(jeez y) Where is the Internet?

Answers from YAHOO! ANSWERS :

(G. Whilikers) Out there.

(Mike) the way I see it, the "internet" has to be somewhere. a router collects the internet to my house sure. but somewhere on earth there HAS to be some where that the internet originates by sending out the first signal so others can collect it. where is that place? also is it like... a box or something? I'm so confused. one more for those who want to answer, if some maintenance guy were to spill coffee on it would the worlds internet crash? like... everywhere?

(tp5com) Everywhere. It's not really a physical being.

(evaohell) The internet is stored in a little black box with a blinking light:

And what does it look like?

?

cambodia4kidsorg: ”What does the Internet look like?” flickr from YAHOO!

The Network Geography of the Internet*

Péter Hága
Eötvös Loránd University
Budapest, Hungary

Spotter
GEOLOCATION SERVICE

* *Spotter: A Model Based Active Geolocation Service,* INFOCOM 2011
* *On the Network Geography of the Internet*, INFOCOM 2011
* *On the Spatial Properties of Internet Routes*, submitted to Journal on Selected Areas of Communications


Geolocation in general

Geolocation in general

- passive geolocation
  - extracting location information from domain names (DNS)
  - registries, Whois databases, commercial databases
  - large and geographically dispersed IP blocks are allocated to a single entity prohibiting the study of several geographic aspects

---

- active geolocation
  - active probing (delay, topology)
  - measurement nodes with known location
  - constraint based techniques

Whois based location estimation

example for passive geoloc

LOCATIONS OF NORTH AMERICAN IPV4 ADDRESSES ASSIGNED TO GOOGLE INC.

Registered location	number of distinct IPv4 address
Mountain View, CA, US	222340
Plano, TX, US	331
Chicago, IL, US	46
Irvine, CA, US	30
Waterloo, ON, CA	22
San Francisco, CA, US	21
Atlanta, GA, US	15
Phoenix, AZ, US	15
Southfield, MI, US	15
Mequon, WI, US	14
Pittsburgh, PA, US	14
Gladwyne, PA, US	7
Richardson, TX, US	7
Total	222877

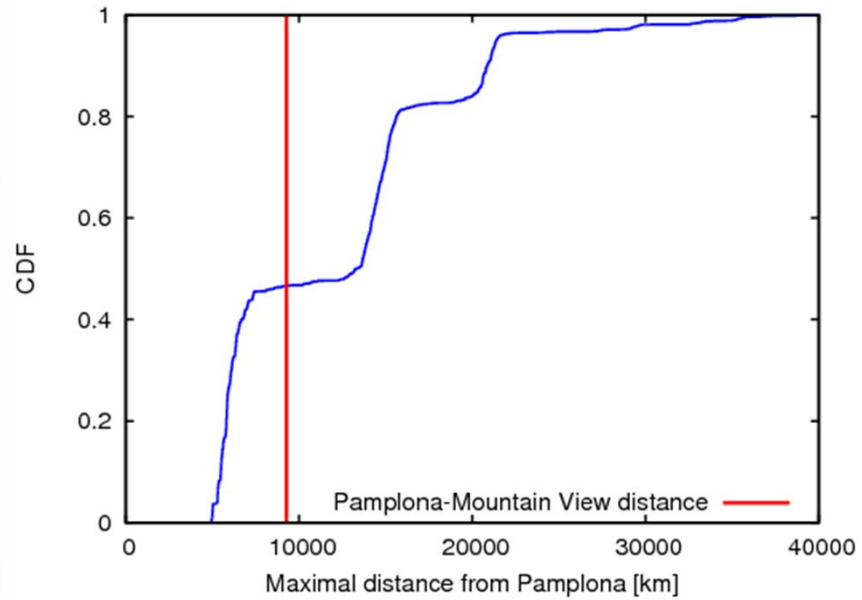

Cumulative distribution of the maximal distances from Pamplona, Spain to 4000 Google IPs. The maximal distances are calculated from the network delays assuming 200000 km/sec signal propagation speed. The vertical line represents the real geographical distance between Pamplona and Mountain View, CA, showing that 47% of the nodes must be closer to Pamplona than Mountain View.


Spotter – theoretical background

Spotter – theoretical background

- roundtrip travel time measured between the landmark (L) and the target (T) node by ping
- RTT is transformed to geographic distance
- ”triangulation” from multiple landmarks

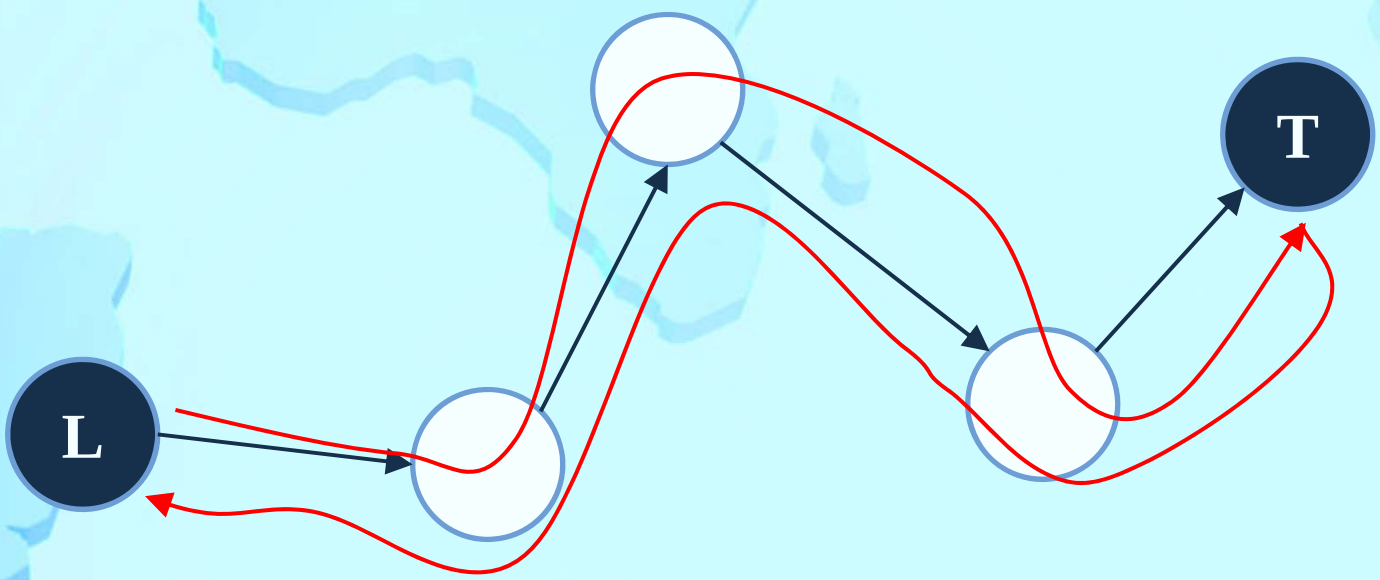

Travel time – distance relation

- reference dataset (nodes with known location)
- distance between the source and destination
- measured RTT

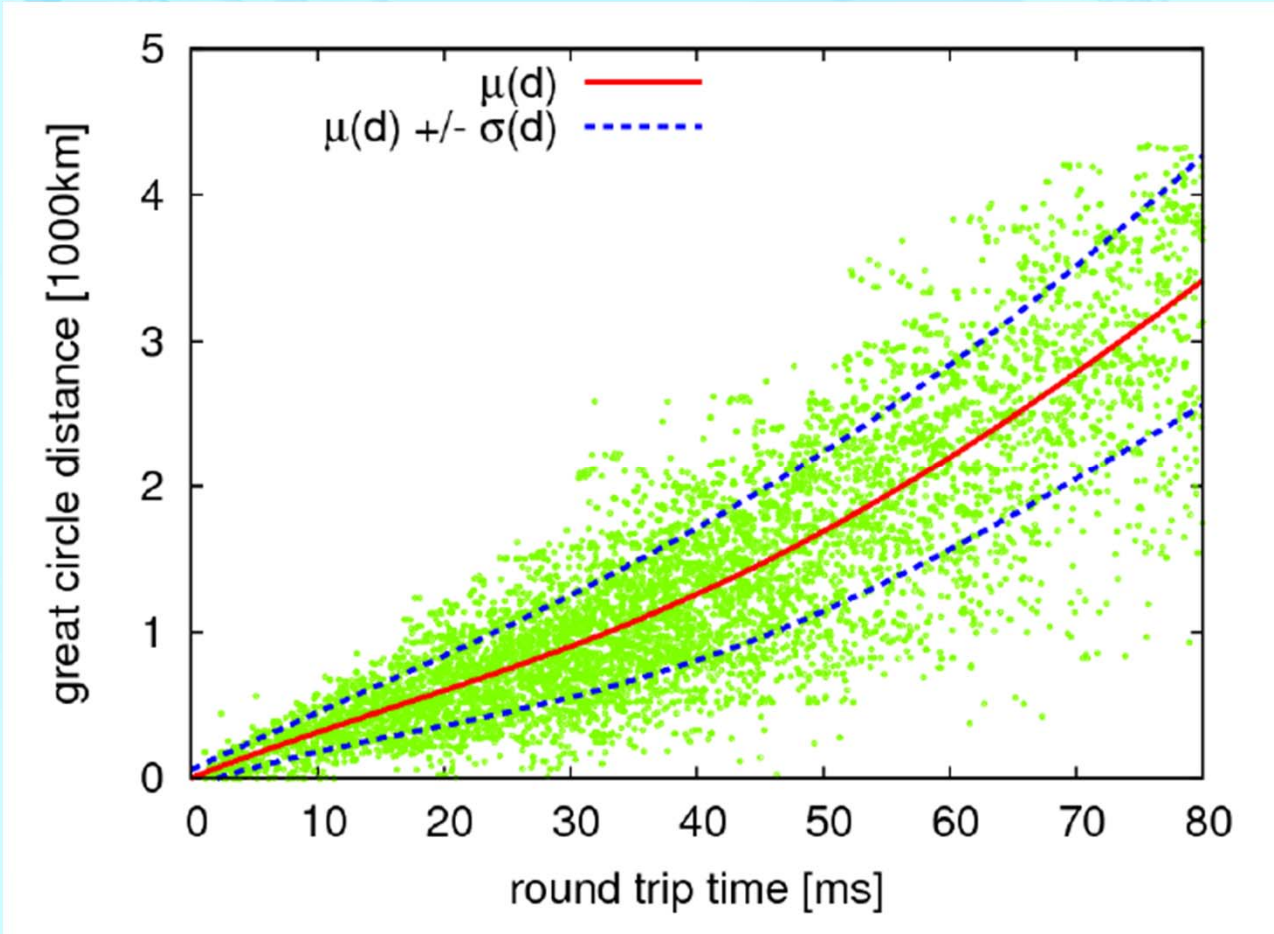

Standardized travel time – distance distribution

- Standardized values, fitted normal distribution

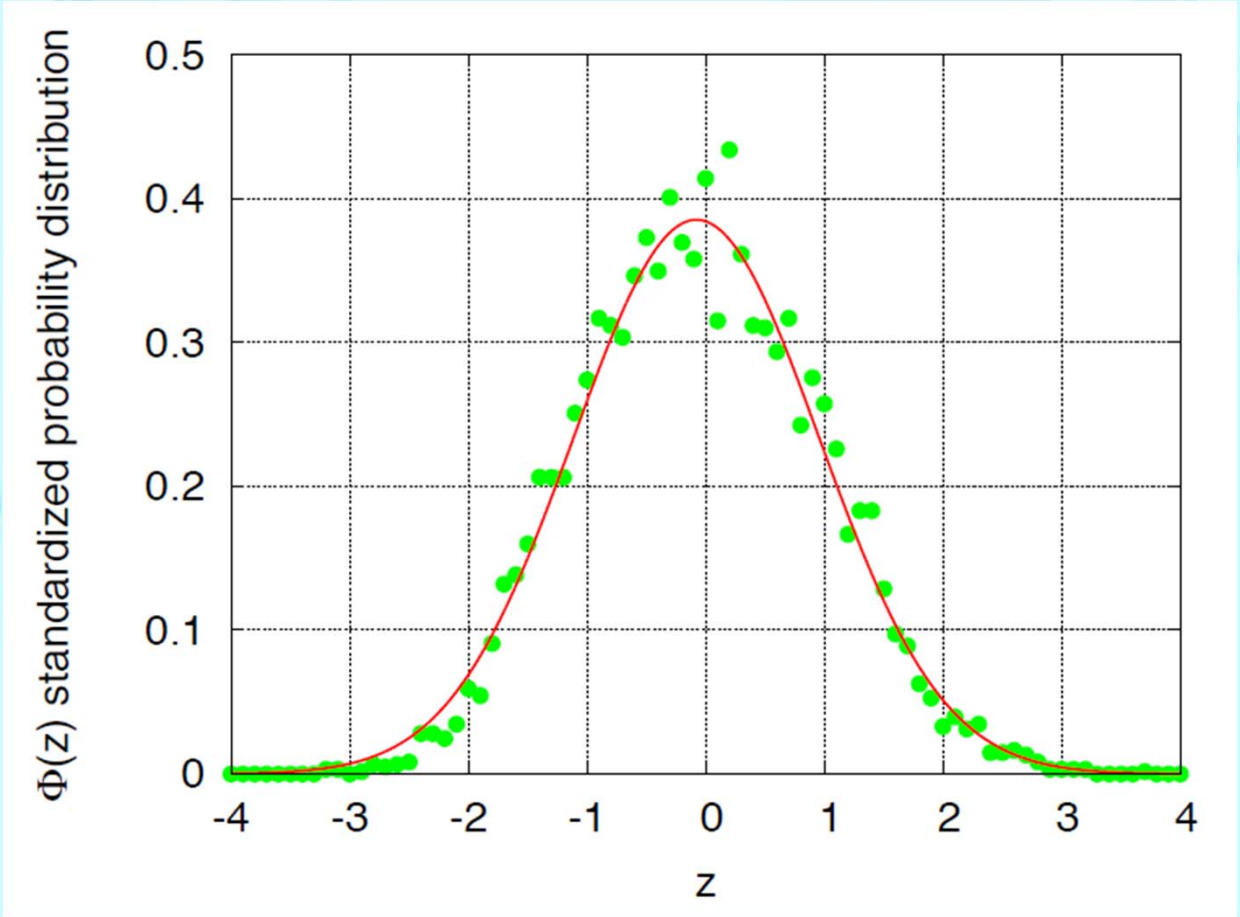

Evaluation process

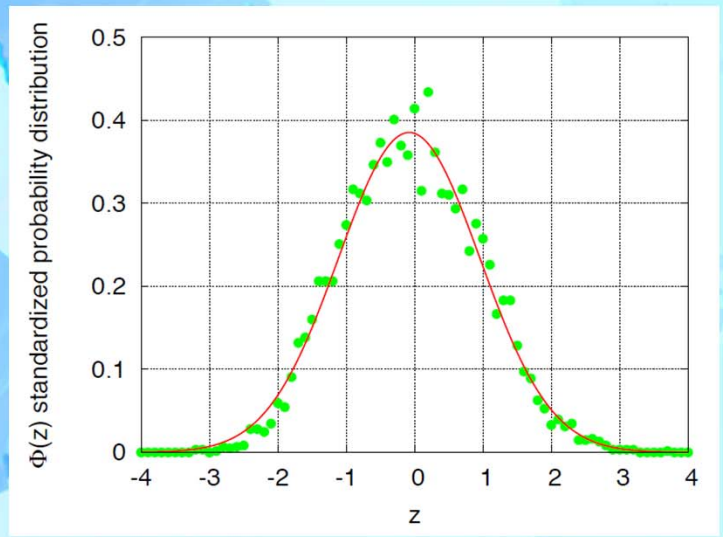

The distances are normally distributed for a given RTT

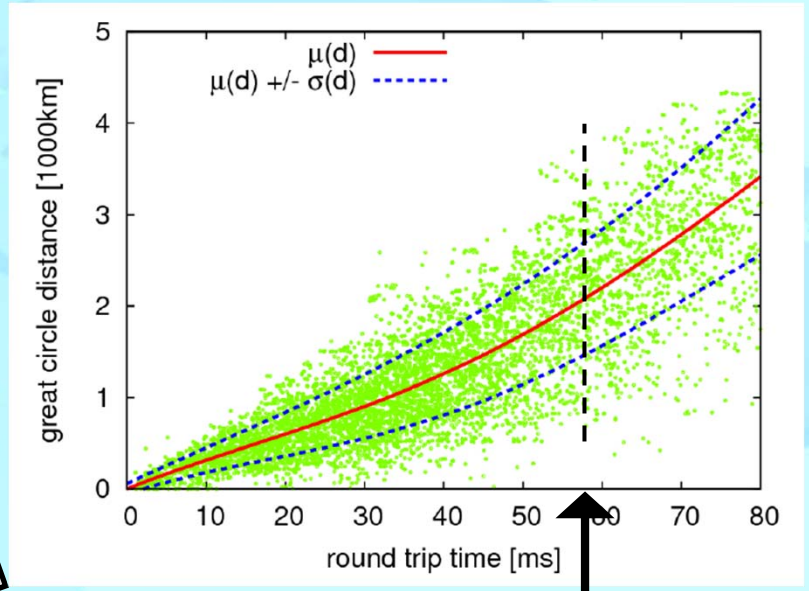

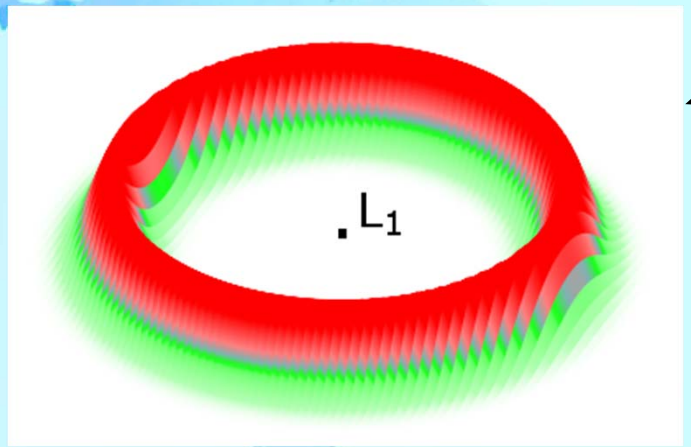

each RTT point out a radial profile with $\mu(d)$ and $\sigma(d)$ => which defines the spatial probability distribution of the target

”Triangulation”

The target’s location must satisfy all the individual spatial constraints.

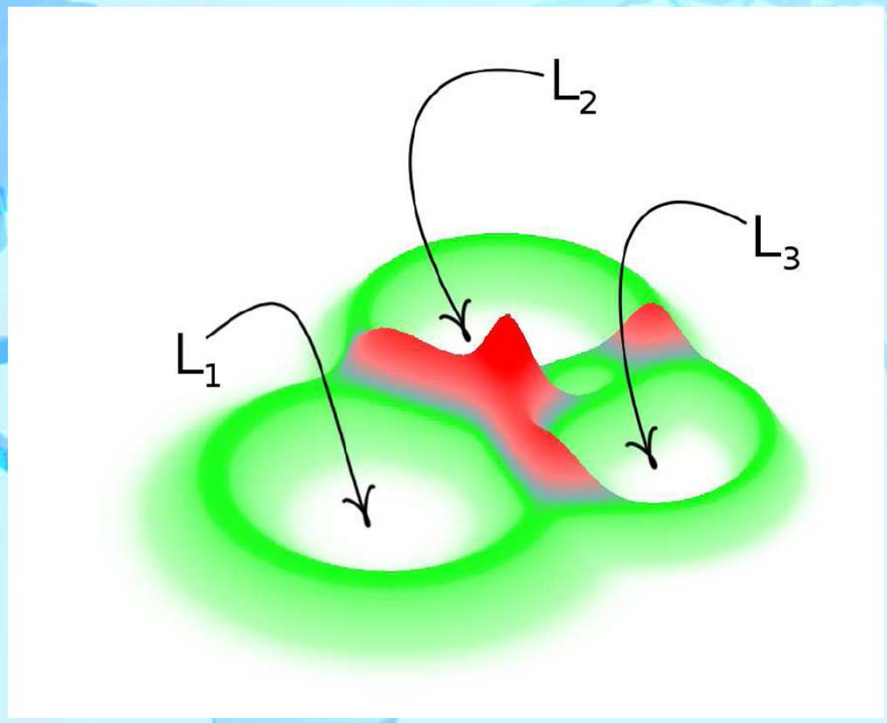

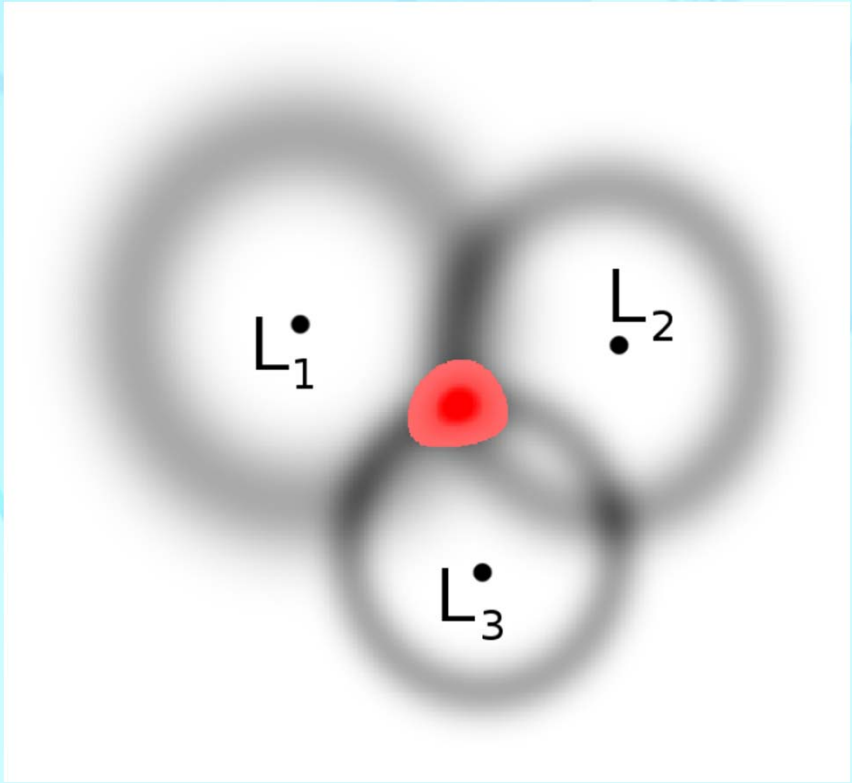

Spotter - online geolocation service

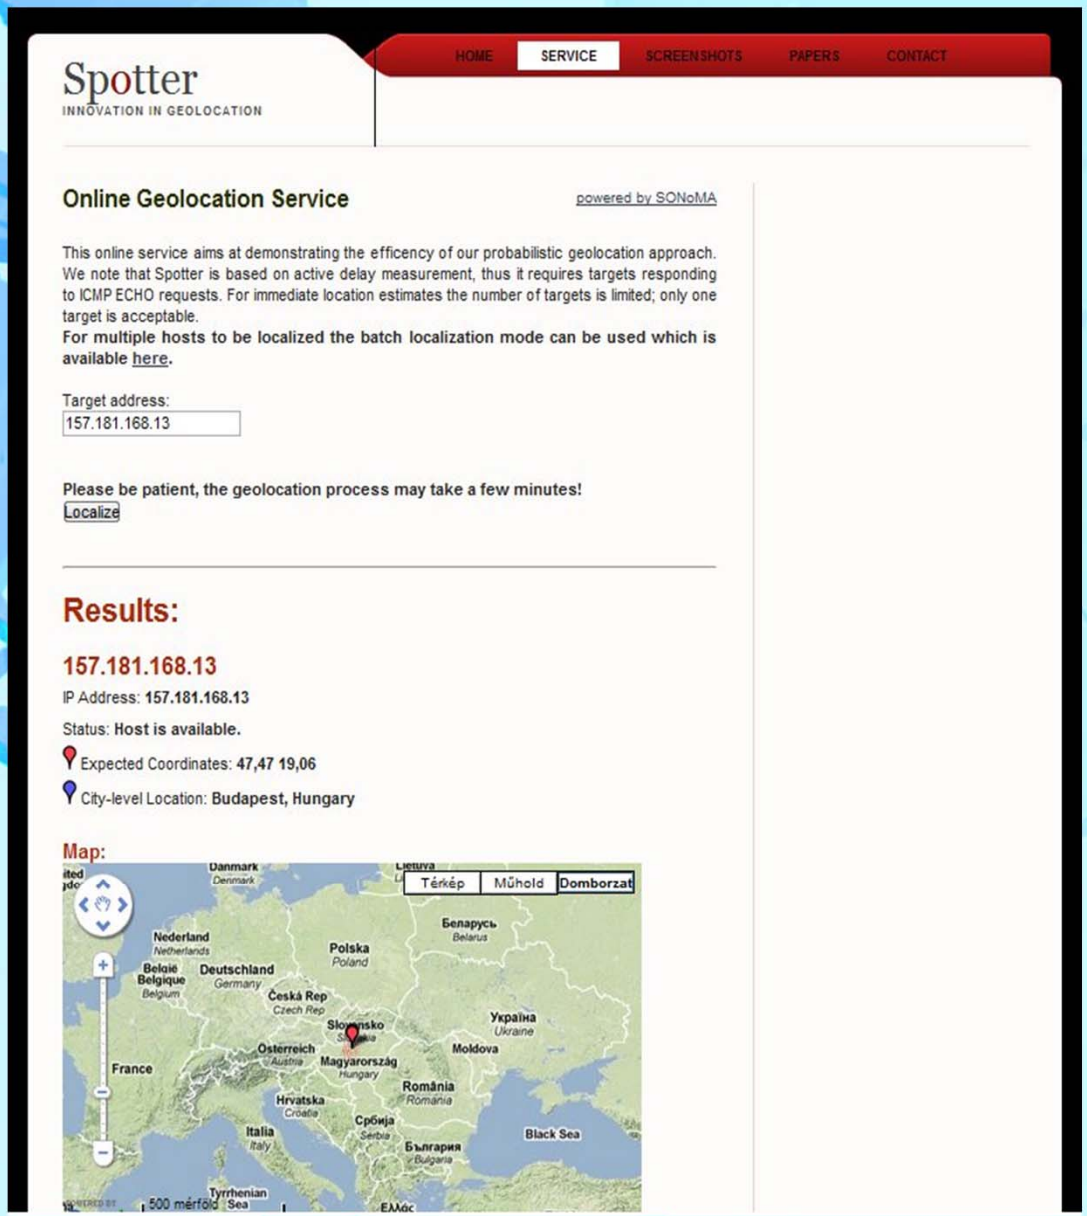

- online: free, easy to use
- offline: batch mode
  ~15k addresses/day
- to fix: DNS resolve problem

visit and use!

http://spotter.etomic.org

Where is the WikiLeaks site (46.59.1.2)?

hostip.info

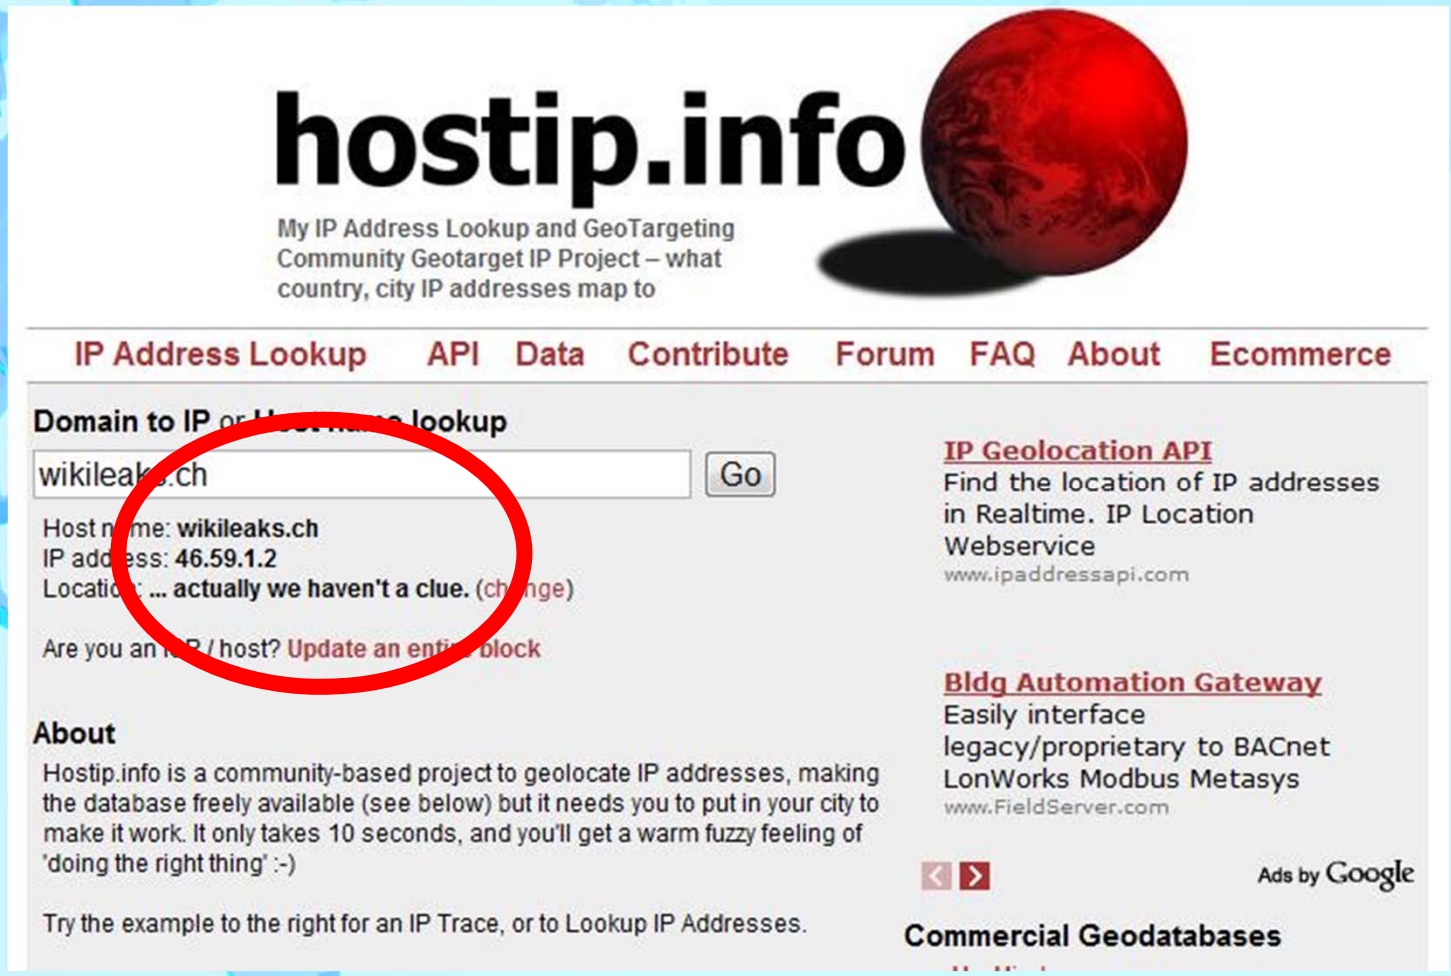

Where is the WikiLeaks site (46.59.1.2)?

IPligence

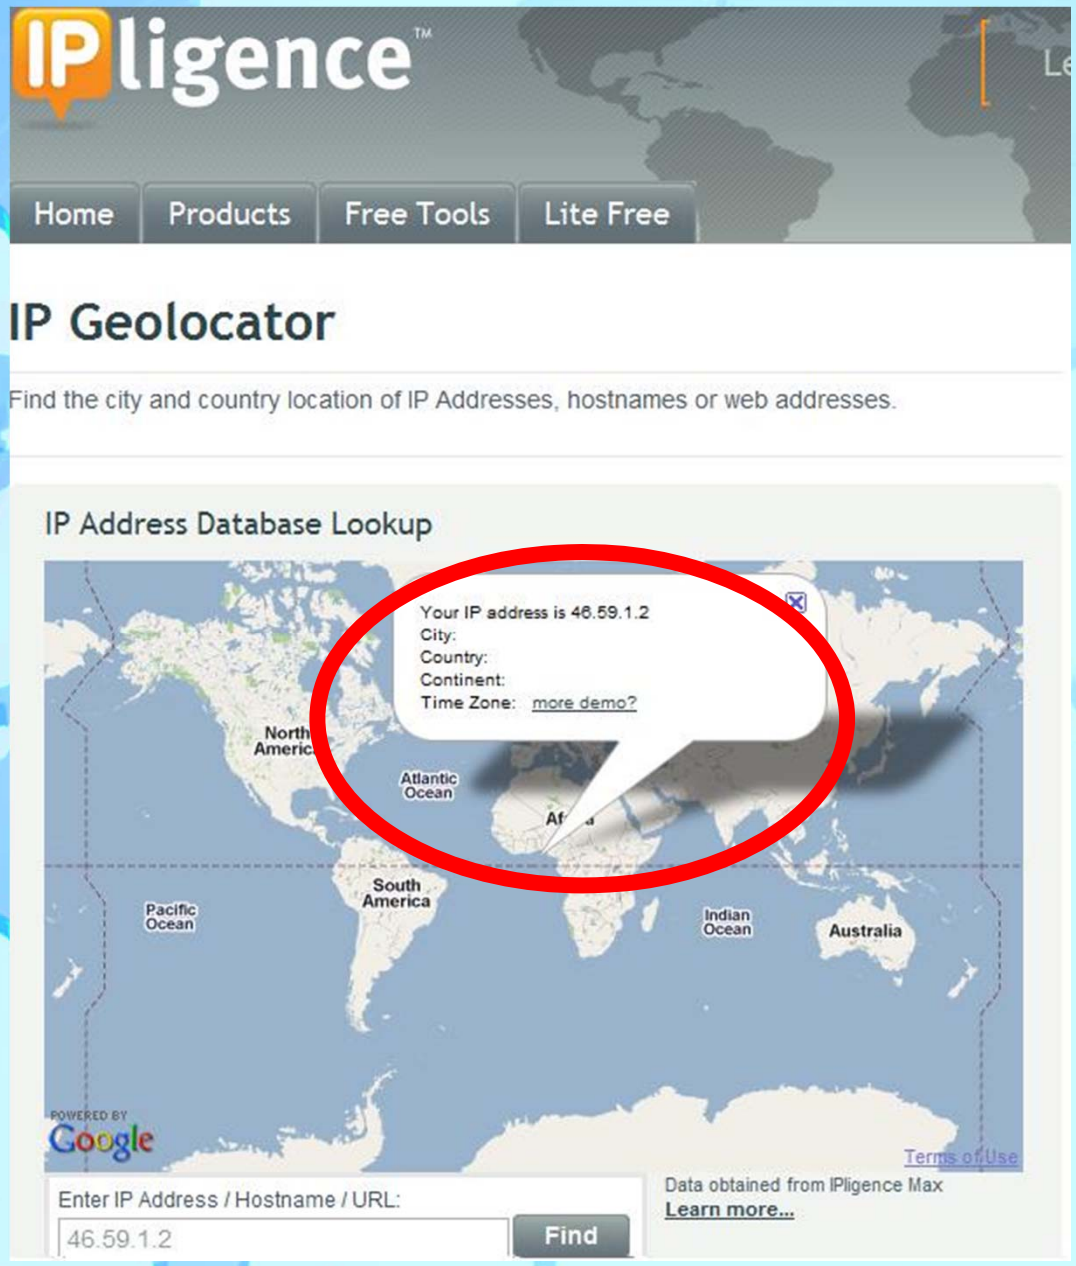

Where is the WikiLeaks site (46.59.1.2)?

MaxMind

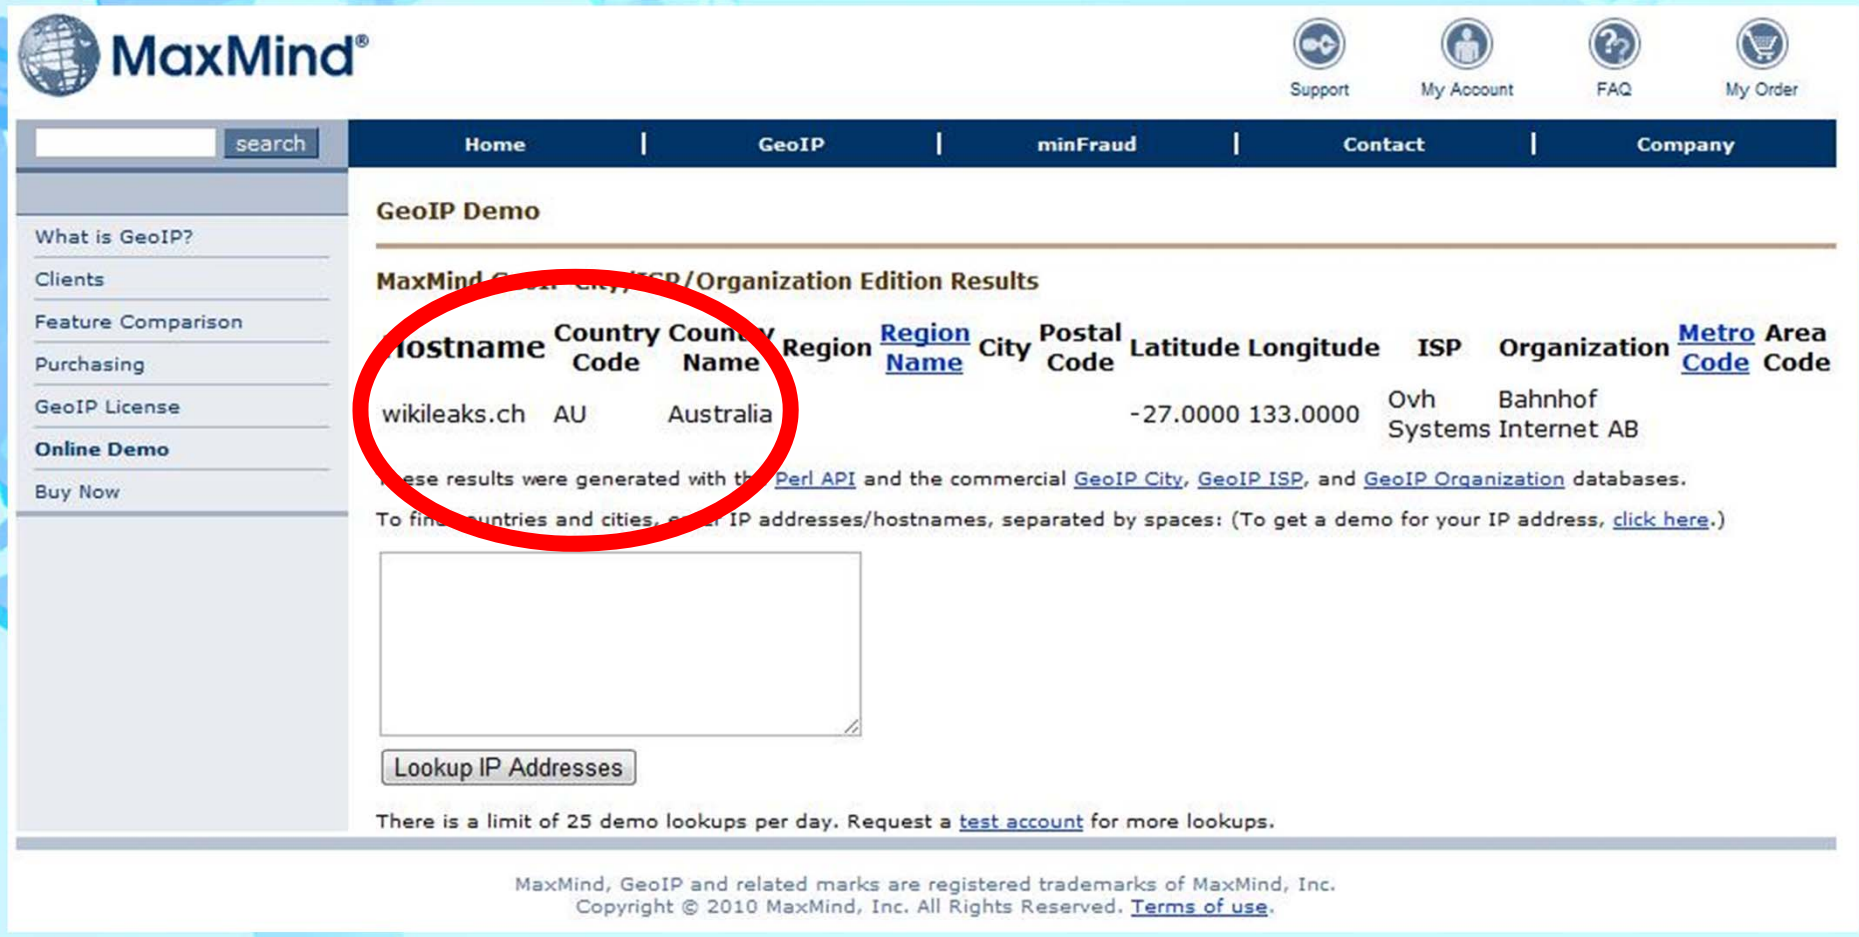

Where is the WikiLeaks site (46.59.1.2)?

Spotter

Results:

wikileaks.ch

IP Address: **46.59.1.2**

Status: **Host is available.**

Expected Coordinates: **59,35 17,94**

City-level Location: **Stockholm, Sweden**

Map:

WikiLeaks

Where is the WikiLeaks site (46.59.1.2)?

http://www.youtube.com/watch?v=qwlATf9xse4

WikiLeaks

Spotter's accuracy

- Estimation accuracy for reference node set

(thanks to Bradley@CAIDA)

- Spotter – „pure probabilistic method"
- Spotter city – population density as extra constraint

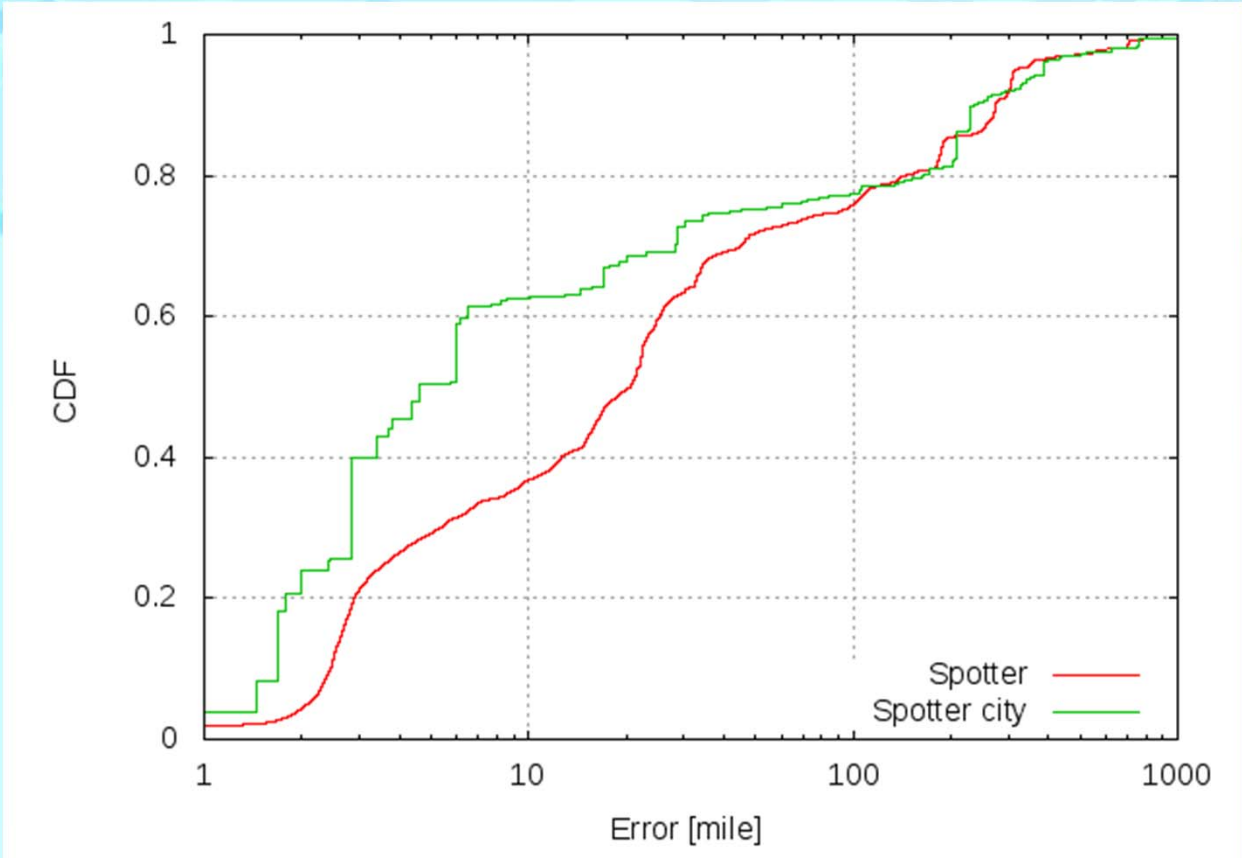

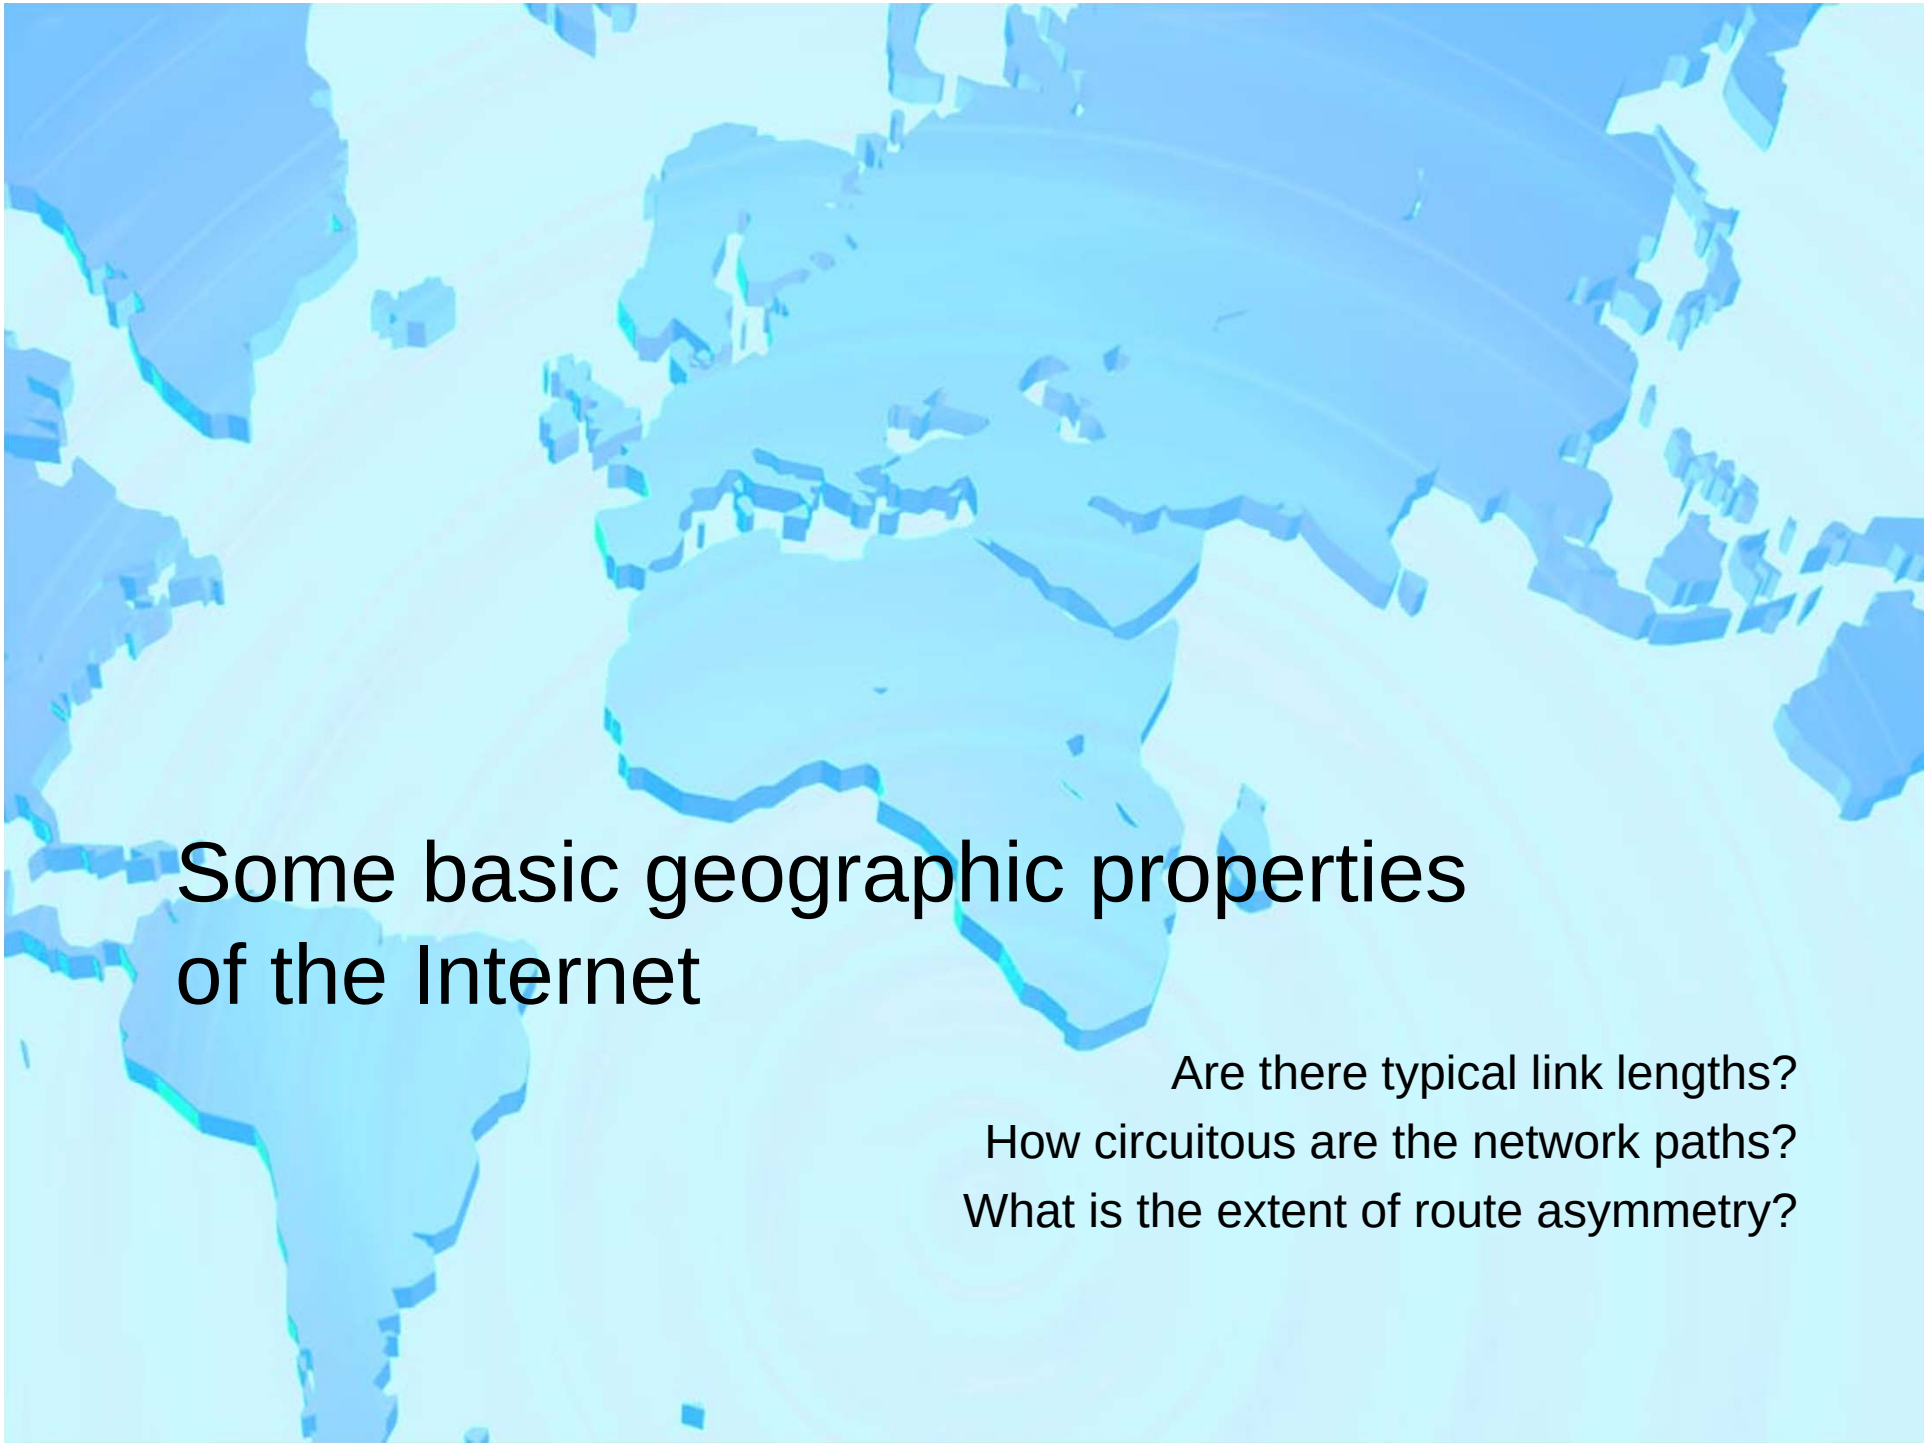

Some basic geographic properties of the Internet

Are there typical link lengths?

How circuitous are the network paths?

What is the extent of route asymmetry?

Router density visualization

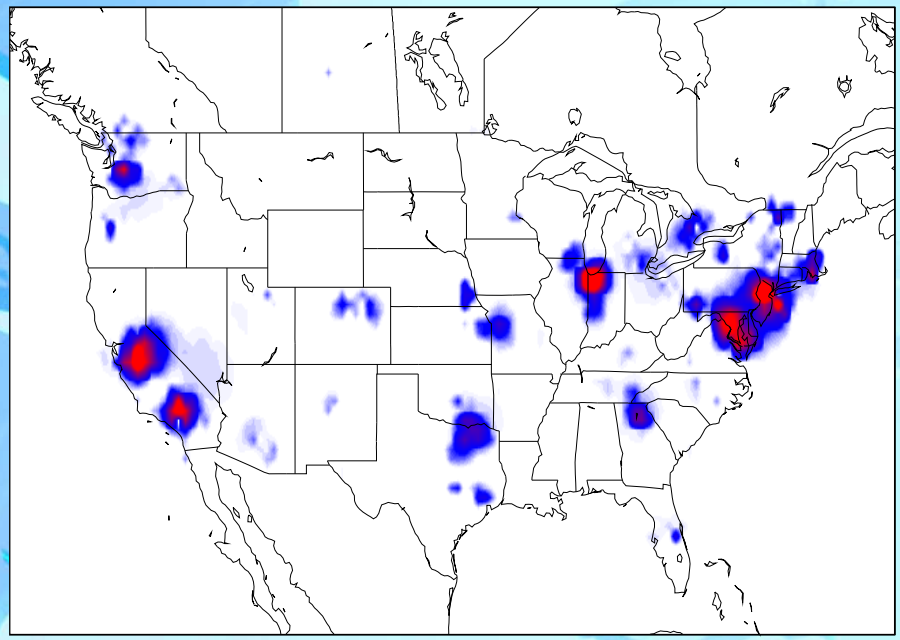

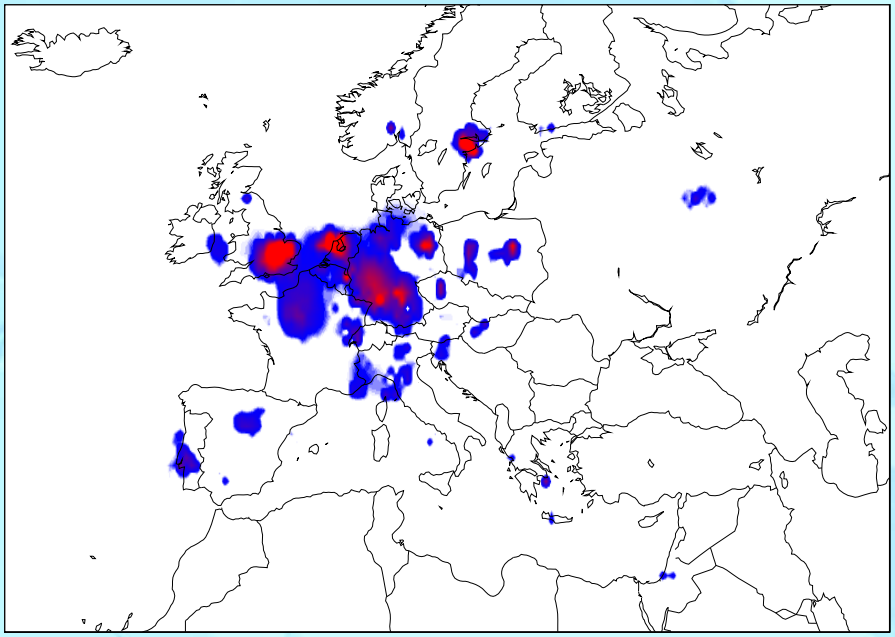

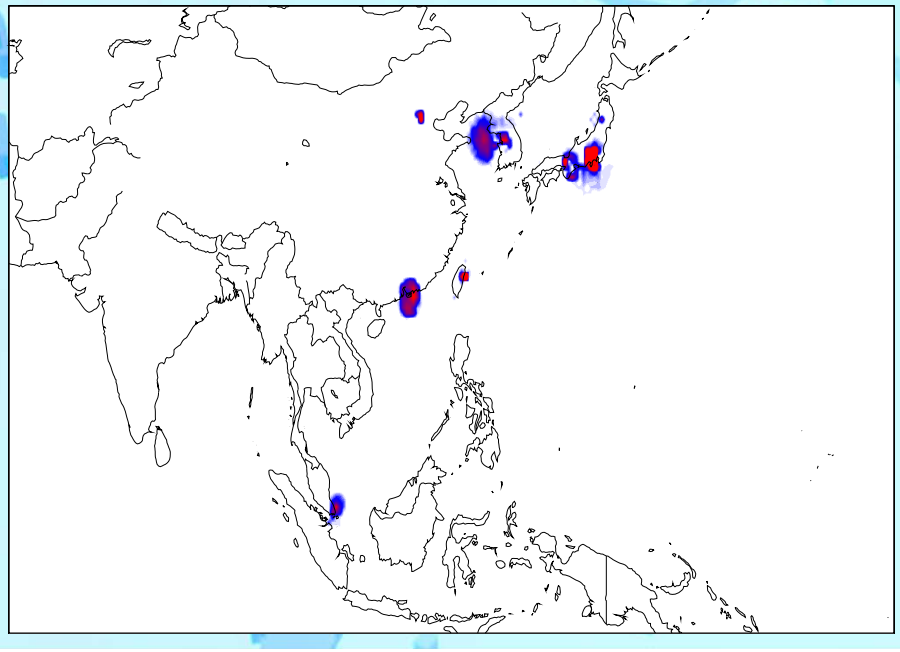

- routers of the full-mesh PlanetLab topology (700x700)
- discovered by traceroute
- both academic and commercial ISPs
- IP level network
- approx. 16.000 addresses


Link length frequency & distribution

Link length frequency

of identified IP level links: approx. 44.000

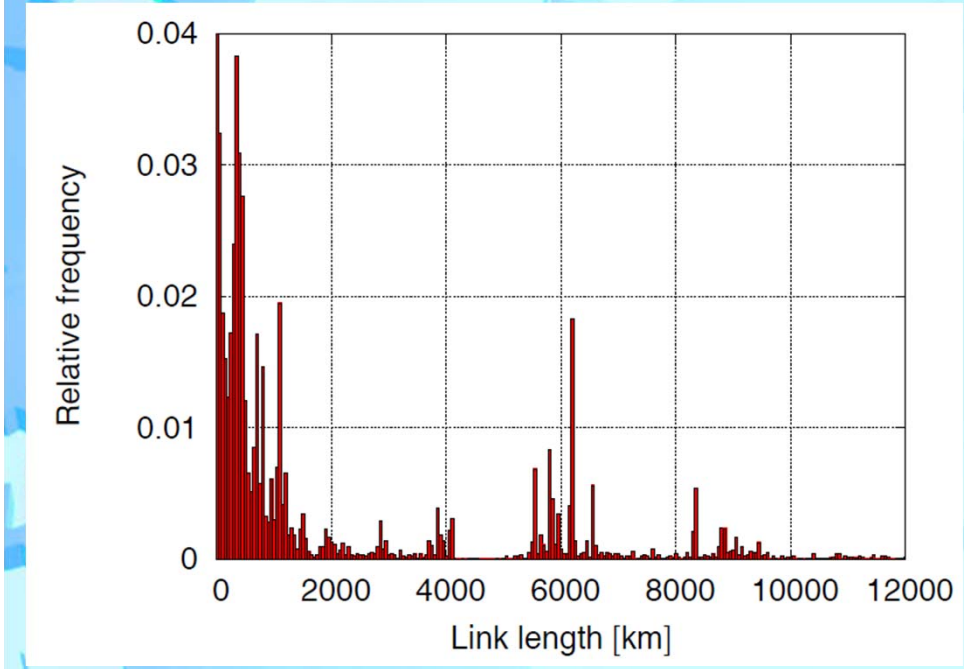

each unique link is represented once

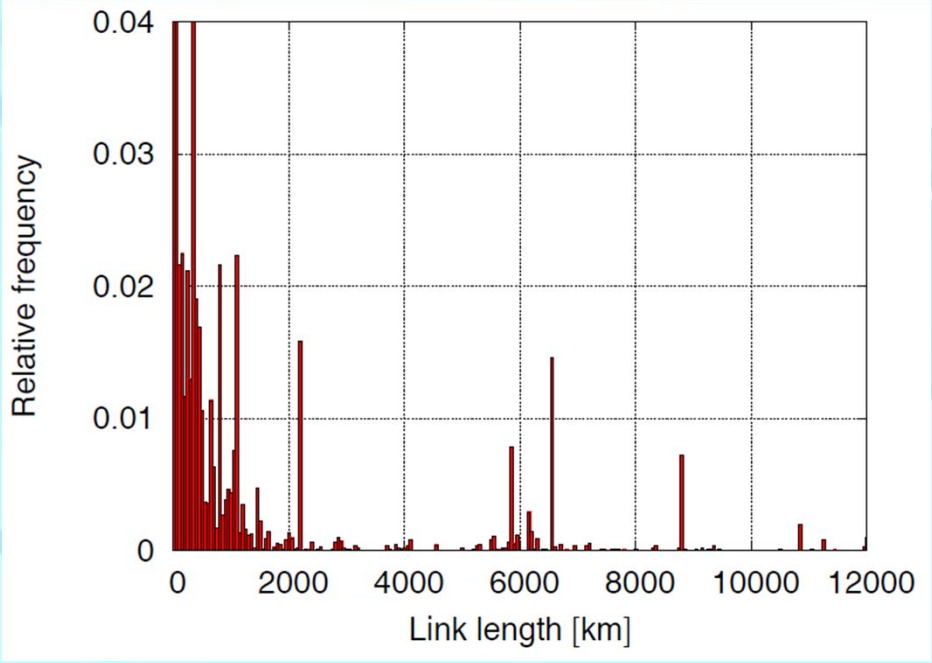

links are weighted up with their prevalence in the traceroute data set collected at one time

Link length distribution

- power law, where $P(d) \propto d^{-\delta}$?
- logarithmic relation, where $P(d) \propto 1/d$?
- model behind the observation?

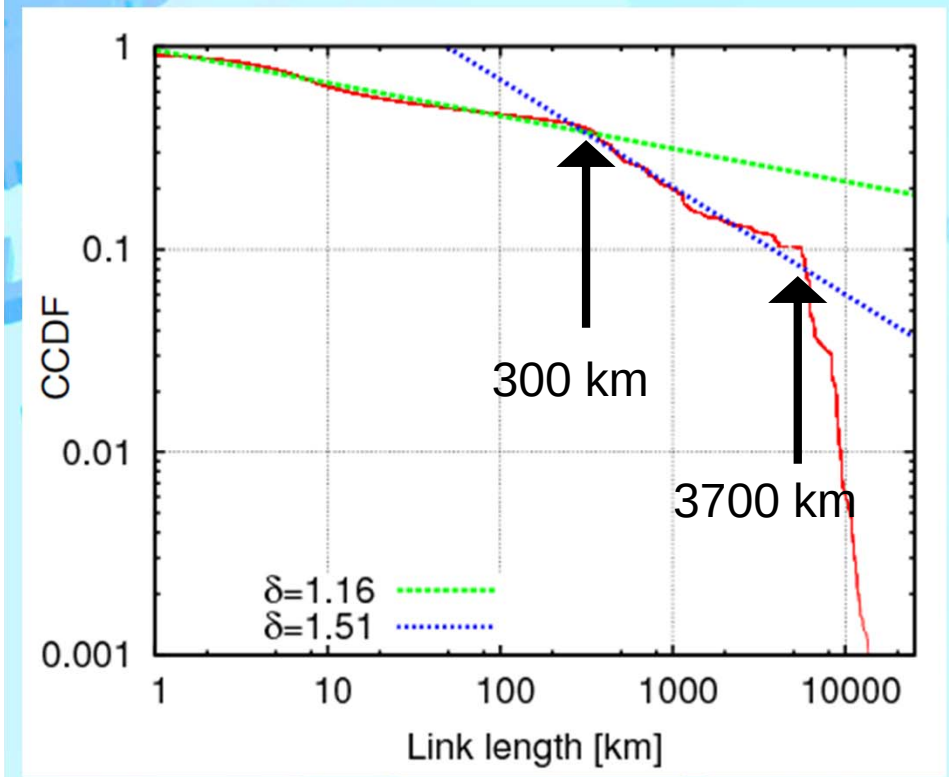

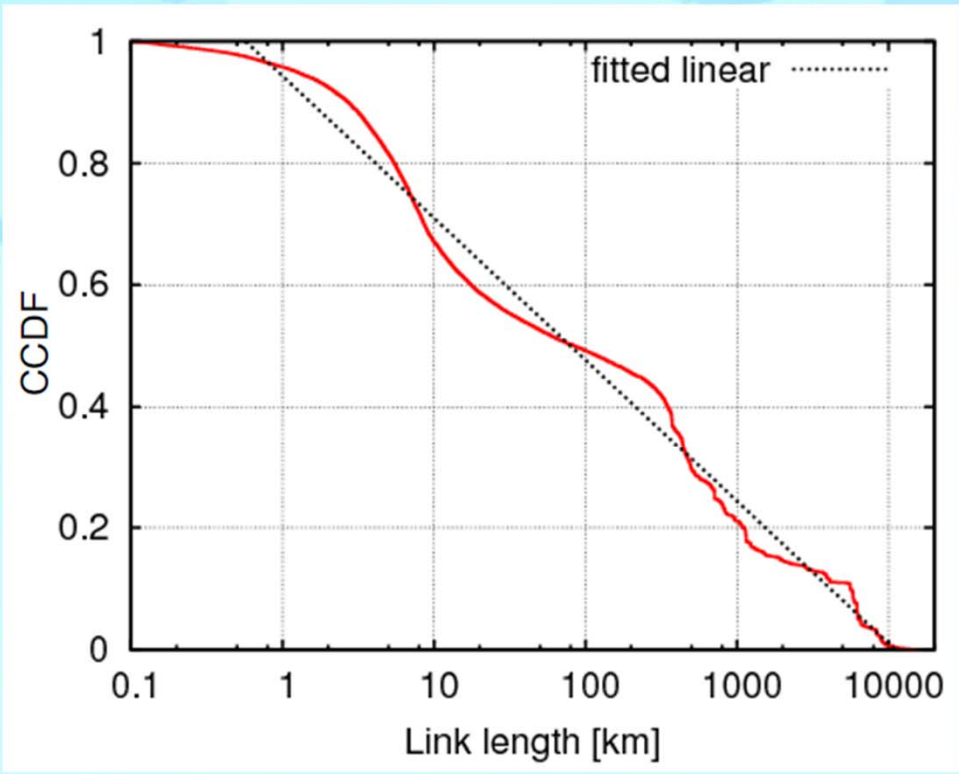


Circuitousness of network paths

Circuitousness of network paths

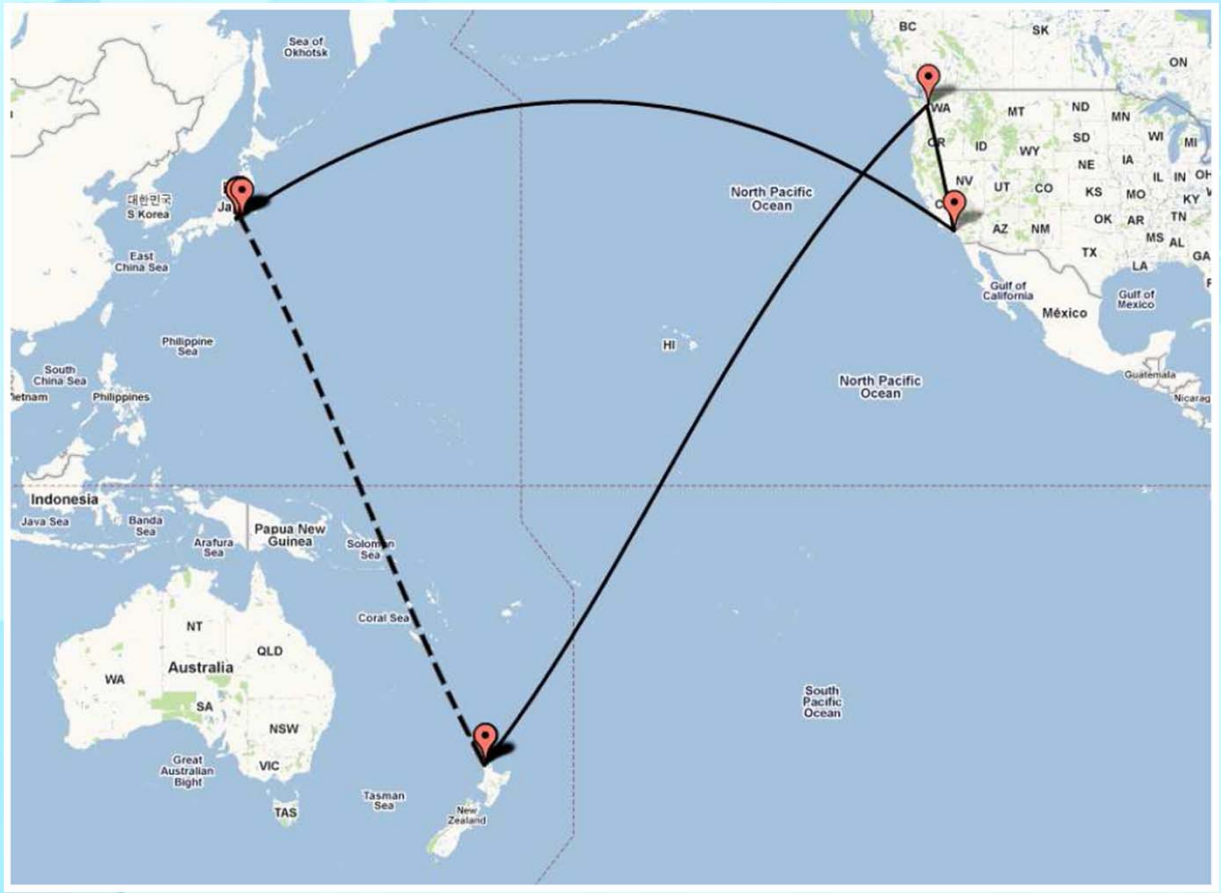

$$\text{relative circuitousness} = \frac{\text{aggregated link length}}{\text{spherical distance of the endpoints}}$$

Circuitousness of network paths

Intra-continental routes: significant circuitousness

Inter-continental routes: determined by the intercontinental gateways

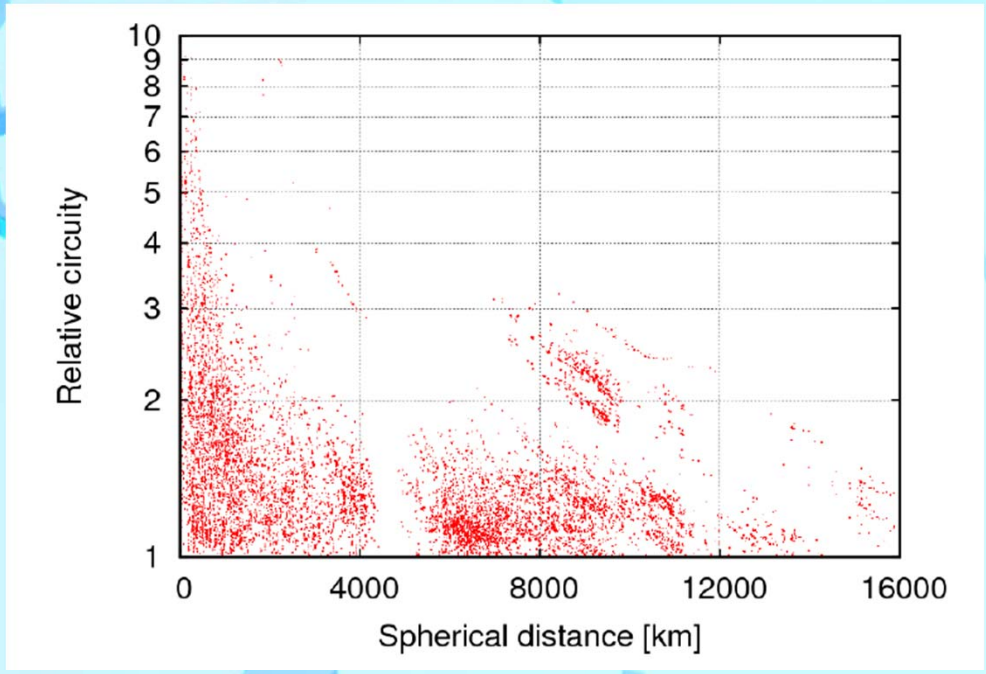

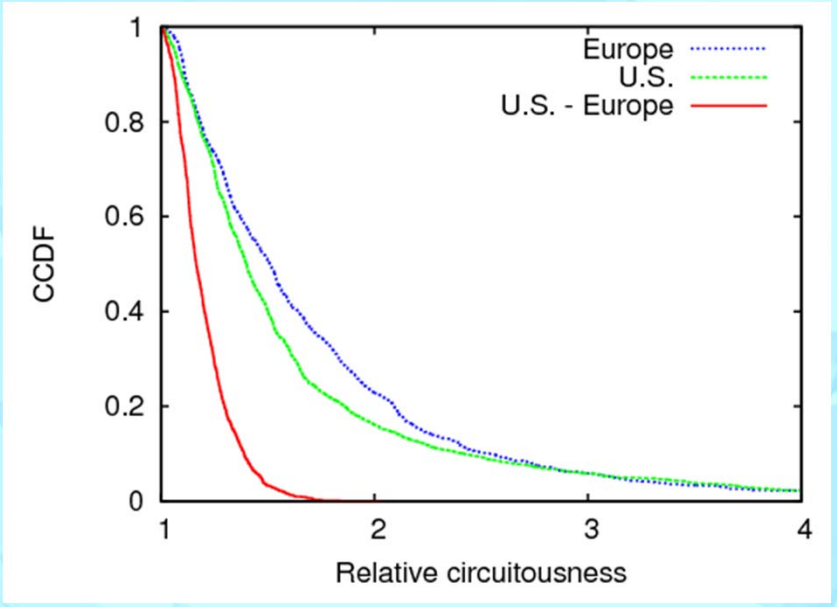

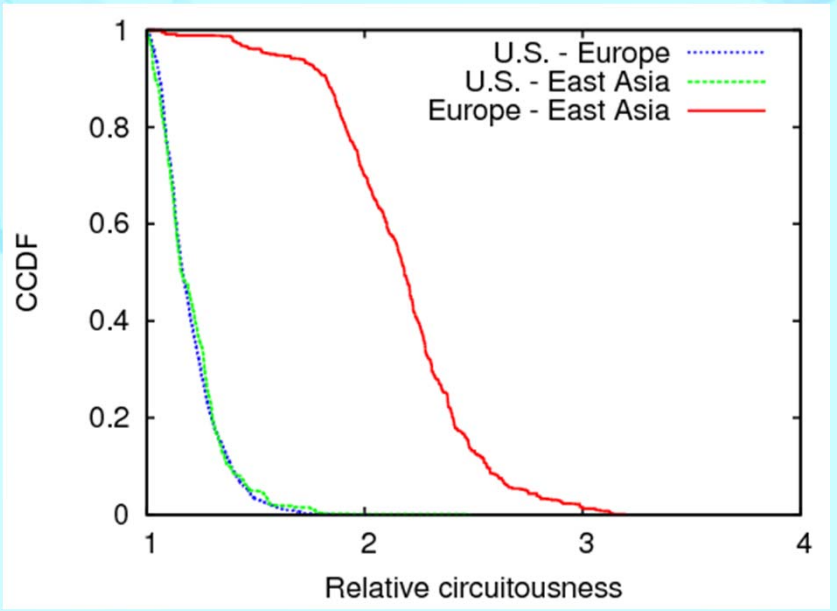

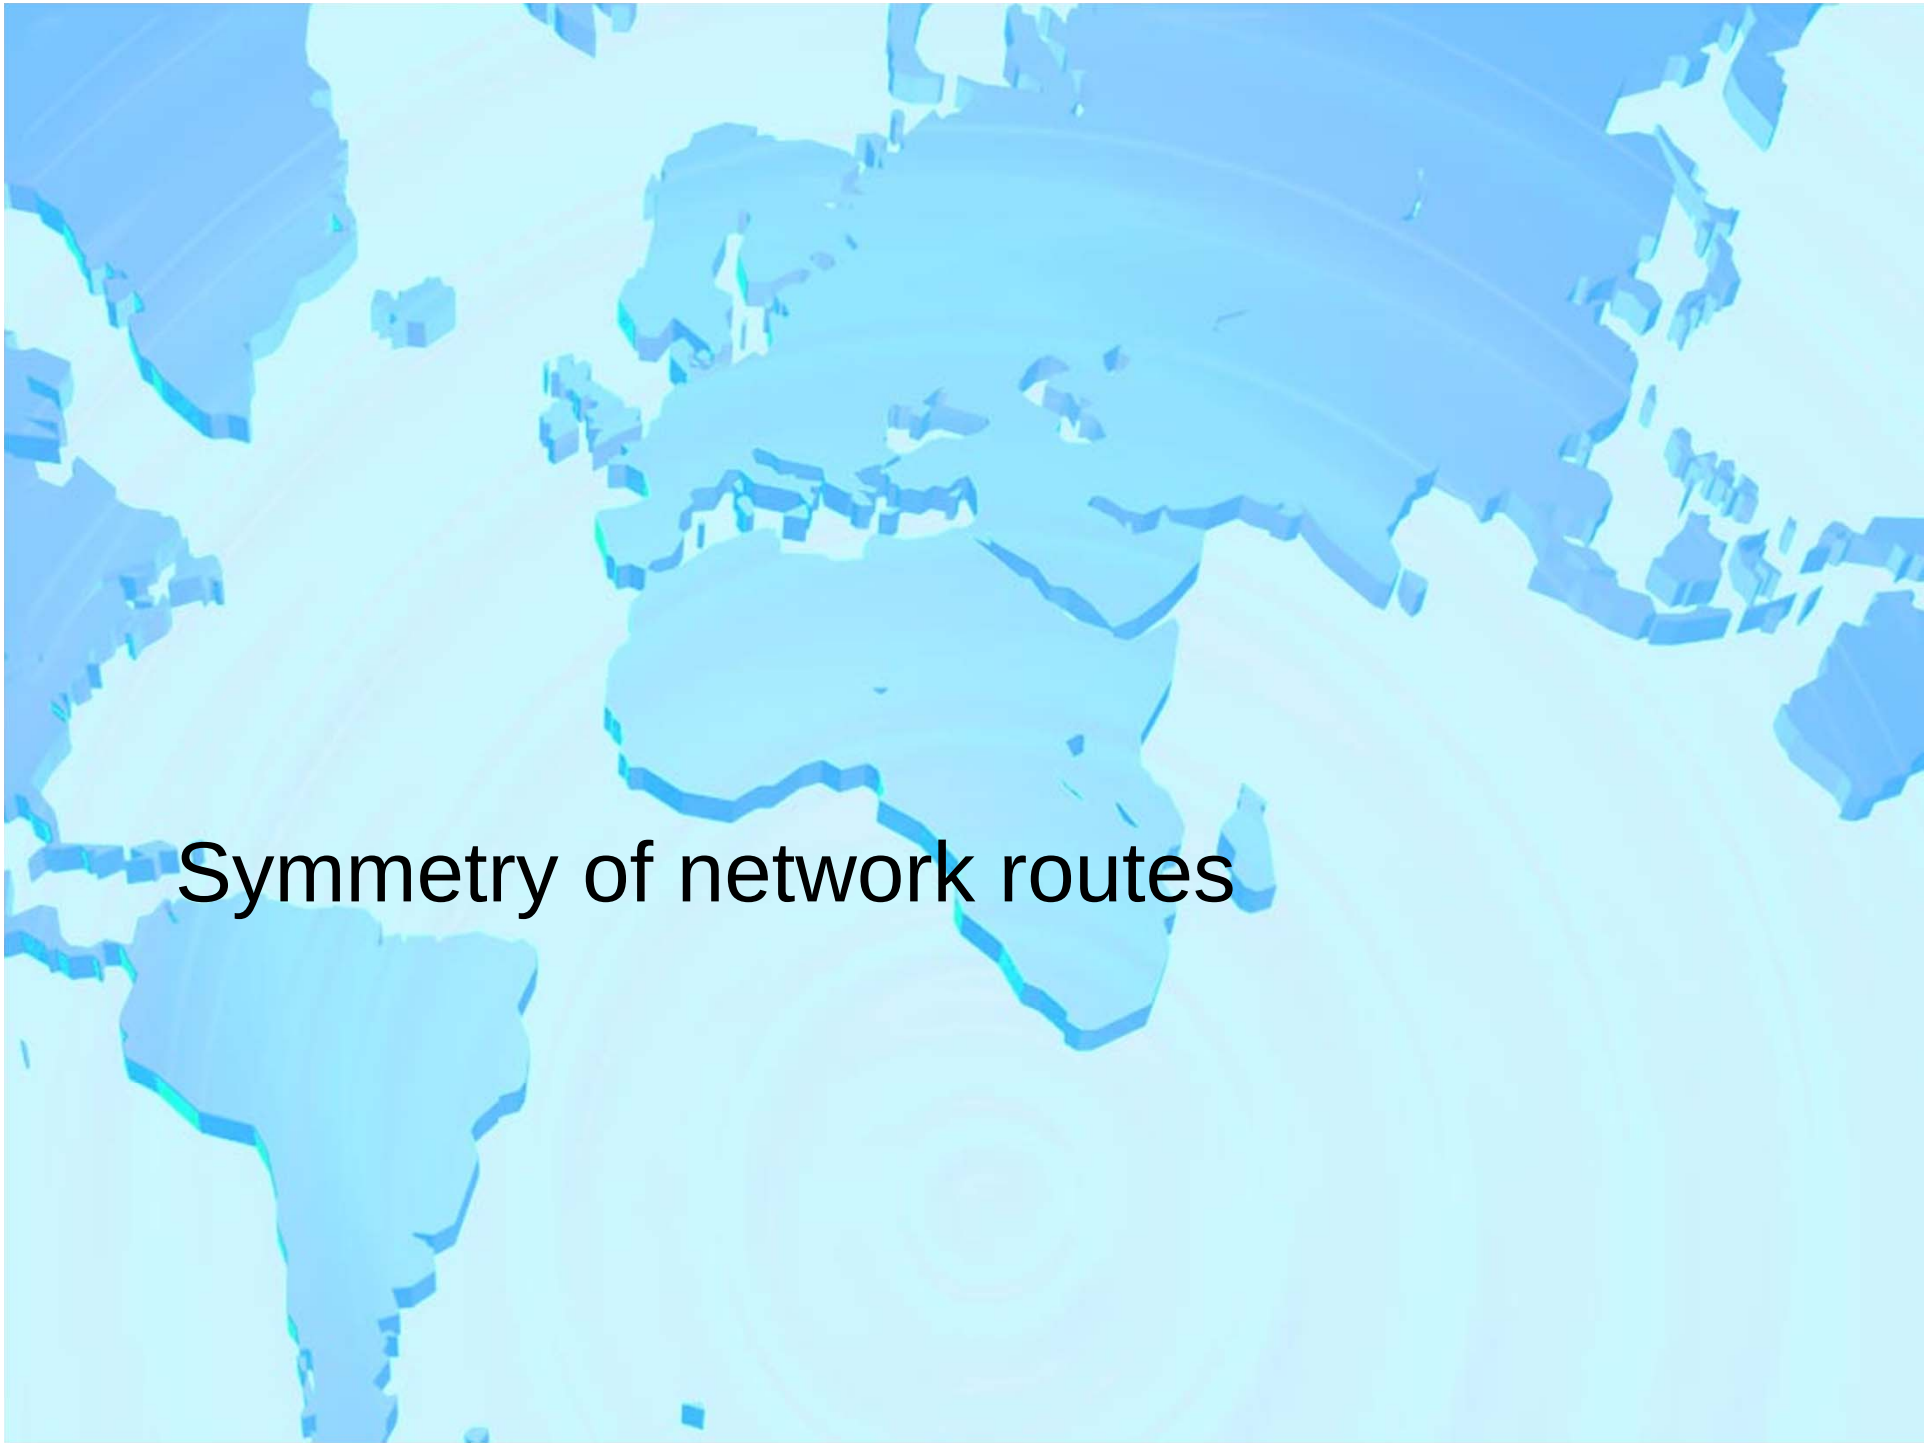
Symmetry of network routes

Symmetry of network routes

Ratio of the route lengths of the forward and the backward directions.

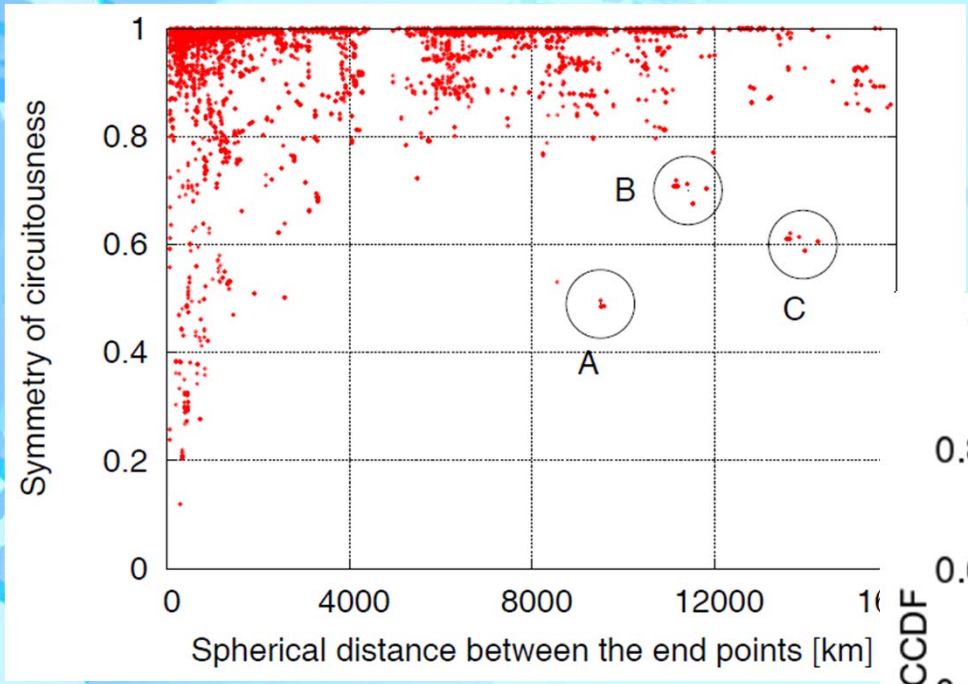

A: United Kingdom – Hong Kong
B: California, USA – Hong Kong
C: California, USA – Singapore

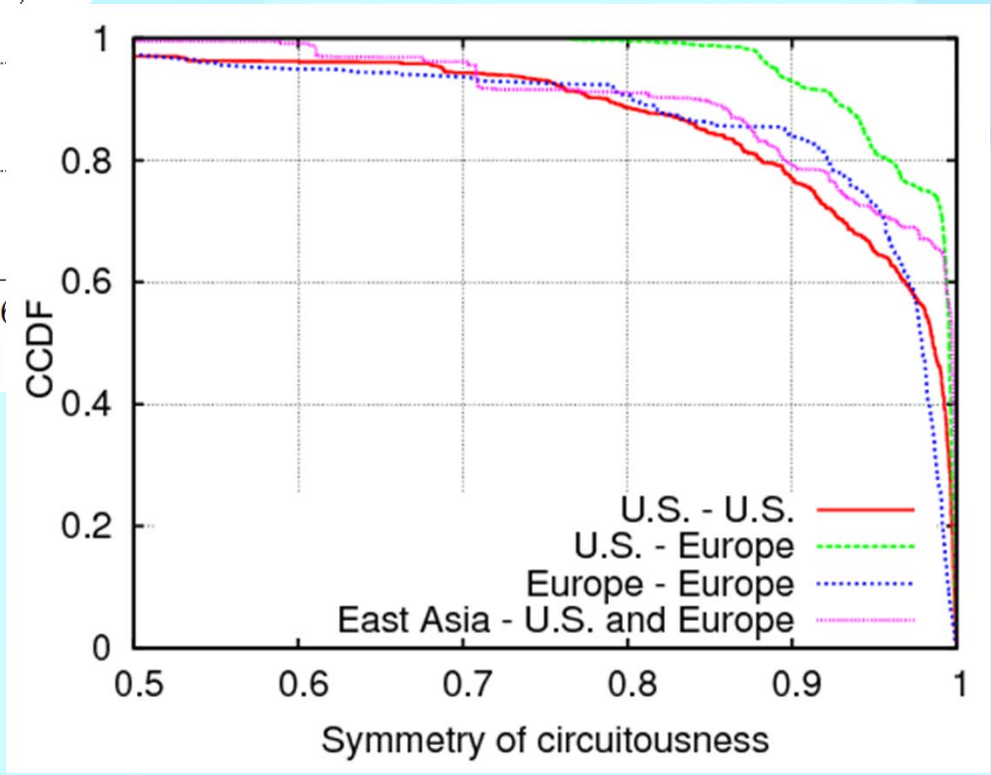

Around the world in 300 ms

from the United Kingdom to Hong Kong:

- forward direction: eastward through Europe and Asia
- backward direction: eastward through the USA

Where is the Internet?*
And what does it look like?*

*as our dataset shows

Thank you for your attention!

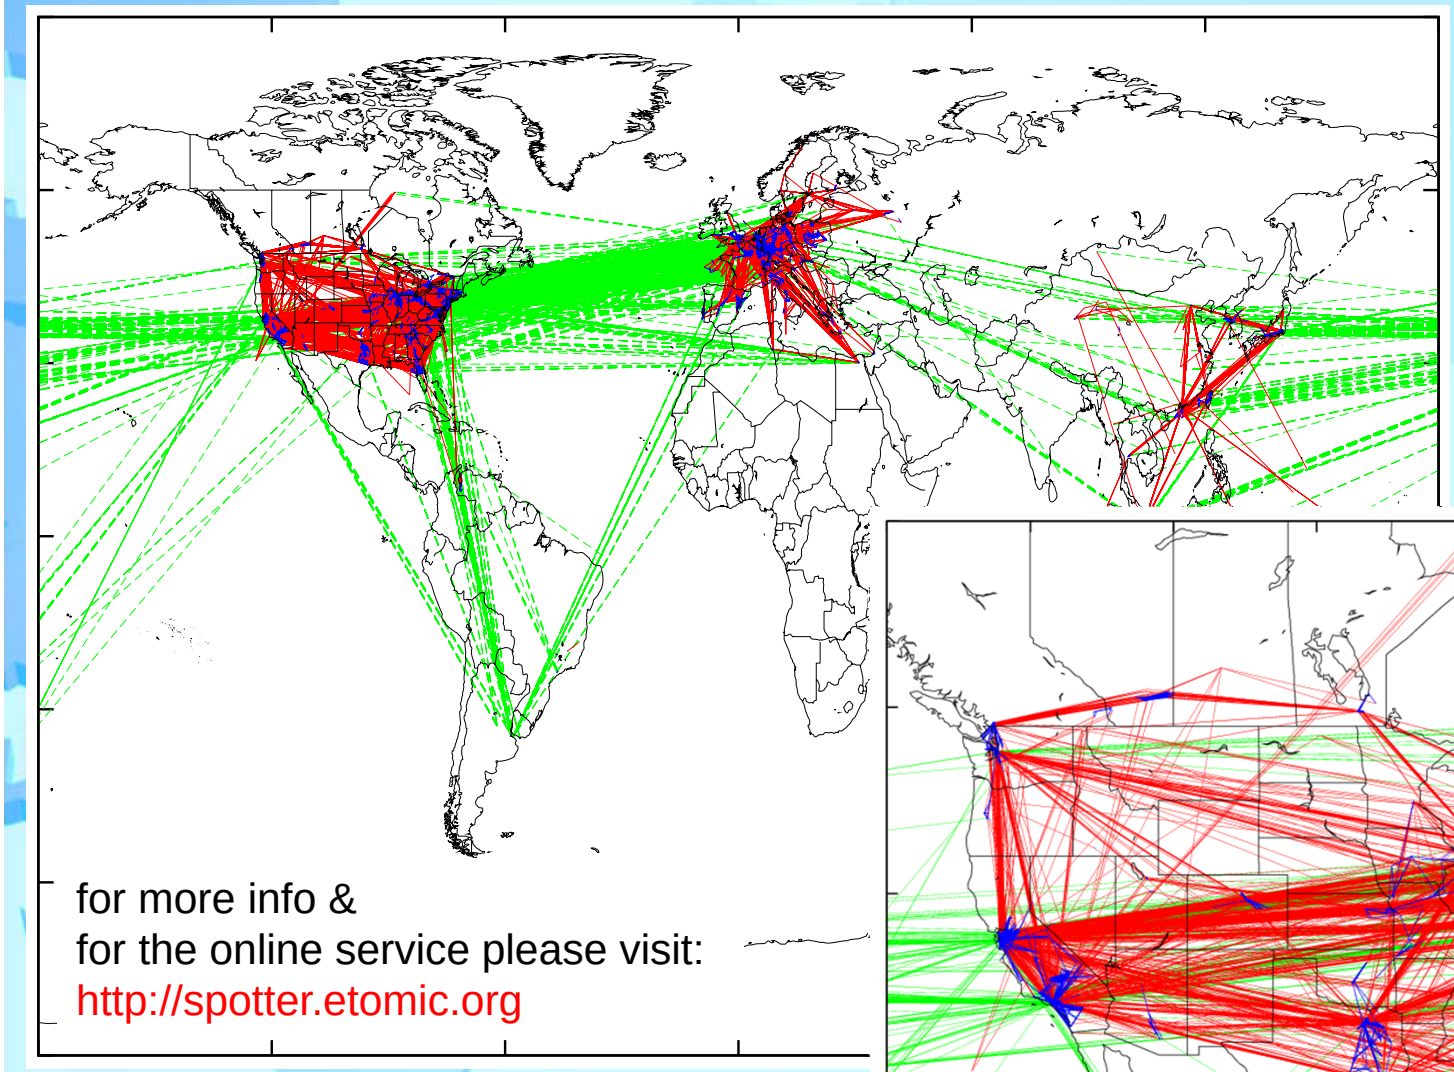

for more info &
for the online service please visit:
http://spotter.etomic.org

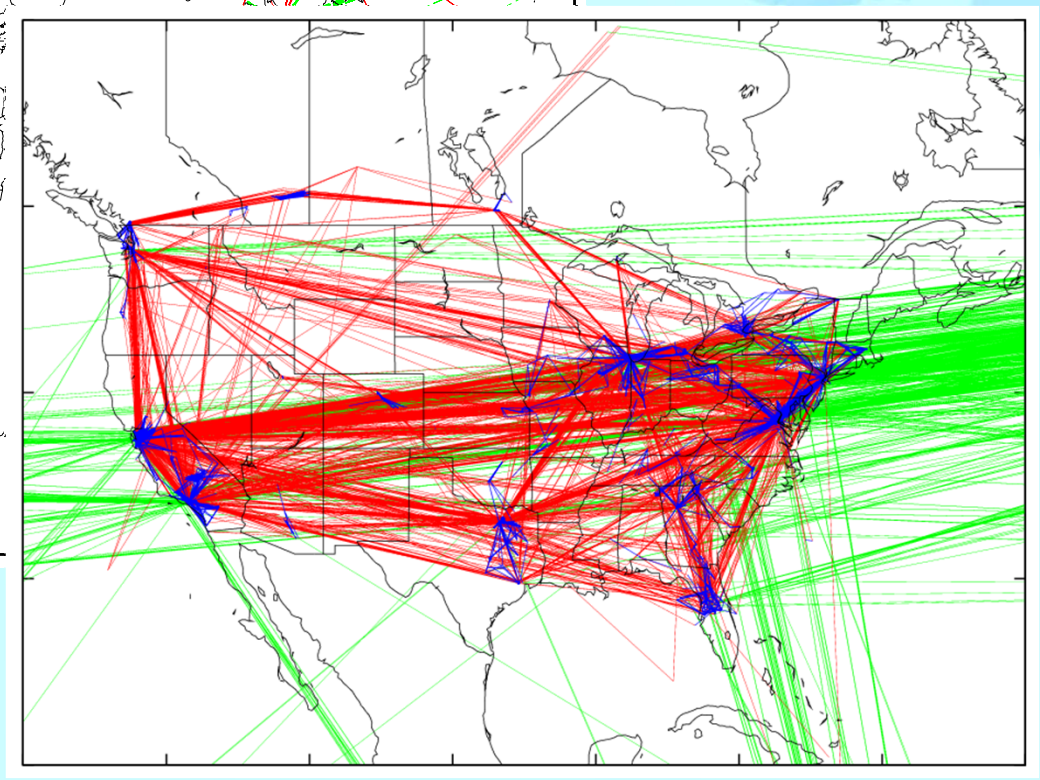

Spotter
GEOLOCATION SERVICE