Toward an Atlas of the Physical Internet

Summer 2014

Paul Barford
Computer Sciences
University of Wisconsin

Motivation

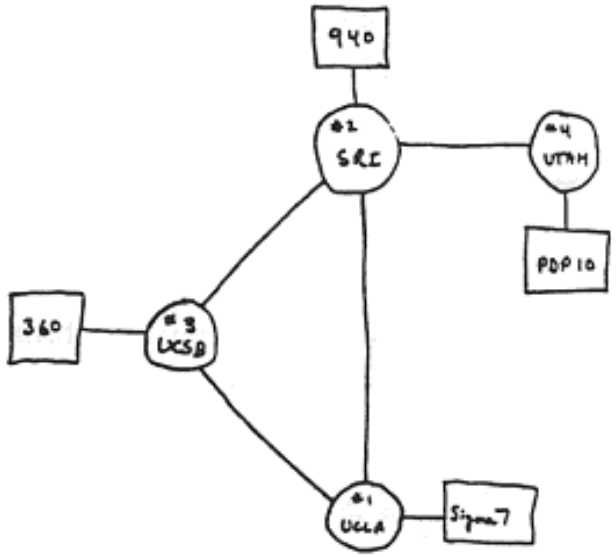

FIGURE 6.2 Drawing of 4 Node Network
(Courtesy of Alex McKenzie)

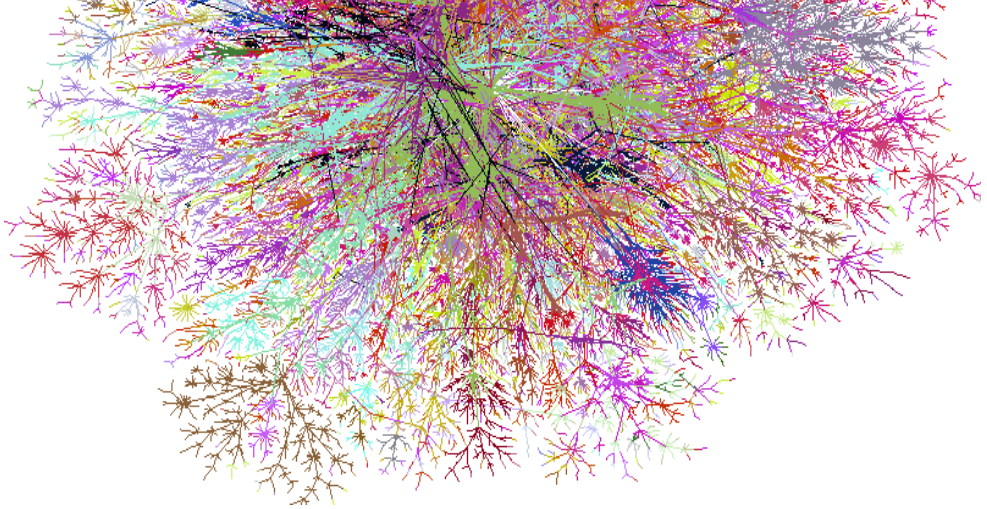

Objectives of our work

- **Create and maintain a comprehensive catalog of the *physical Internet***
  - **Geographic locations of nodes (buildings that house PoPs, IXPs etc.) and links (fiber conduits)**
- **Extend with relevant related data**
  - **Active probes, BGP updates, weather, etc.**
- **Maintain portal for visualization and analysis**
- **Apply maps to problems of interest**
  - **Robustness, performance, security**

Related work

- Many prior Internet mapping efforts
  - S. Gorman studies from early 2000's
  - CAIDA
  - DIMES
  - iPlane
- Commercial activities
  - TeleGeography
  - Renesys/Dyn
  - Lumeta
- Internet Topology Zoo

Compiling a physical repository

- **Step #1: Identification**
  - **Utilize *search* to find maps of physical locations**
- **Step #2: Transcription**
  - **Multiple methods to automate data entry**
- **Step #3: Verification**
  - **Ensure that data reflects latest network maps**
- **Our hypothesis is that physical sites are limited in number and fixed in location**
  - **But the raw number is still large!**

Example: Telstra world wide

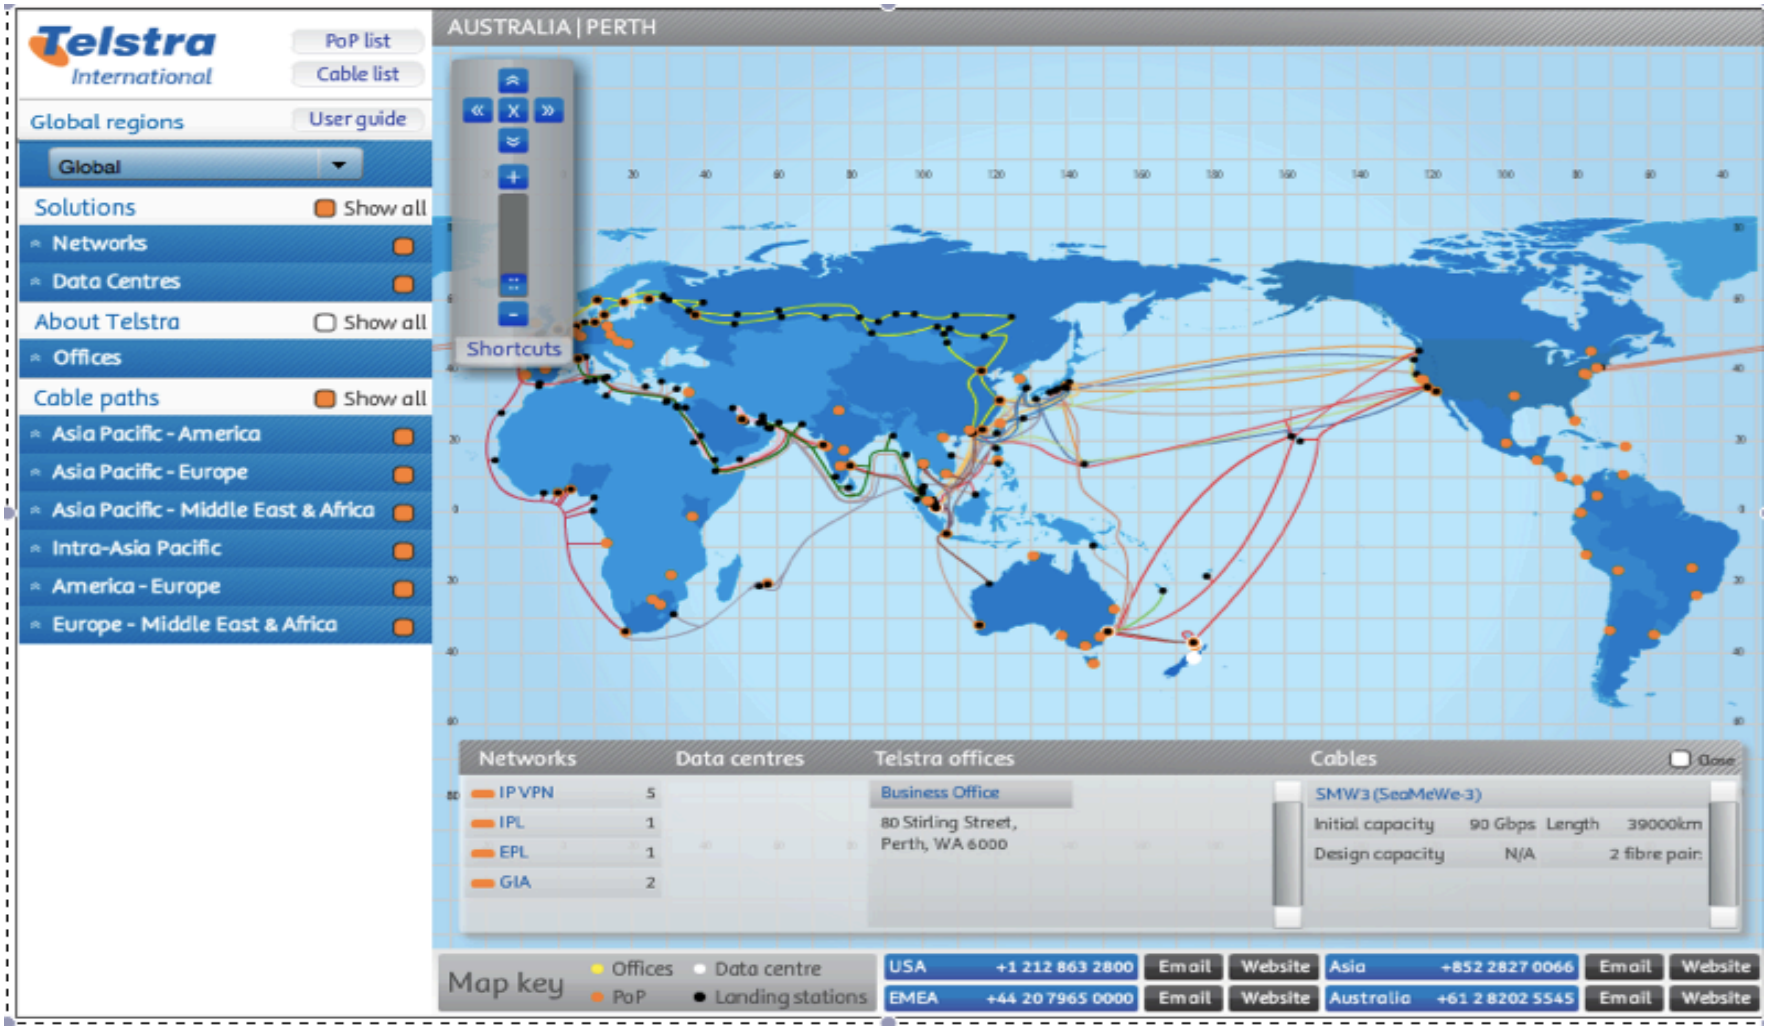

Example: Sprint IP network (US)

Example: Regional fiber

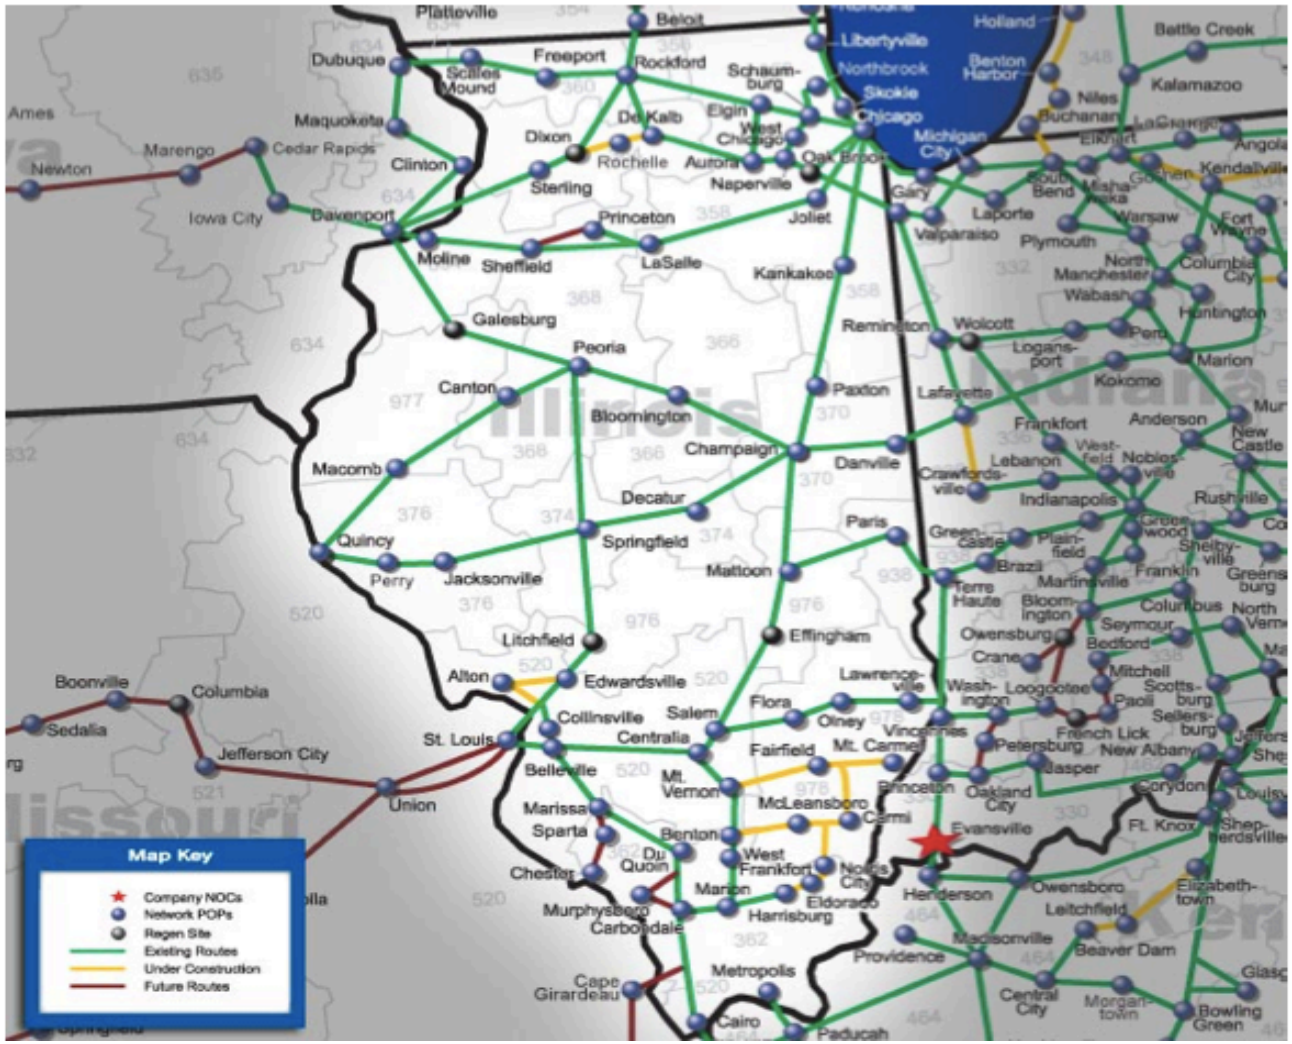

Illinois POP List

ALTON
Address: 1805 Washington Ave Zip: 62002
Type: CO Status: FUTURE
CLLI: ALTNILAK

BELLEVILLE
Address: 211 Kretschmer Ave Zip: 62220
Type: CO Status: ACTIVE
CLLI: BLVLILAD

BLOOMINGTON
Address: 110 E Monroe St Zip: 61701
Type: CO Status: ACTIVE
CLLI: BLTNILXD

Address: 110 E Monroe St Zip: 61701
Type: CO Status: DOUBLE
CLLI: BLTNILXD

CAIRO
Address: 221 15th St Zip: 62914
Type: CO Status: ACTIVE
CLLI: CAIRILCF

CANTON
Address: 75 W Pine St Zip: 61520
Type: CO Status: ACTIVE
CLLI: CNTNILCN

CARBONDALE
Address: 208 W Monroe St Zip: 62901
Type: CO Status: ACTIVE
CLLI: CRDLILXE

CARMI
Address: 200 W Cherry St Zip: 62821

Example: Metro fiber maps

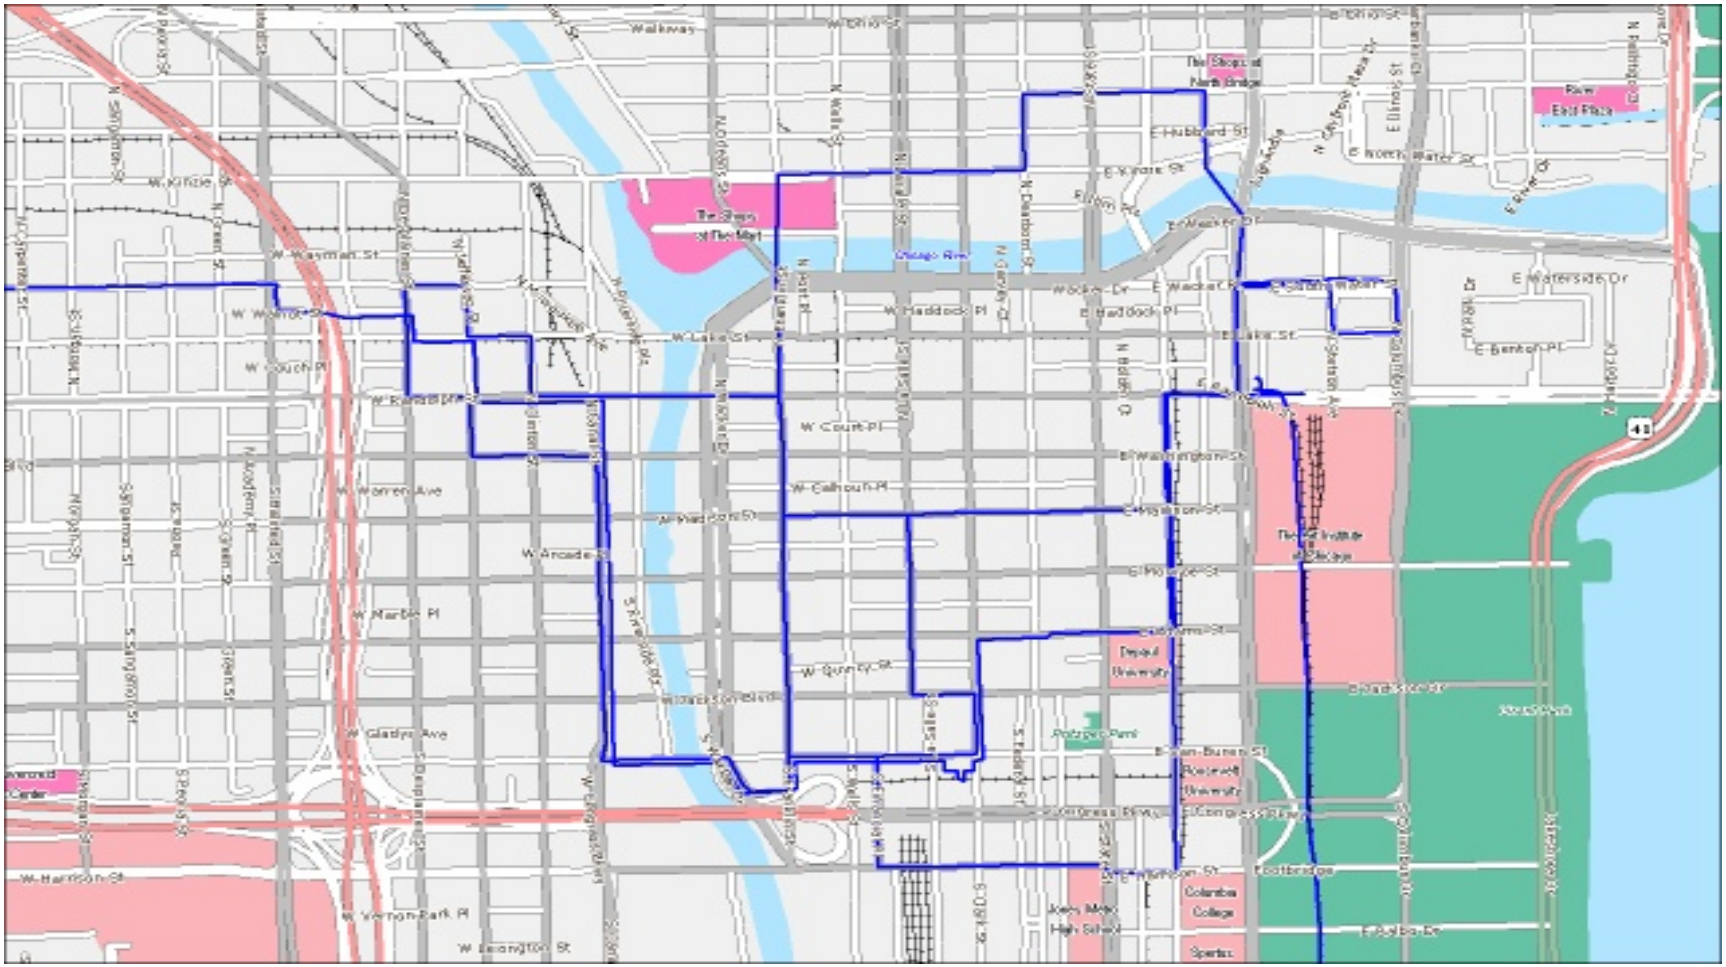

Internet Atlas @ UW

- **Effort began in September ’11**
  - **Capture everything from maps discovered by search**
  - **Use all relevant data sources (ISP maps, colocation, data centers, NTP, traceroute, etc.)**
- **Data extraction and verification tools**
- **Comprehensive database**
- **Interactive web portal**
  - **Includes ArcGIS for visualization and analysis**
- **Paper in ACM SIGCOMM HotPlanet WS ‘13**

Current DB

- Number of networks: 389
- Number of tier 1 networks: 10
- Number of data centers: 2,232
- Number of NTP servers: 744
- Number of traceroute servers: 221
- Number and type of other nodes: IXP (358), DNS root (282)
- Total number of nodes: 14,827
- Number of unique locations of nodes: 7,988
- Maximum overlap at any one node: 92
- Total number of links: 13,861
- Peering DB facility locations: 1058
- WiGLE Wireless SSID locations: 5202
- Antenna locations from FCC: 5786

Internet Atlas – Full View

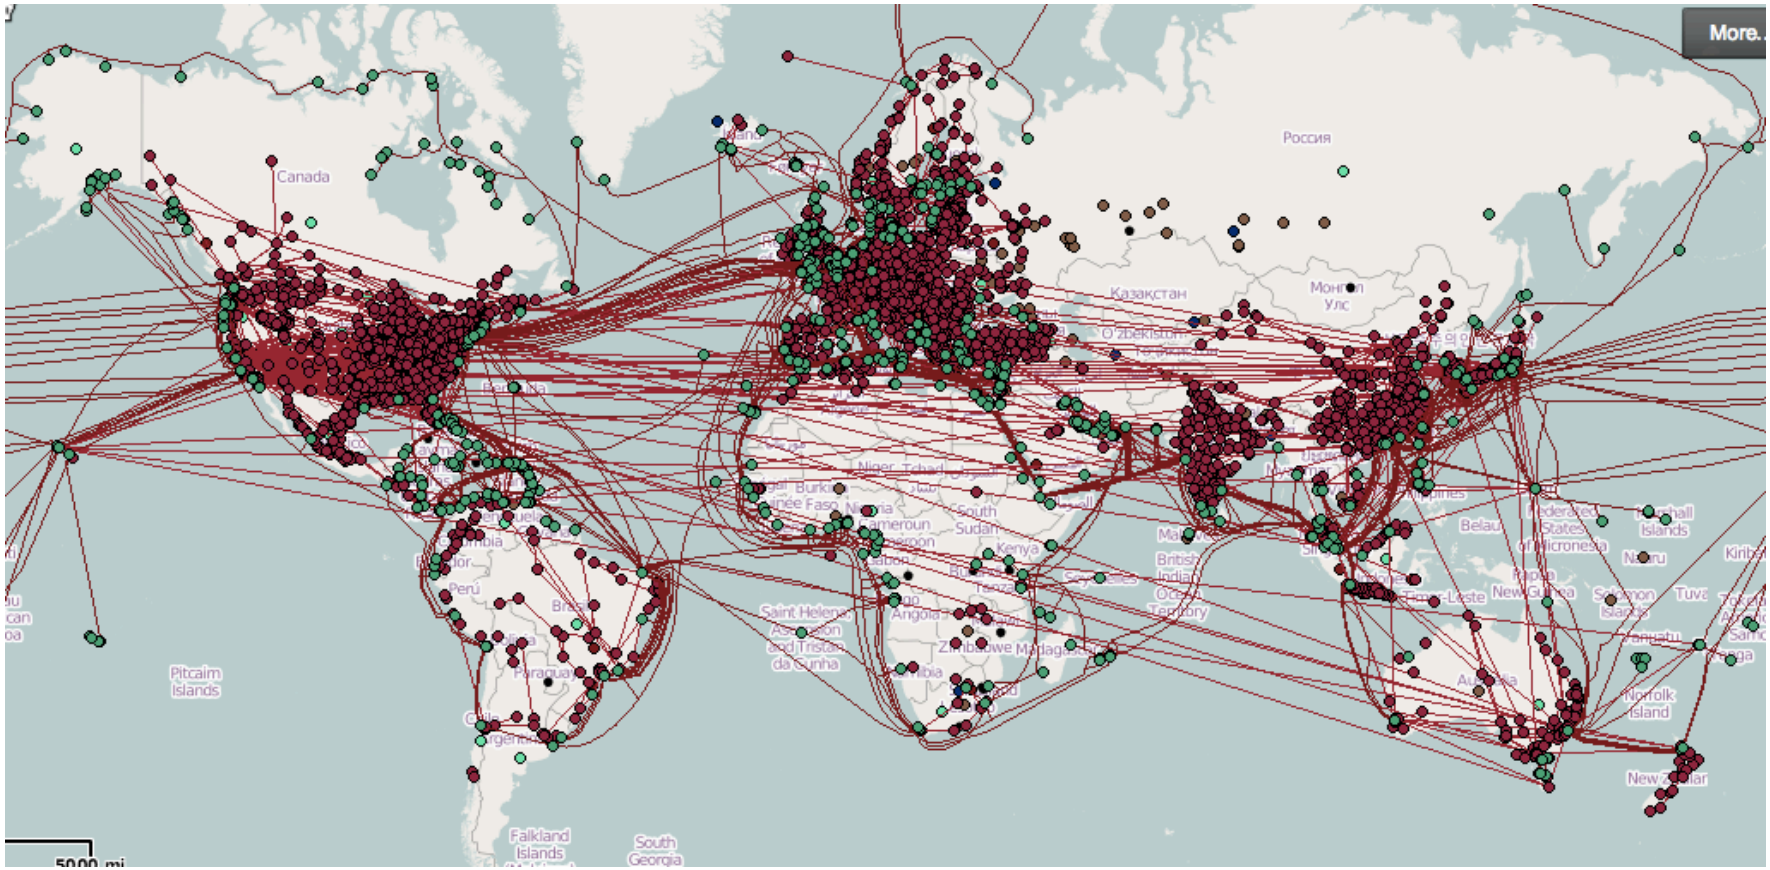

Case study: RiskRoute

Consider Internet physical infrastructure:

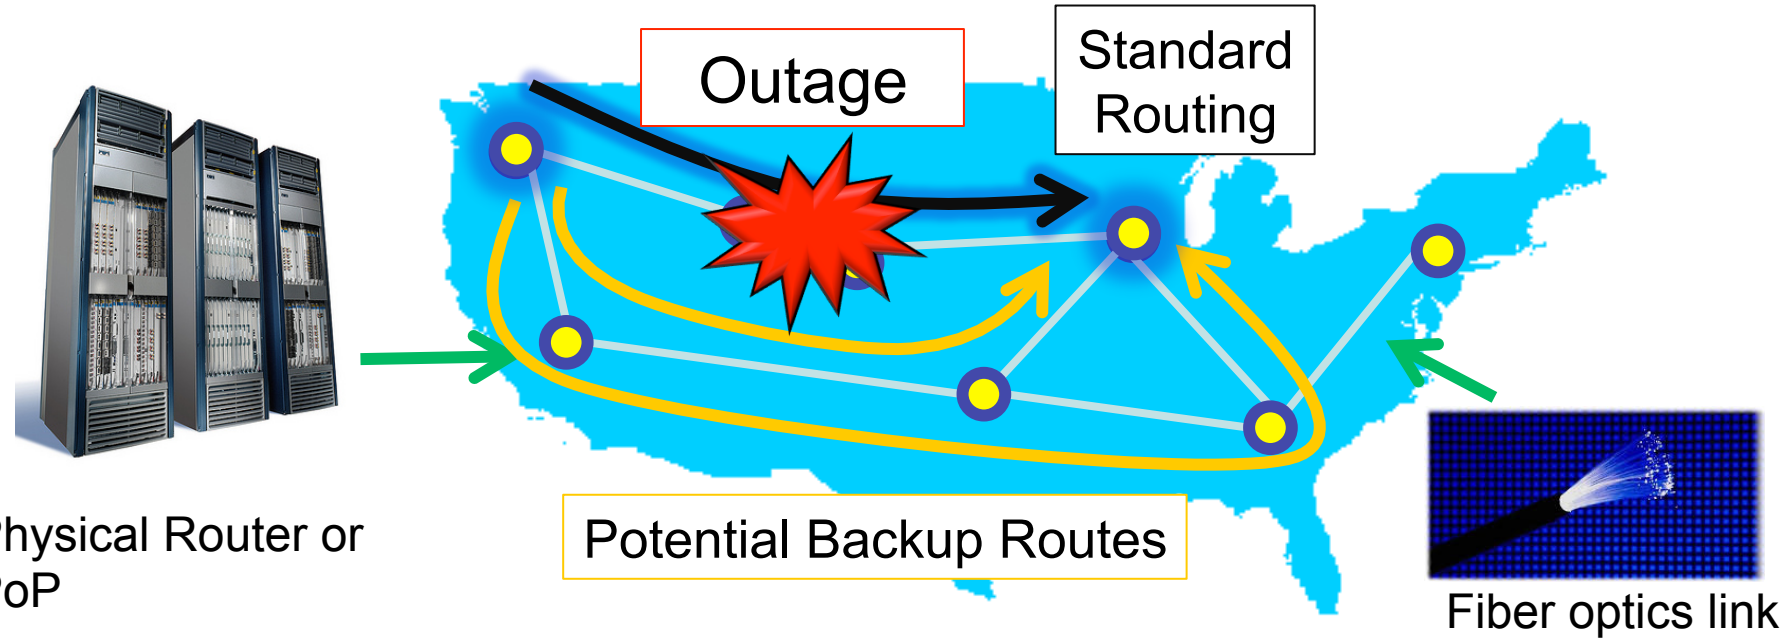

- **Can we automatically adjust routes to avoid outages before they happen?**
- **Can we identify the best backup routes?**

Bit-risk miles metric

- **The idea of bit-miles motivates the introduction of *bit-risk miles***

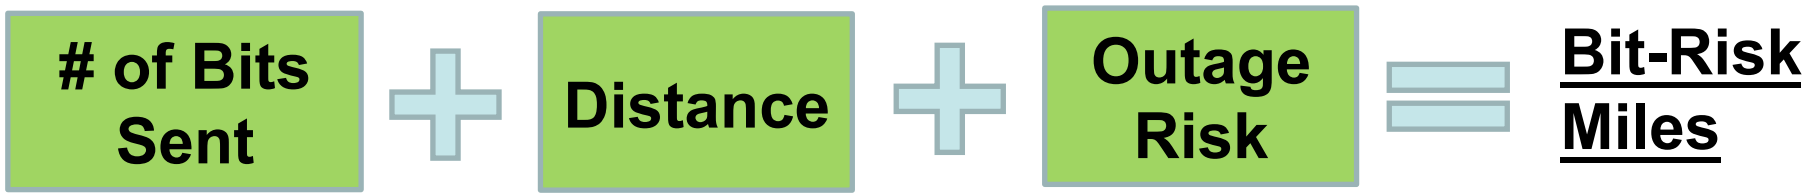

  - **Consider a network path:**

$$\mathbf{p} = \{p_1, p_2, \ldots, p_K\}$$

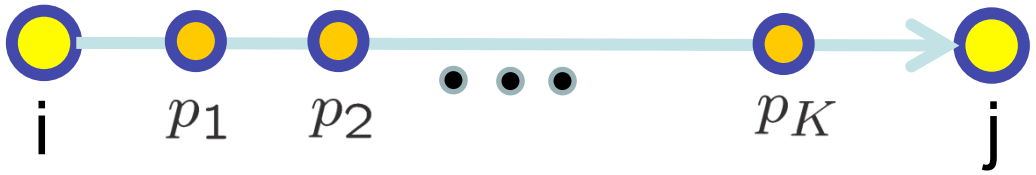

  - **The bit-risk miles of the routing path is defined as:**

$$r_{i,j}(\mathbf{p}) = \sum_{x=1}^{K} \left(d_{p_x, p_{x+1}} + \gamma_{i,j} \left(\lambda_h o_h(p_x) + \lambda_f o_f(p_x) \right) \right)$$

Bit miles **Outage risk**

Utility of bit-risk miles

- **Quantifies the trade-offs between:**

Short geographic routing paths with **high** outage risk

vs.

Long geographic routing paths with **low** outage risk

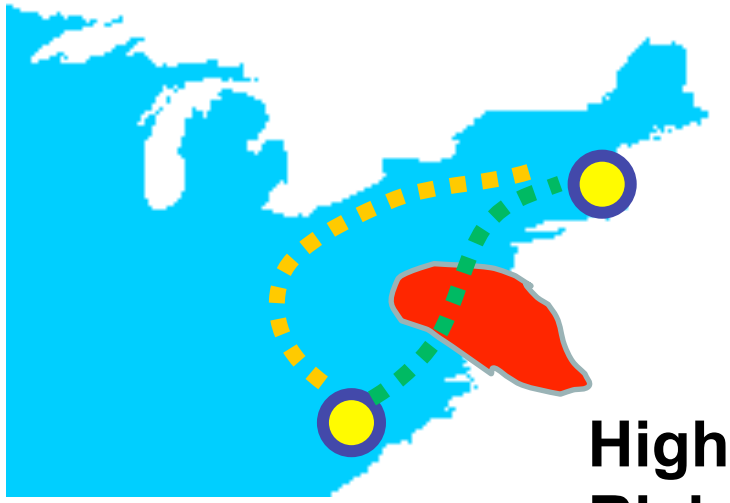

Defining terms

- **Gamma: what is the cost of an outage between the source and destination?**
  - **To approximate this, we use the fraction of population affected clustered to nearest PoP**
- **$Lambda_h$: what is historical outage probability at a PoP location?**
  - **We use corpus of events from 1970 to 2010 (29,865 FEMA emergency declarations and over 145,000 NOAA severe weather events**
- **$Lambda_f$: what is forecasted outage probability at a PoP location?**
  - **Based on reported information from NWS, NHC, etc.**

RiskRoute methodology

How do we choose which backup path has the smallest bit-risk miles?

- **Current techniques: Storing only one backup path (*e.g.*, Fast Reroute) is fragile to large-scale outages**
- **Storing all the backup paths is combinatorial**
- **<u>RiskRoute Framework</u>: Using shortest path techniques, continuously recalculate all paths with the smallest bit-risk miles**

Analysis

- **Real-World Network**
  - **7 Tier-1 ISPs, 16 regional networks**
- **Intra-domain Routing**
  - **Routing inside a specified network**
- **Inter-domain Routing**
  - **Routing between networks**
- **Performance Metrics:**

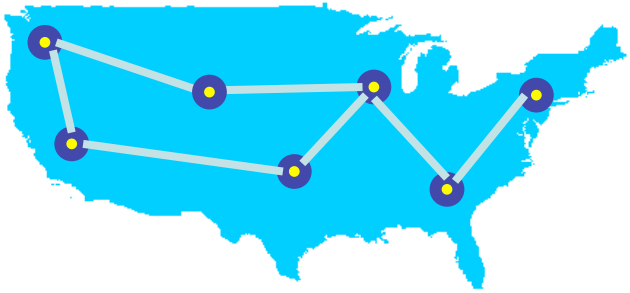

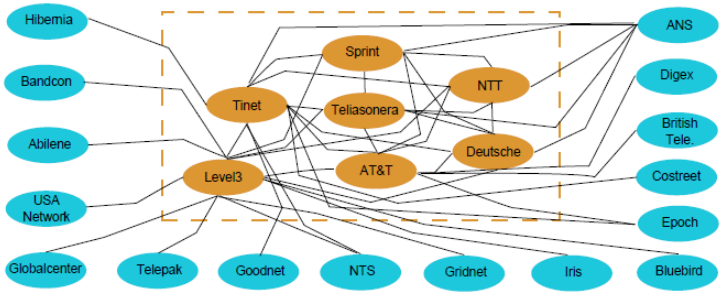

1 **Risk Ratio** – The average reduction in bit-risk miles using RiskRoute compared with shortest path routing

$$r_r = 1 - \frac{1}{N^2}\sum_{i=1}^{N}\sum_{j=1}^{N}\frac{r\left(\mathbf{p}_{i,j}^{rr}\right)}{r\left(\mathbf{p}_{i,j}^{shortest}\right)}$$

2 **Distance Ratio** – The average increase in bit-miles using RiskRoute compared with shortest path routing

$$d_r = \frac{1}{N^2}\sum_{i=1}^{N}\sum_{j=1}^{N}\frac{d\left(\mathbf{p}_{i,j}^{rr}\right)}{d\left(\mathbf{p}_{i,j}^{shortest}\right)} - 1$$

Intradomain results

- **Why is RiskRoute more advantageous to some networks?**

Reduction in Bit-Risk Miles (Risk Ratio)

Geographic footprint

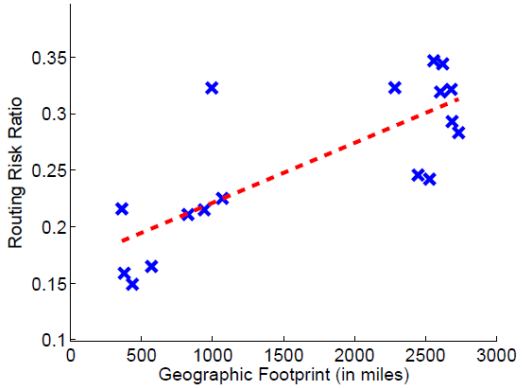

Avg. router outdegree

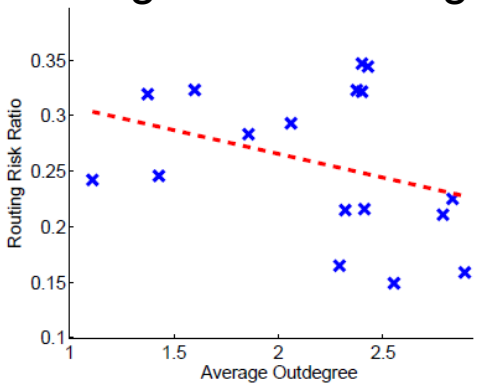

Increase in Bit-Miles (Distance Ratio)

Number of PoPs

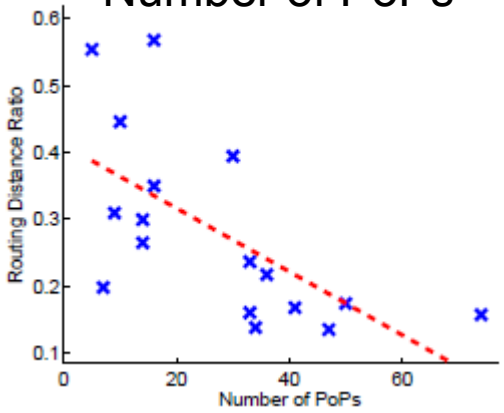

Number of peers

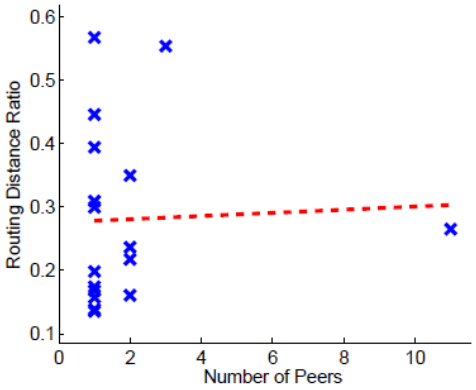

Robustness results

- **Can all networks decrease risk via the new link infrastructure?**

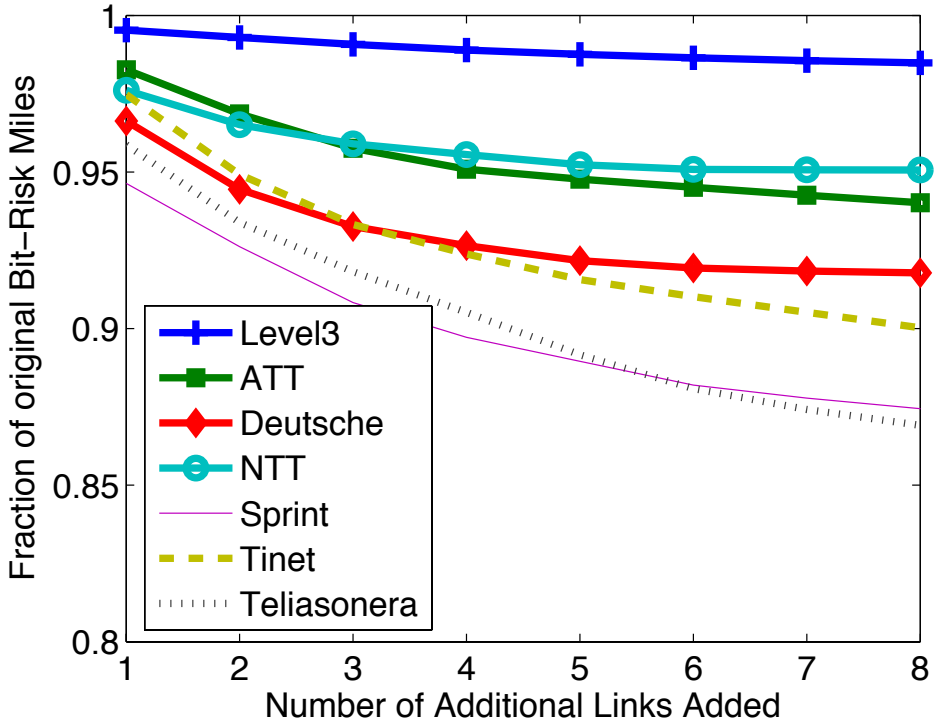

Hurricane Katrina and Level3

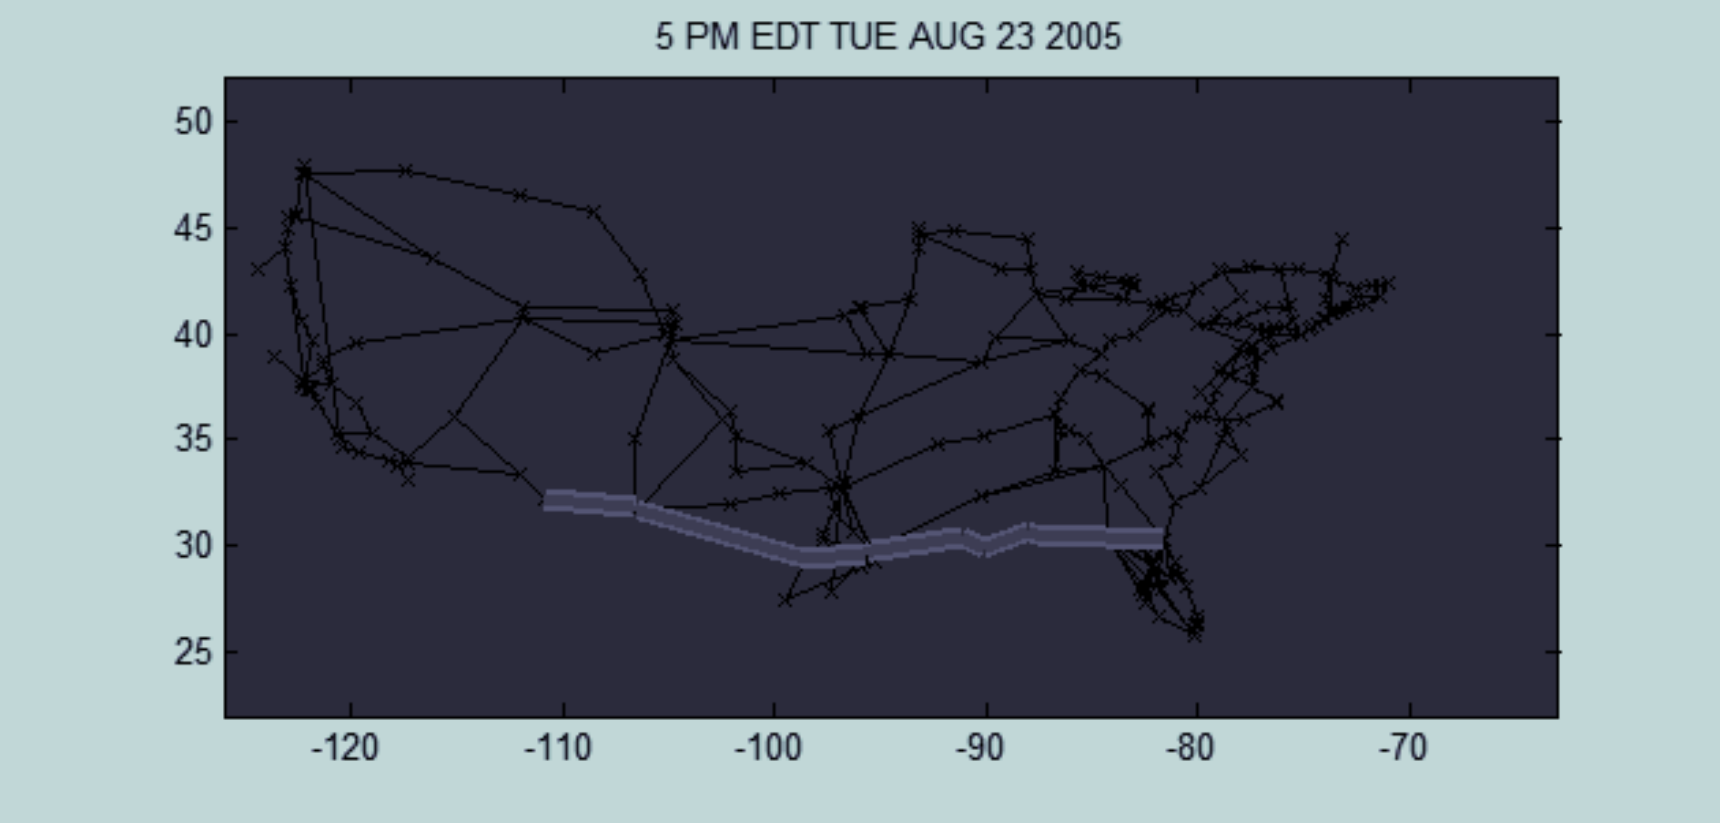

Next steps

- Continue to populate DB
  - Goal = 500 networks by December, ‘14
- Continue to enhance web portal
  - Expanded analytic capability
- Add related data for physical sites
  - PoPs, routers, IP addresses, peering, etc.
- Expanded active probing capability
  - IP geolocation is the key
- Expand focus for target applications
  - Shared infrastructure risk

Thank you!

- Ram Durairajan
- Brian Eriksson
- Xin Tang
- Subhadip Ghosh

Portal

http://internetatlas.org