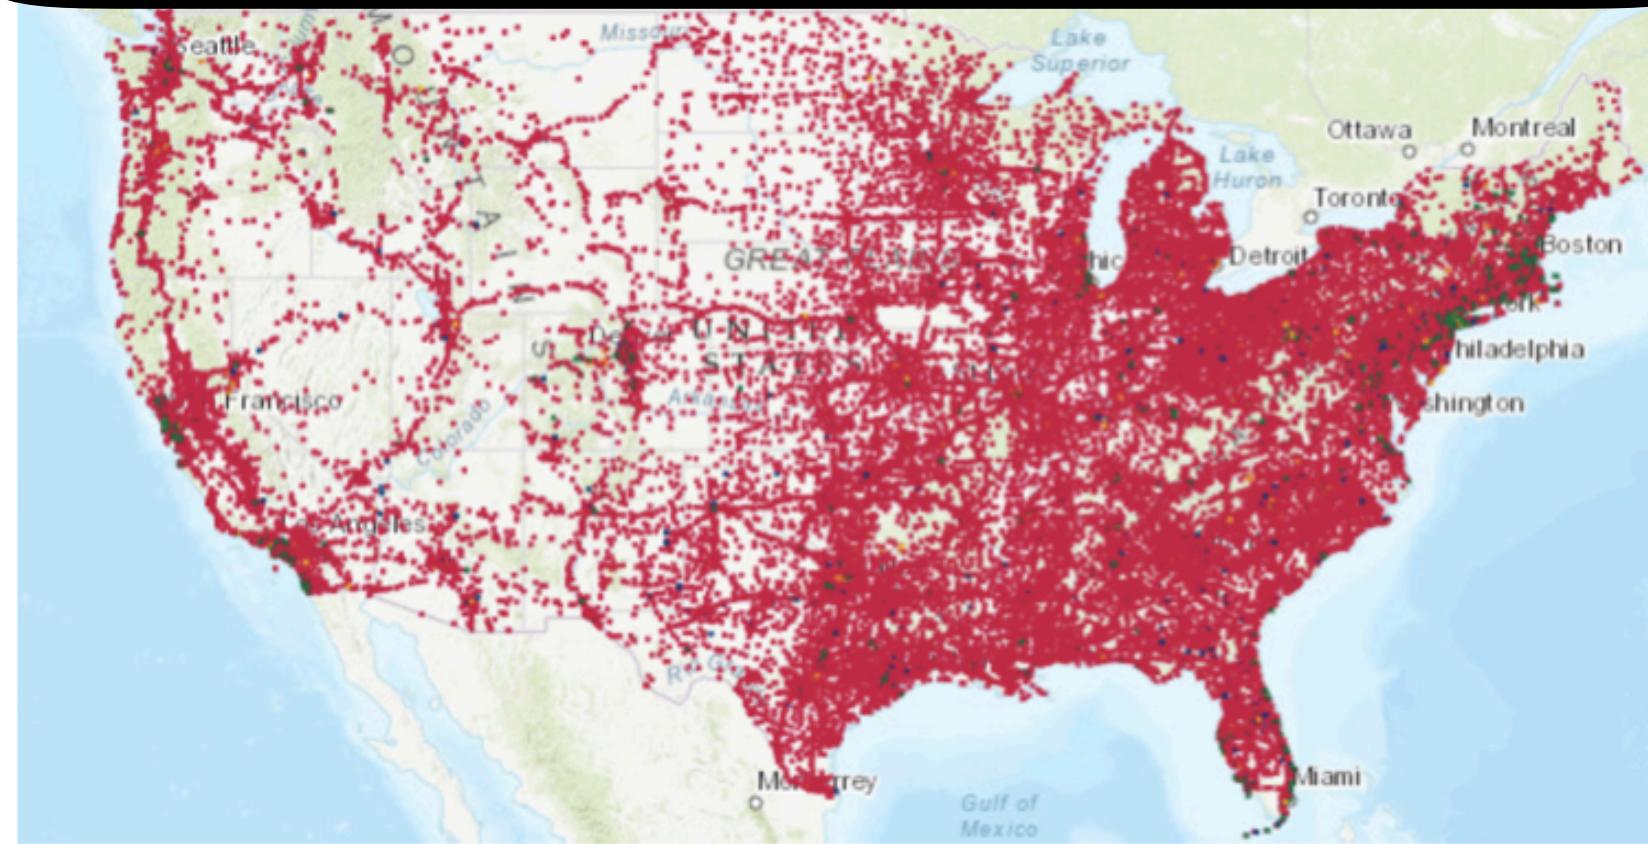
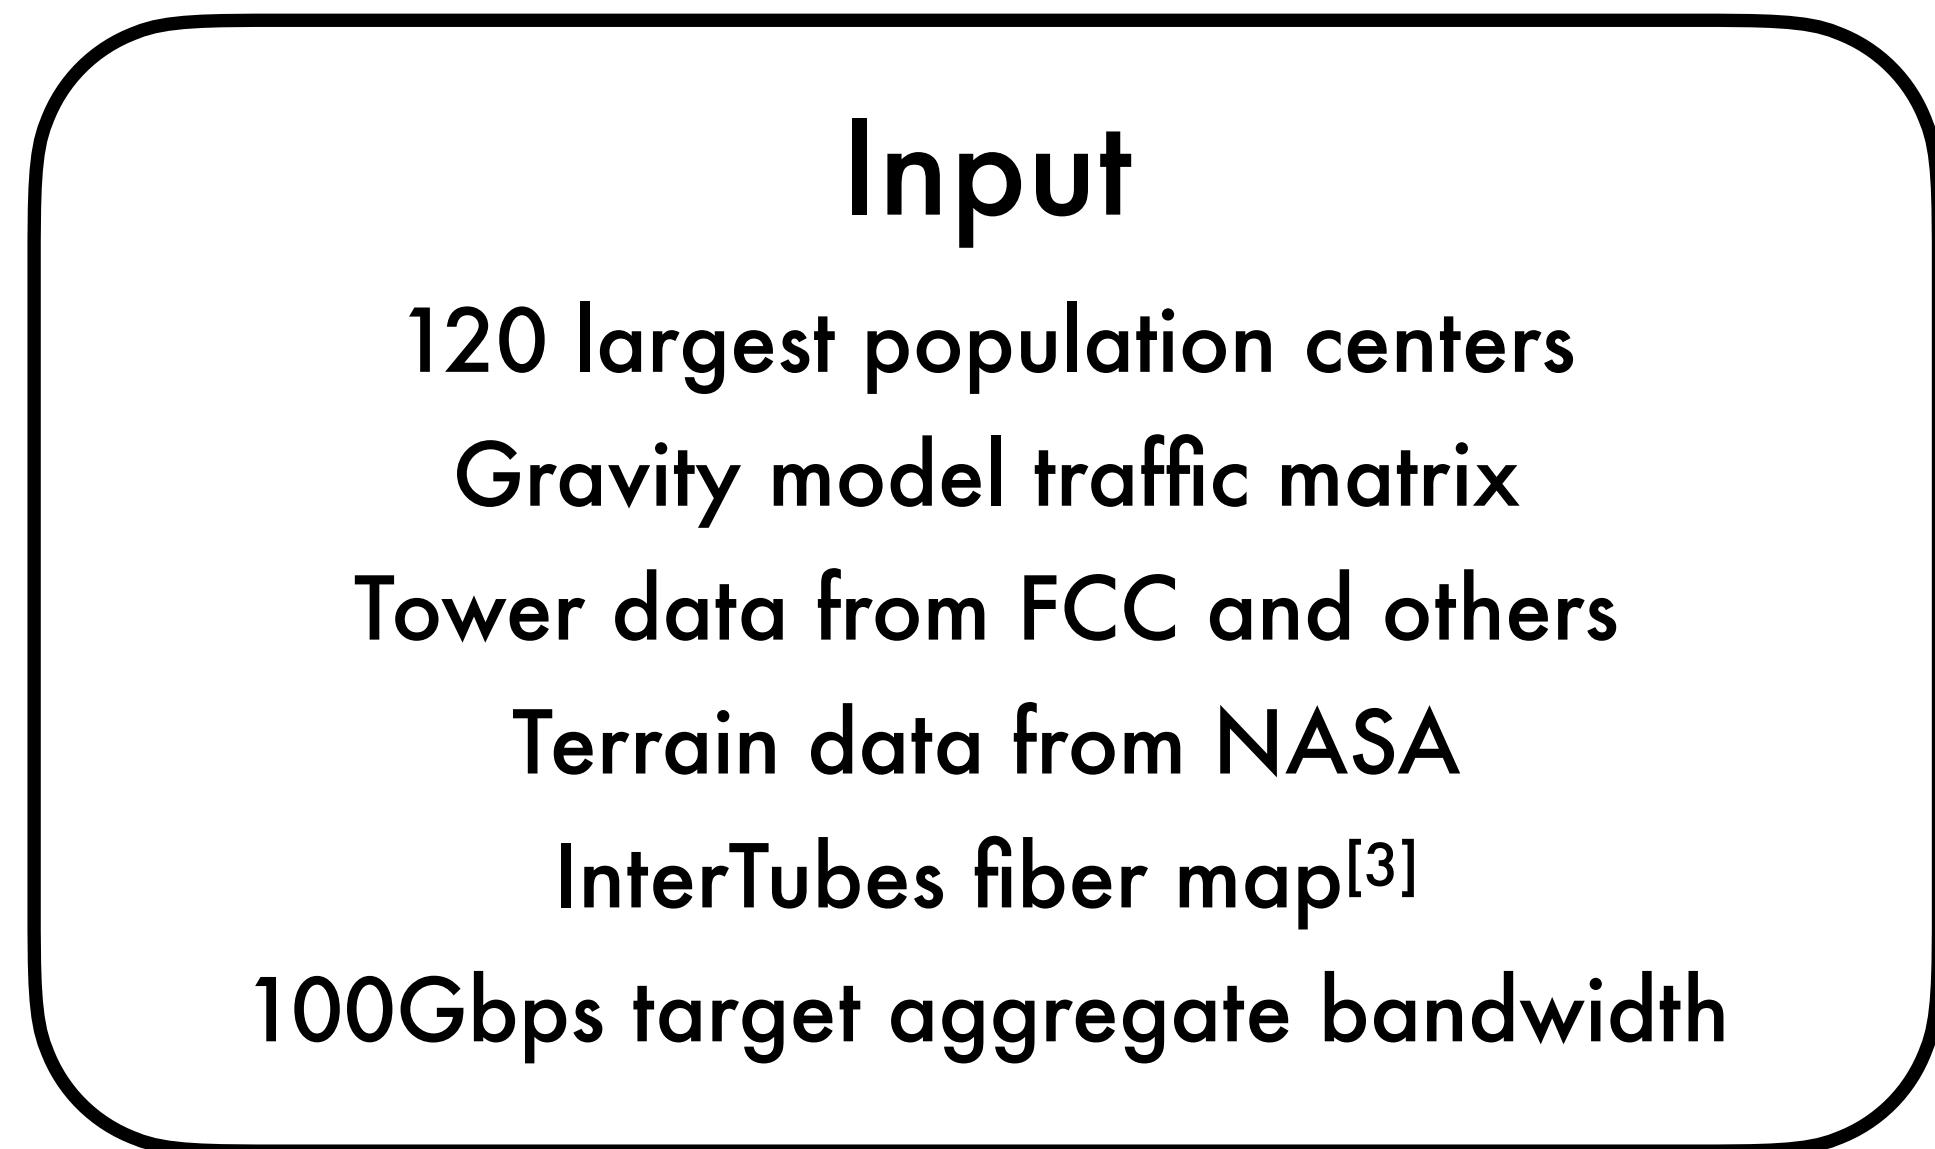
Optimizer

$$\min \sum_{s,t} \frac{h_{st}}{d_{st}} \sum_{i,j} (o_{i,j} f_{stij,o} + m_{i,j} f_{stij,m})$$

[3] Durairajan et al. "InterTubes: A study of the US long-haul fiber-optic infrastructure." ACM SIGCOMM. 2015.

cISP for the United States

Budget = 3000 towers

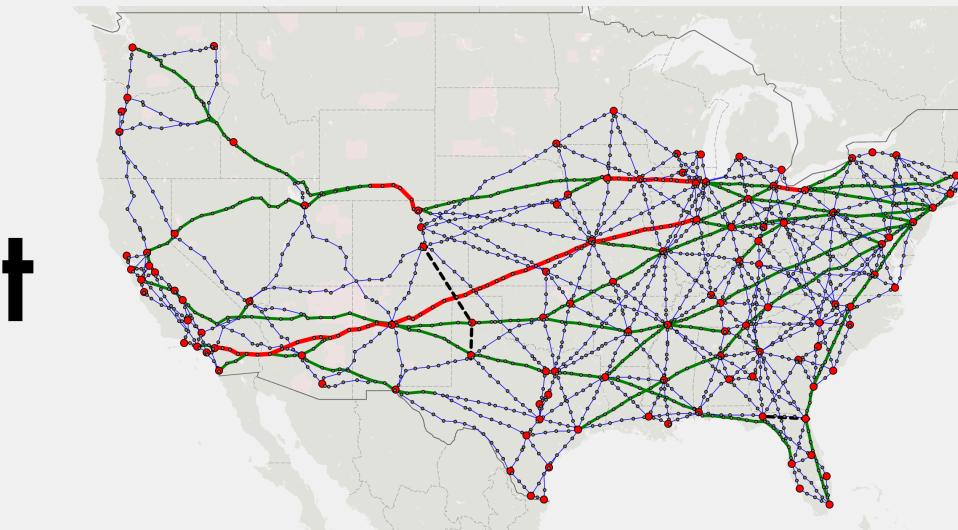


Identify feasible hops
Select shortest links

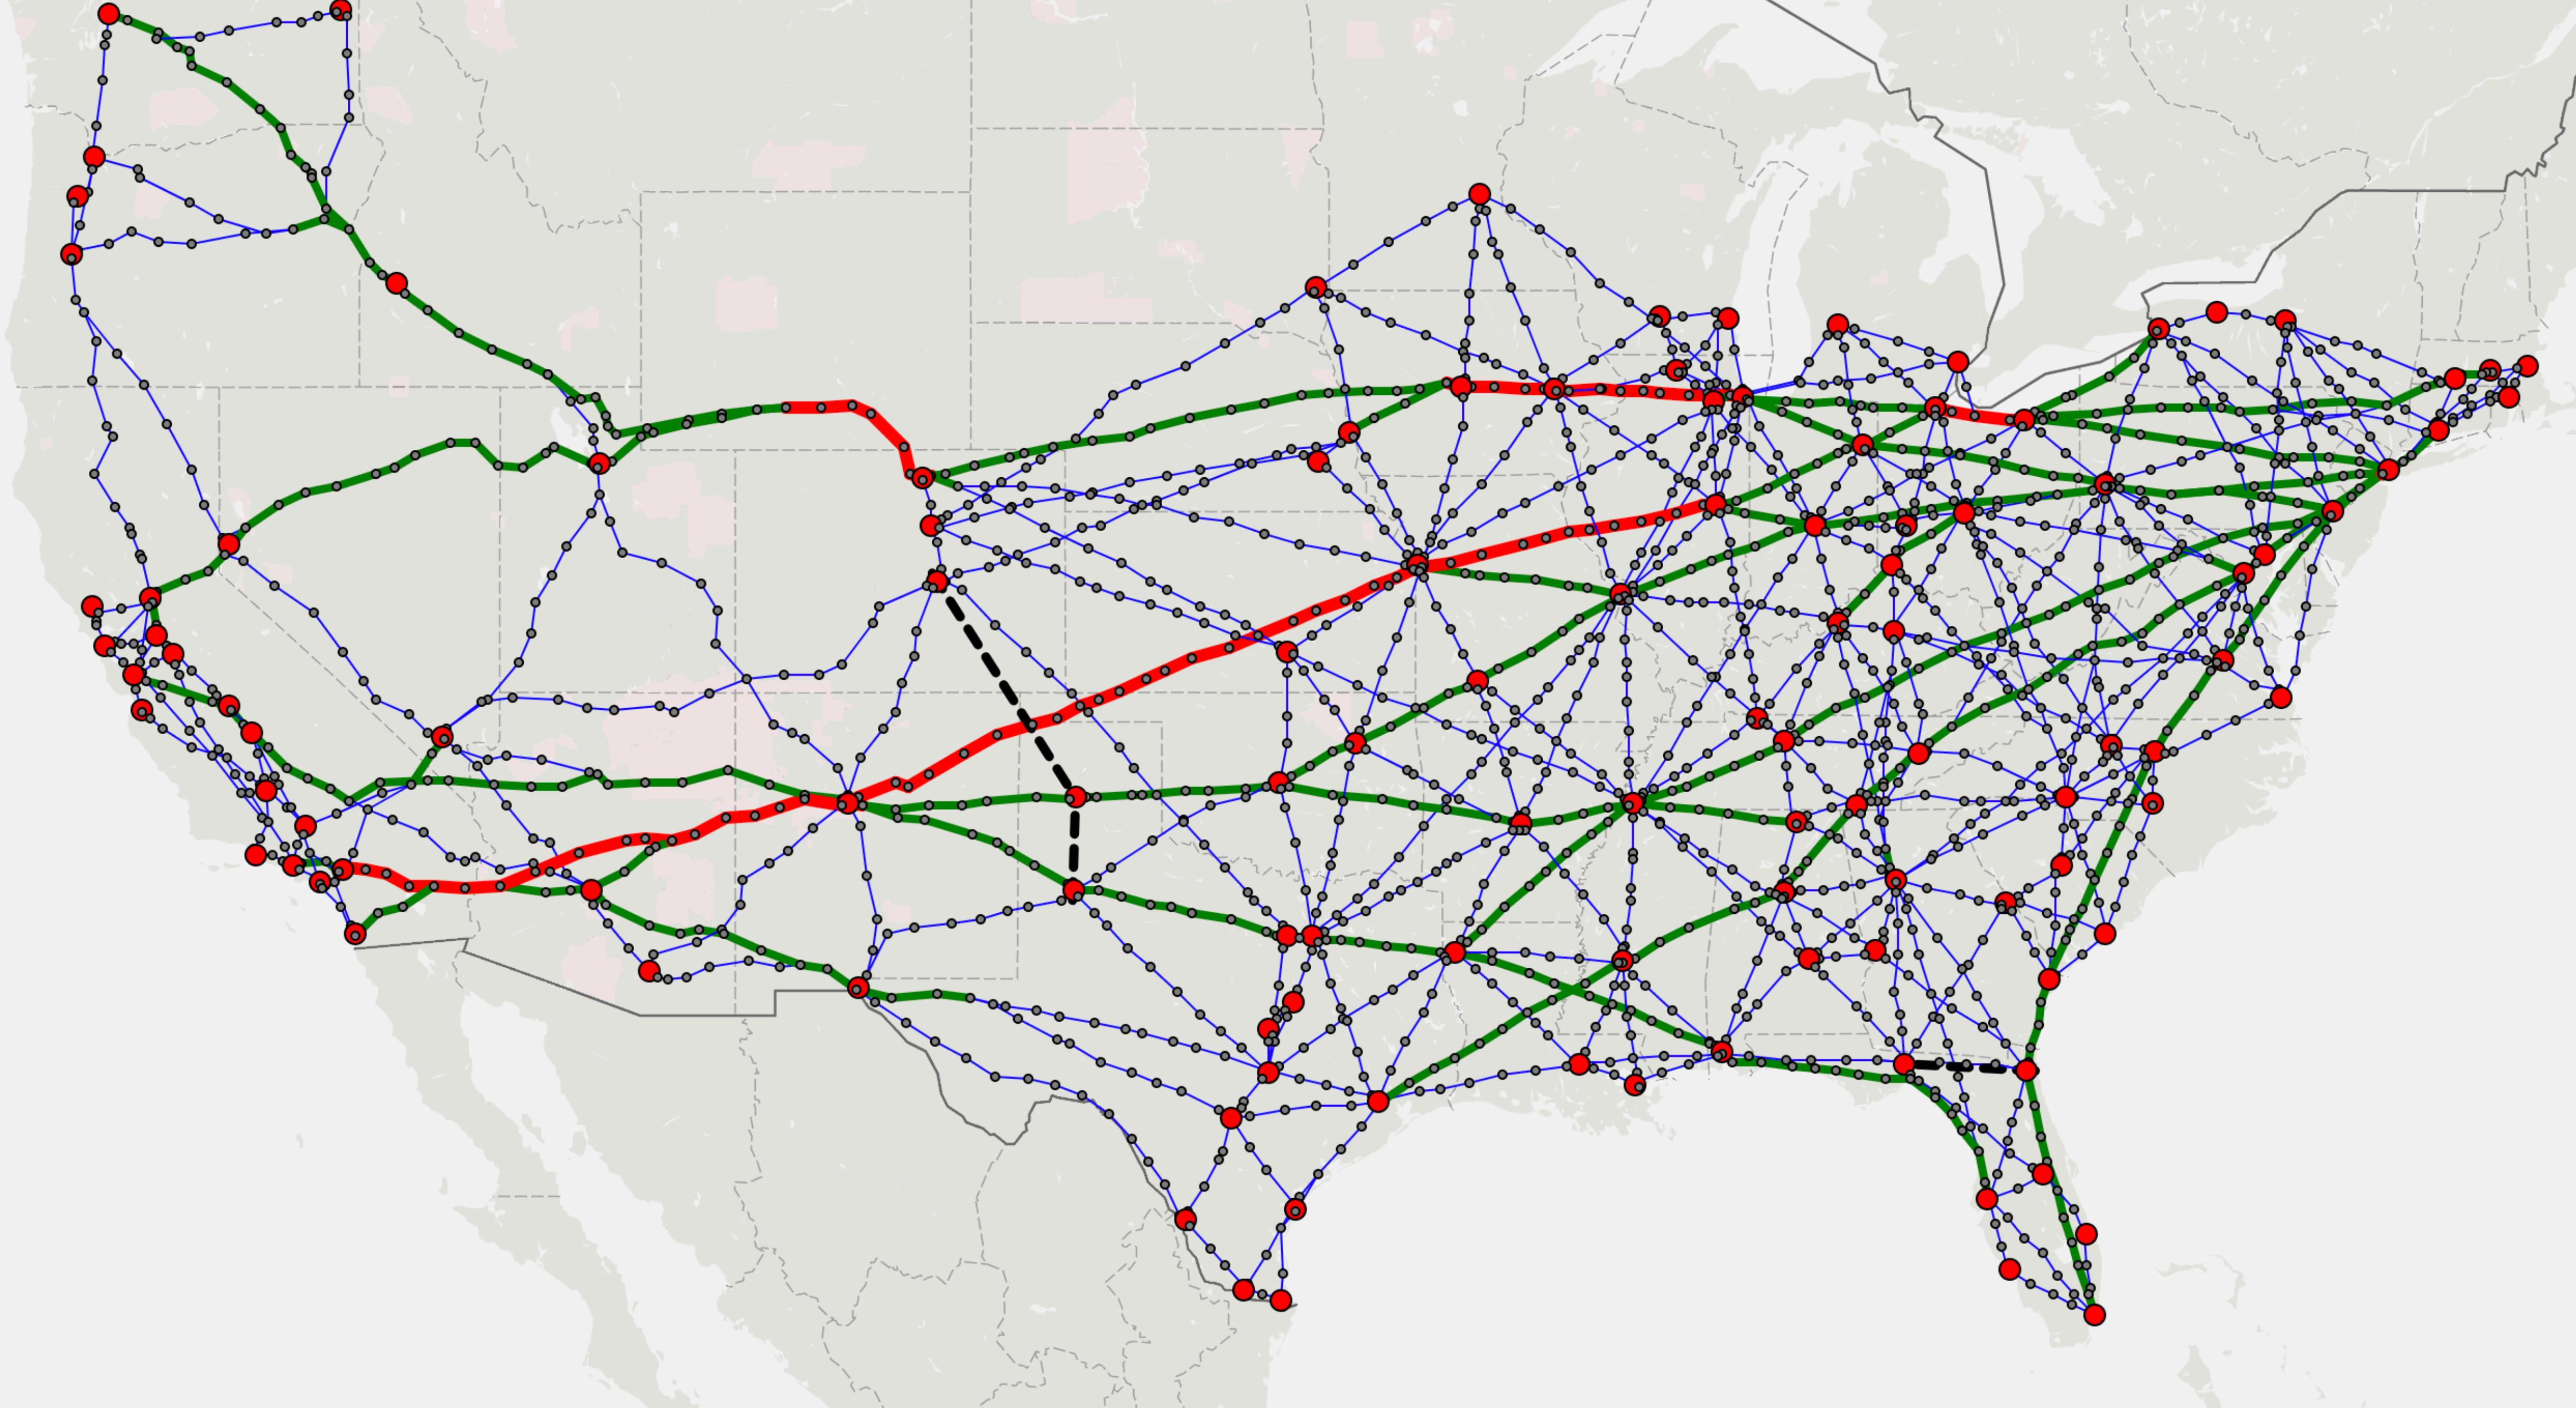
Optimizer

$$\min \sum_{s,t} \frac{h_{st}}{d_{st}} \sum_{i,j} (o_{i,j} f_{stij,o} + m_{i,j} f_{stij,m})$$

Output



[3] Durairajan et al. "InterTubes: A study of the US long-haul fiber-optic infrastructure." ACM SIGCOMM. 2015.

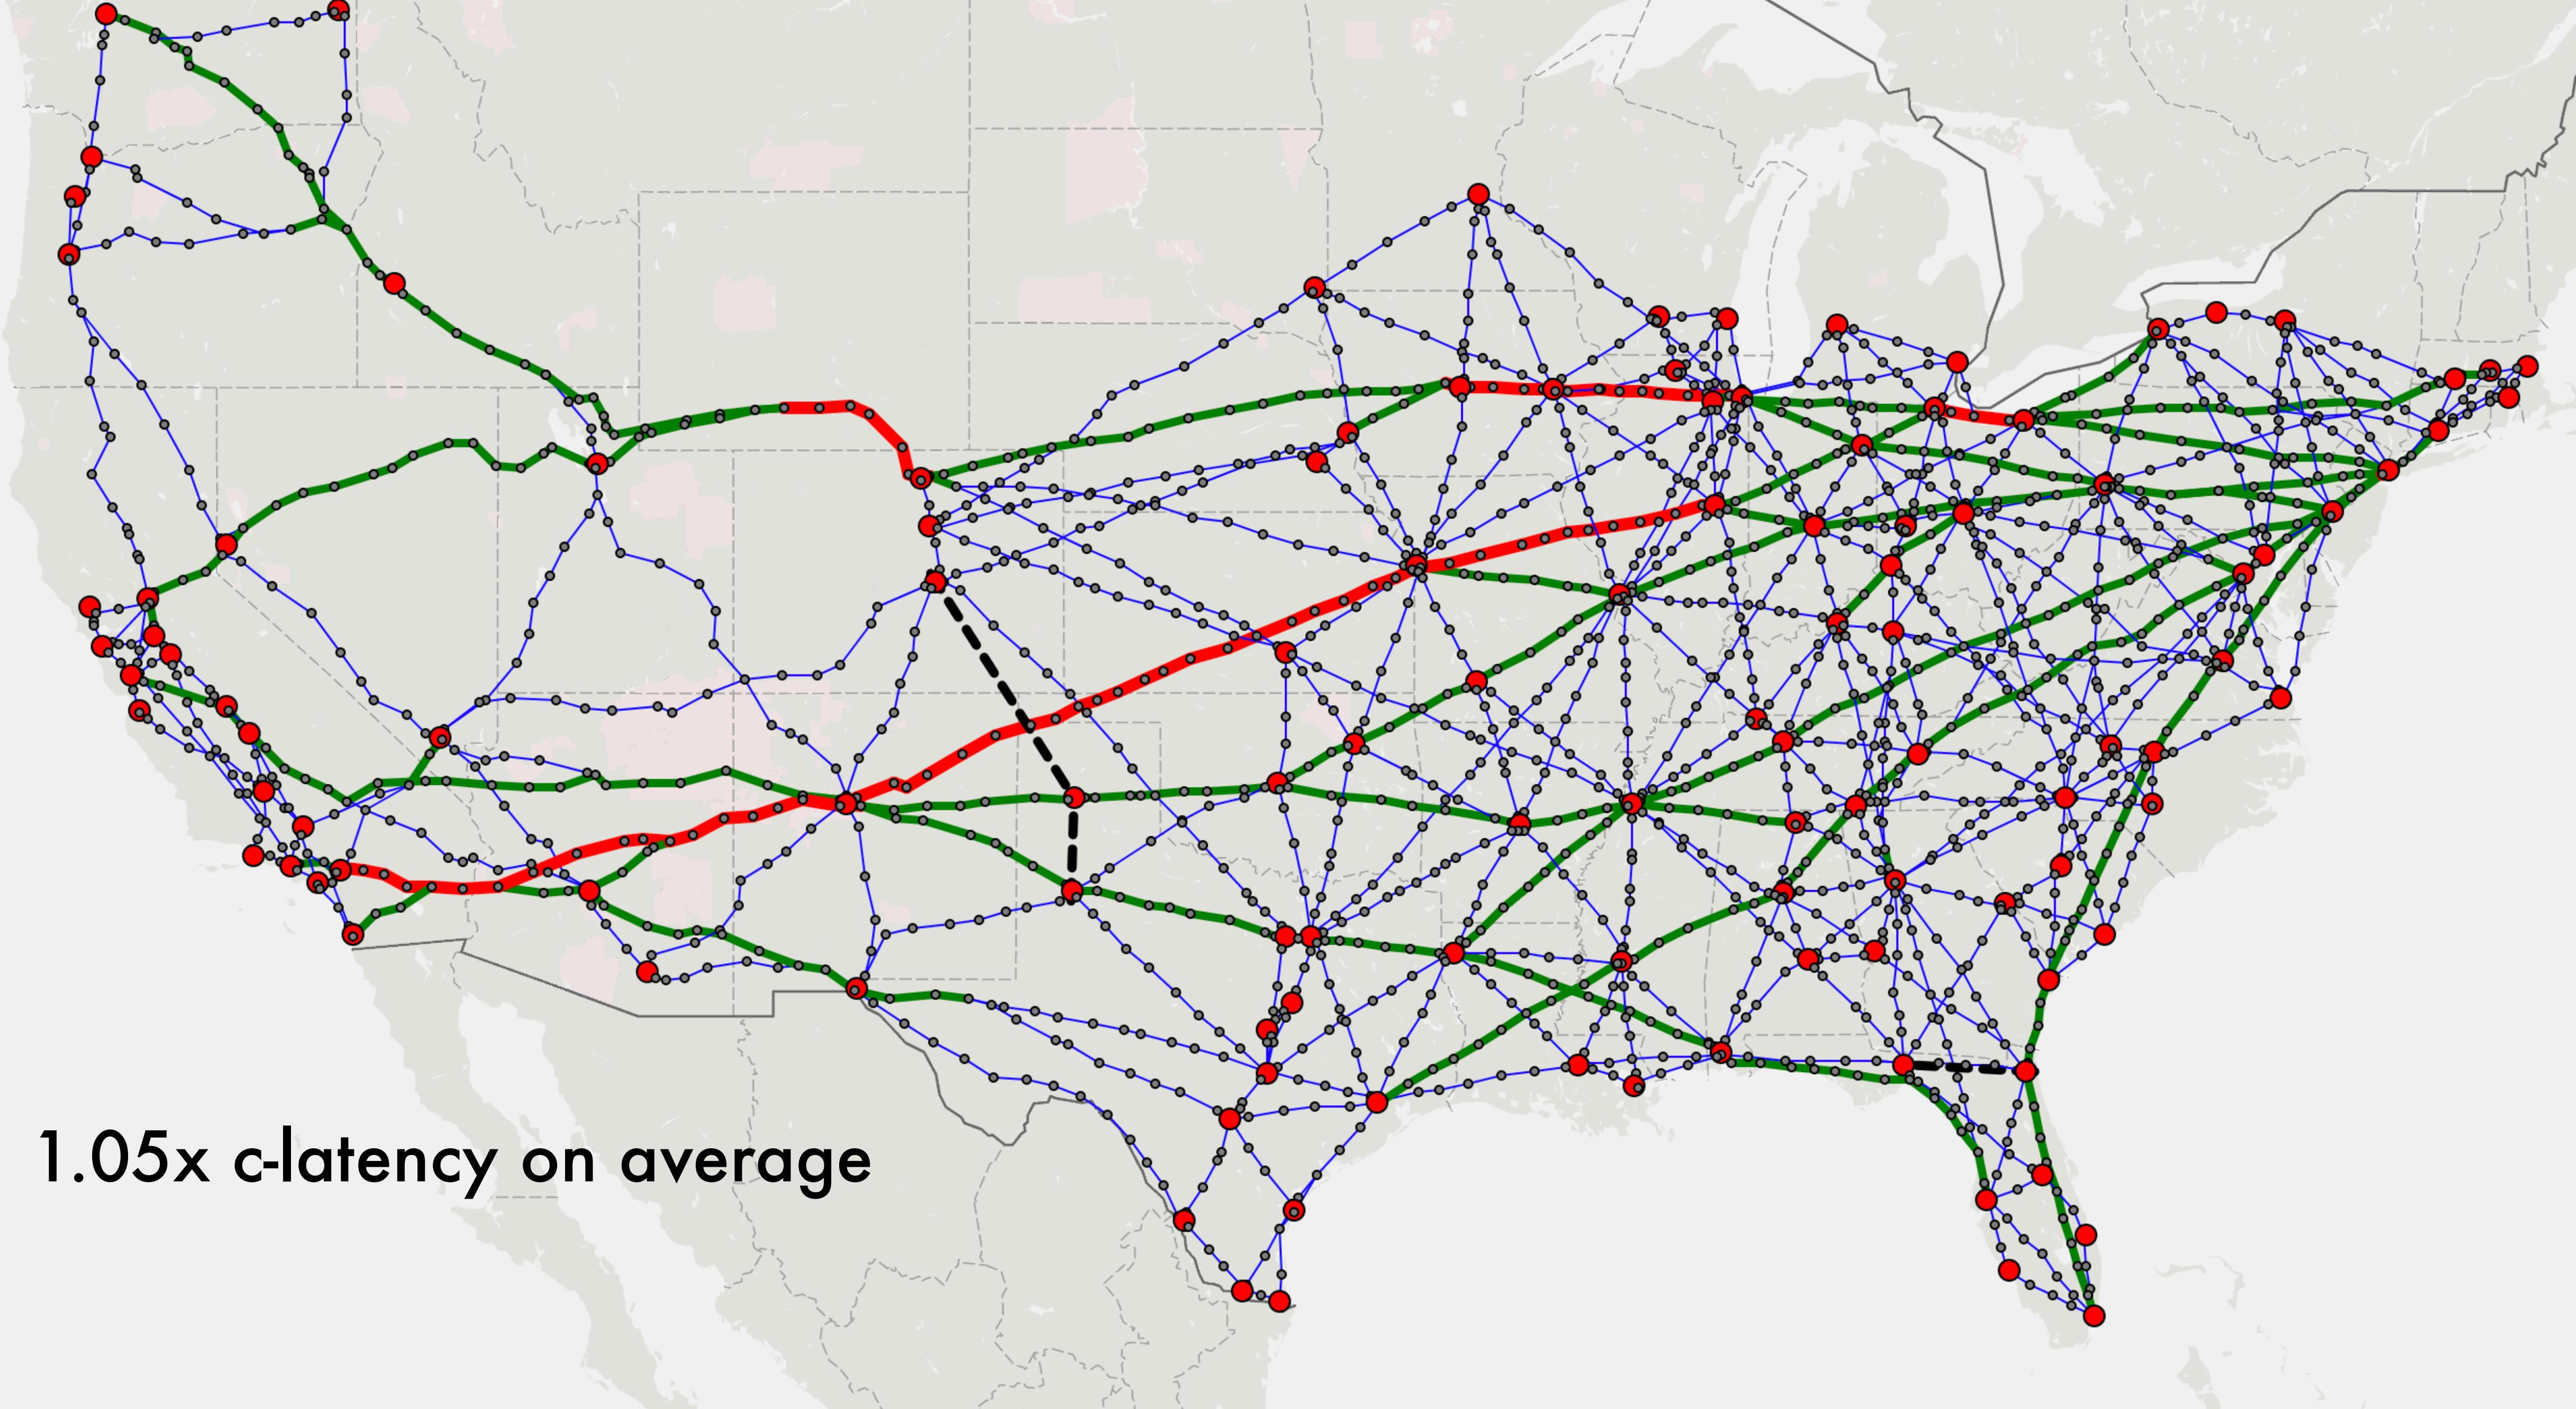


1Gb/s link

4Gb/s link

9Gb/s link

fiber

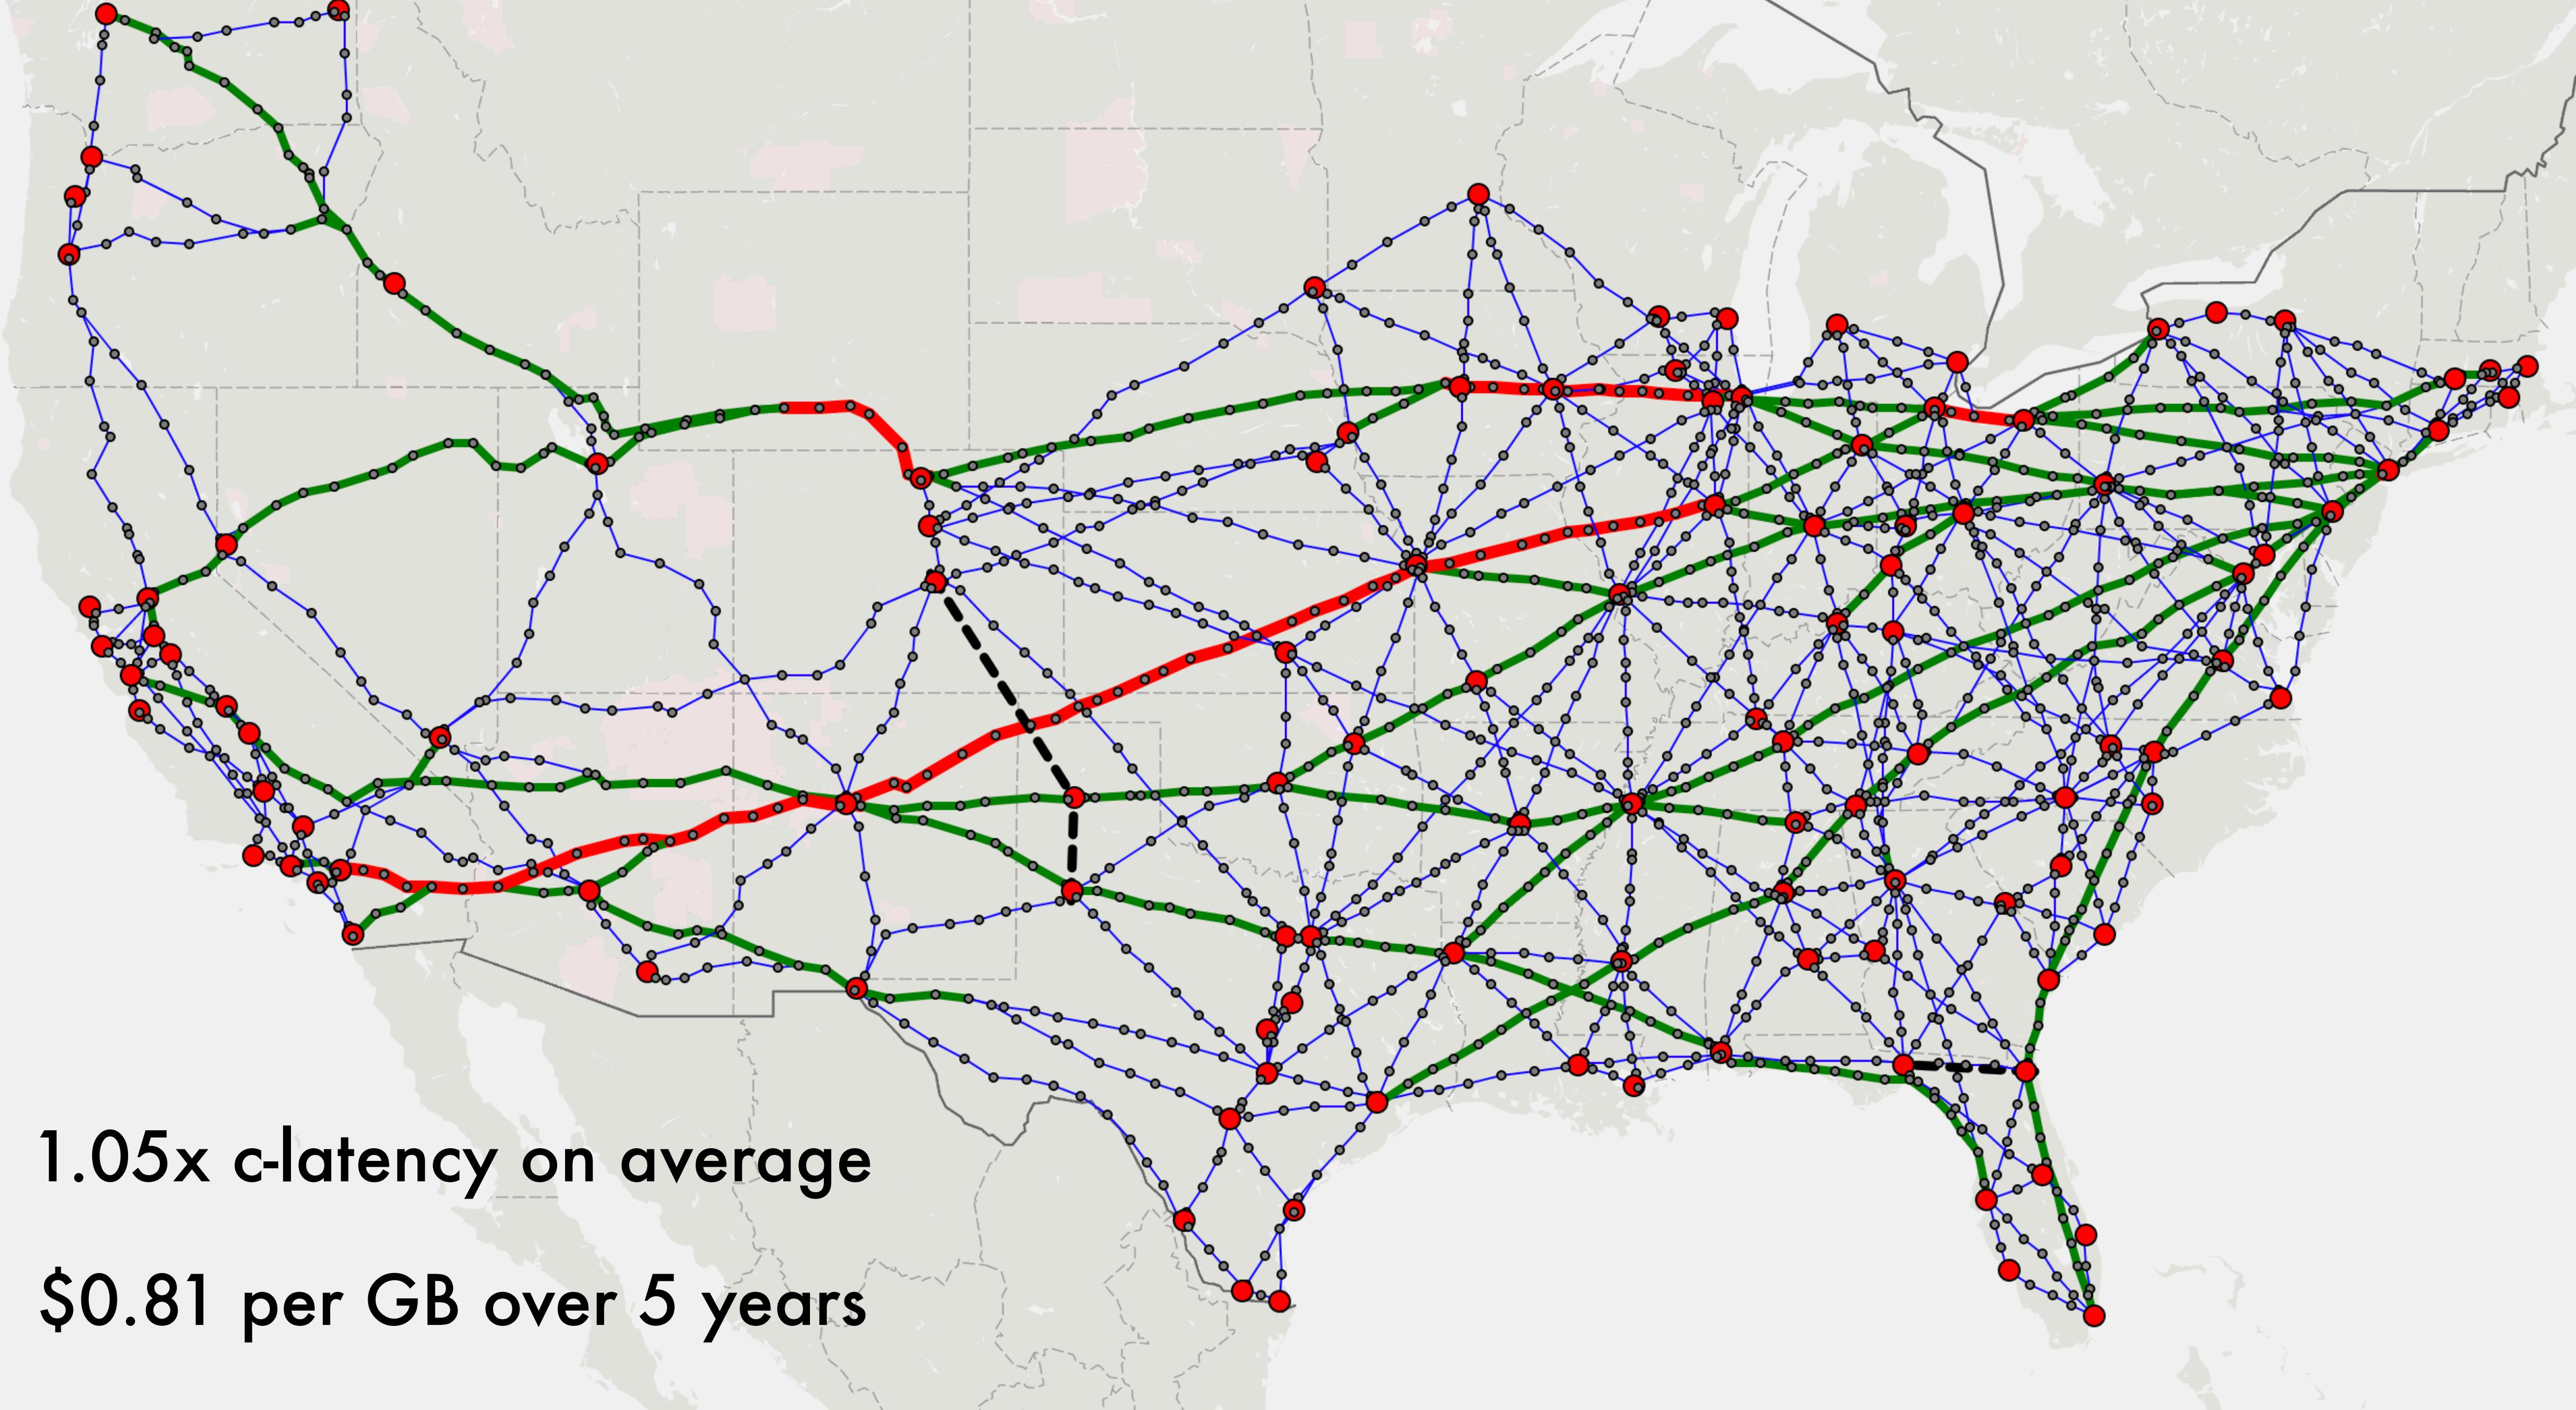


— 1Gb/s link

— 4Gb/s link

— 9Gb/s link

---- fiber



1.05x c-latency on average

\$0.81 per GB over 5 years

— 1Gb/s link

— 4Gb/s link

— 9Gb/s link

---- fiber

But wait, will this really work in practice?

How does a real MW network perform?

McKay Brothers network performance

How does a real MW network perform?

McKay Brothers network performance

- Machines connected over 1140 km long microwave network

How does a real MW network perform?

McKay Brothers network performance

- Machines connected over 1140 km long microwave network
- Round-trip time of 7.7ms (1.3% inflation at c-latency=7.6ms)

How does a real MW network perform?

McKay Brothers network performance

- Machines connected over 1140 km long microwave network
- Round-trip time of 7.7ms (1.3% inflation at c-latency=7.6ms)
- At least 120Mbps Bandwidth

How does a real MW network perform?

McKay Brothers network performance

- Machines connected over 1140 km long microwave network
- Round-trip time of 7.7ms (1.3% inflation at c-latency=7.6ms)
- At least 120Mbps Bandwidth
- Bit errors and packet loss
 - High packet loss rate with no error correction

How does a real MW network perform?

McKay Brothers network performance

- Machines connected over 1140 km long microwave network
- Round-trip time of 7.7ms (1.3% inflation at c-latency=7.6ms)
- At least 120Mbps Bandwidth
- Bit errors and packet loss
 - High packet loss rate with no error correction
 - Median underlying bit error rate of 10^{-5}

How does a real MW network perform?

McKay Brothers network performance

- Machines connected over 1140 km long microwave network
- Round-trip time of 7.7ms (1.3% inflation at c-latency=7.6ms)
- At least 120Mbps Bandwidth
- Bit errors and packet loss
 - High packet loss rate with no error correction
 - Median underlying bit error rate of 10^{-5}
 - Can be corrected with lightweight FEC

Is cISP resilient to weather?

Will rain bring cISP down?

Is cISP resilient to weather?

Will rain bring cISP down?

- Rain causes the most attenuation at 11GHz

Is cISP resilient to weather?

Will rain bring cISP down?

- Rain causes the most attenuation at 11GHz
- Consider link down when attenuation exceeds 1.5 dB/km

Is cISP resilient to weather?

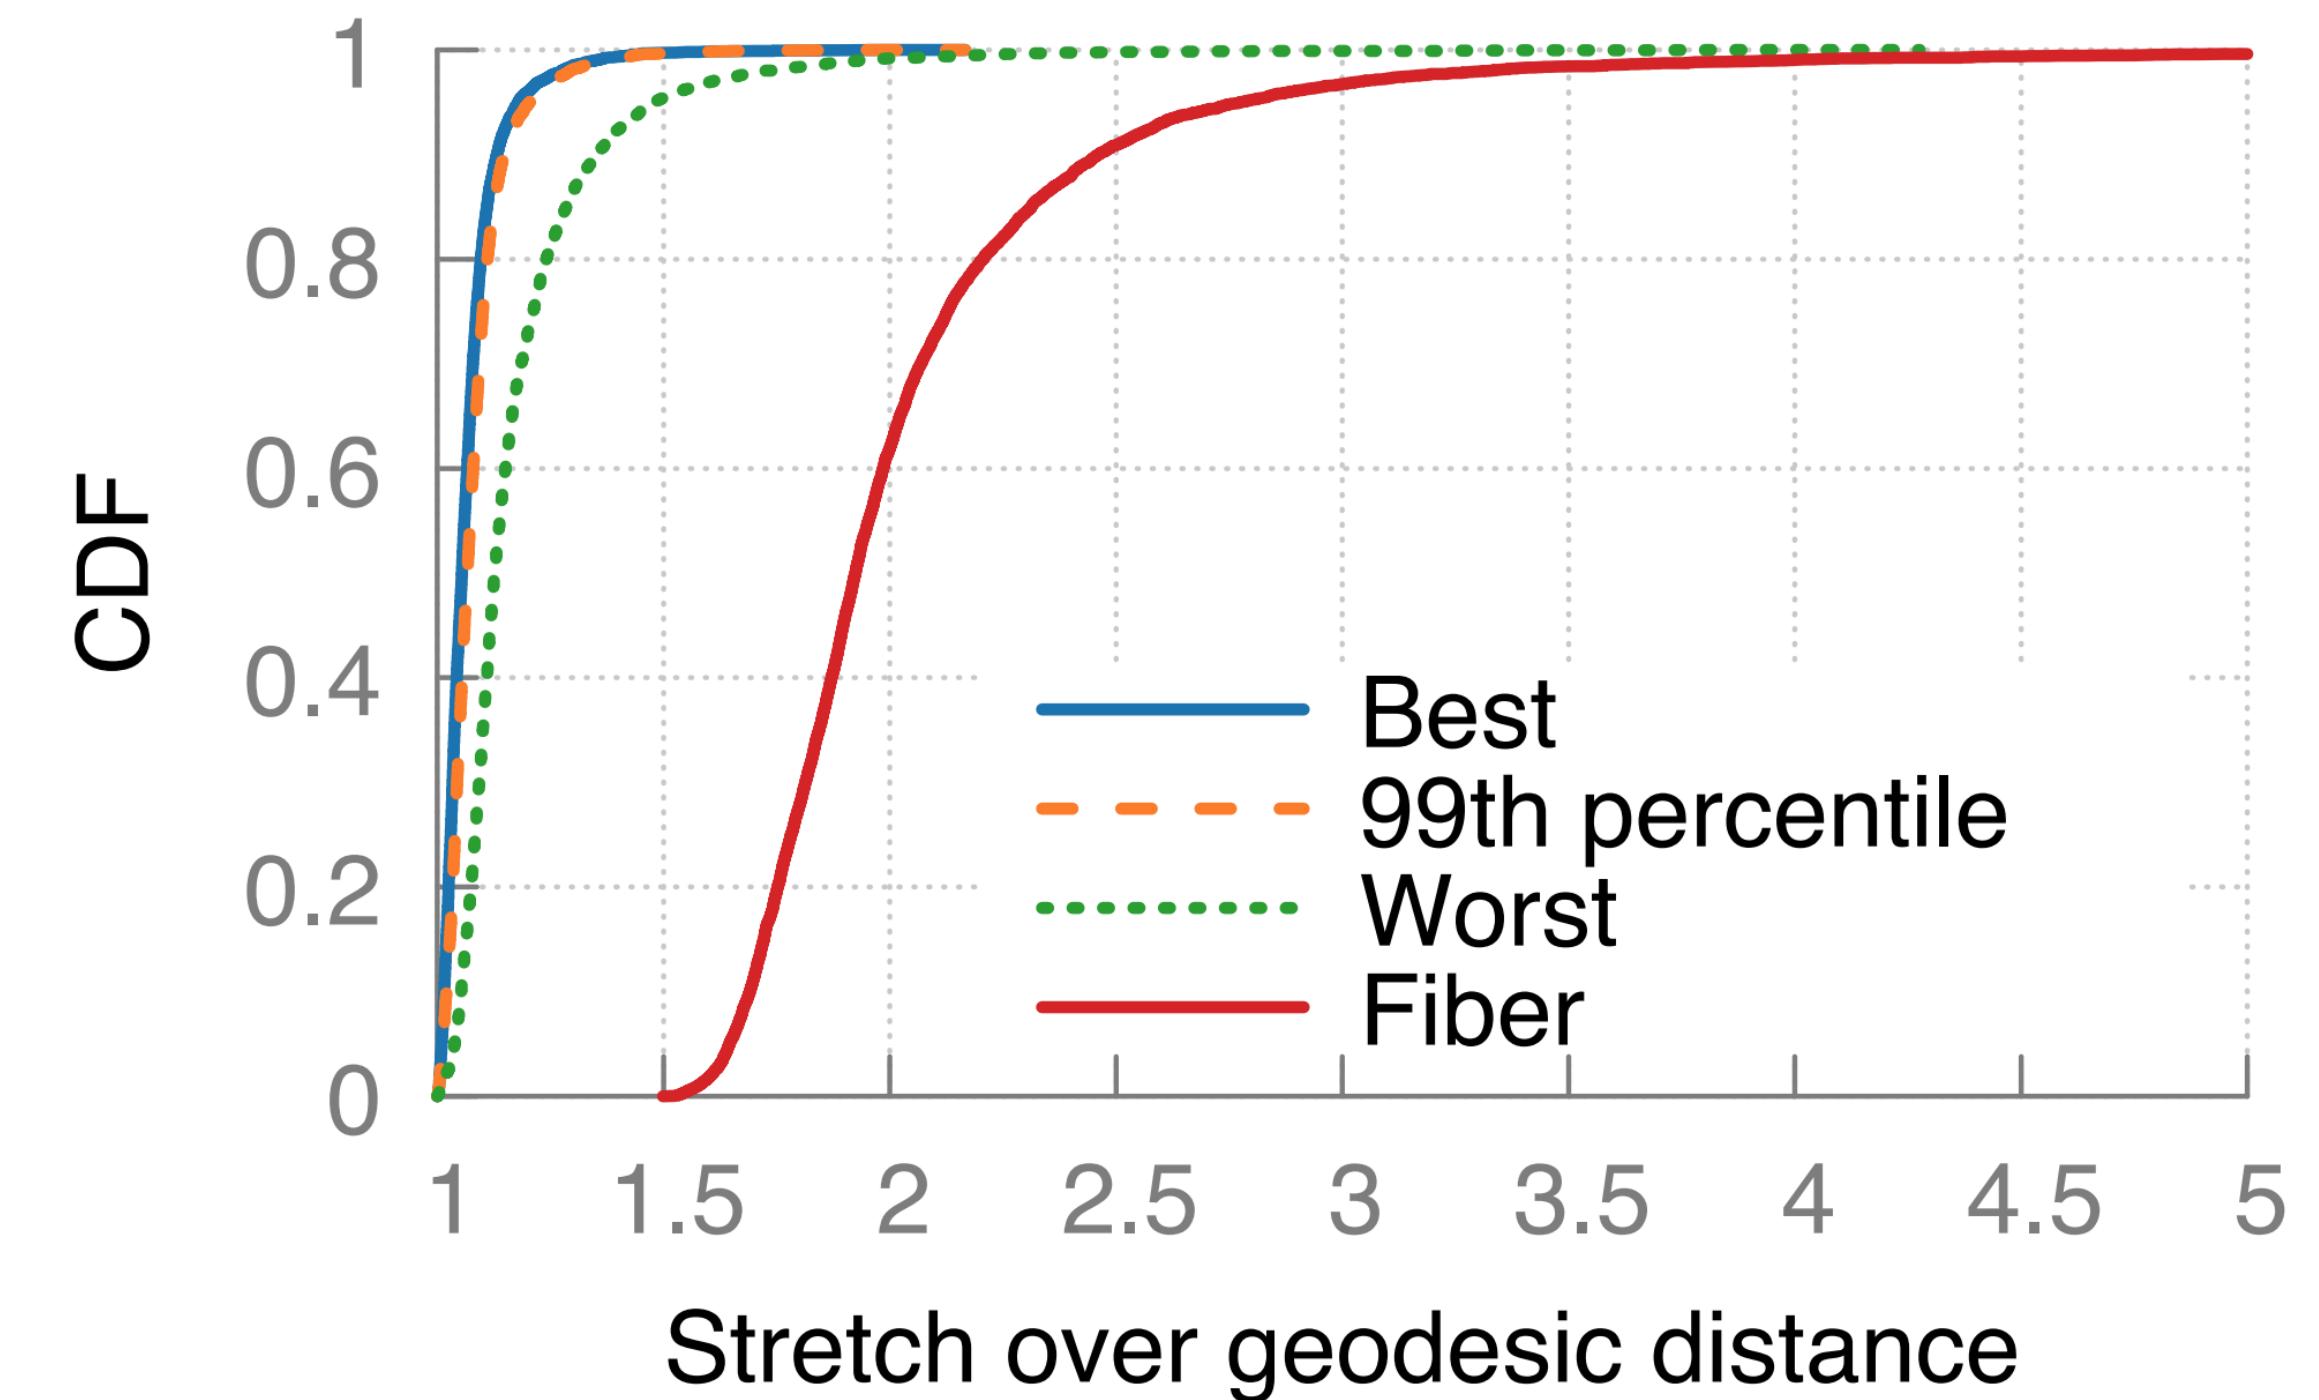
Will rain bring cISP down?

- Rain causes the most attenuation at 11GHz
- Consider link down when attenuation exceeds 1.5 dB/km
- 99th-percentile latencies remain same as best fair-weather latencies

Is cISP resilient to weather?

Will rain bring cISP down?

- Rain causes the most attenuation at 11GHz
- Consider link down when attenuation exceeds 1.5 dB/km
- 99th-percentile latencies remain same as best fair-weather latencies



Is MW resilient to weather in practice?

What the markets show

Is MW resilient to weather in practice?

What the markets show

- Events in Chicago (stimulus) trigger trades in New York (response)

Is MW resilient to weather in practice?

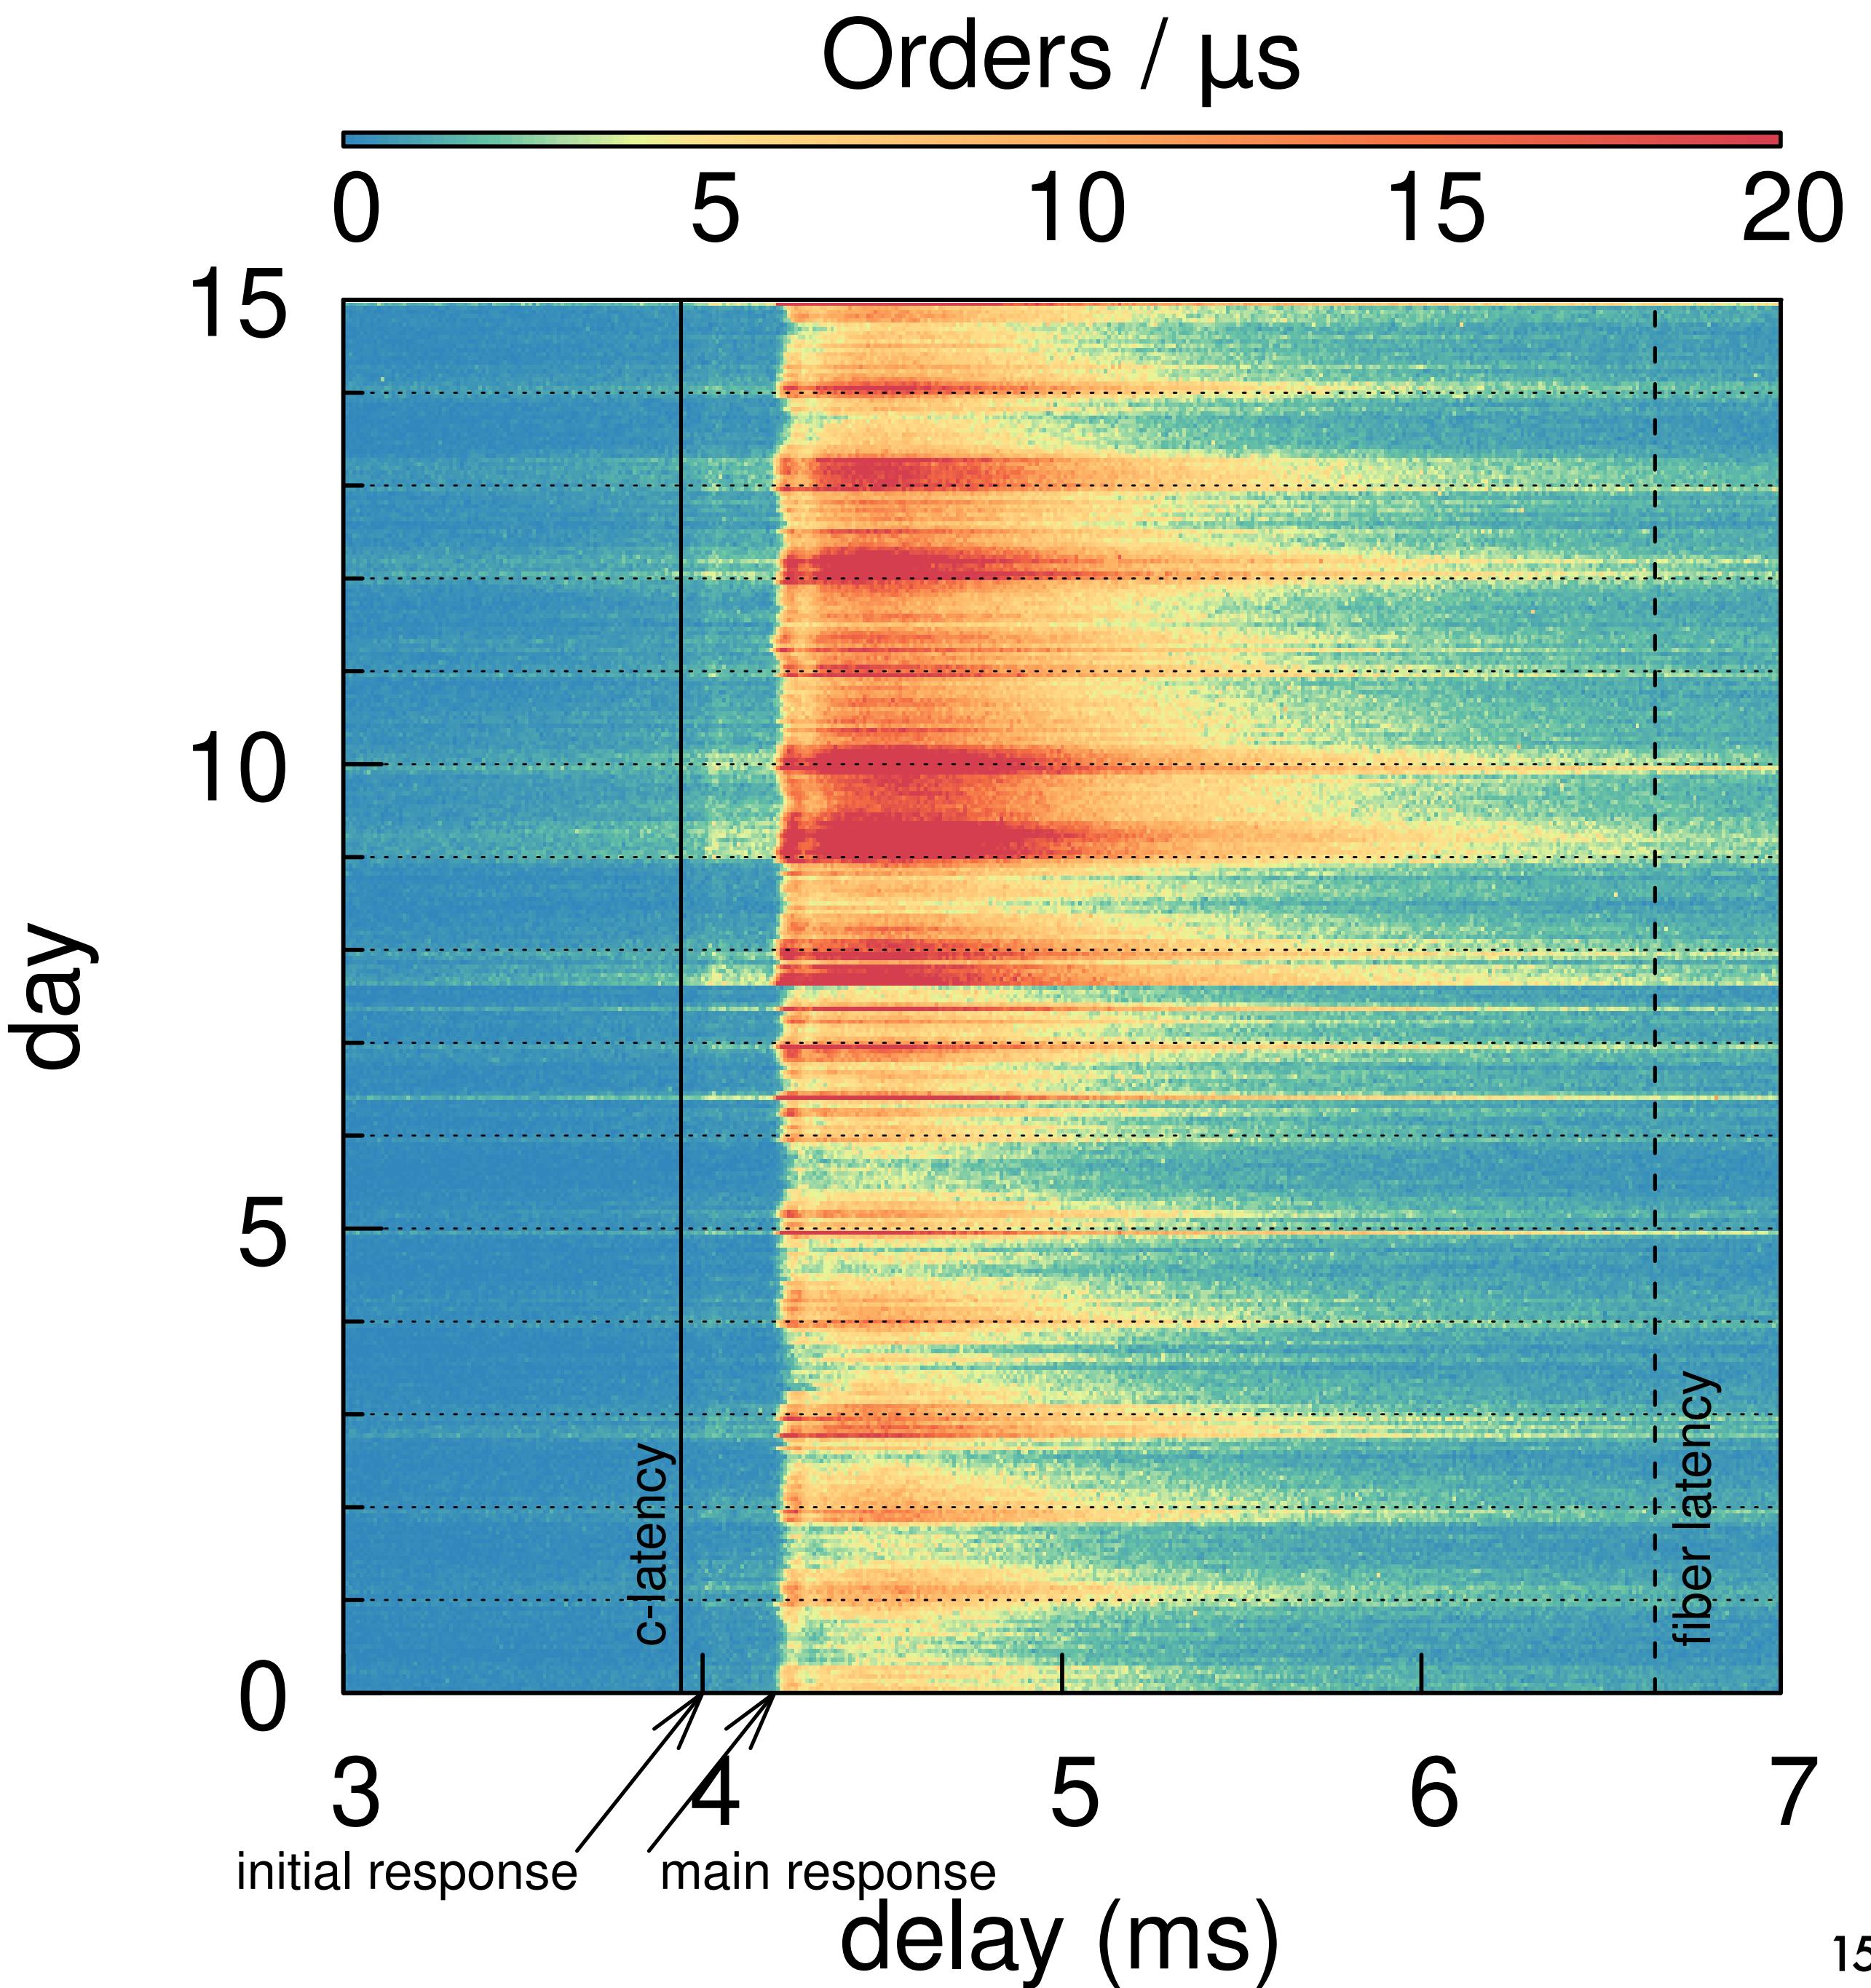
What the markets show

- Events in Chicago (stimulus) trigger trades in New York (response)
- Latency is lag between stimulus and response

Is MW resilient to weather in practice?

What the markets show

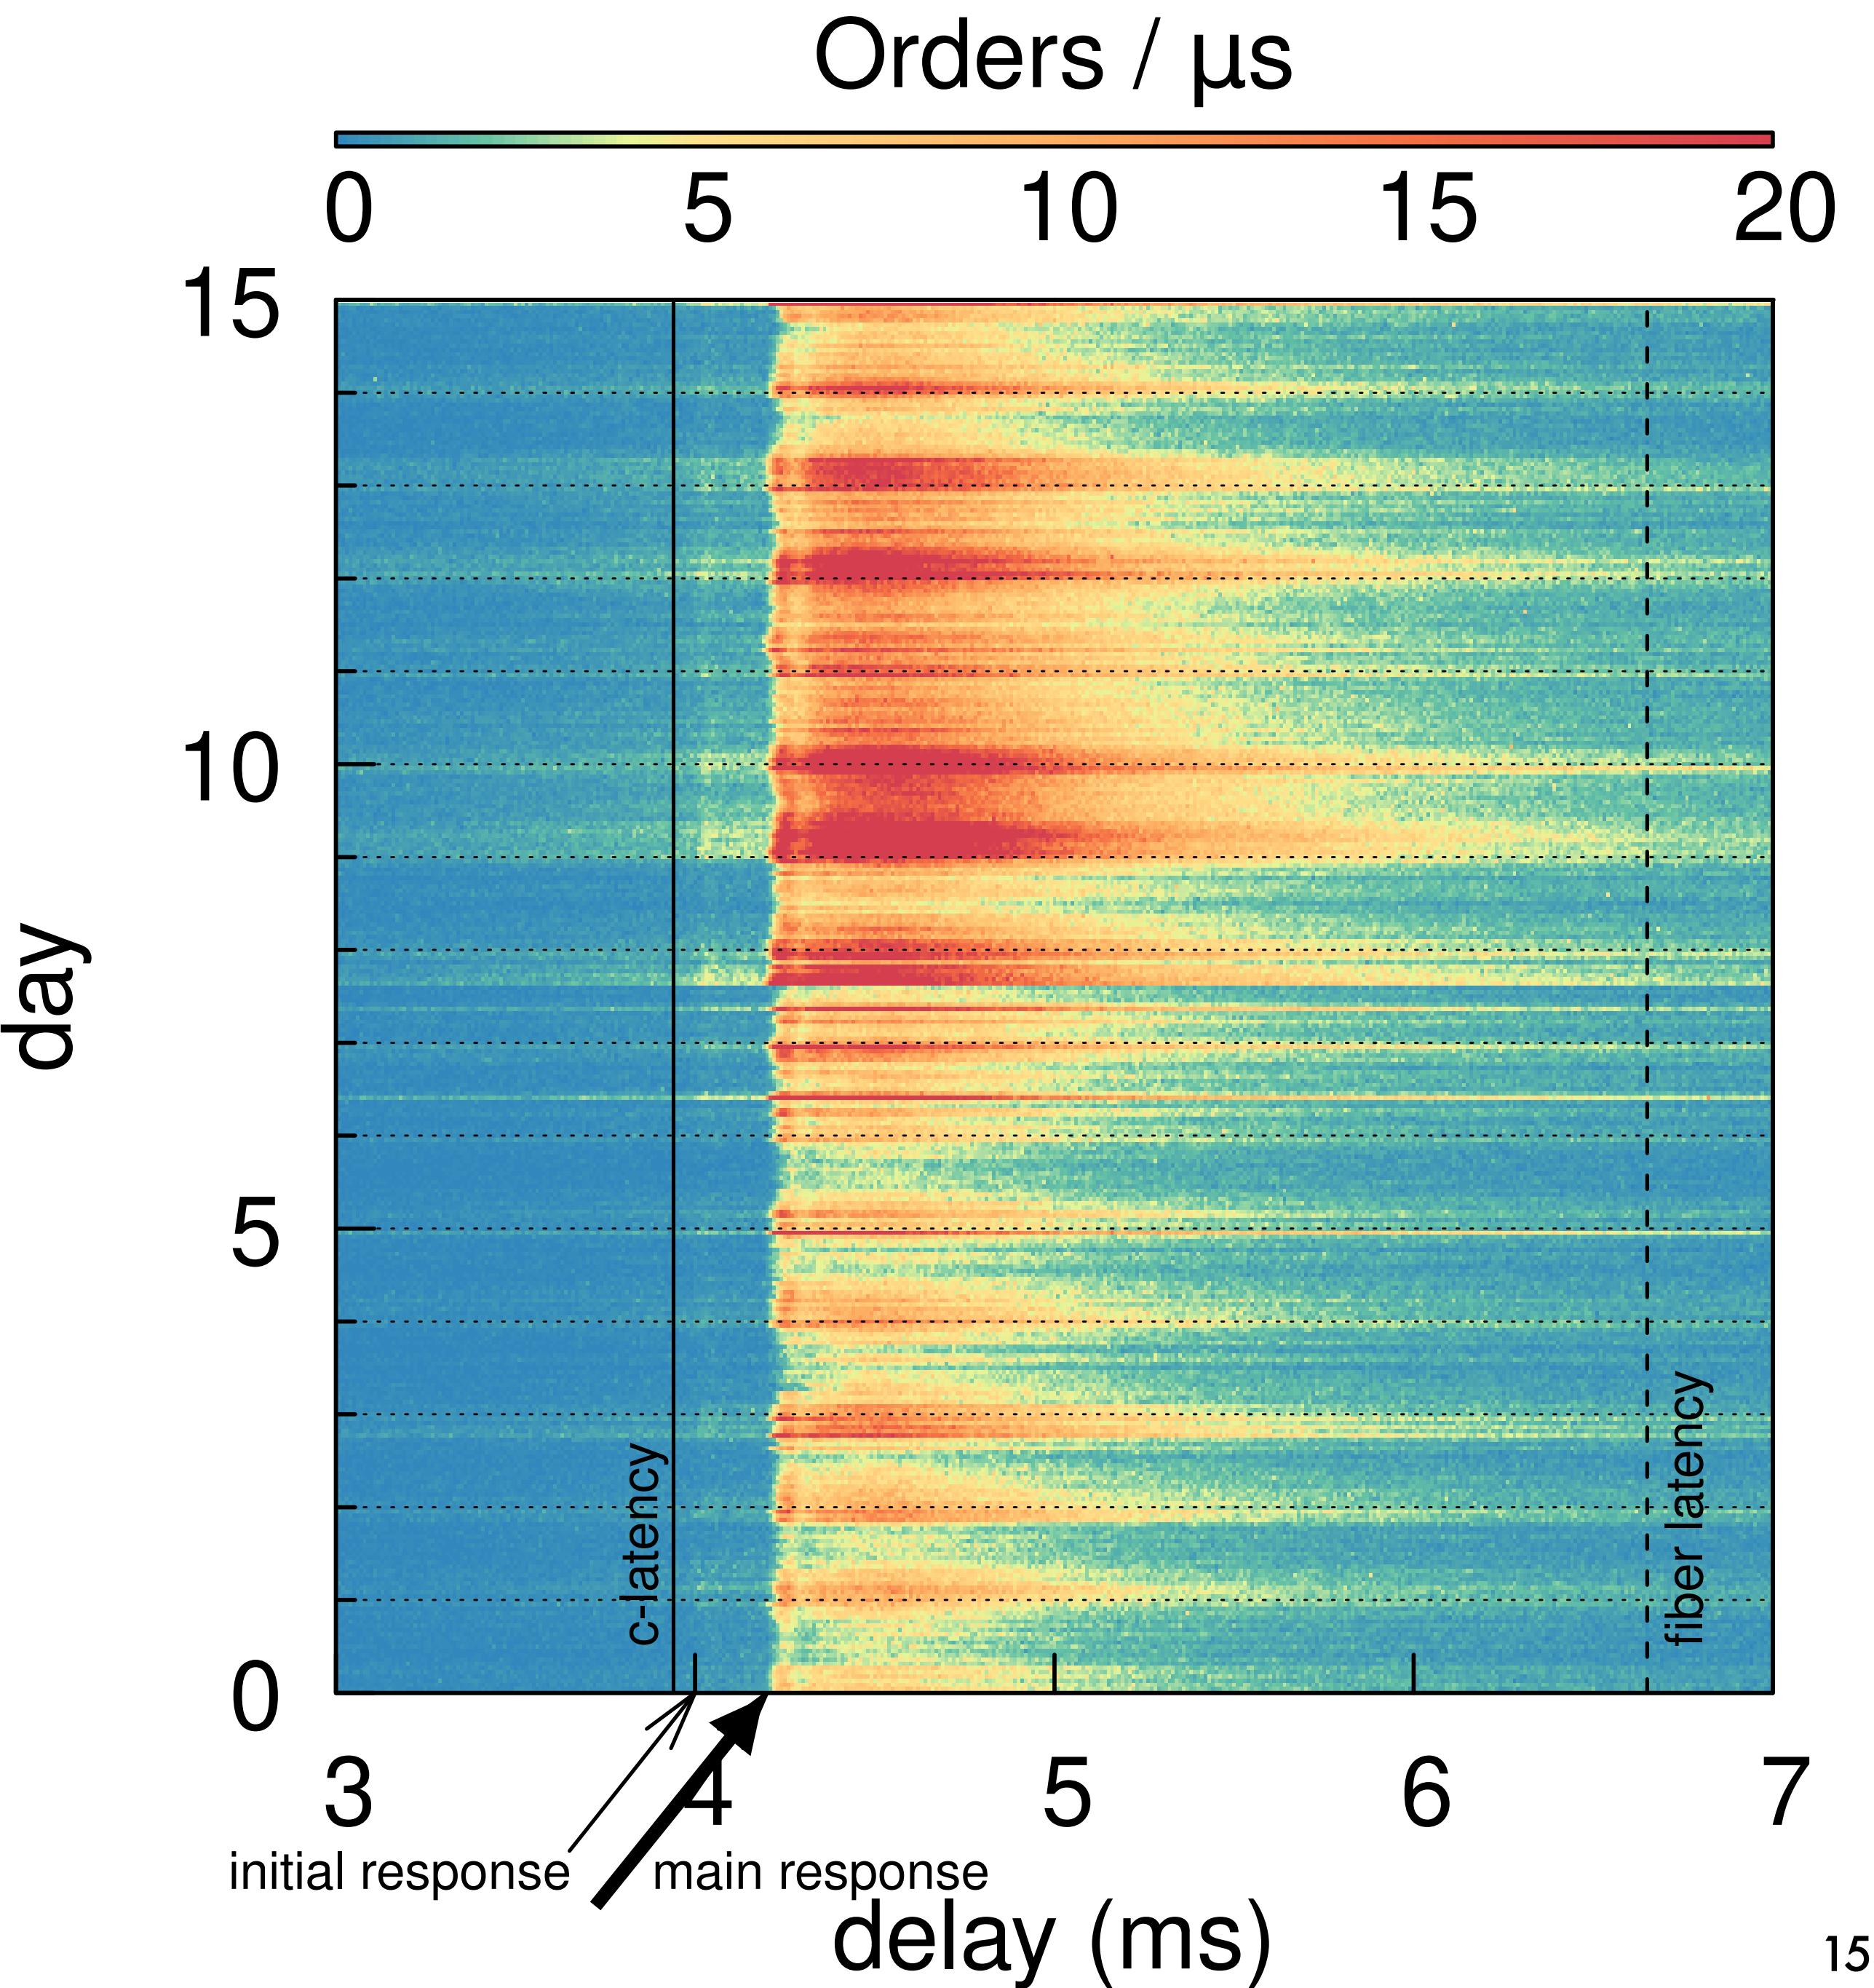
- Events in Chicago (stimulus) trigger trades in New York (response)
- Latency is lag between stimulus and response



Is MW resilient to weather in practice?

What the markets show

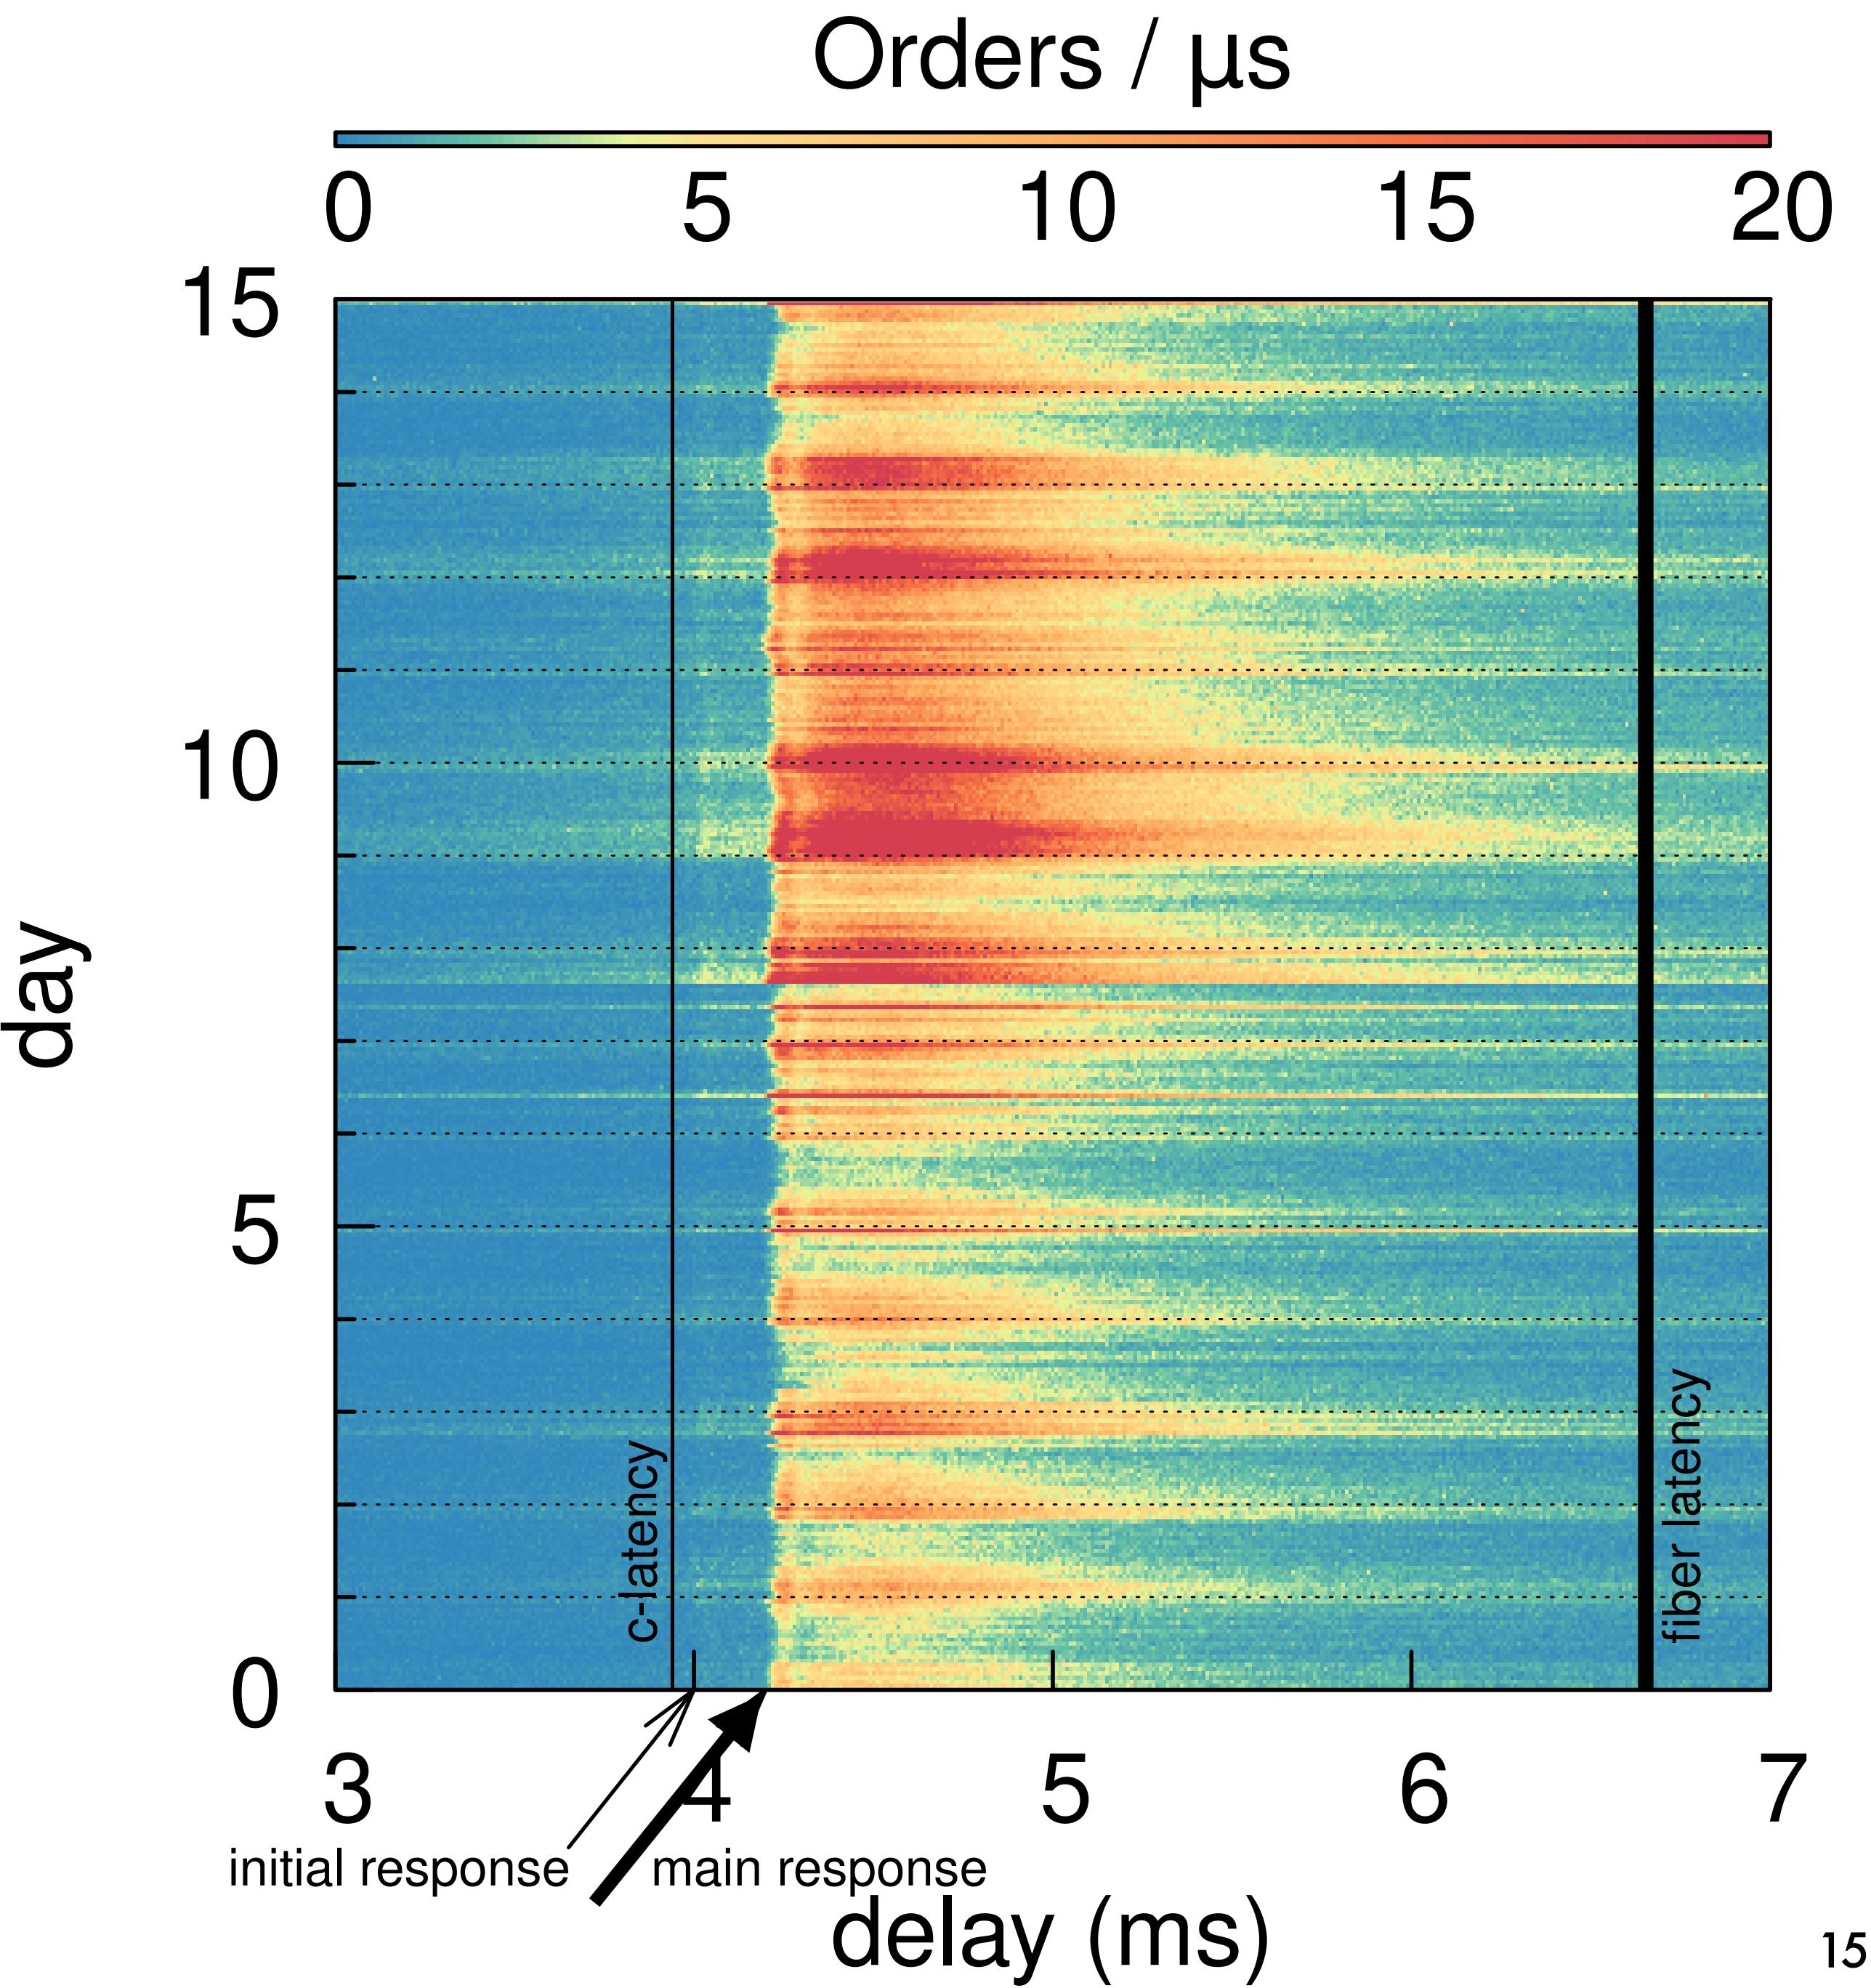
- Events in Chicago (stimulus) trigger trades in New York (response)
 - Latency is lag between stimulus and response



Is MW resilient to weather in practice?

What the markets show

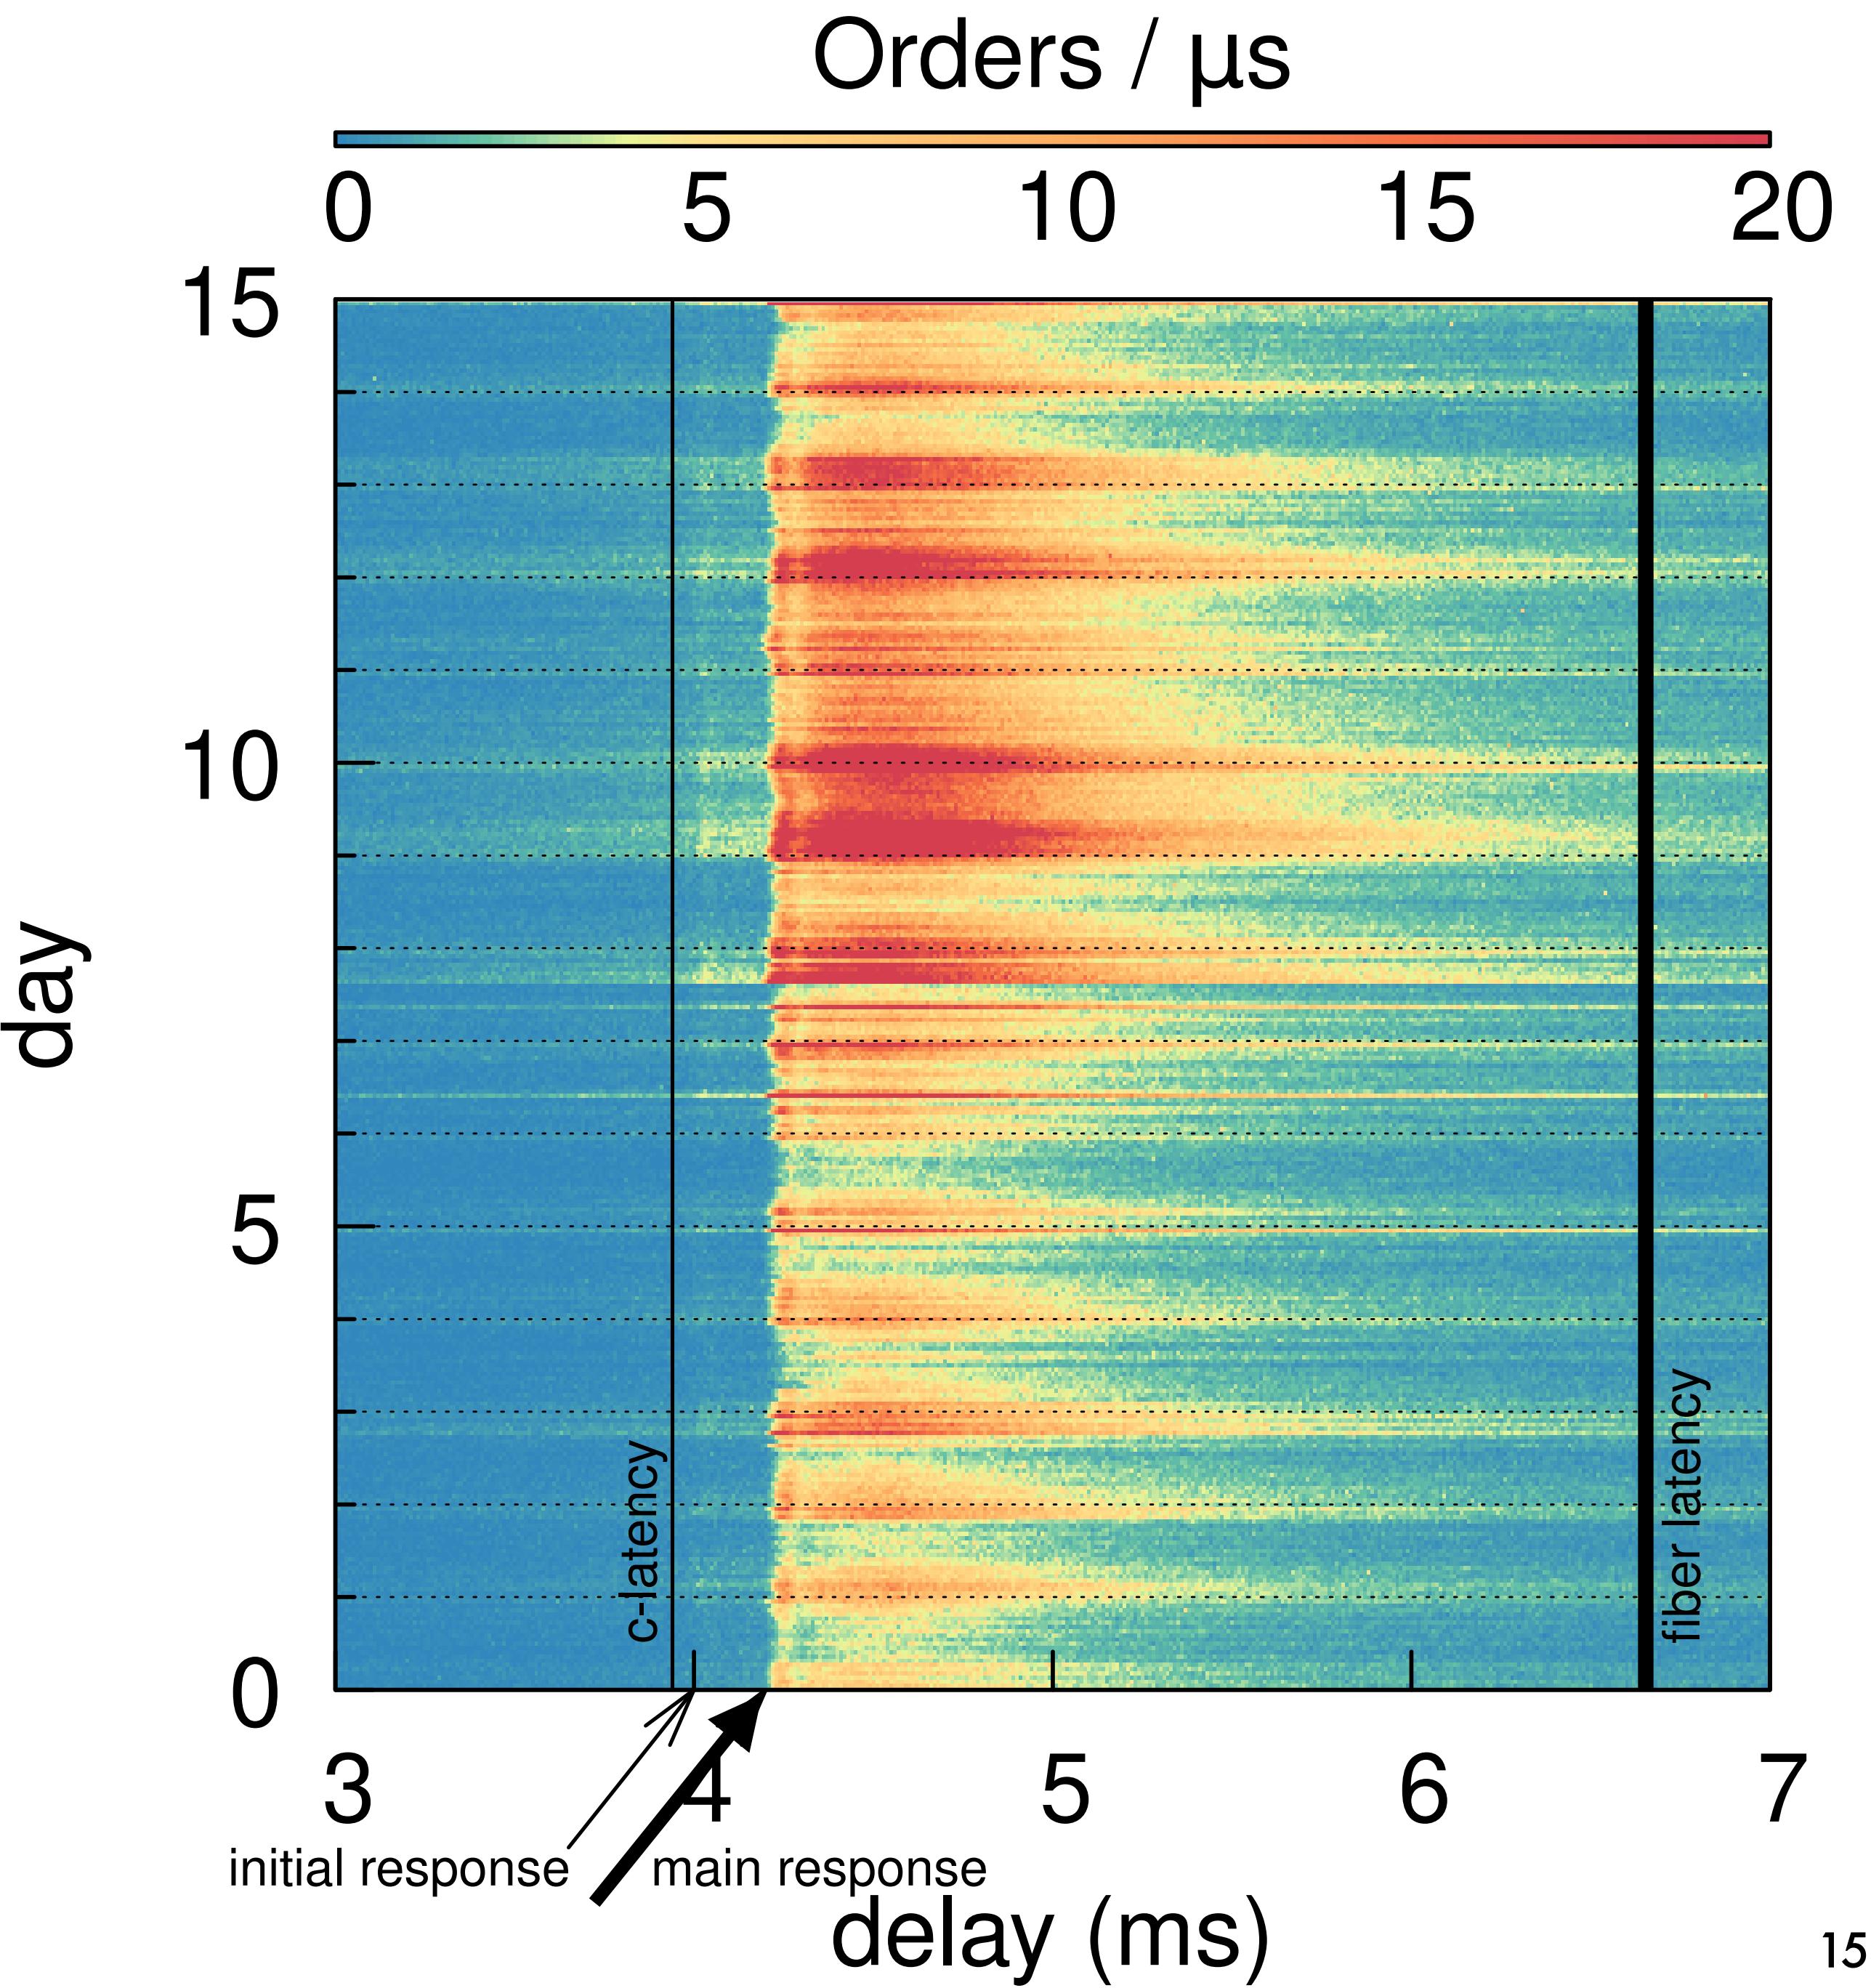
- Events in Chicago (stimulus) trigger trades in New York (response)
- Latency is lag between stimulus and response



Is MW resilient to weather in practice?

What the markets show

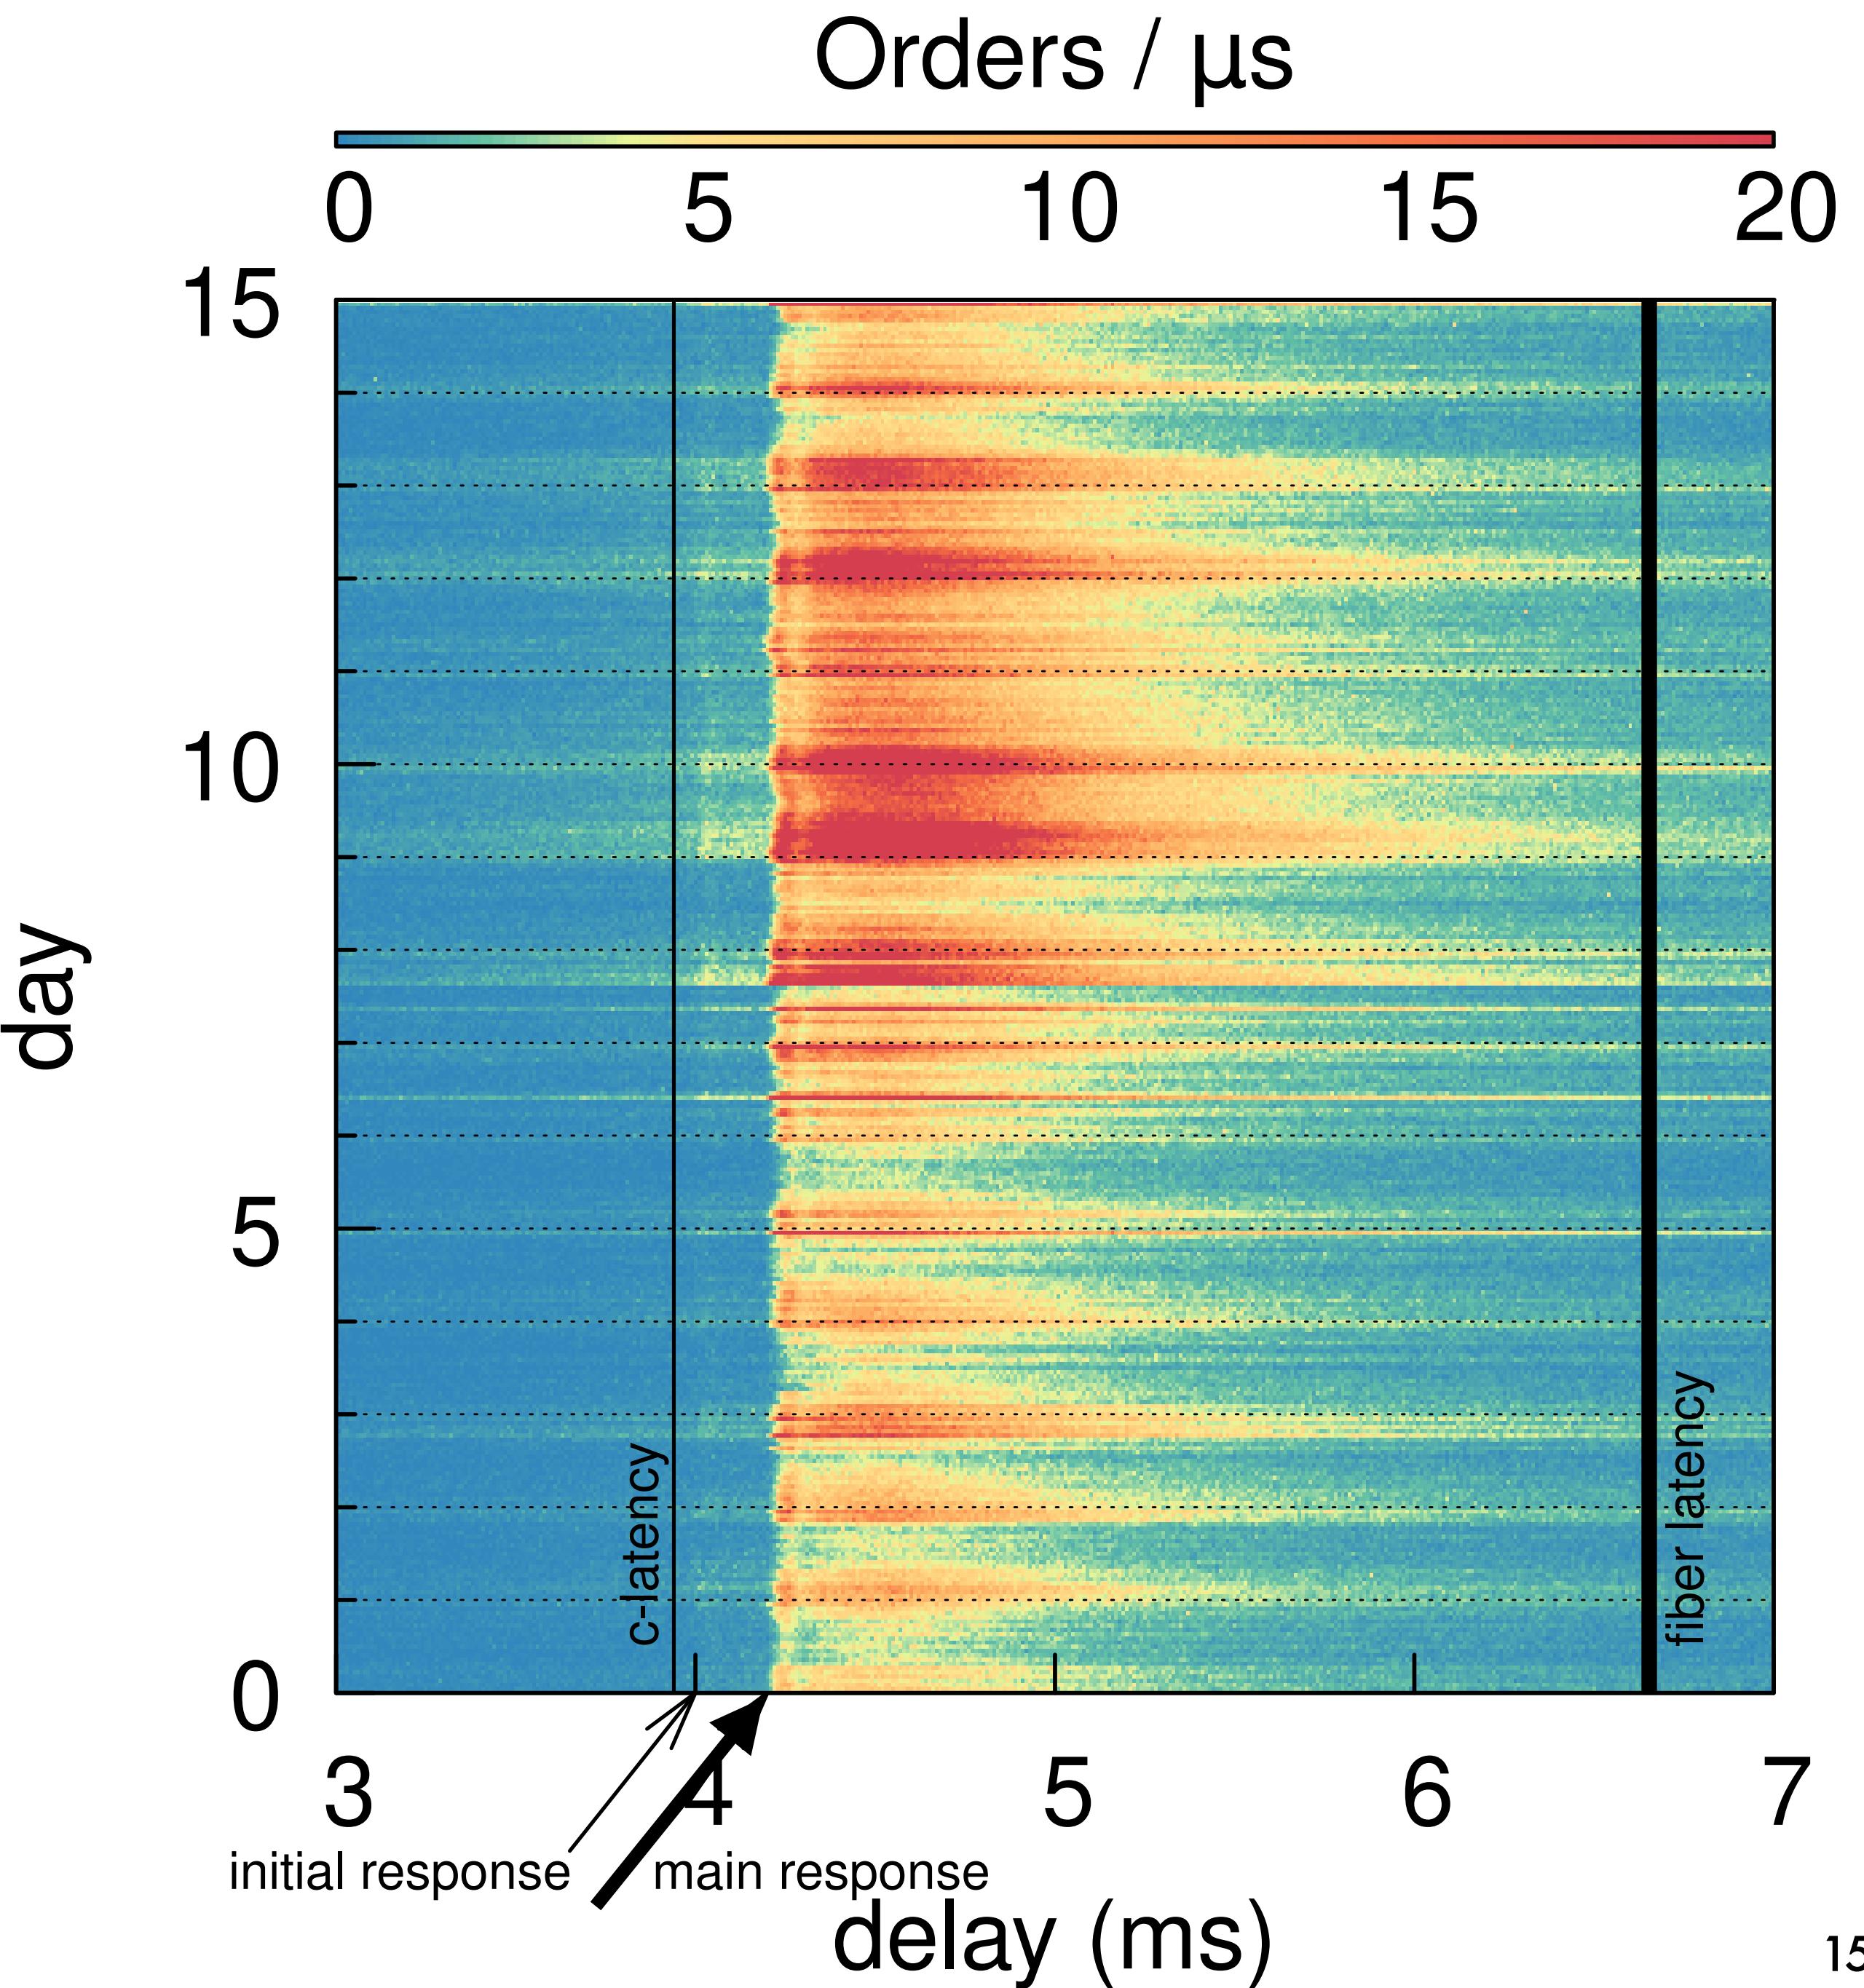
- Events in Chicago (stimulus) trigger trades in New York (response)
- Latency is lag between stimulus and response
- Some MW networks were always up



Is MW resilient to weather in practice?

What the markets show

- Events in Chicago (stimulus) trigger trades in New York (response)
- Latency is lag between stimulus and response
- Some MW networks were always up
- Bad weather correlated with latency increase



Does cISP provide cost-effective benefits to applications?

Speedup for web browsing

Browsing the web with cISP

Speedup for web browsing

Browsing the web with cISP

- Measured page load time (PLT) with CDN cache emulation

Speedup for web browsing

Browsing the web with cISP

- Measured page load time (PLT) with CDN cache emulation
- For traffic over cISP, latency reduced to 0.33x

Speedup for web browsing

Browsing the web with cISP

- Measured page load time (PLT) with CDN cache emulation
- For traffic over cISP, latency reduced to 0.33x
- PLT reduction by sending traffic over cISP: **42%**

Speedup for web browsing

Browsing the web with cISP

- Measured page load time (PLT) with CDN cache emulation
- For traffic over cISP, latency reduced to 0.33x
- PLT reduction by sending traffic over cISP: **42%**
- Idea: split traffic over cISP and traditional Internet

Speedup for web browsing

Browsing the web with cISP

- Measured page load time (PLT) with CDN cache emulation
- For traffic over cISP, latency reduced to 0.33x
- PLT reduction by sending traffic over cISP: **42%**
- Idea: split traffic over cISP and traditional Internet
 - Upload only (10% bytes): **21%**

Speedup for web browsing

Browsing the web with cISP

- Measured page load time (PLT) with CDN cache emulation
- For traffic over cISP, latency reduced to 0.33x
- PLT reduction by sending traffic over cISP: **42%**
- Idea: split traffic over cISP and traditional Internet
 - Upload only (10% bytes): **21%**
 - Small packets only (10% bytes): **28%**

Cost and benefit per GB

At 3000 towers and 100Gbps, cISP cost estimate is **\$0.81** per GB

Cost and benefit per GB

At 3000 towers and 100Gbps, cISP cost estimate is **\$0.81** per GB

Web search

Numbers from
Google

\$1.84 per GB

Cost and benefit per GB

At 3000 towers and 100Gbps, cISP cost estimate is **\$0.81** per GB

Web search

E-commerce

Numbers from
Google

Numbers from
Amazon

\$1.84 per GB

\$6.8 per GB

Cost and benefit per GB

At 3000 towers and 100Gbps, cISP cost estimate is **\$0.81** per GB

Web search	E-commerce	Gaming
Numbers from Google	Numbers from Amazon	Numbers from VPN services
\$1.84 per GB	\$6.8 per GB	\$3.7 per GB

In the paper...

More goodies



In the paper...

More goodies

- Issues of integrating cISP into the Internet



In the paper...

More goodies

- Issues of integrating cISP into the Internet
- On speed mismatch and pacing of regular vs. cISP traffic



In the paper...

More goodies

- Issues of integrating cISP into the Internet
- On speed mismatch and pacing of regular vs. cISP traffic
- Alternative deployment models



In the paper...

More goodies

- Issues of integrating cISP into the Internet
- On speed mismatch and pacing of regular vs. cISP traffic
- Alternative deployment models
- cISP design for continental Europe



In the paper...

More goodies

- Issues of integrating cISP into the Internet
- On speed mismatch and pacing of regular vs. cISP traffic
- Alternative deployment models
- cISP design for continental Europe
- and more...



Thank you

To coauthors, collaborators, reviewers, our shepherd George Porter, and you!

See more at cspeed.net

Questions?