Archipelago

Measurement Infrastructure

kc claffy
CAIDA

AAF Workshop
Nov 23, 2009

Outline

* Focus and Architecture
* Monitor Deployment
* Measurements
* Future Work

Introduction

* Archipelago (Ark) is CAIDA's next-generation active measurement infrastructure
  * evolution of the skitter infrastructure
* in production since Sep 12, 2007

Focus

* easy development and rapid prototyping
  * lower barriers => implement better measurements faster with lower cost
    - measurement infrastructures notoriously lack funding
  * raise level of abstraction with high-level API and scripting language
    - inspiration from Scriptroute, Metasploit, Scapy, Racket

Focus

* dynamic and coordinated measurements
  * take advantage of multiple distributed measurement nodes in sophisticated ways
    - one measurement triggers another measurement
    - use multiple nodes to divide and conquer
    - synchronize measurements
* for example: Doubletree; tomography; Rocketfuel-like targeted discovery of a single network's topology

Focus

* measurement services
  * build upon the work of others; share services between measurement activities
    - for example, on-demand traceroute/ping service; IP-to-AS mapping service
* similiar in goal to service-oriented architecture (SOA) but at finer granularity and without the complexity

Architecture

- Ark is composed of measurement nodes (machines) located in various networks worldwide
  - many thanks to the organizations hosting Ark boxes
  - please contact us if you want to host an Ark box
- Ark employs a tuple space to enable communication and coordination
  - a tuple space is a distributed shared memory combined with a small number of easy-to-use operations
  - a tuple space stores tuples, which are arrays of simple values (strings and numbers), and clients retrieve tuples by pattern matching

Architecture

* use tuple space for decentralized (that is, peer-to-peer) communication, interaction, and coordination

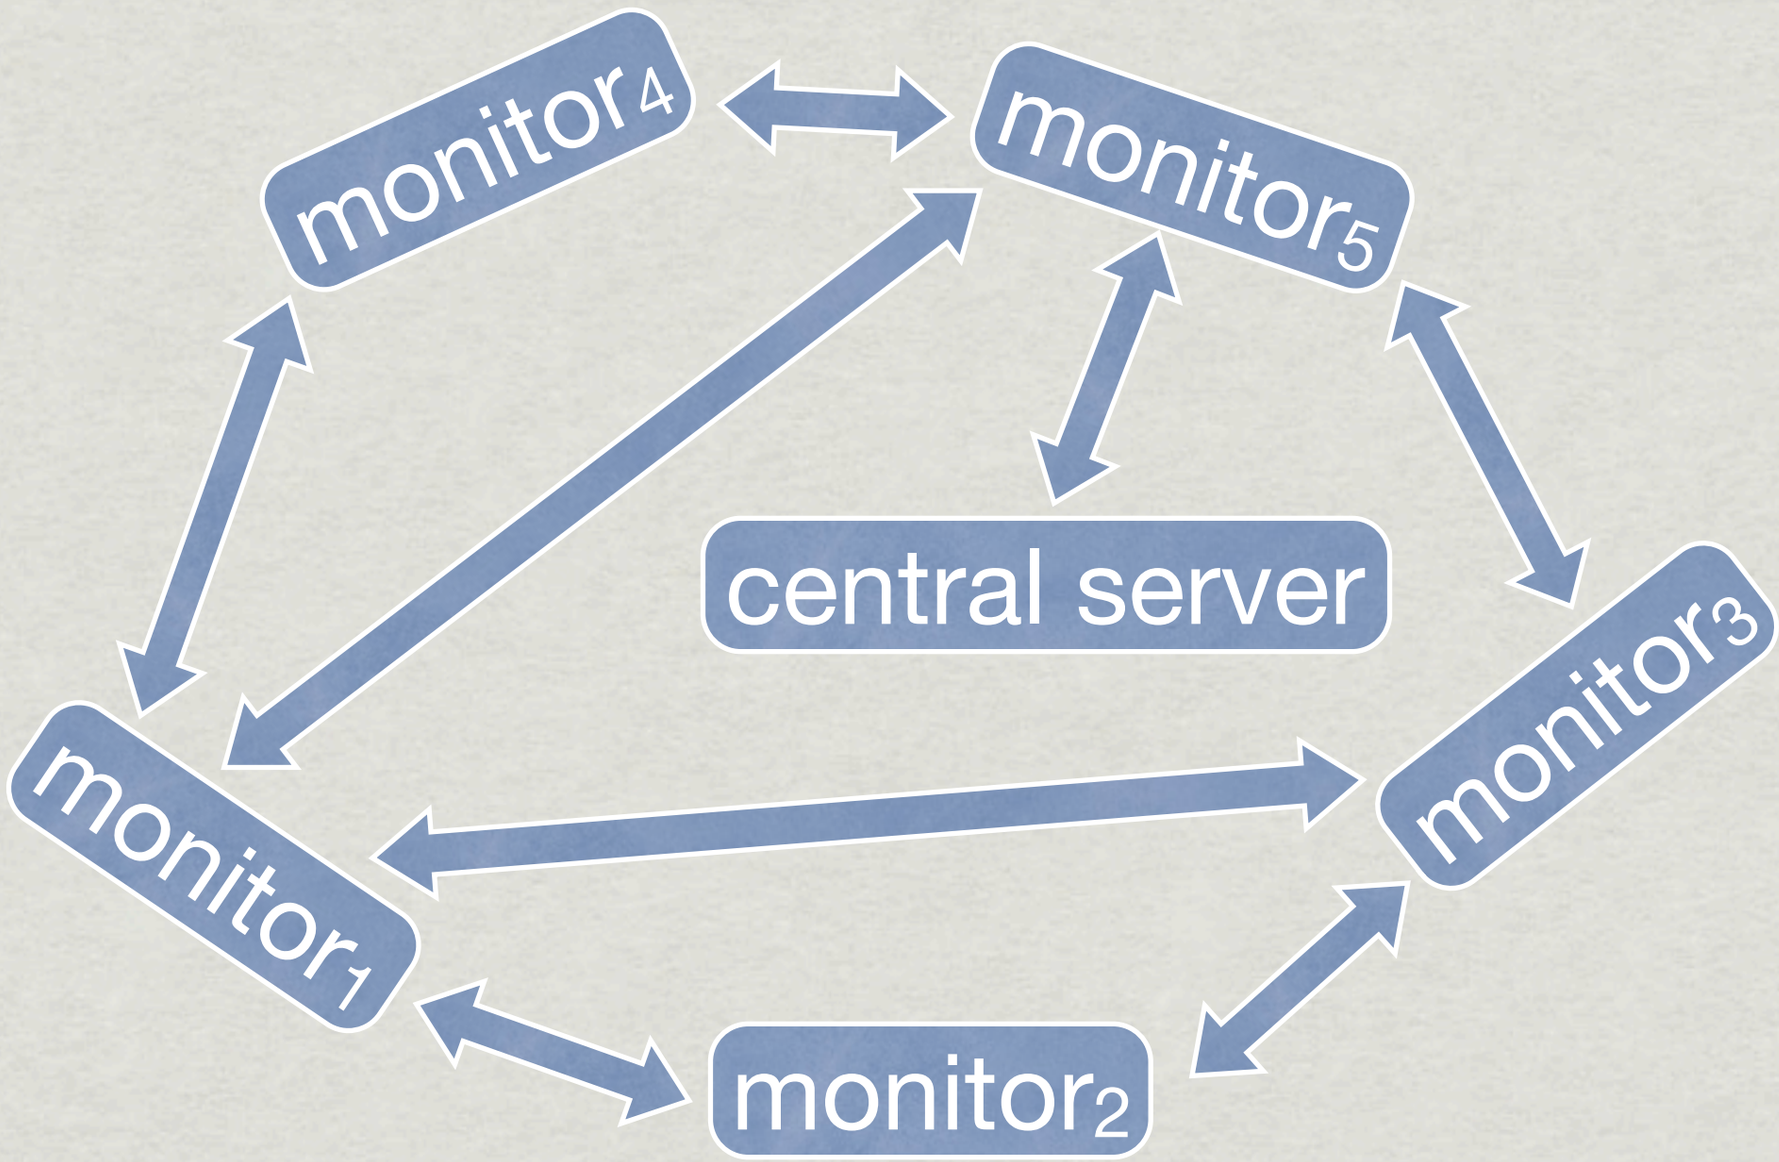

Monitor Deployment

* 38 monitors in 24 countries

	Continent
13	North America
2	South America
12	Europe
1	Africa
5	Asia
3	Oceania

	Organization
20	academic
10	research network
4	network infrastructure
2	commercial network
1	community network
1	military research

Measurements

* IPv4 Routed /24 Topology
* IPv4 Routed /24 AS Links
* IPv**6** Topology
* DNS Names
* DNS Query/Response Traffic
* Spoofer Project Collaboration

IPv4 Routed /24 Topology

- ongoing large-scale topology measurements
  - ICMP Paris traceroute to every routed /24 (7.4 million)
  - running scamper
    - written by Matthew Luckie of WAND, University of Waikato
- group monitors into teams and dynamically divide up the measurement work among team members
  - 13-member team probes every /24 in 48 hrs at 100pps
  - only one monitor probes each /24 per cycle
  - 3 teams active

IPv4 Routed /24 Topology

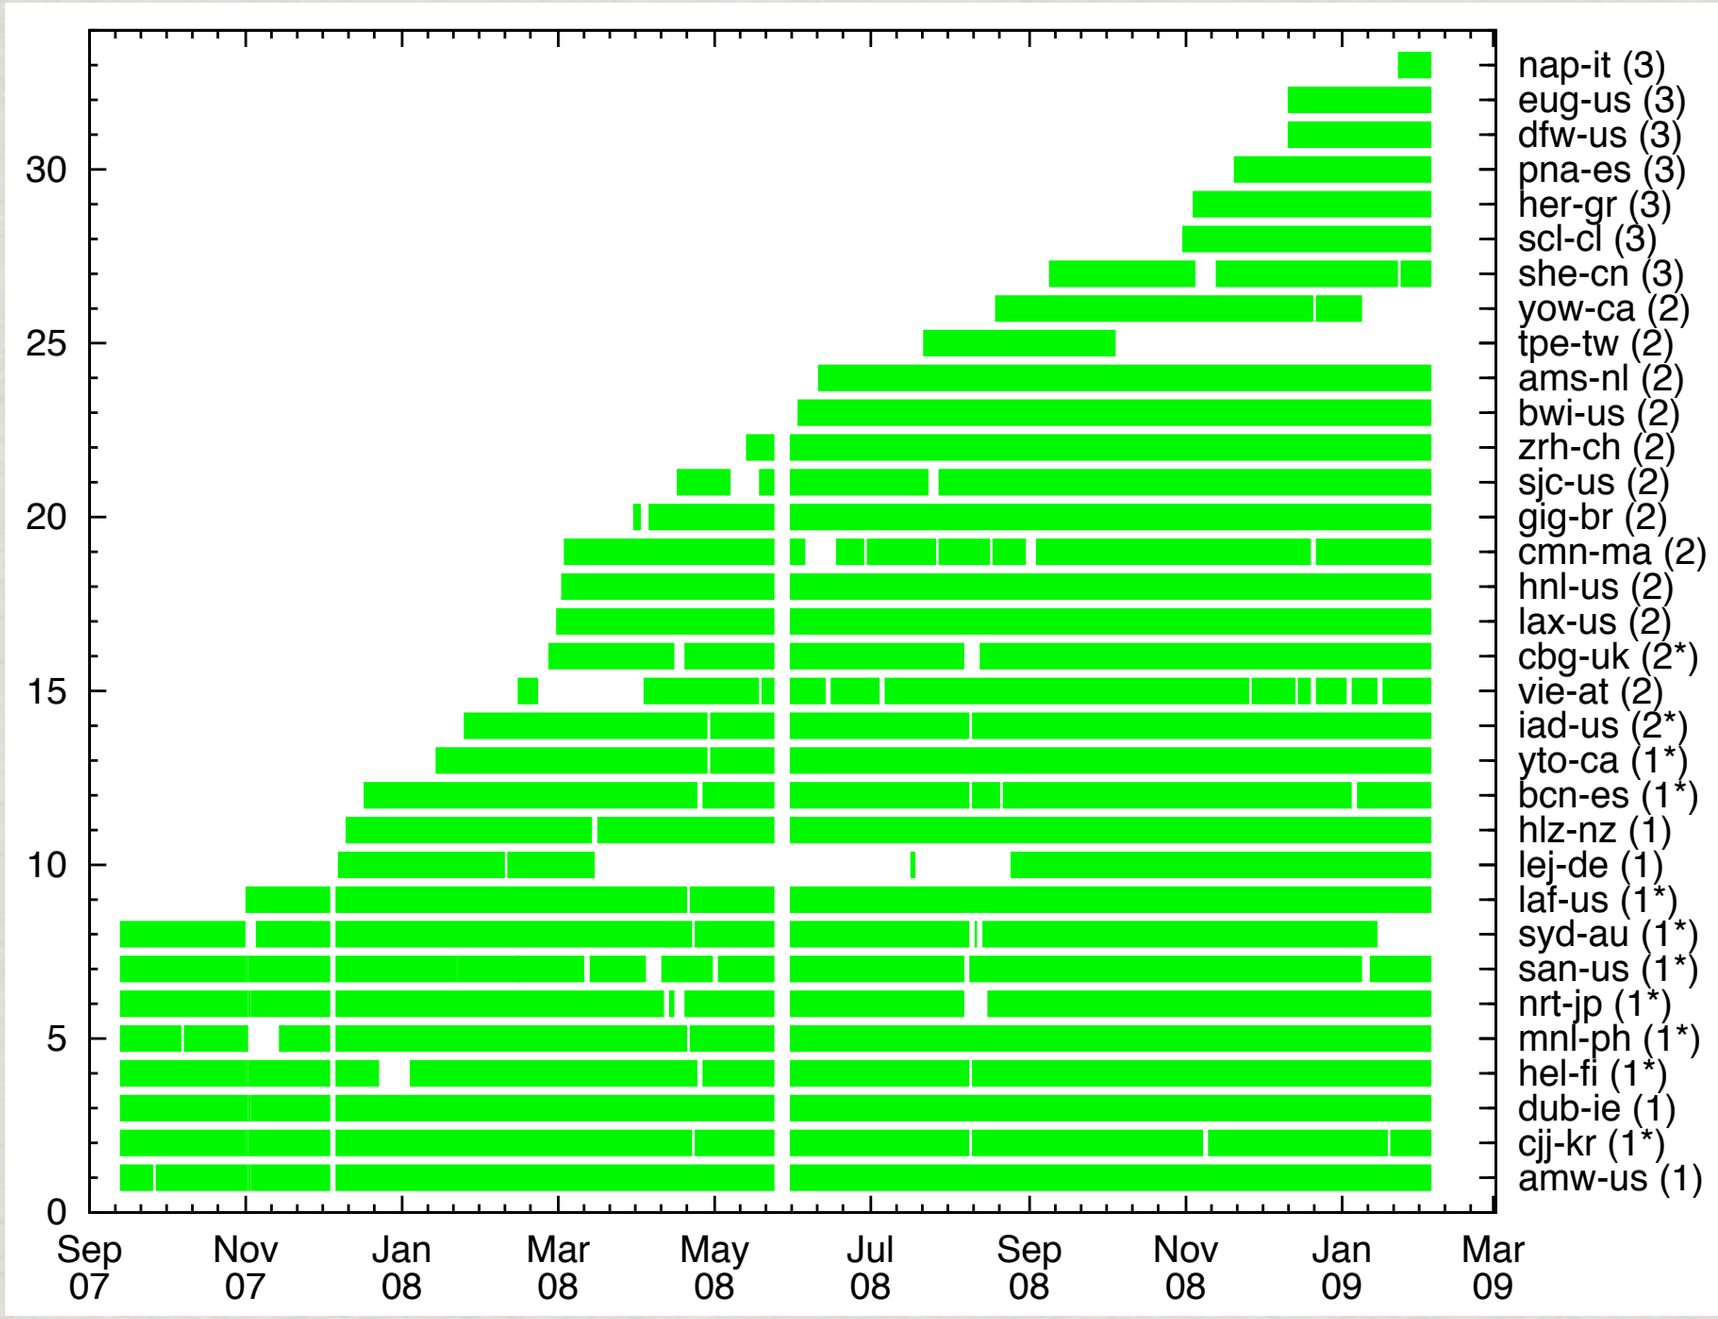

Sep 2007 to Jan 2009 (17 months): 2.5 billion traceroutes; 1.0TB data

IPv4 Routed /24 Topology

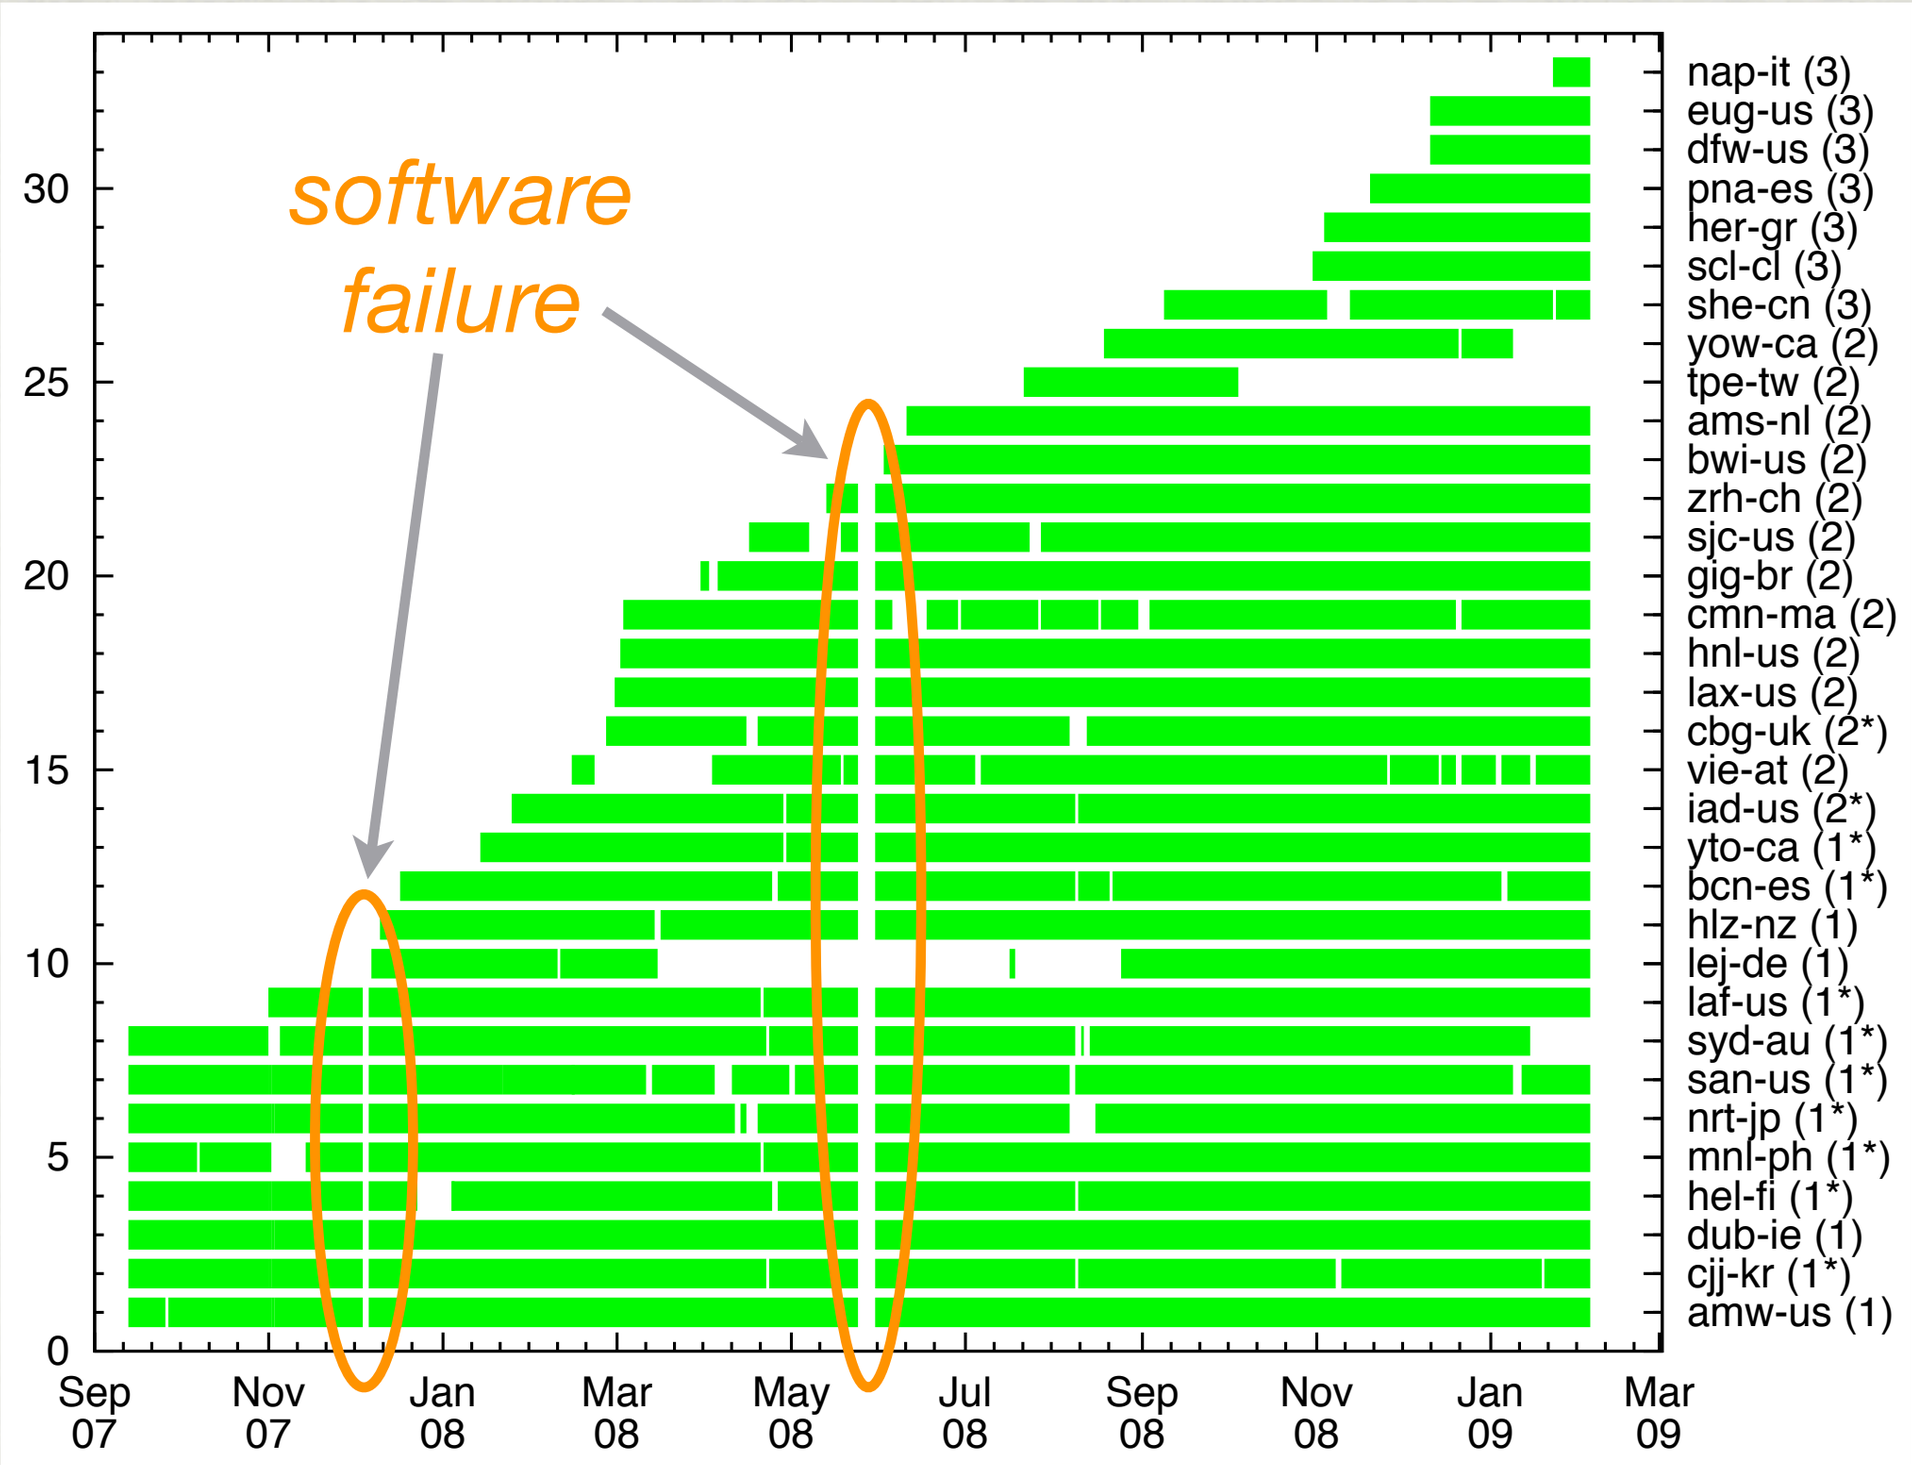

Sep 2007 to Jan 2009 (17 months): 2.5 billion traceroutes; 1.0TB data

IPv4 Routed /24 Topology

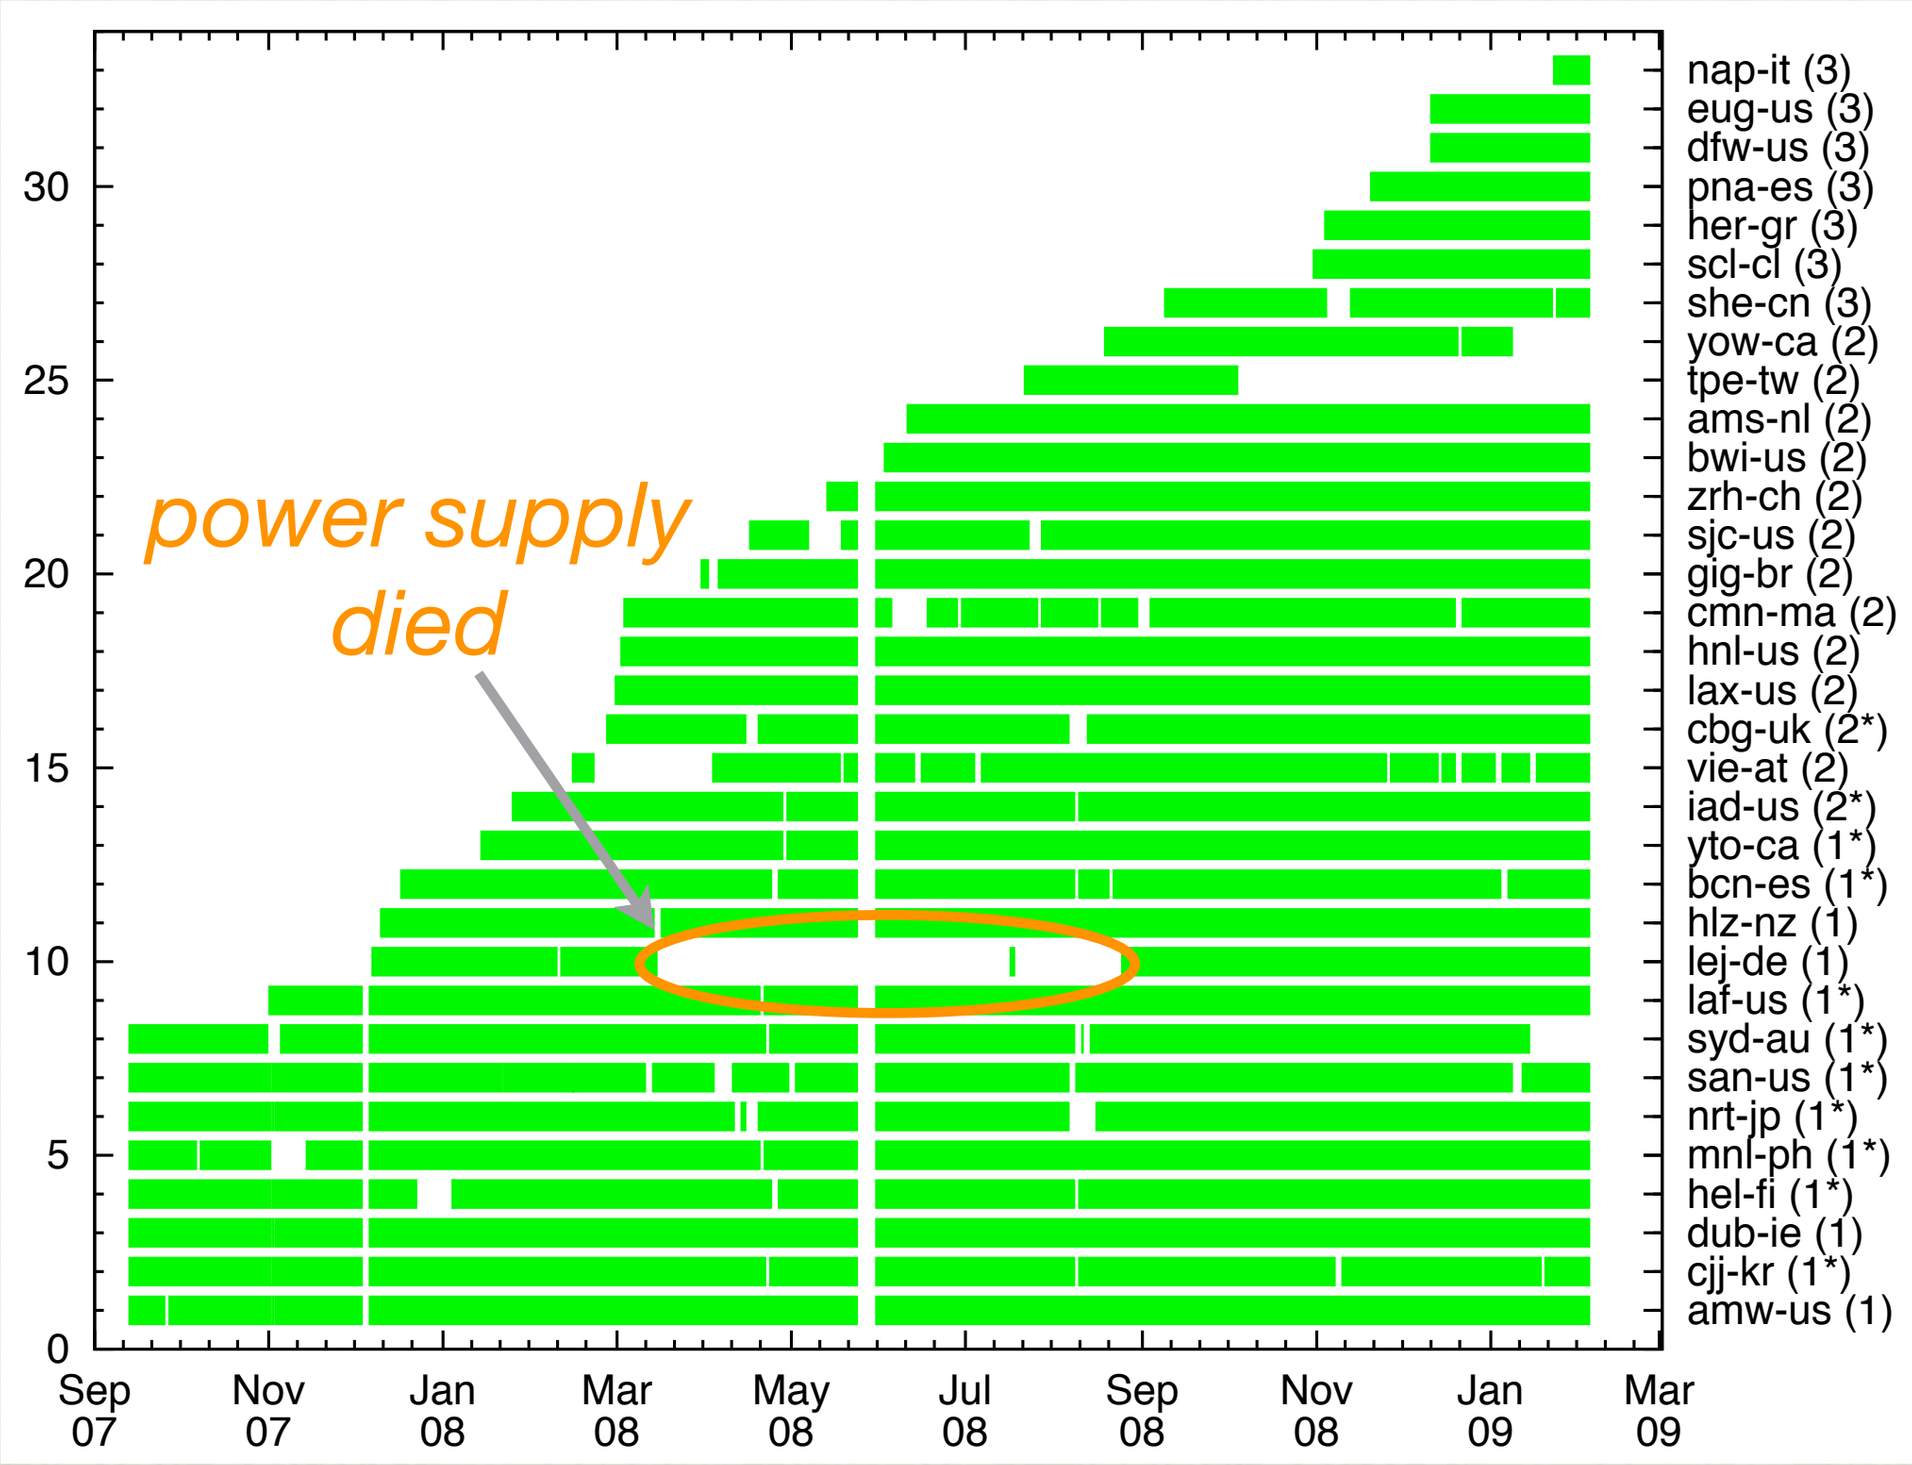

Sep 2007 to Jan 2009 (17 months): 2.5 billion traceroutes; 1.0TB data

IPv4 Routed /24 Topology

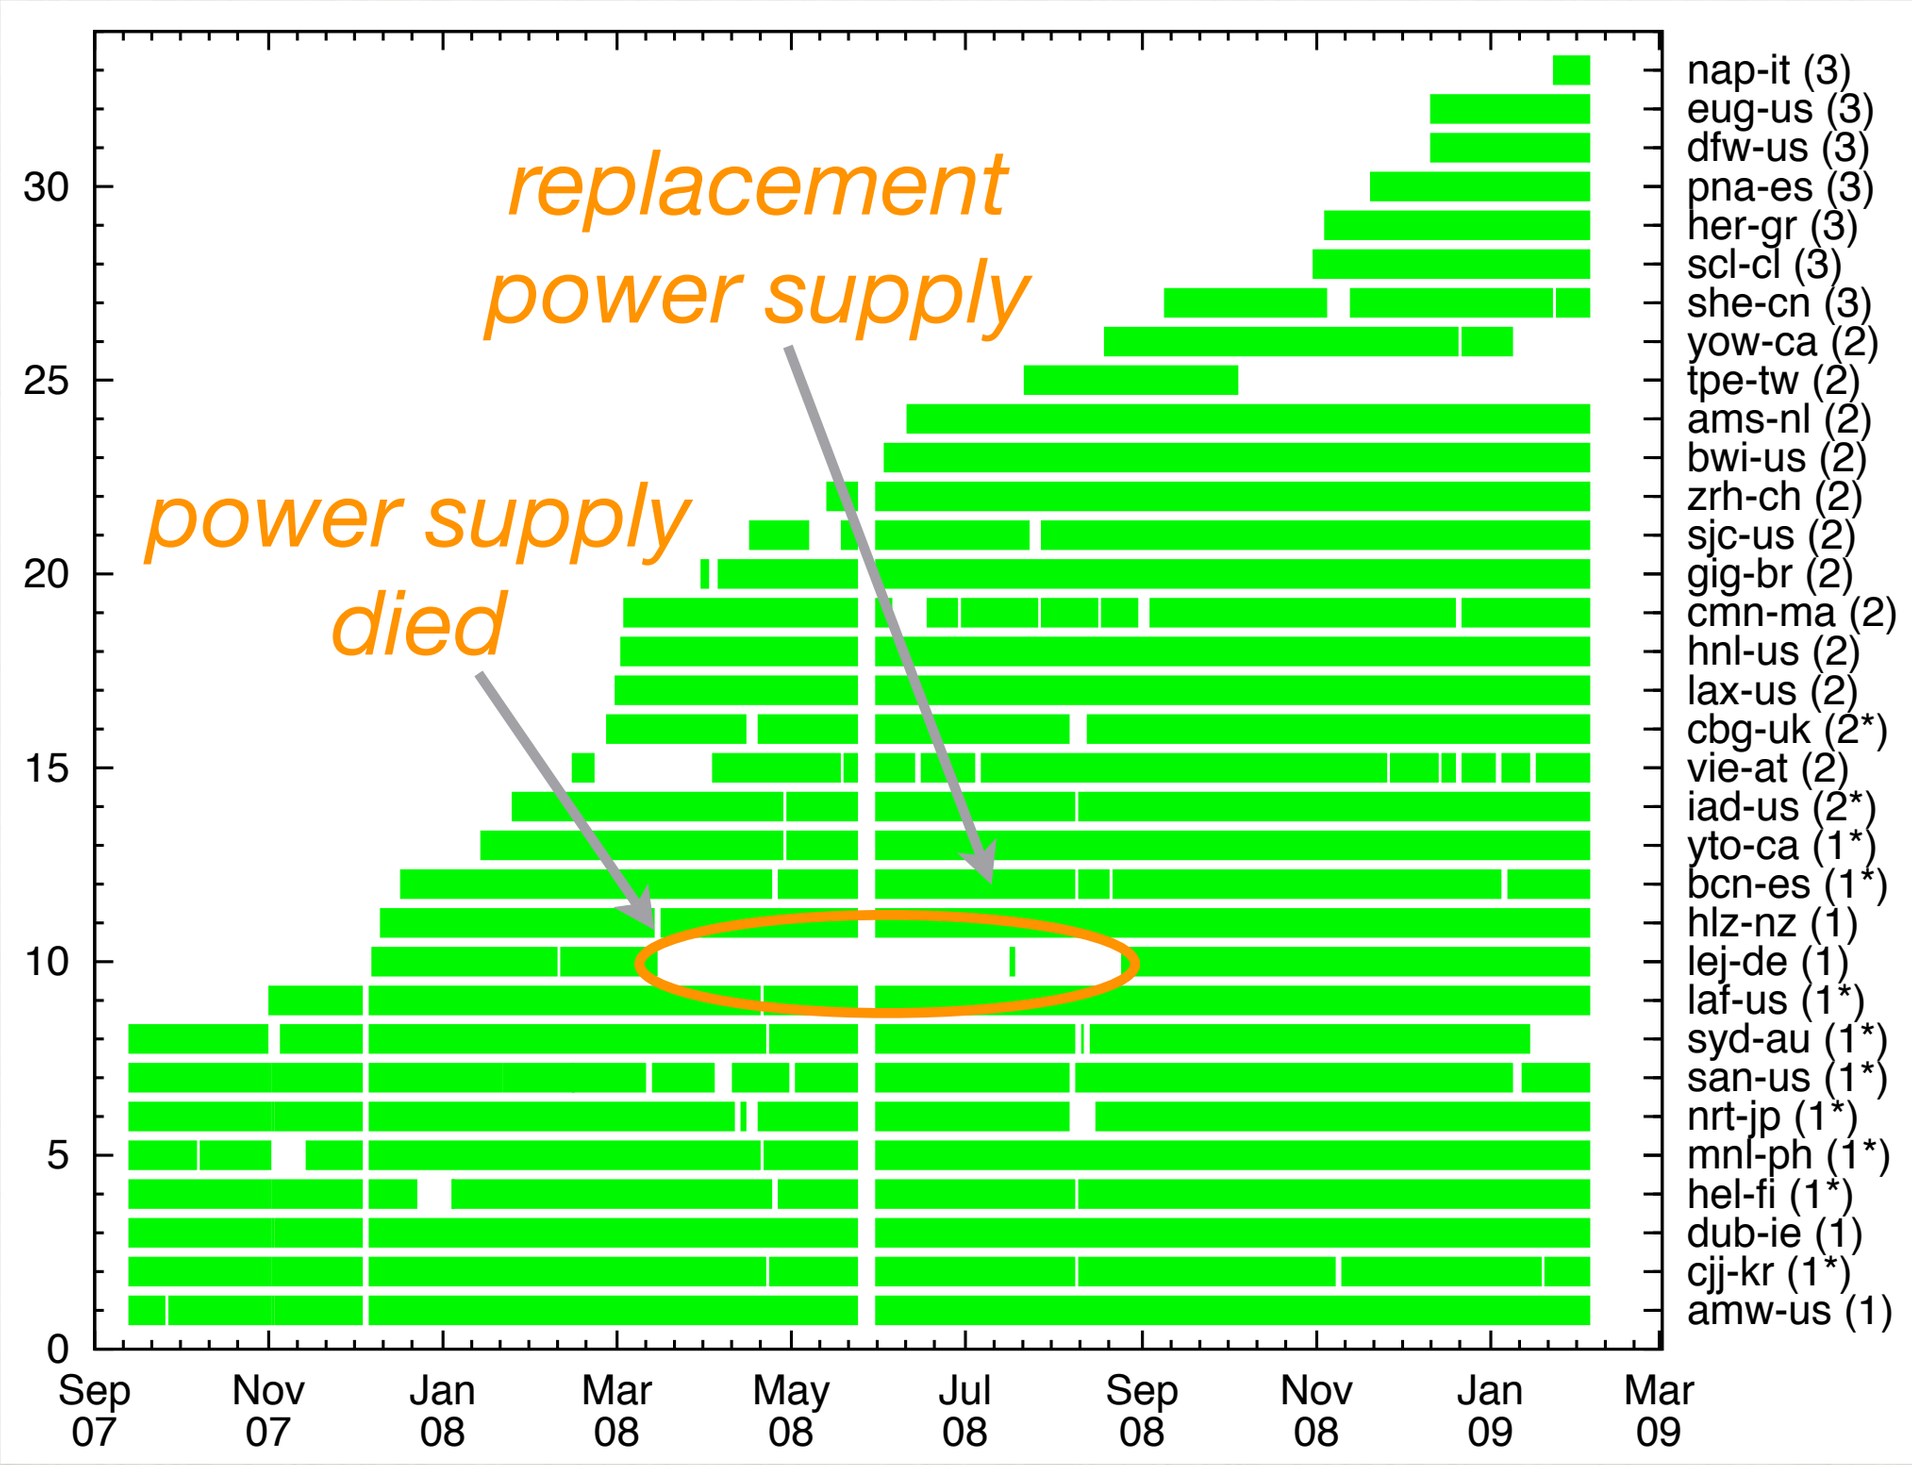

Sep 2007 to Jan 2009 (17 months): 2.5 billion traceroutes; 1.0TB data

IPv4 Routed /24 Topology

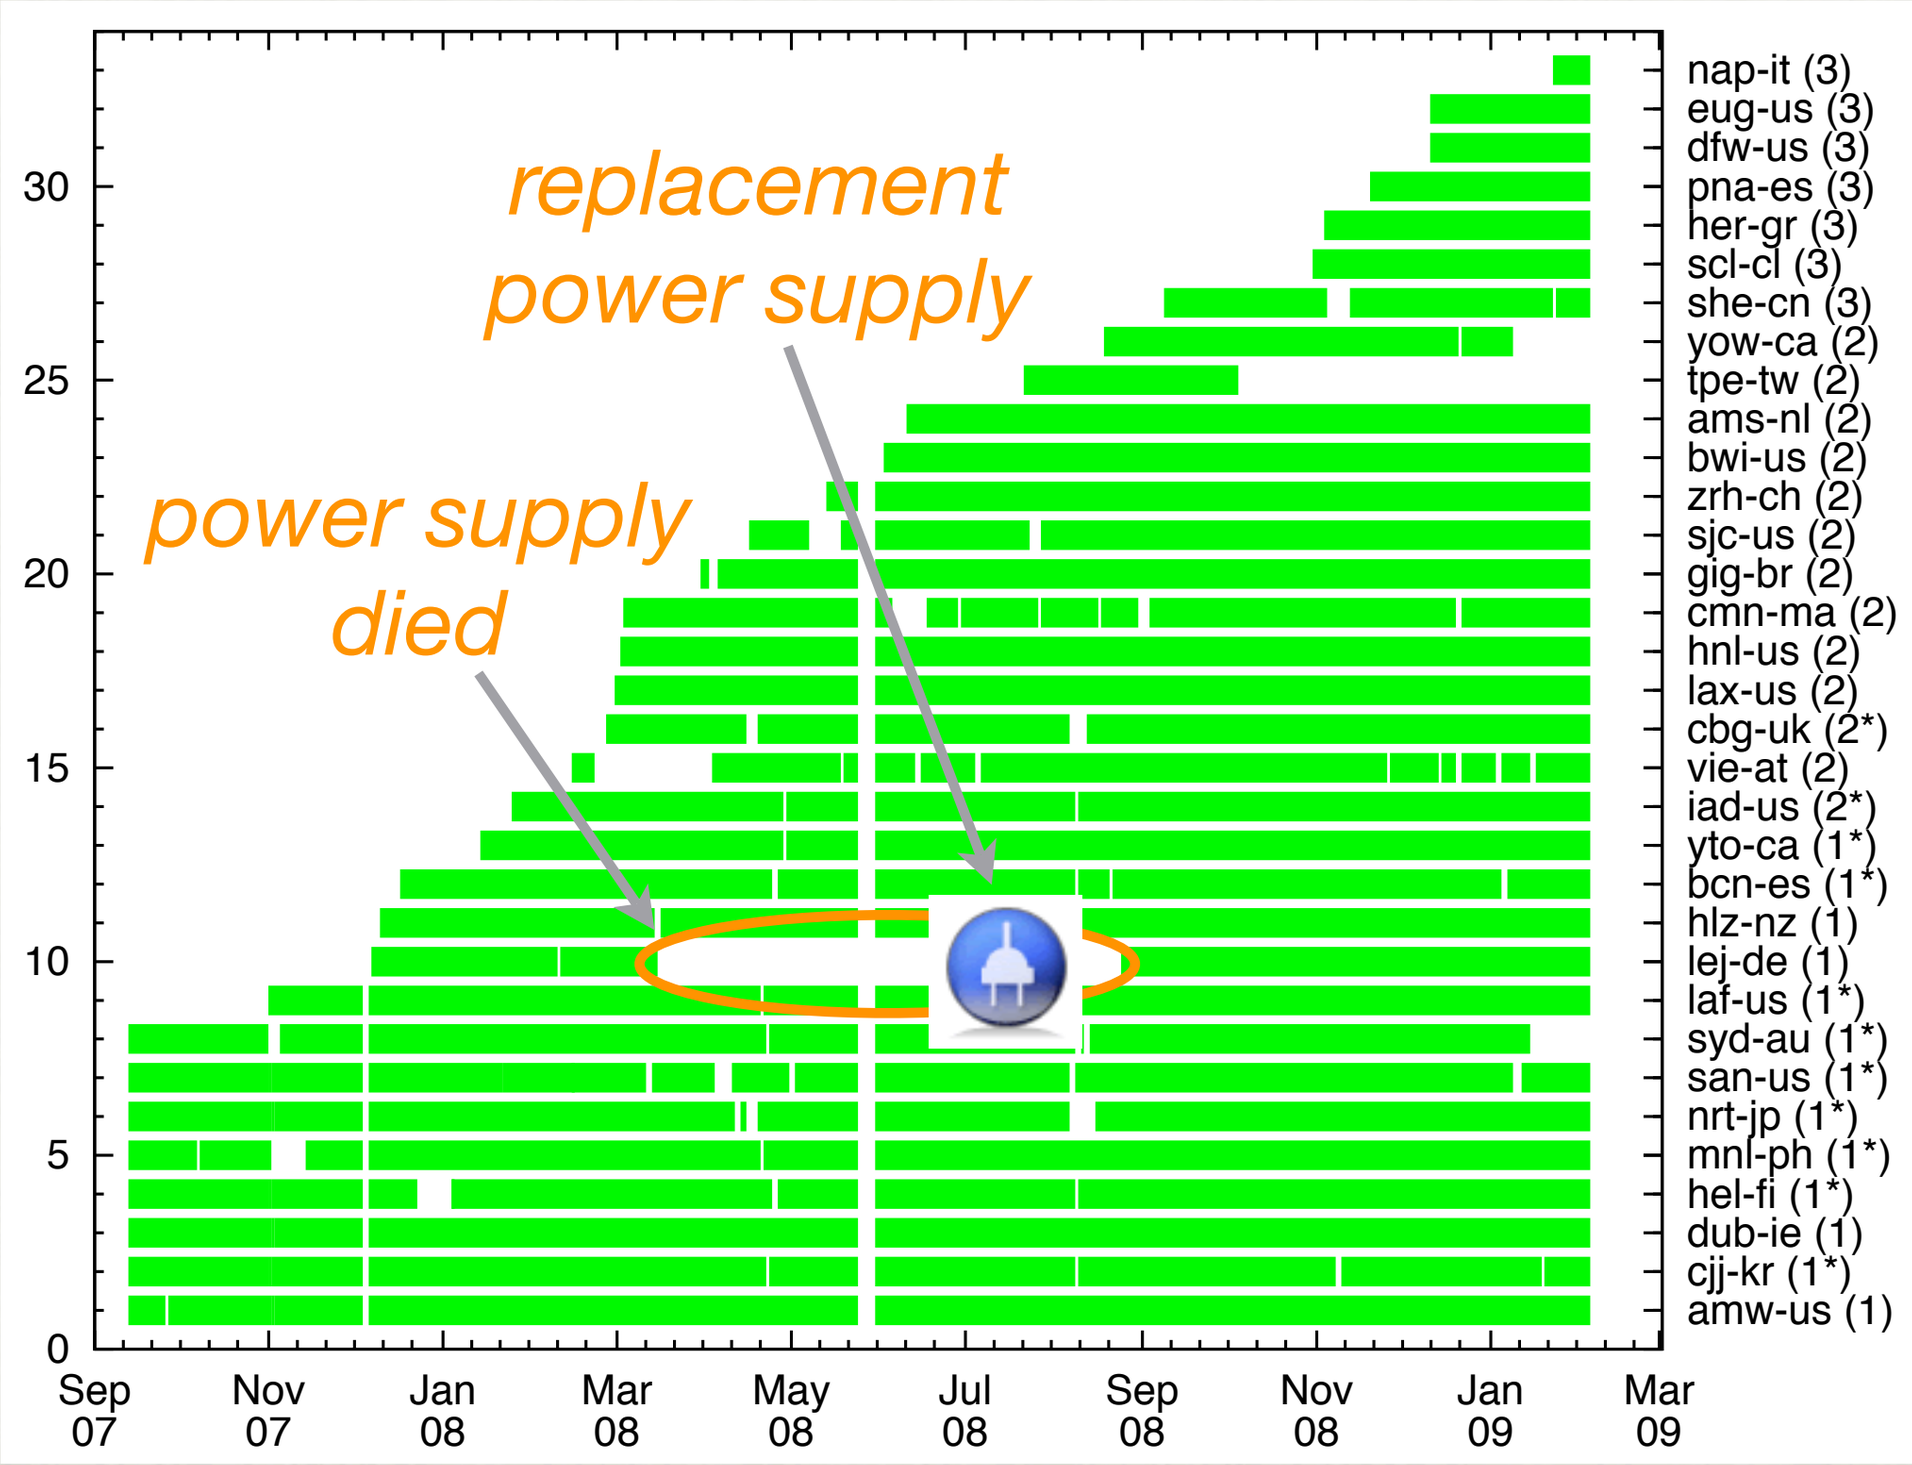

Sep 2007 to Jan 2009 (17 months): 2.5 billion traceroutes; 1.0TB data

IPv4 Routed /24 Topology

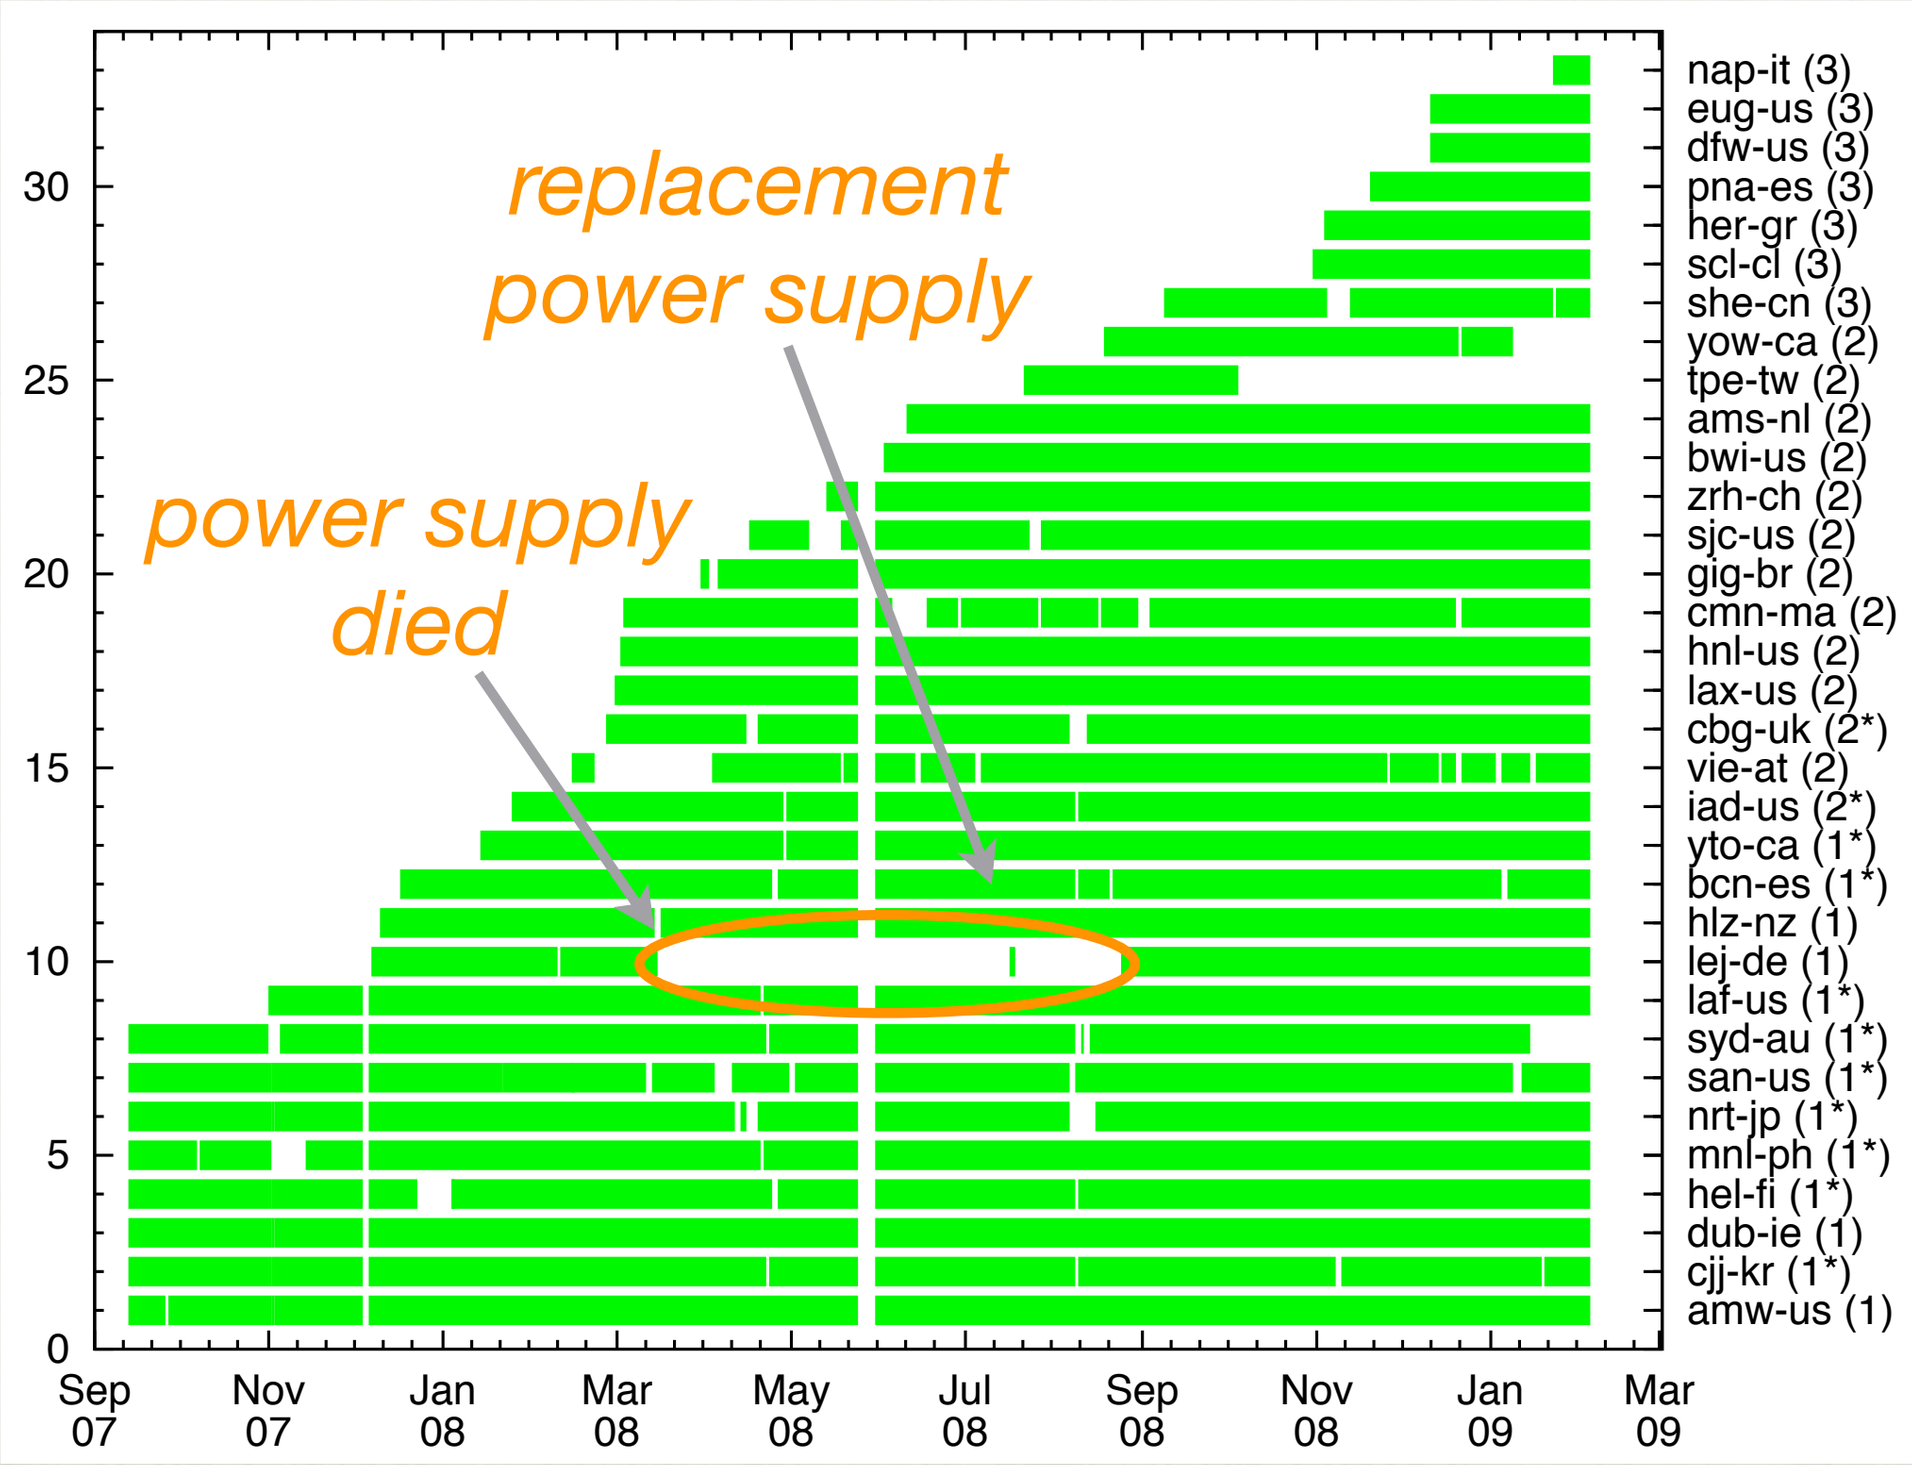

Sep 2007 to Jan 2009 (17 months): 2.5 billion traceroutes; 1.0TB data

IPv4 Routed /24 Topology

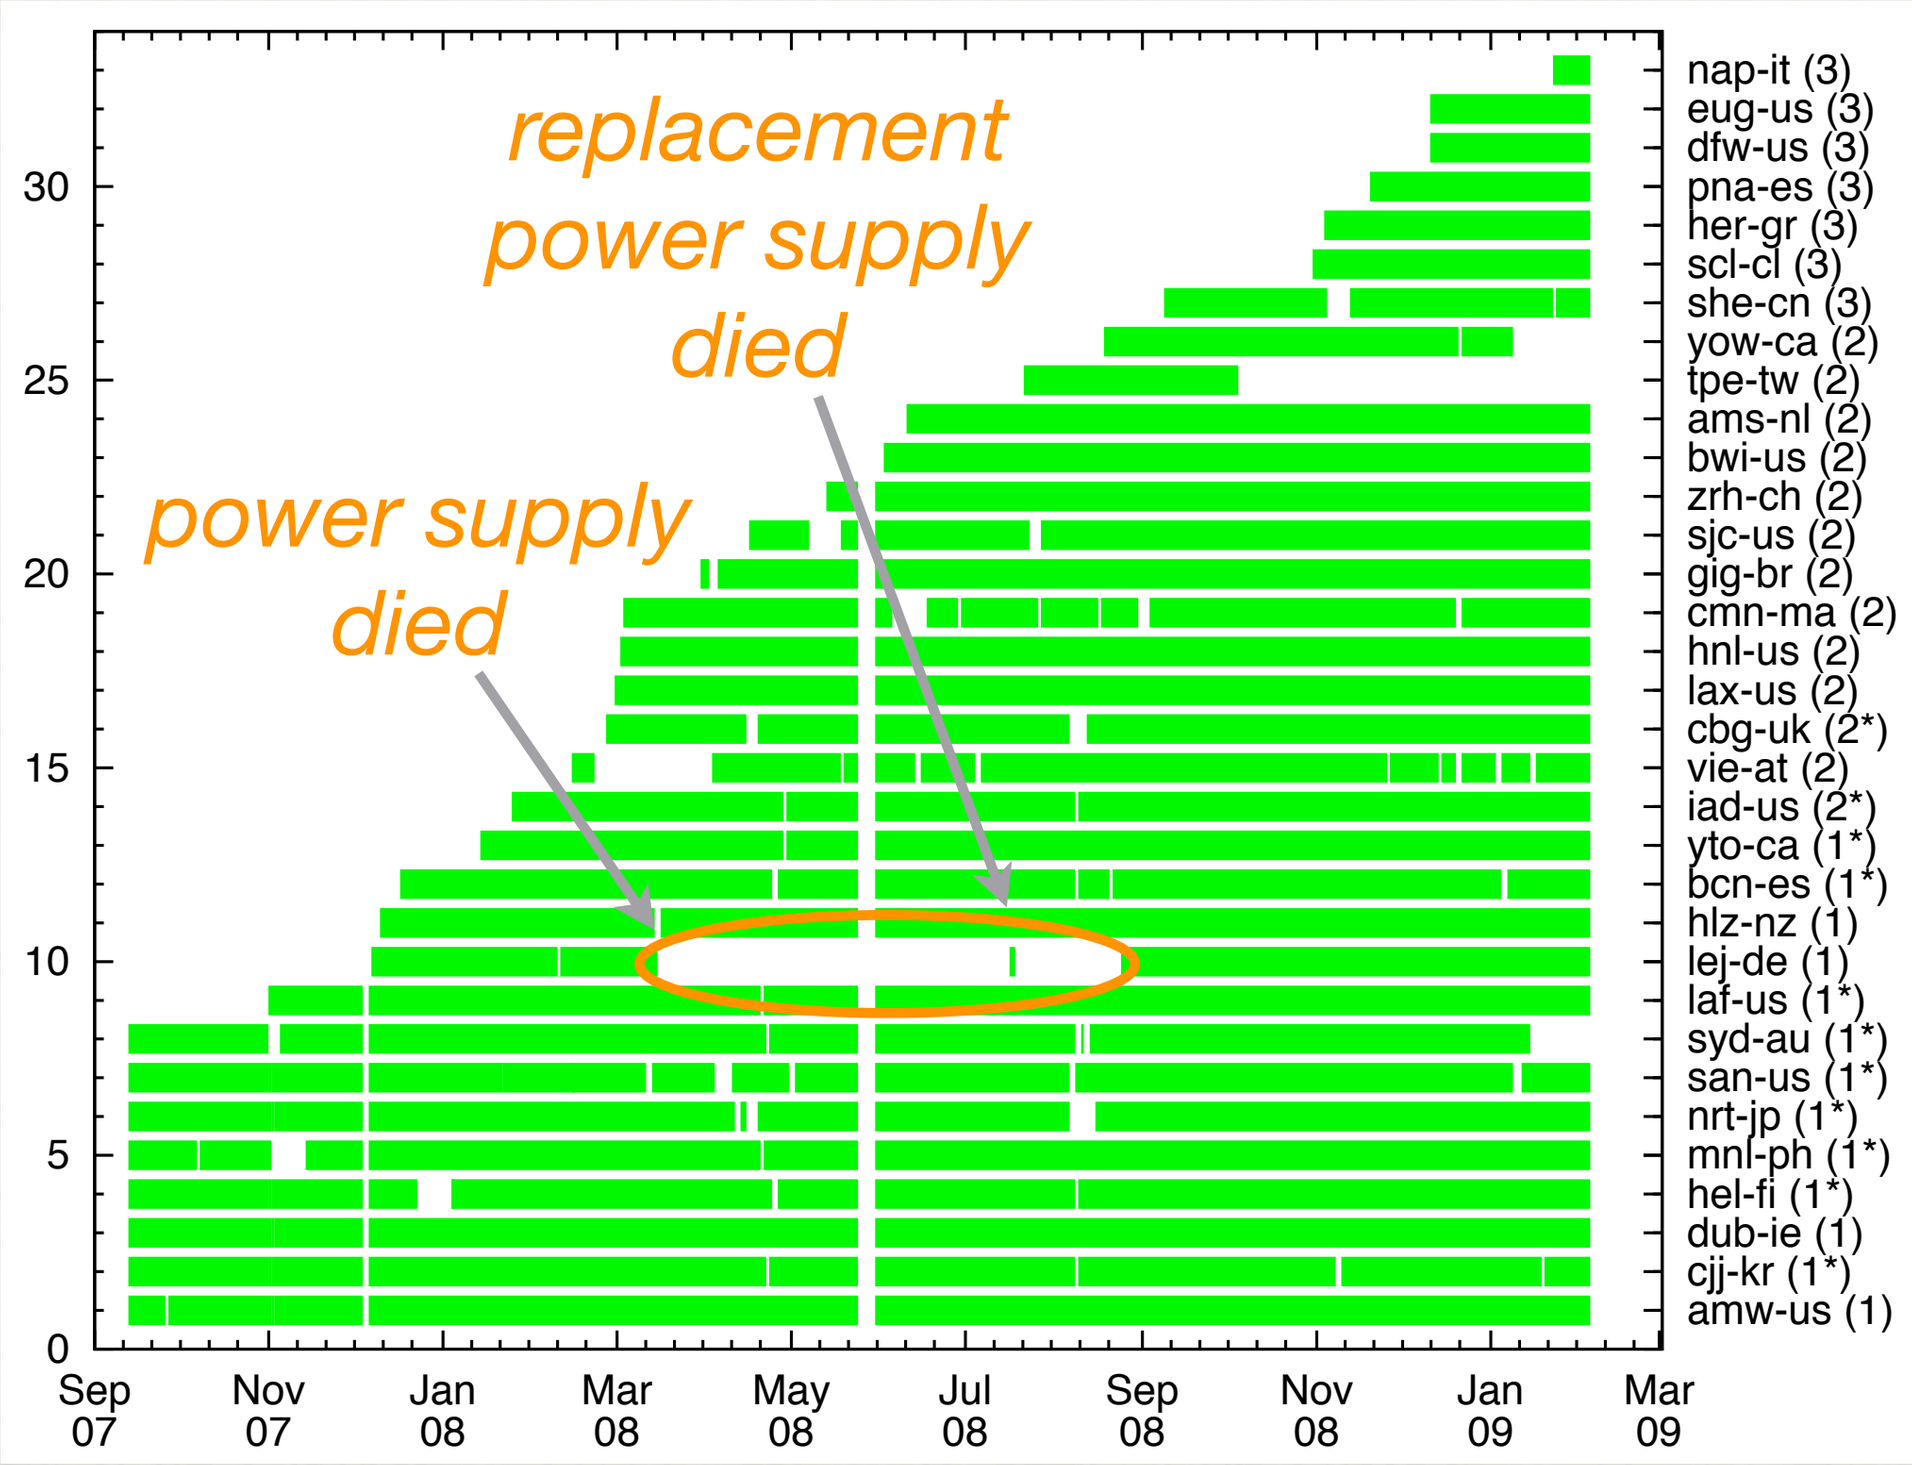

Sep 2007 to Jan 2009 (17 months): 2.5 billion traceroutes; 1.0TB data

IPv4 Routed /24 AS Links

- AS links from Routed /24 Topology traces
  - map IP addresses to ASes with RouteViews BGP table

IPv4 Routed /24 AS Links

* statistics for 1 month of AS links from three sources (Dec 2008):

	nodes	links	max degree	average degree	average neighbor degree	mean clustering
Ark	23,425	56,760	2,509	4.85	467.3	0.354
DIMES	22,995	74,140	3,590	6.45	705.4	0.446
RouteViews (rv2)	30,760	65,775	2,328	4.28	487.2	0.241

* “avg neighbor deg” = avg neighbor degree of the avg *k*-degree node averaged over all *k*
* “mean clustering” = (avg number of links between neighbors of *k*-deg nodes) / (max possible such links for *k*) averaged over all *k*

3 AS Links Sources: 1 Month

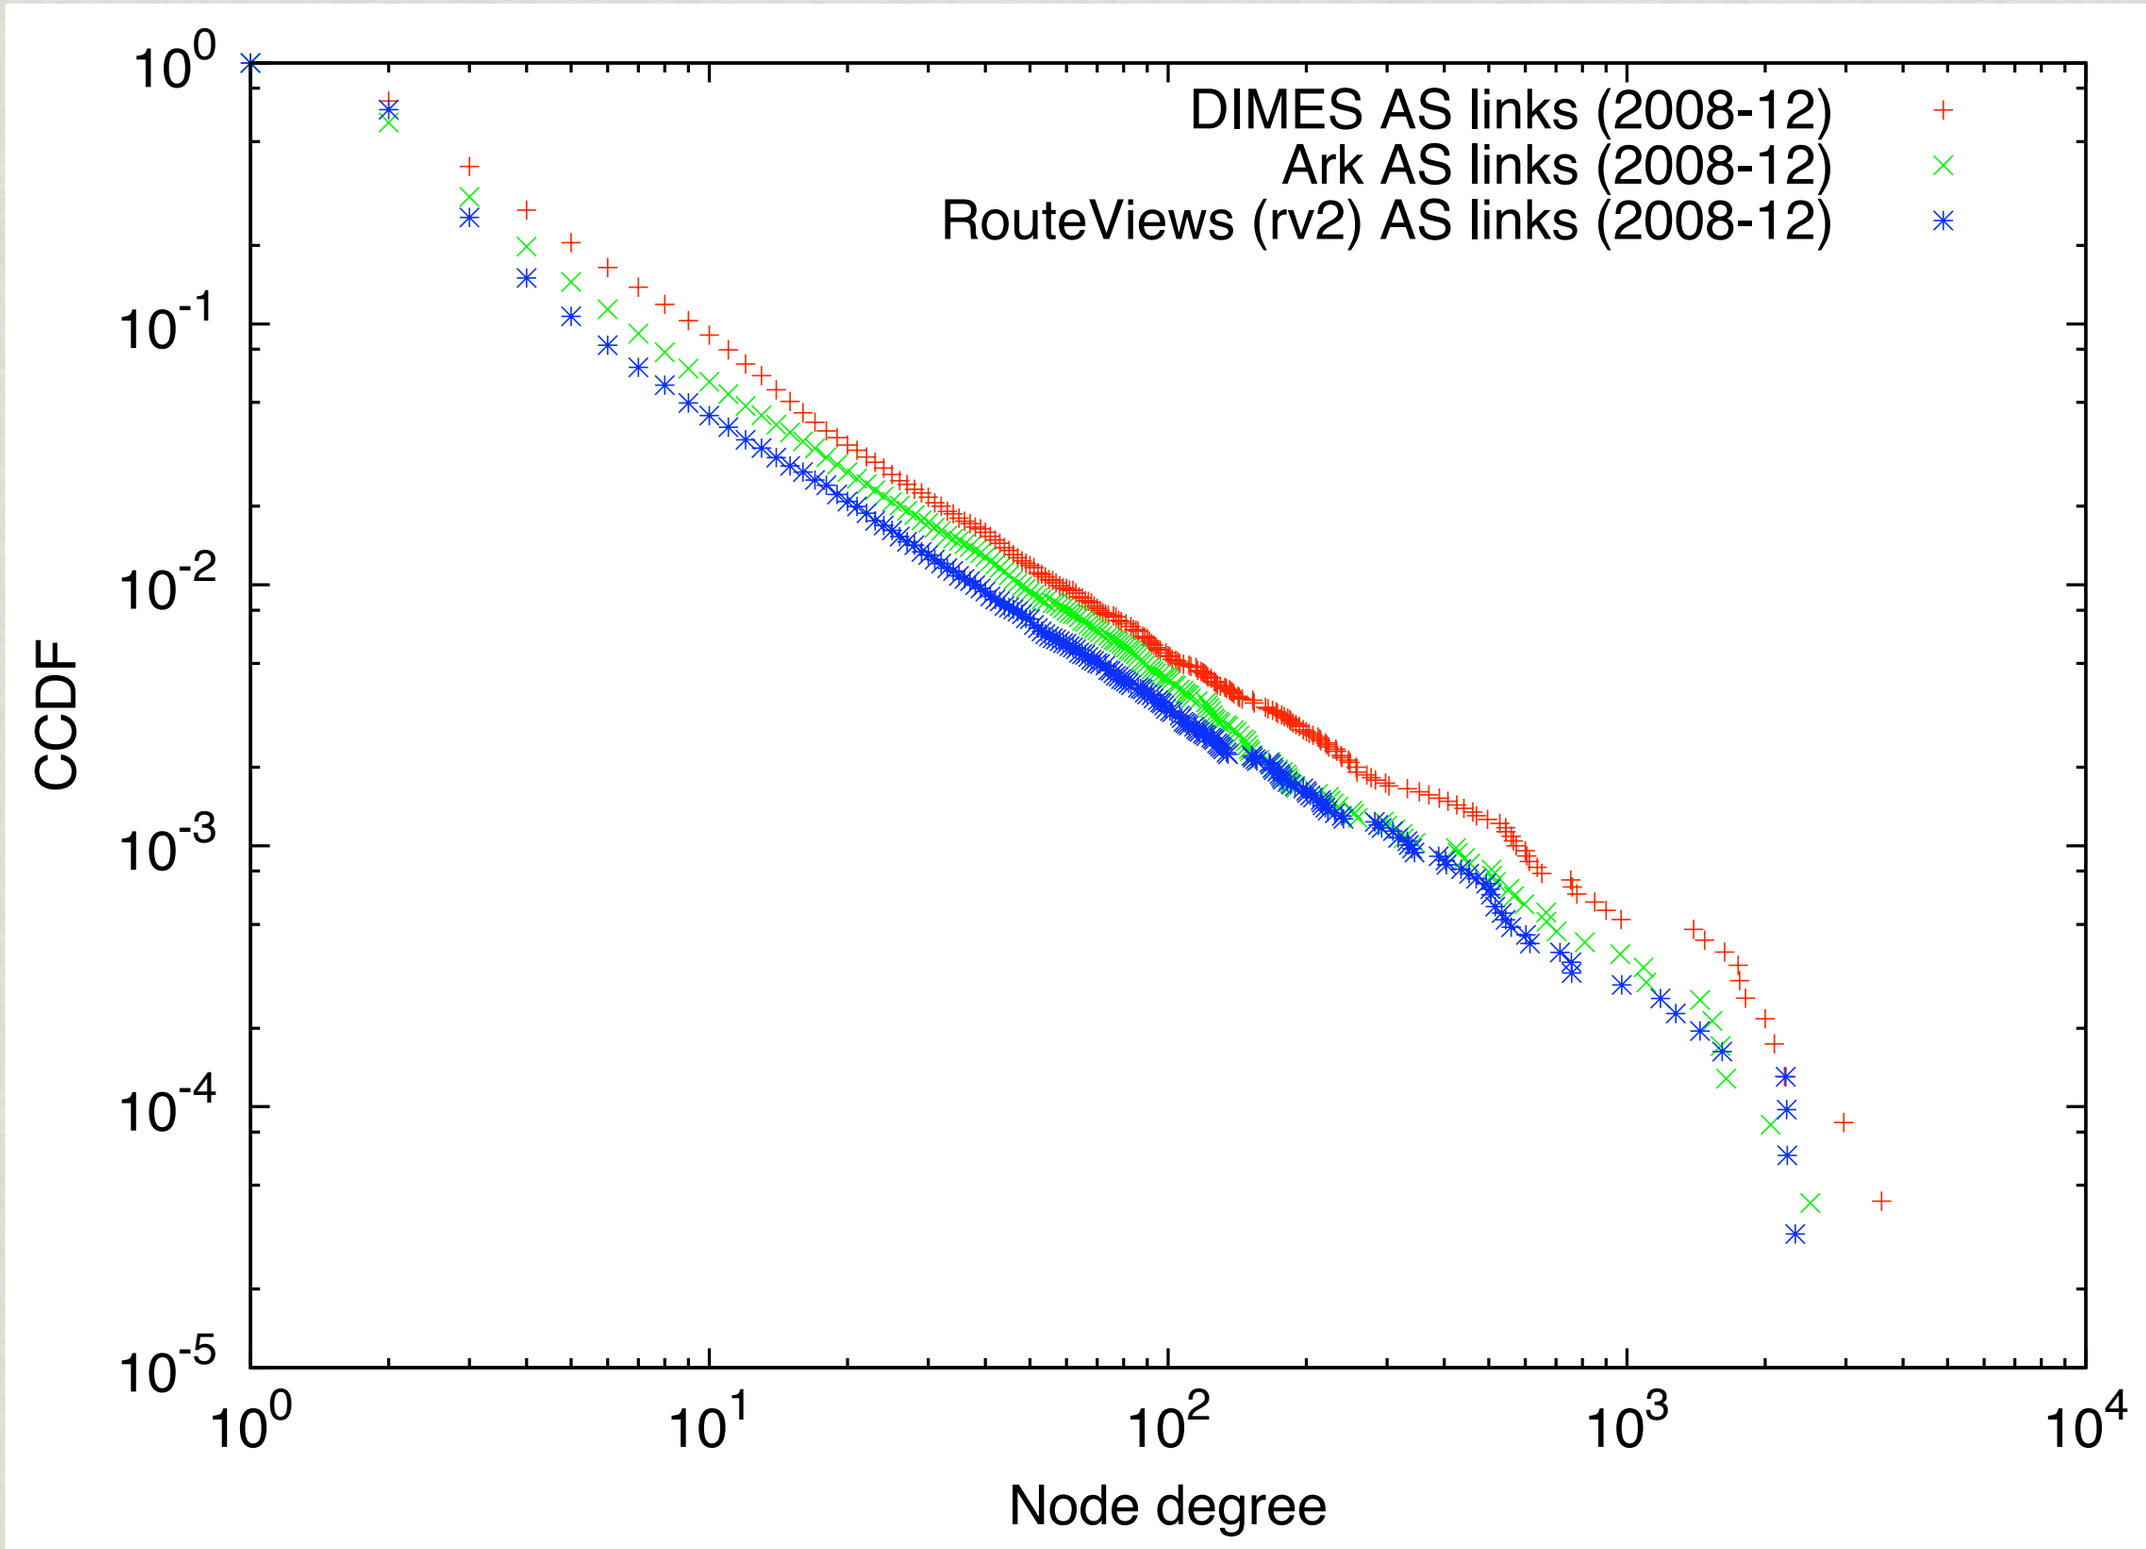

3 AS Links Sources: 1 Month

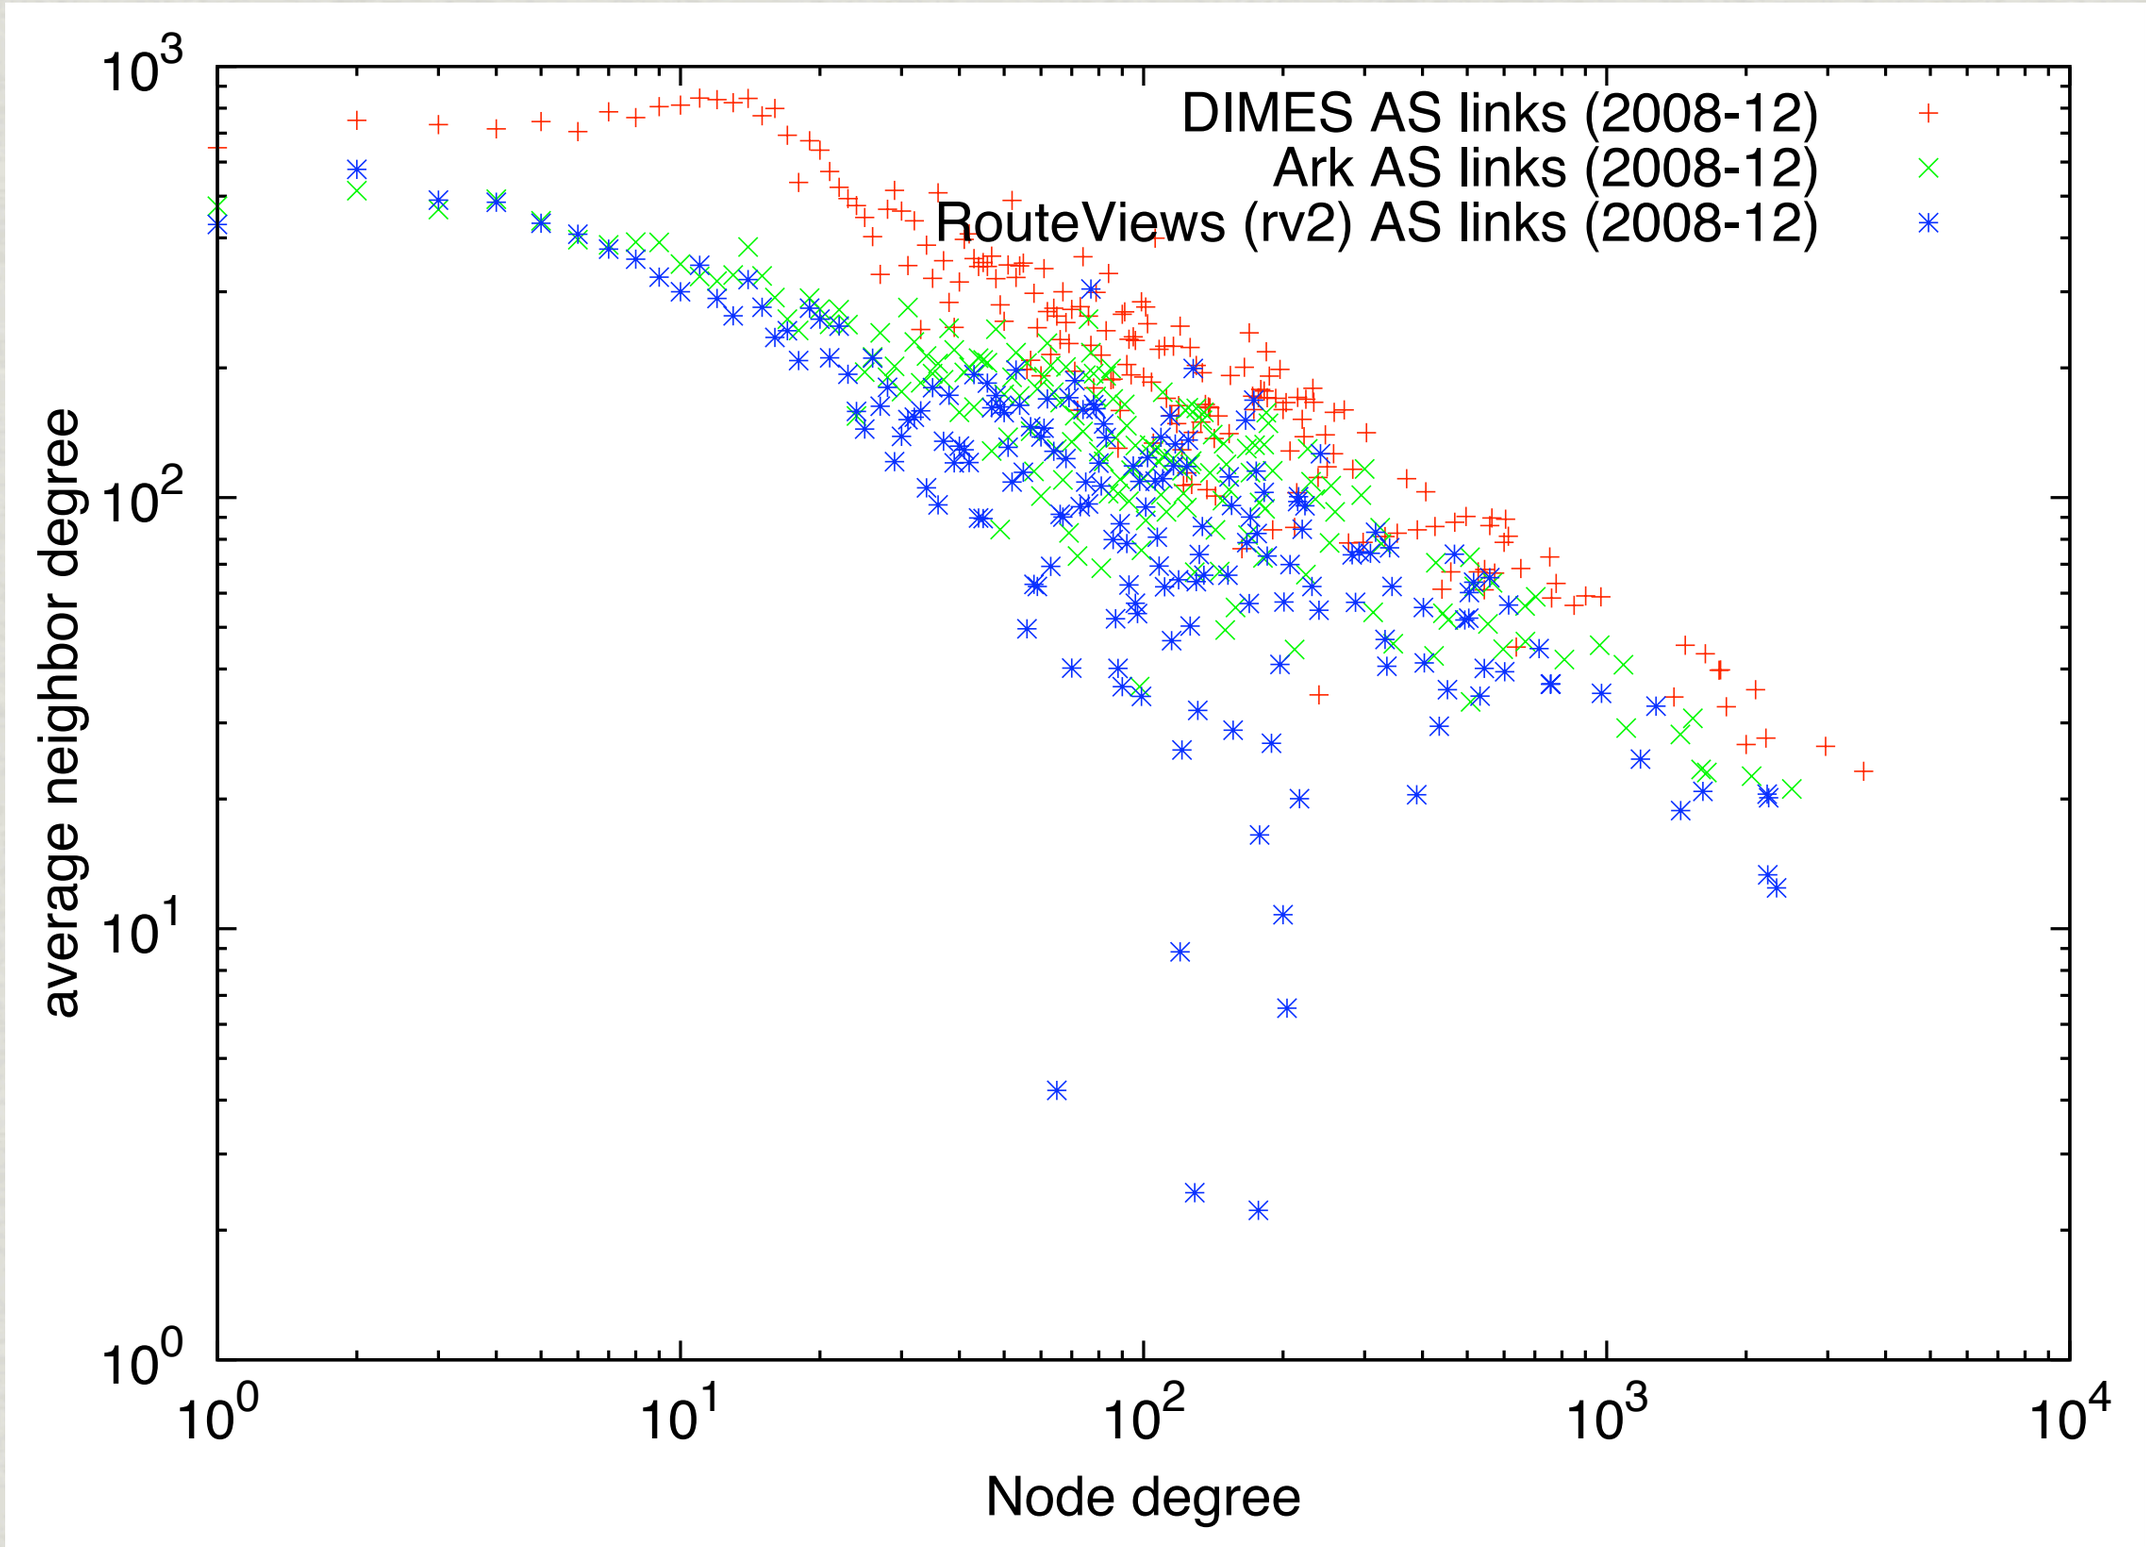

3 AS Links Sources: 1 Month

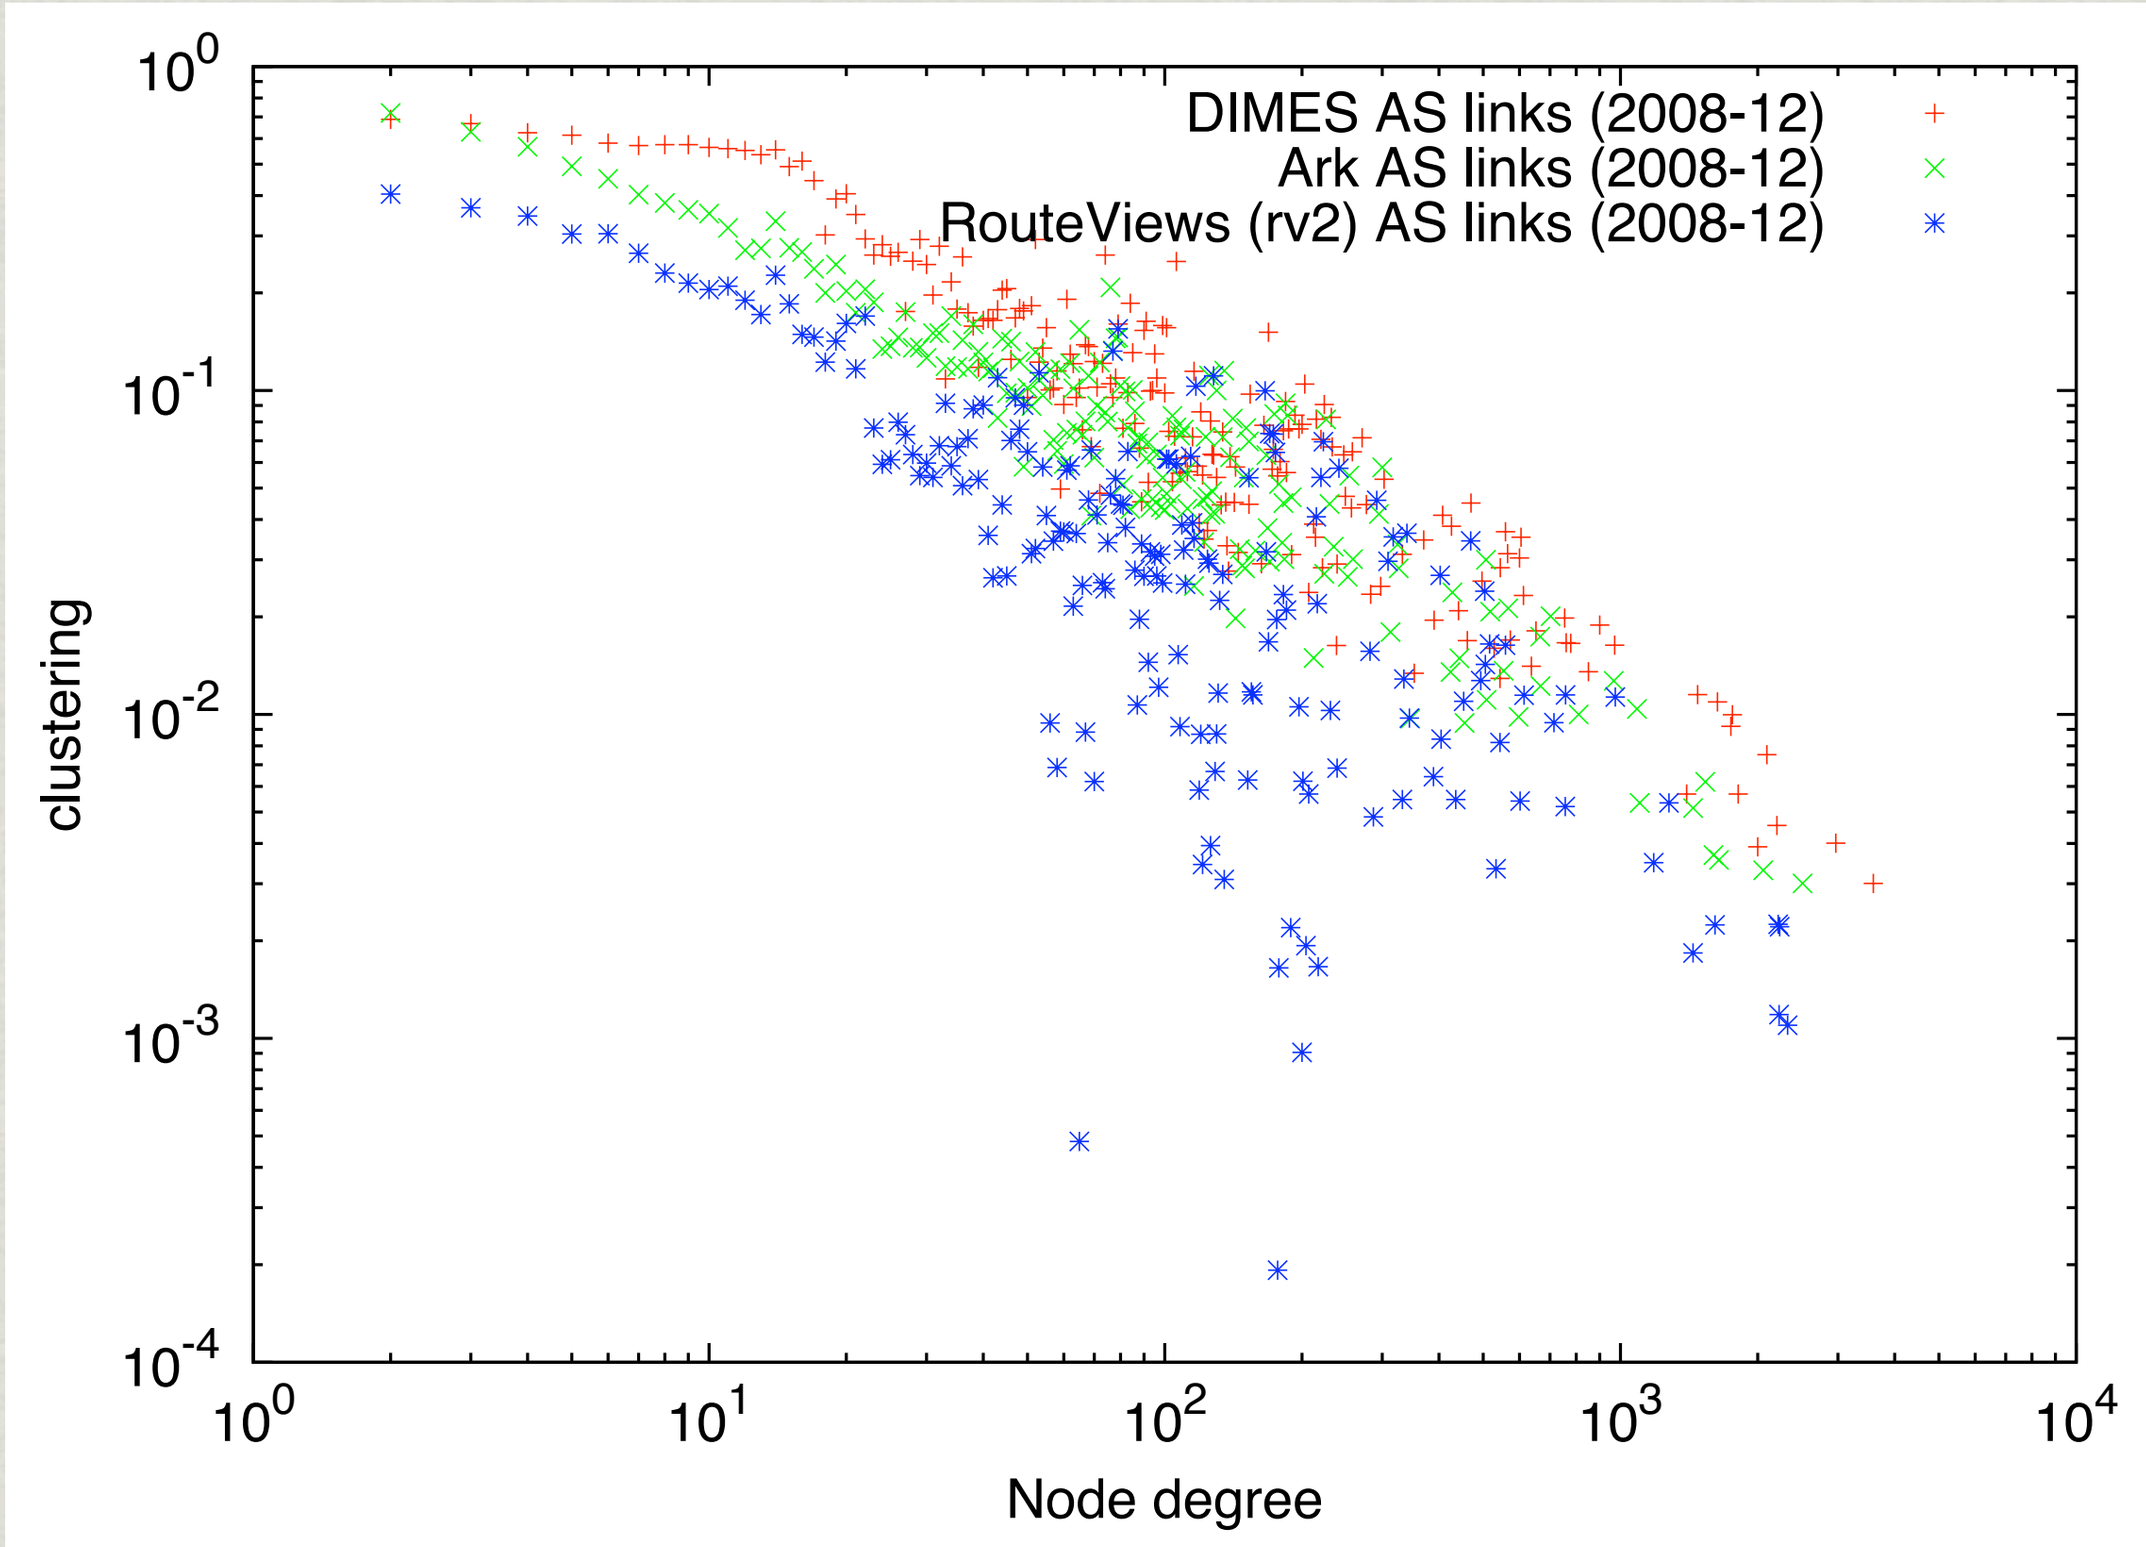

AS Links Growth

- AS links seem to accumulate linearly without bound
  - in skitter, Ark, DIMES; possibly in BGP
  - even with fixed traceroute sources and destination list (which happened with skitter for 4 years)
- AS graph densification: average degree increases
- for example:
  - 1 year of Ark (2008): 104k AS links, 28k ASes
  - 2 years of DIMES: 356k AS links, 29k ASes
  - 7.5 years of skitter: 209k AS links, 27k ASes

AS Links Growth

- hard to determine the “natural” time period to aggregate AS links
  - 1 month? 6 months? years?
  - when do we get a representative AS graph?

Ark AS Links Growth

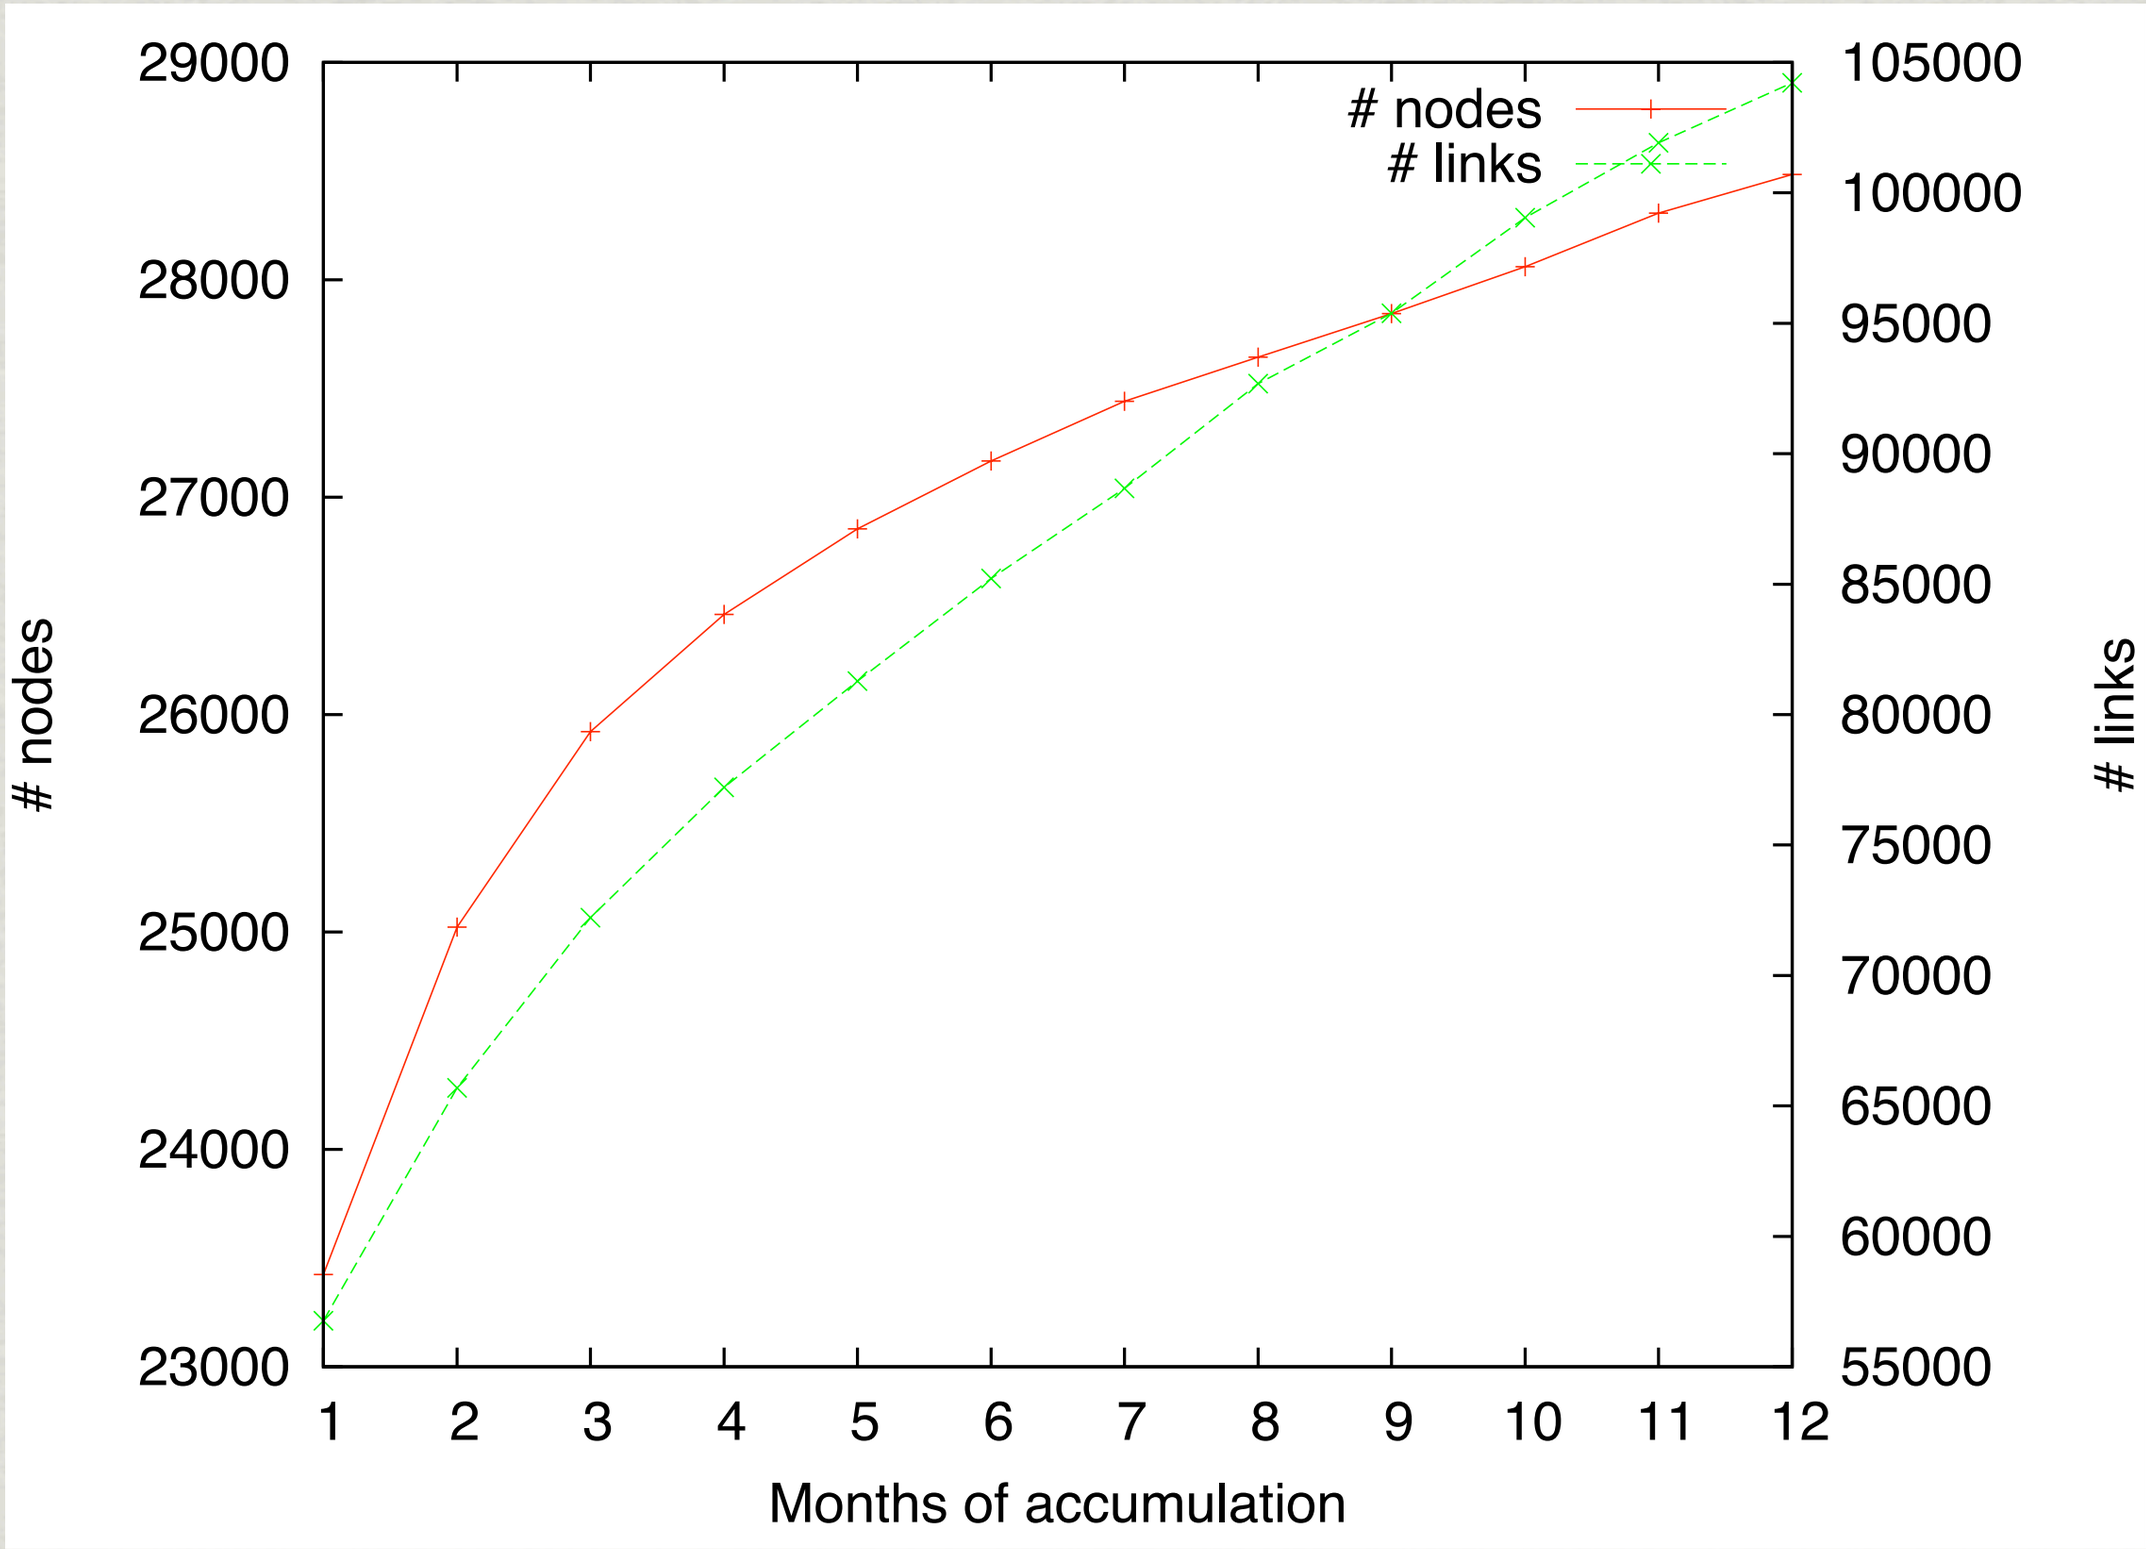

Ark AS Links Growth

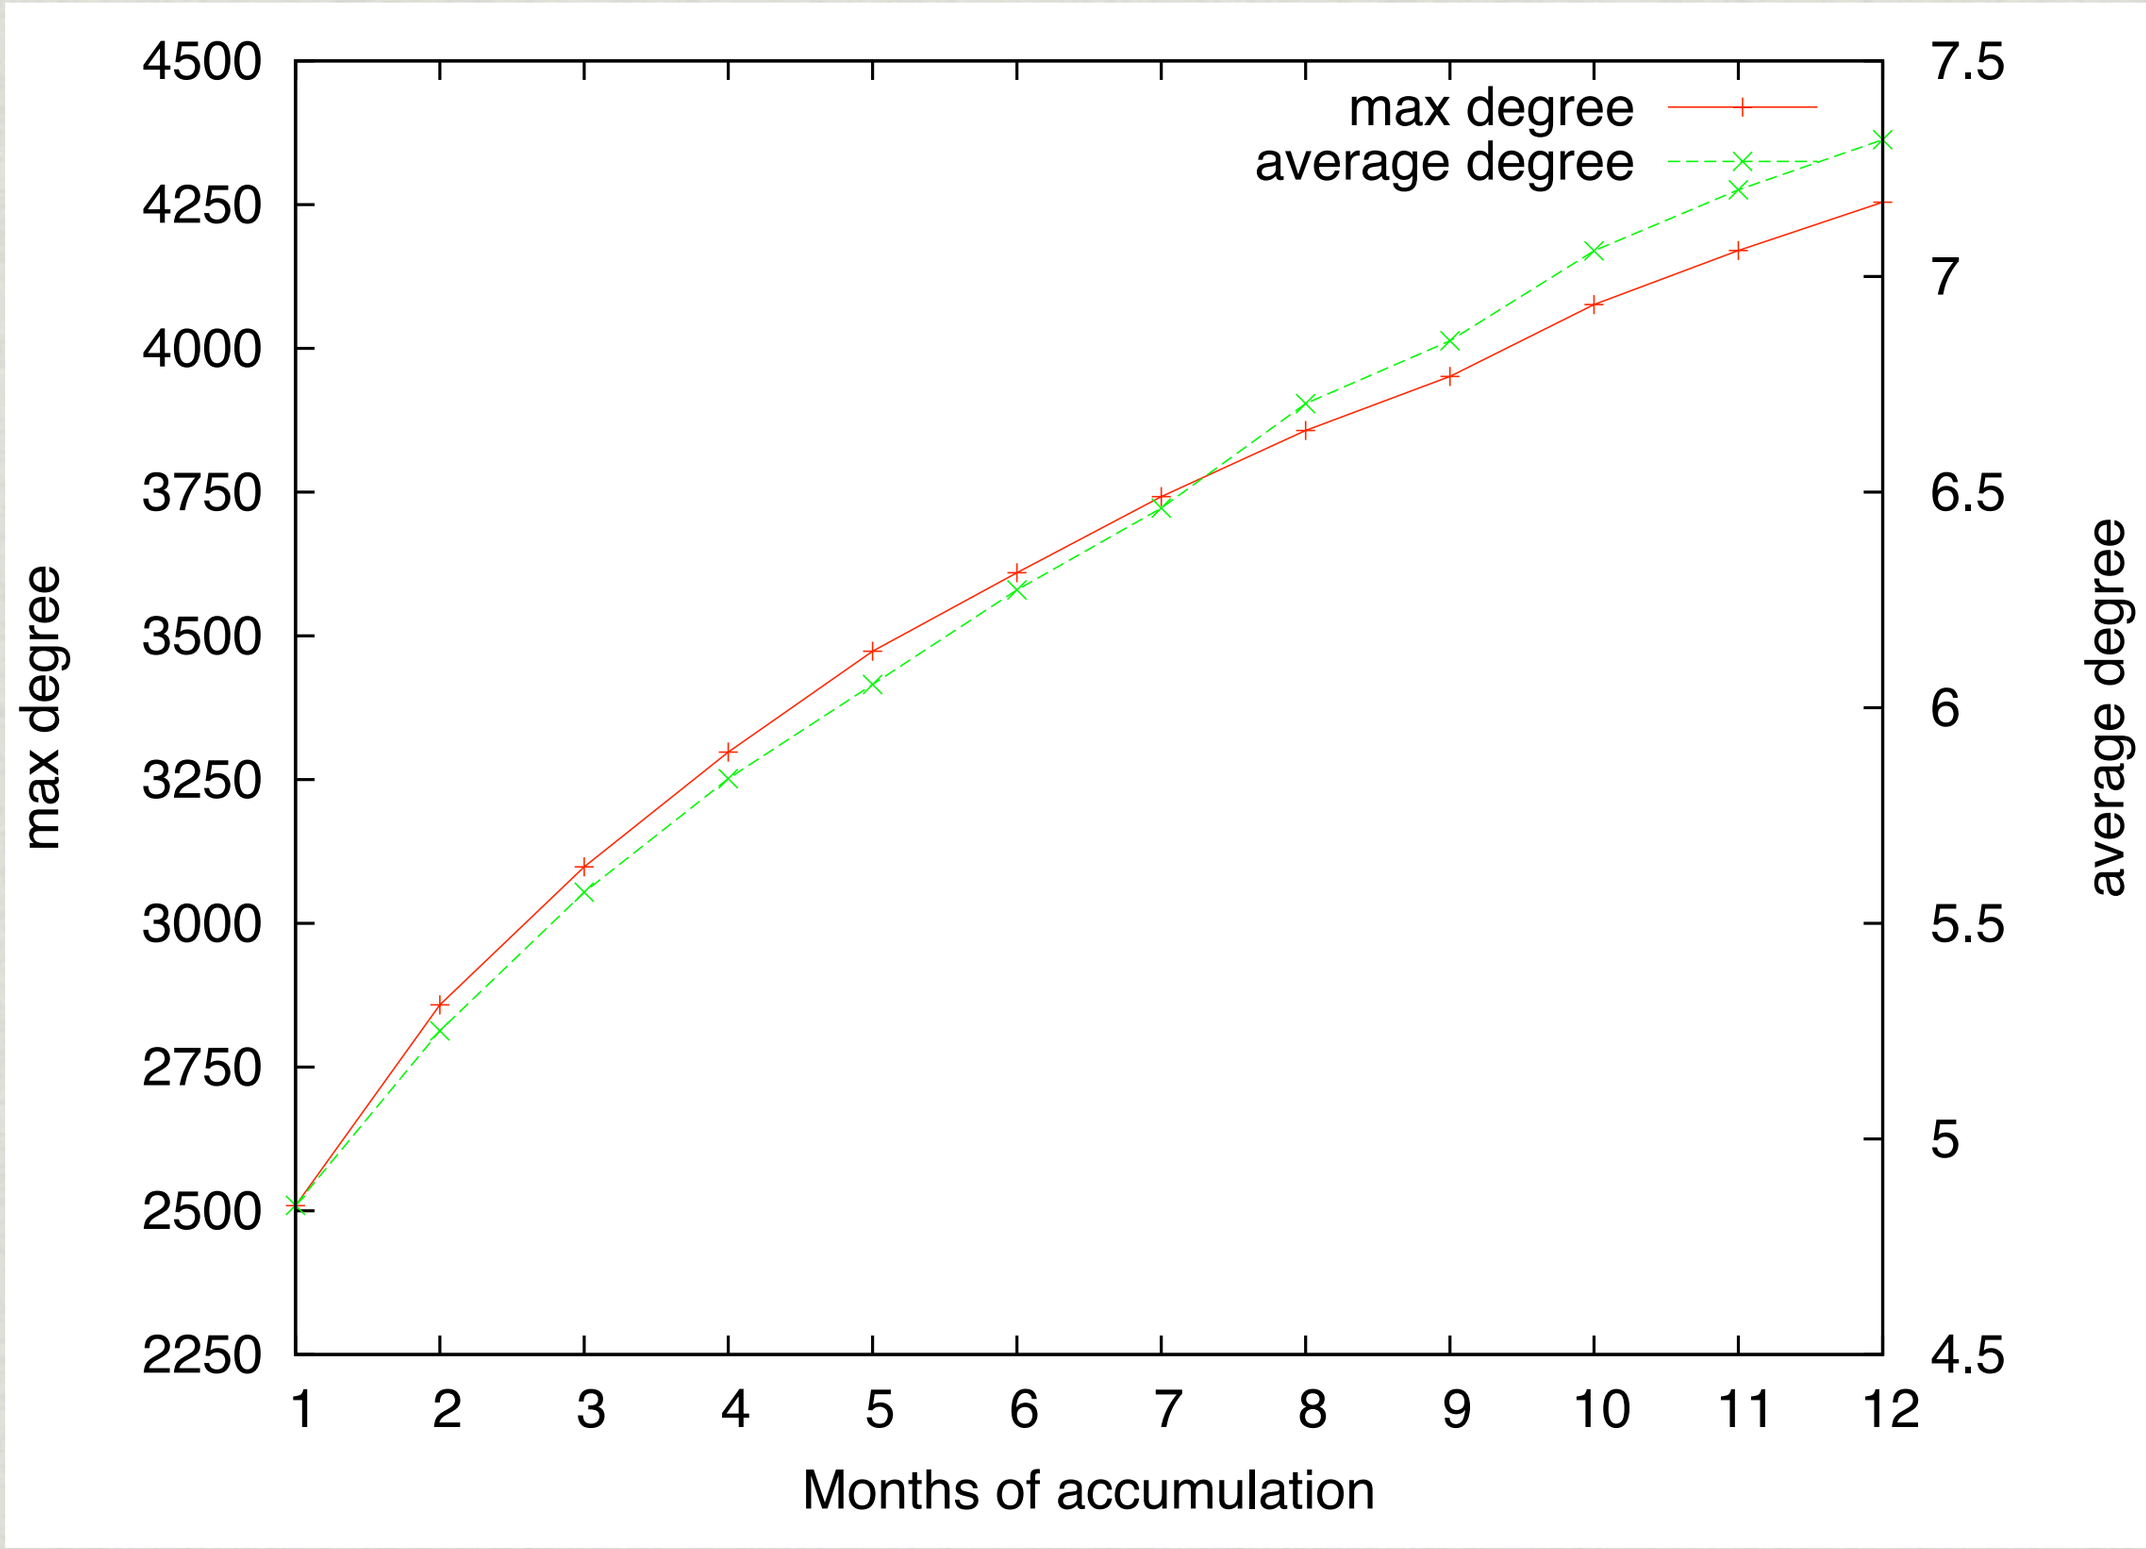

Ark AS Links Growth

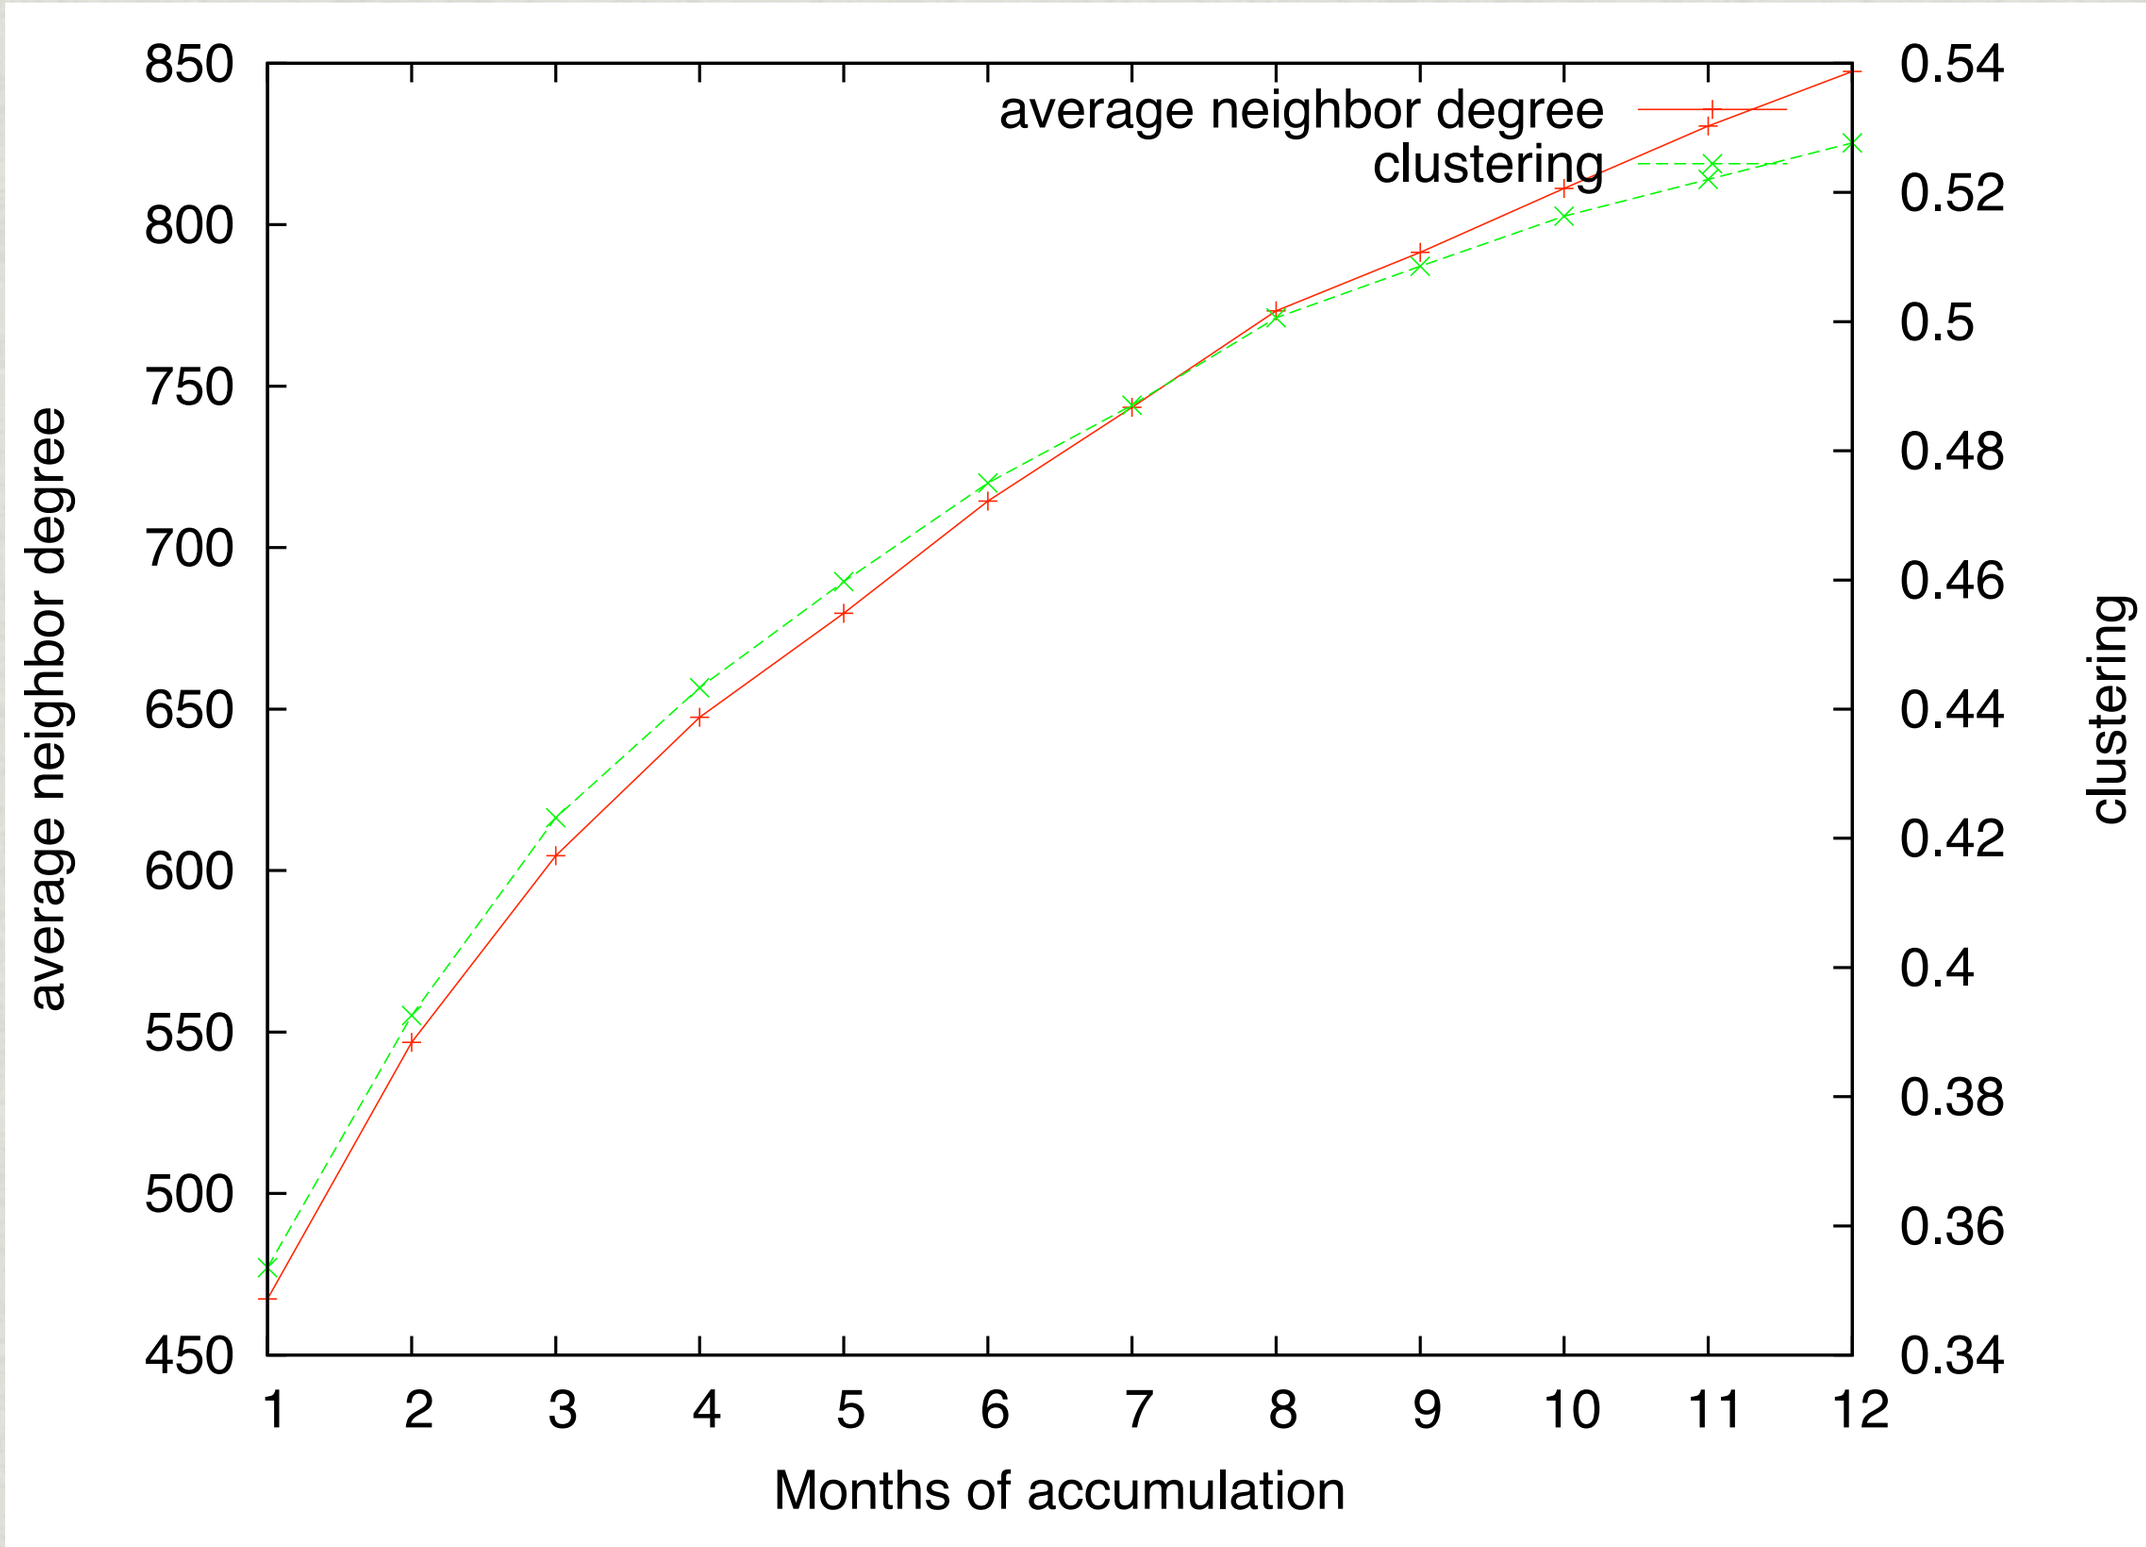

Ark AS Links: 1, 6, 12 Months

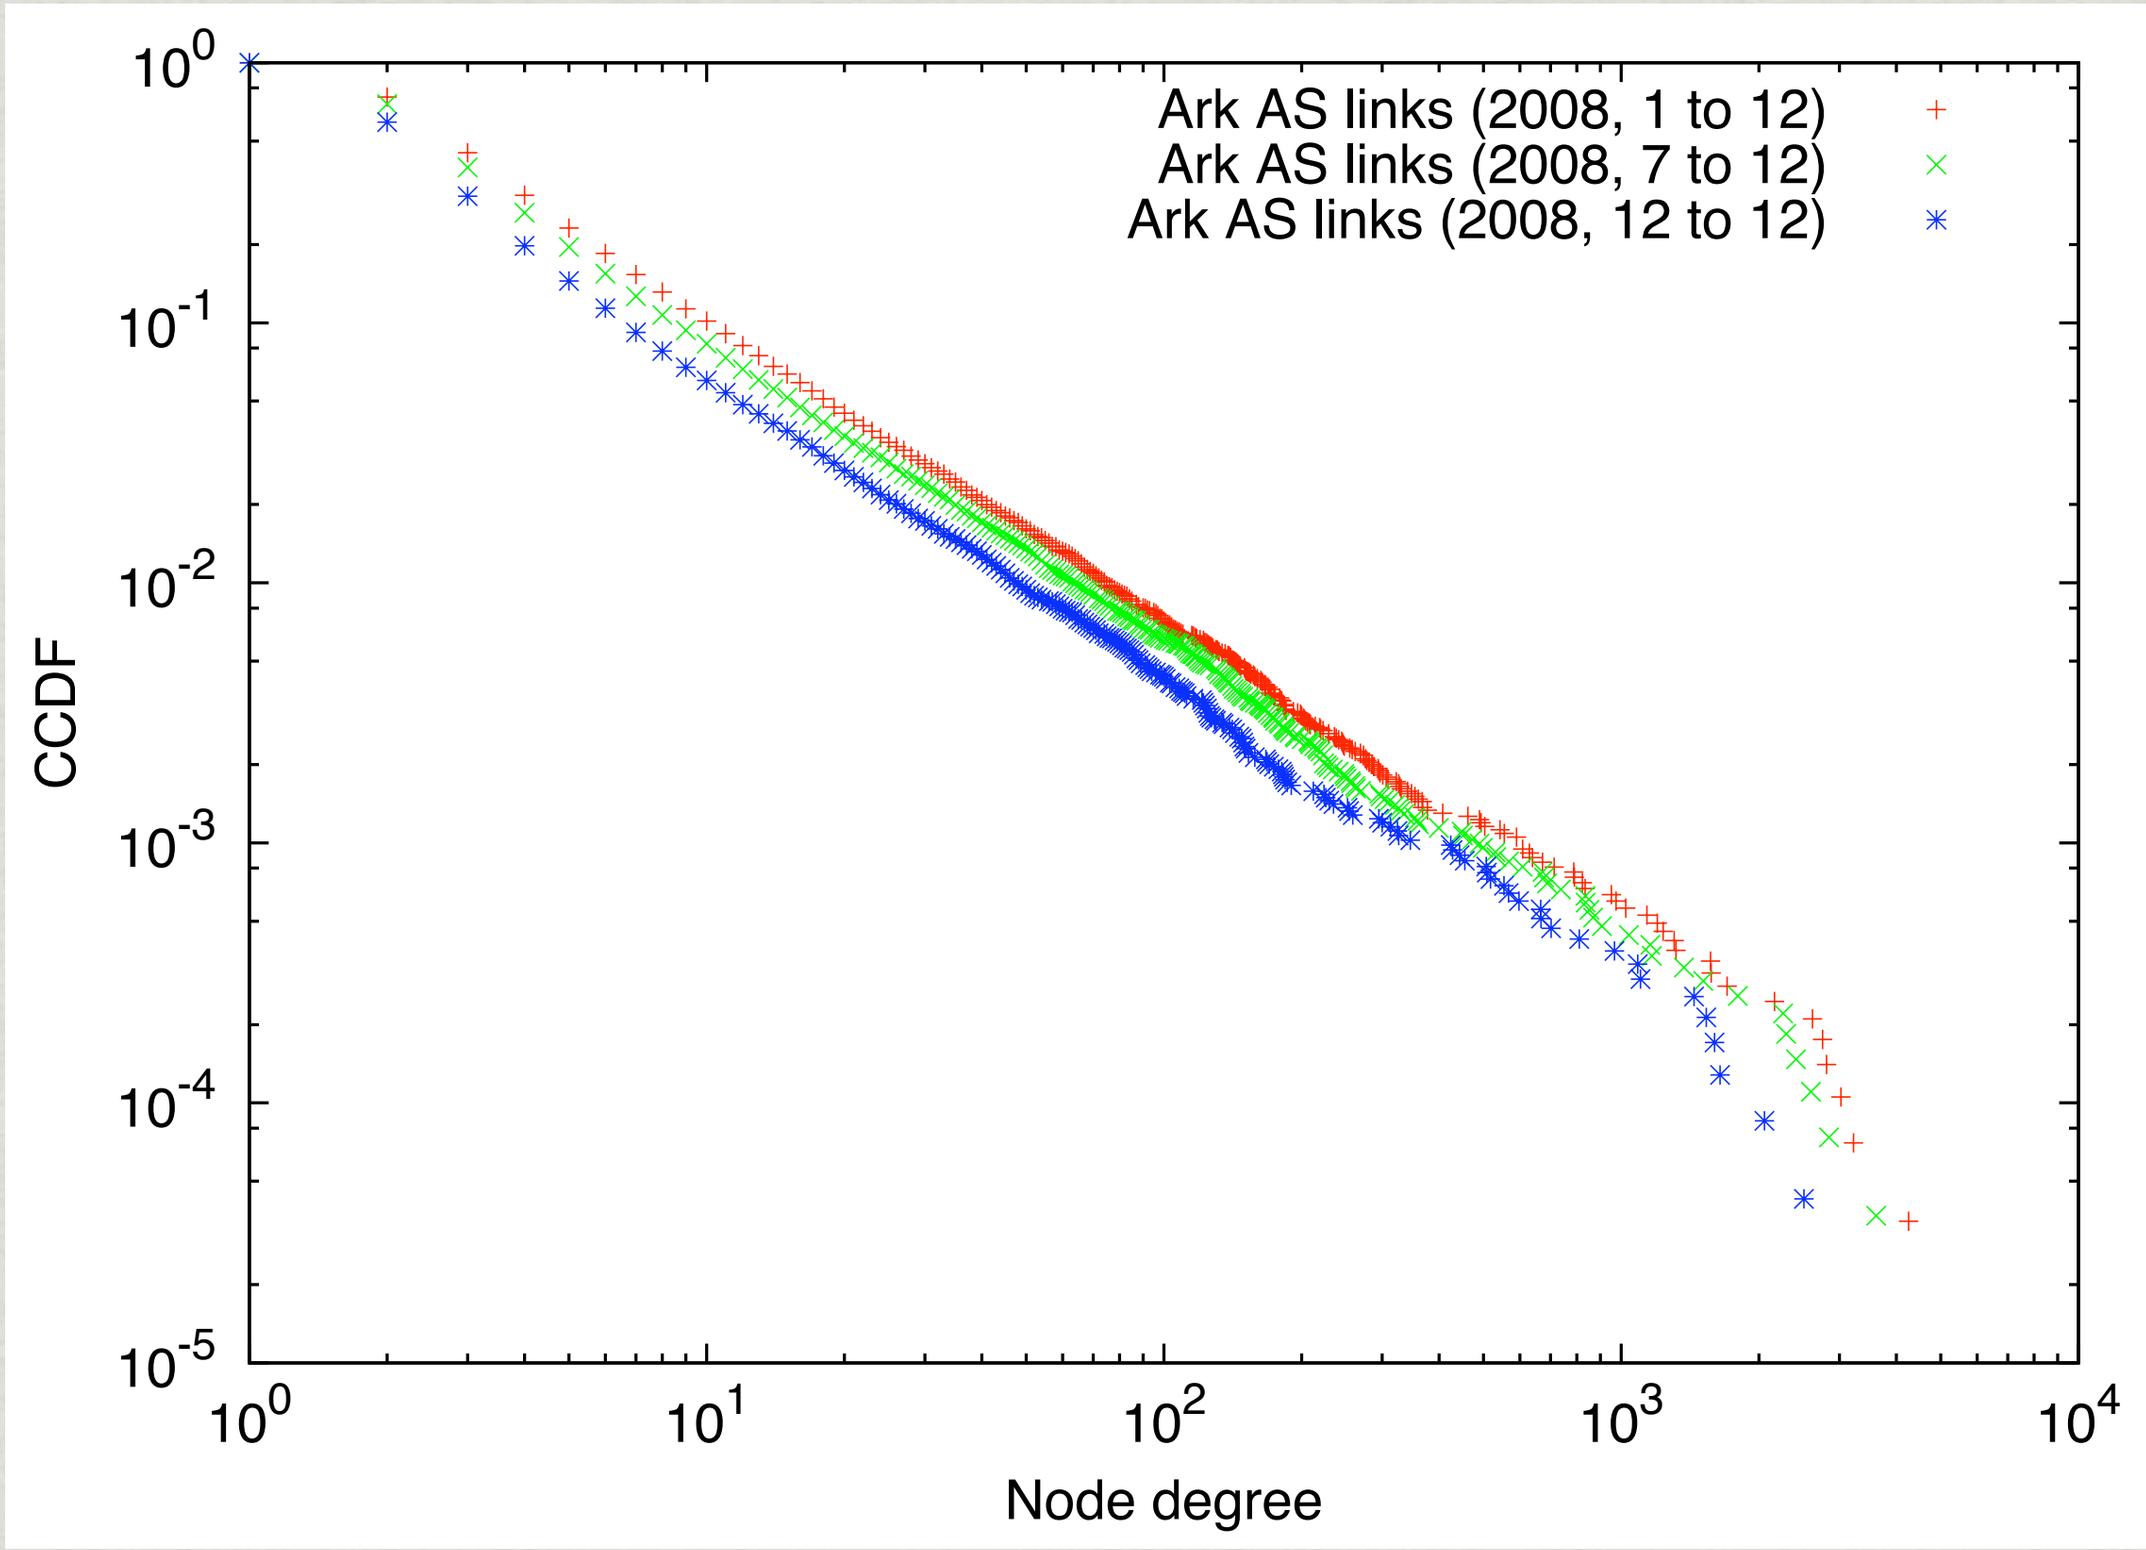

Ark AS Links: 1, 6, 12 Months

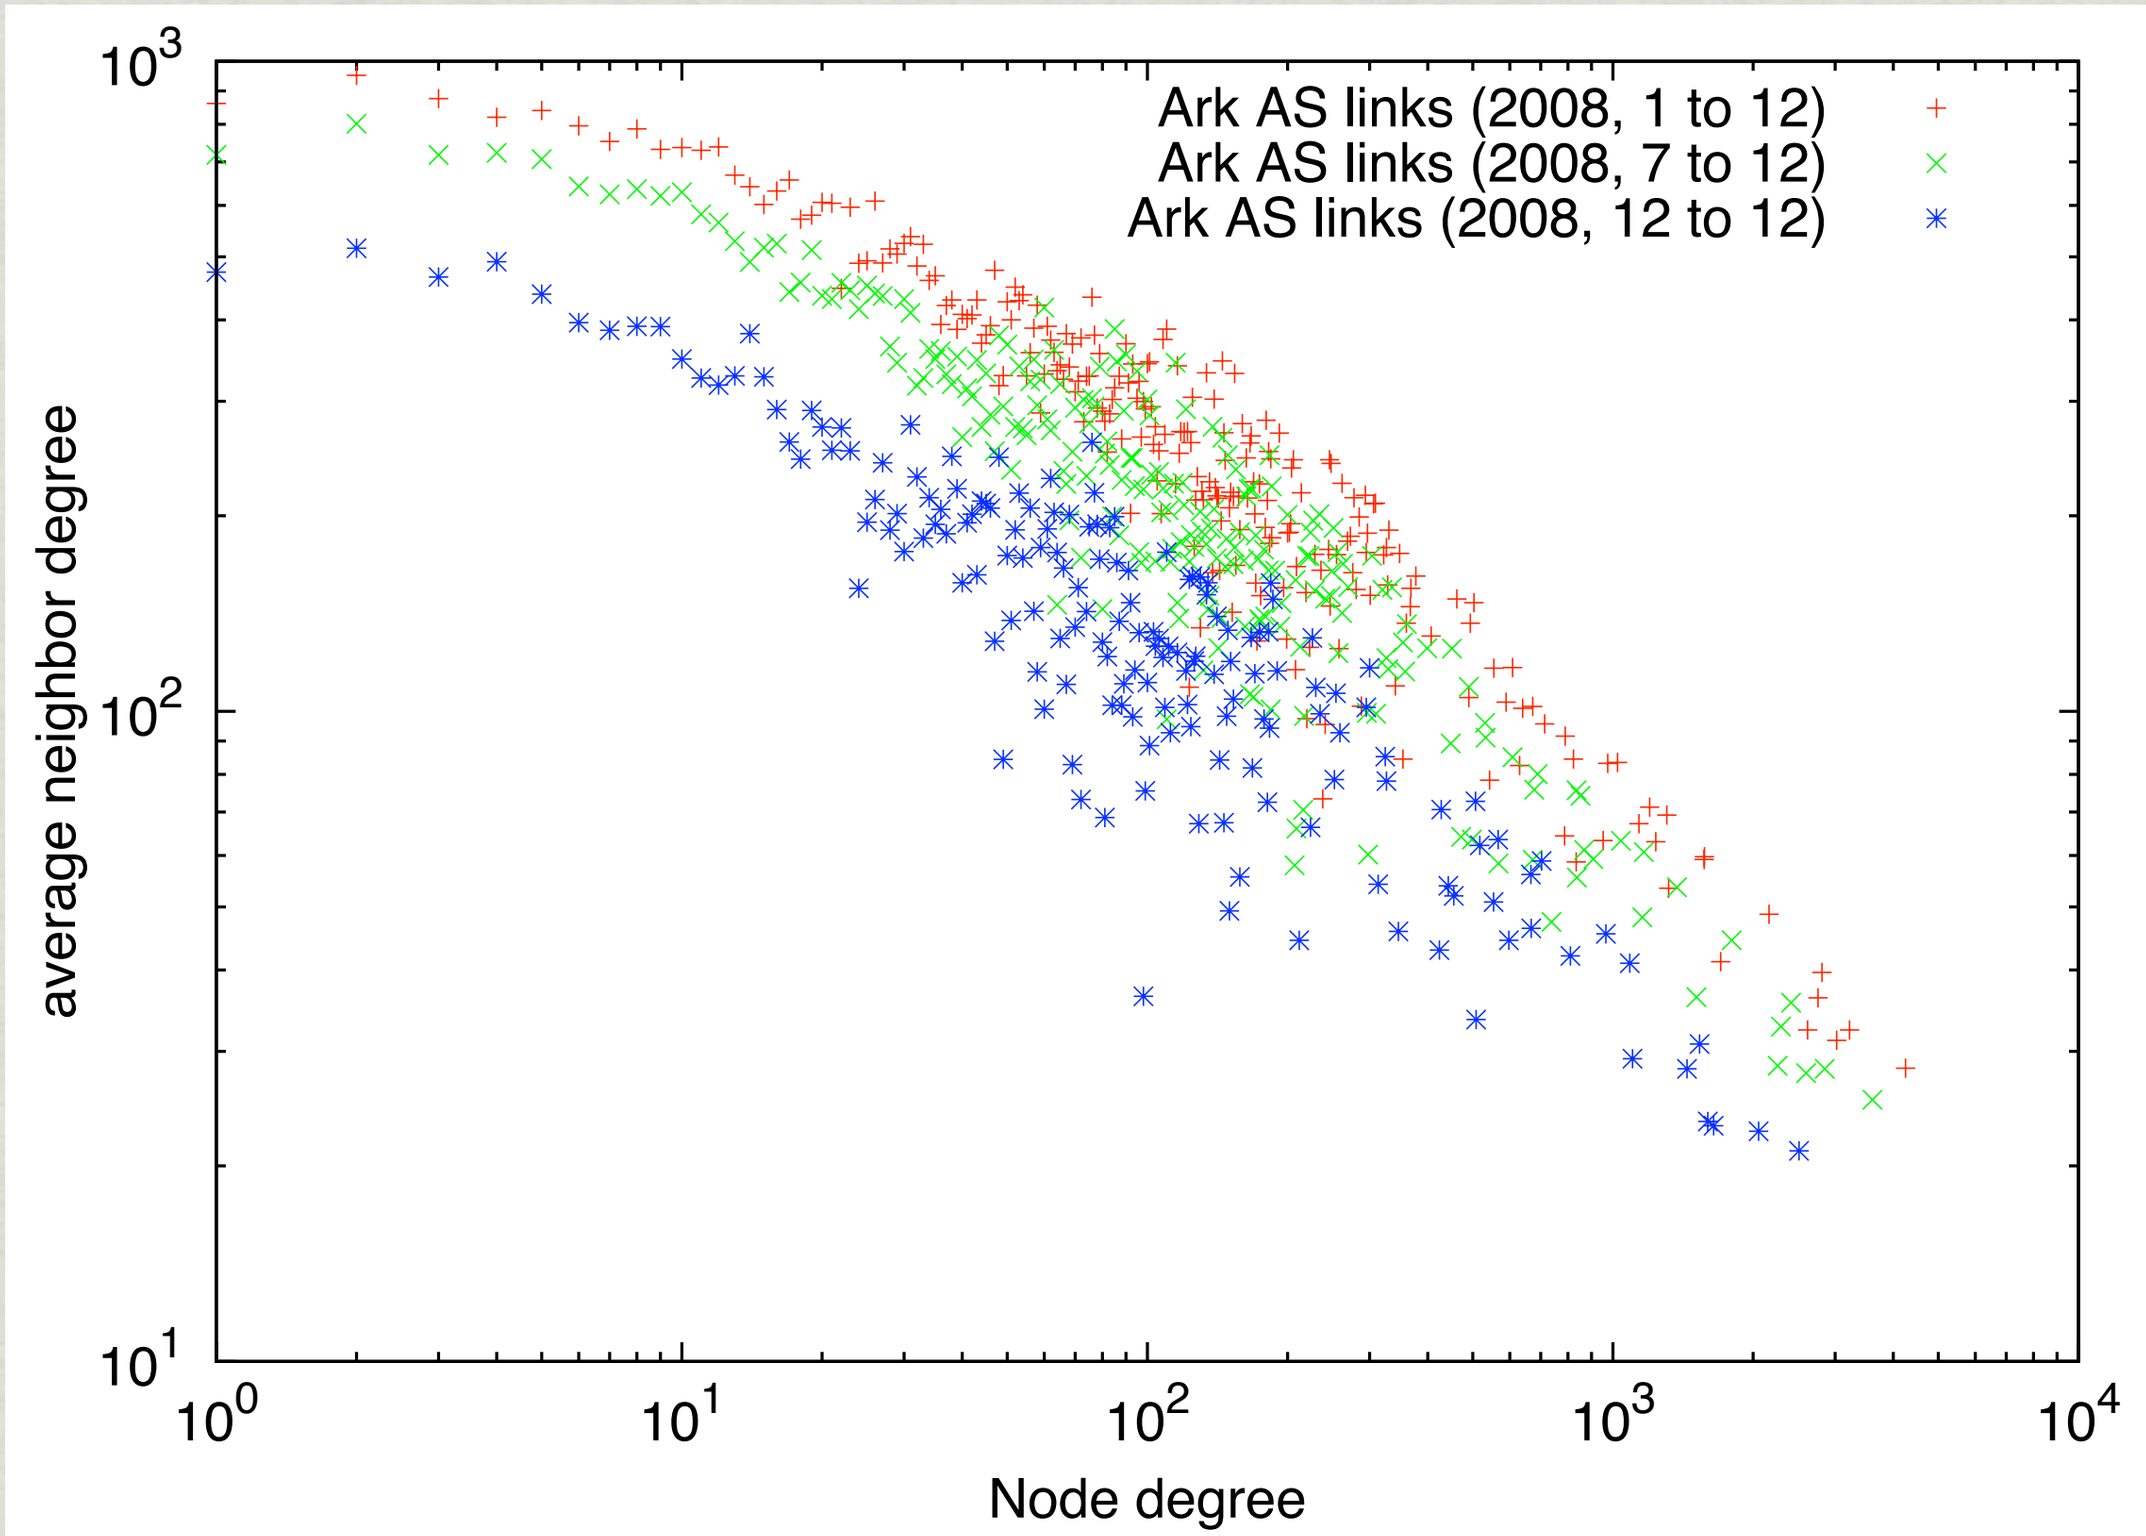

Ark AS Links: 1, 6, 12 Months

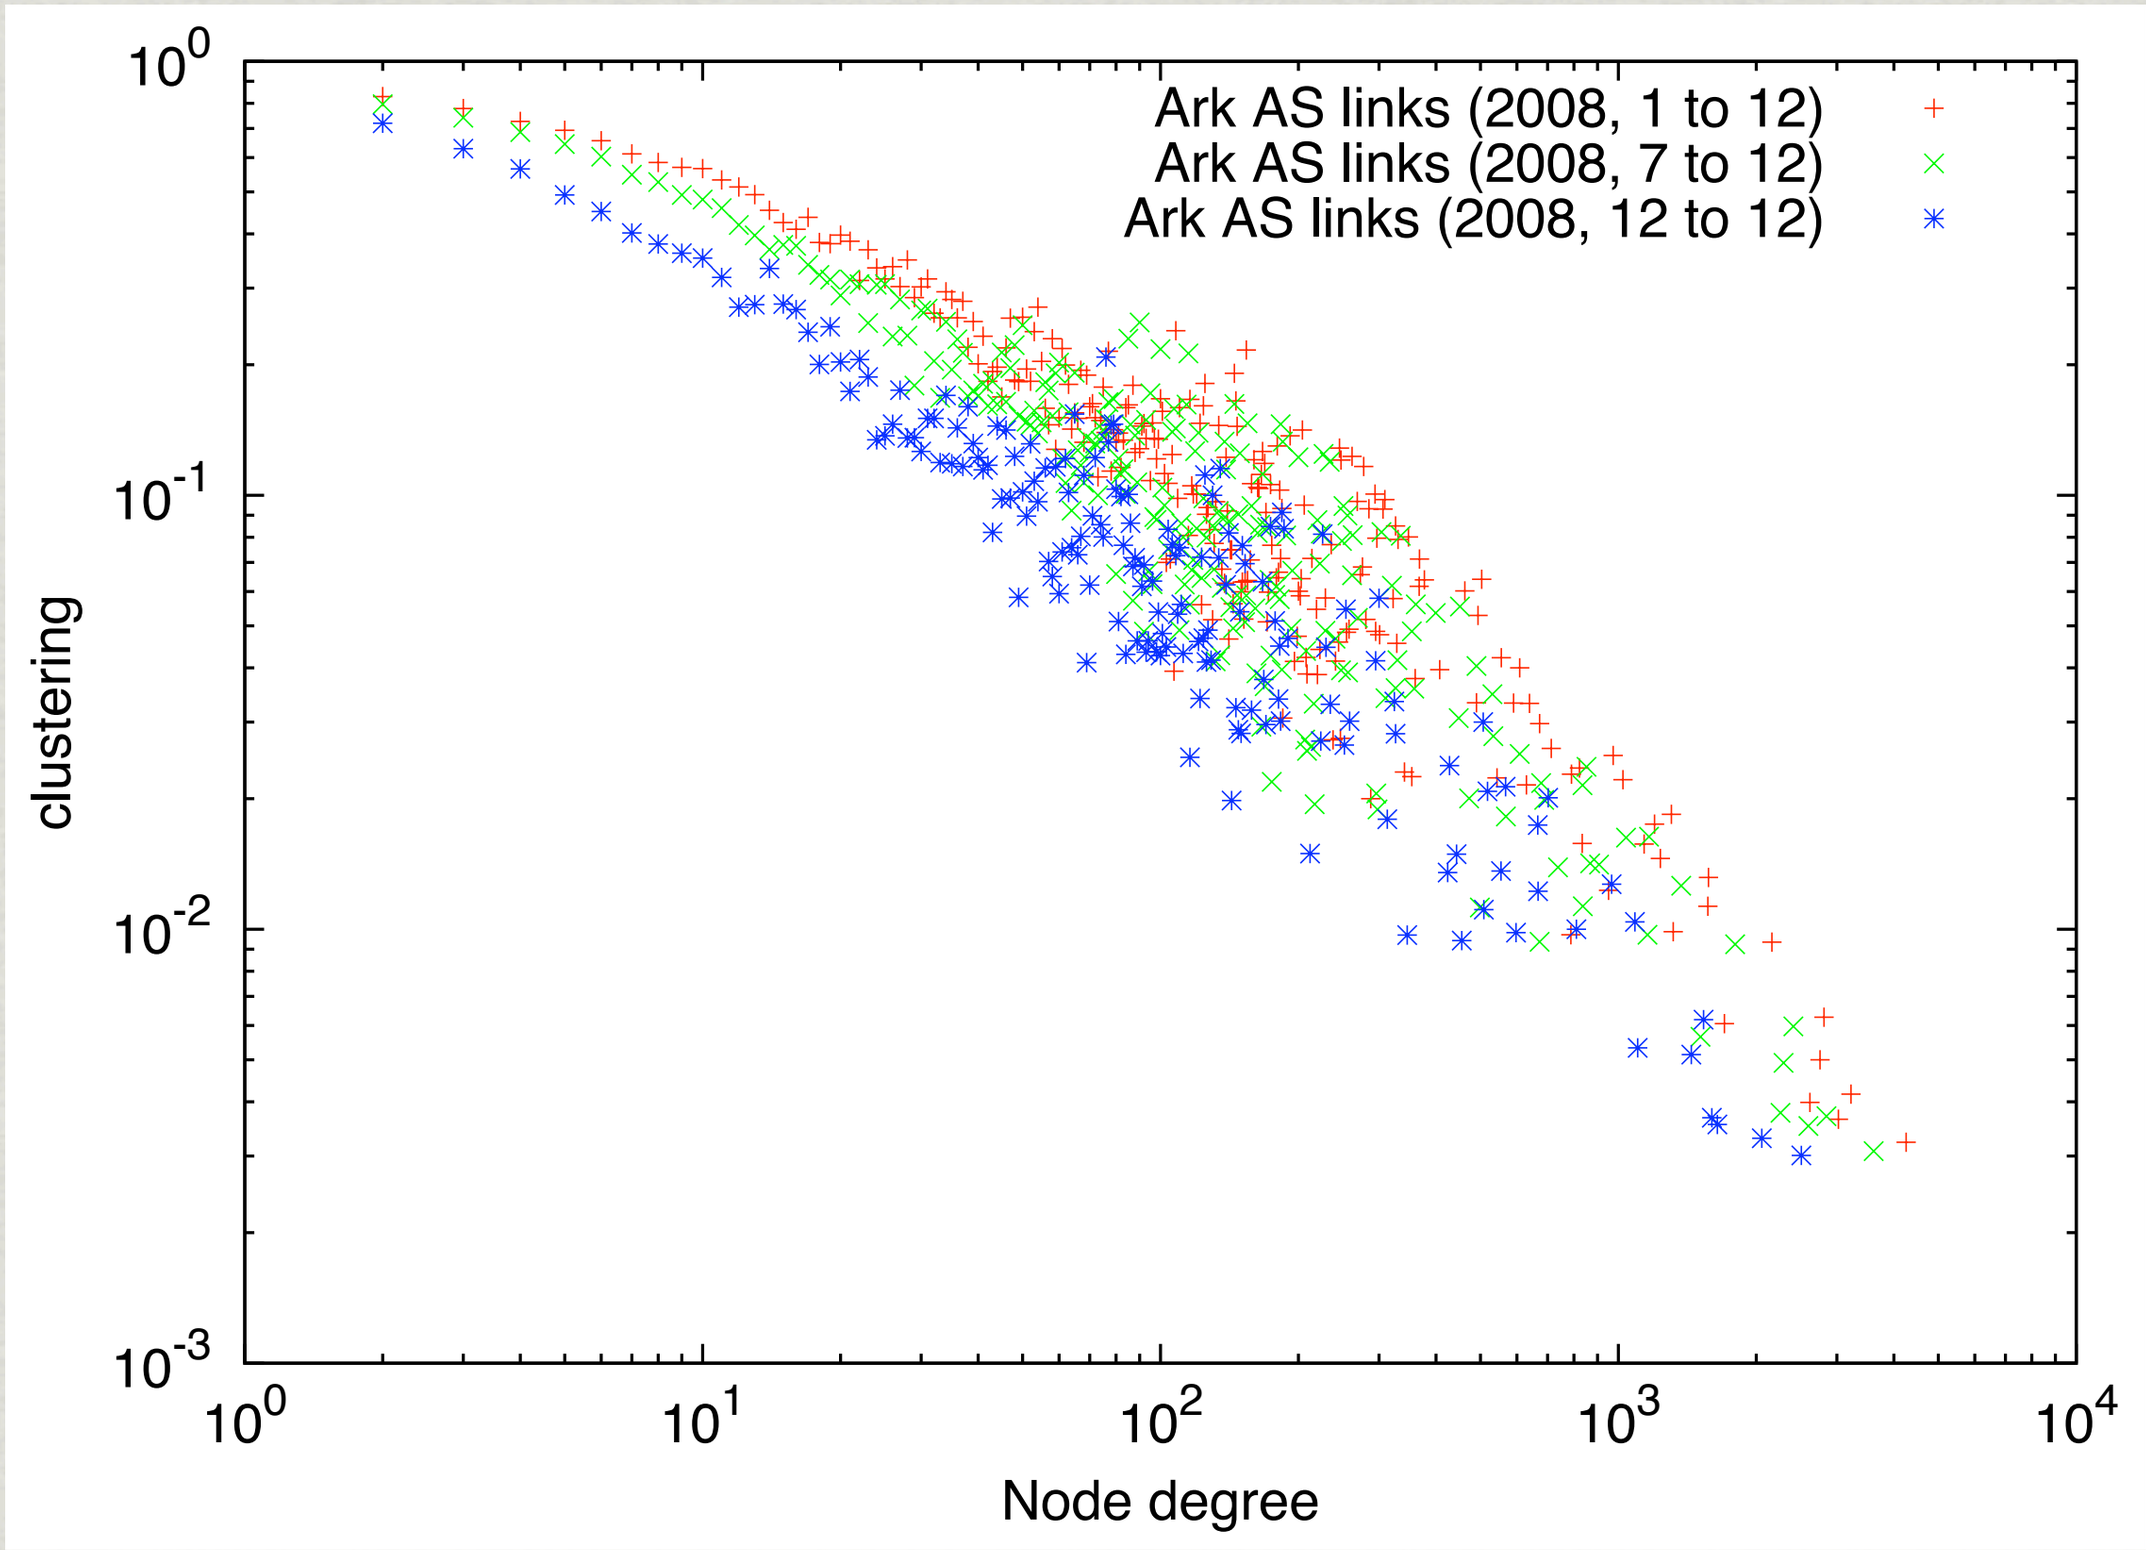

Ark IPv**6** Topology

- ongoing “large-scale” IPv6 measurements since Dec 12, 2008
- 10 monitors: 3 in US, 7 International
  - another IPv6 box coming Real Soon Now
- ICMP Paris traceroute to every routed prefix
  - each monitor probes a random destination in every routed prefix in every cycle; 1,553 prefixes <= /48
  - reduced probing rate to take 2 days per cycle
  - running scamper

Ark IPv**6** Topology

* statistics for 8 weeks of AS links from six sources:
  * Dec 12, 2008 to Feb 7, 2009

	nodes	links	max degree	average degree	average neighbor degree	mean clustering
IPv6 8 weeks	520	1,181	94	4.54	36.3	0.265
IPv4 4 weeks	23,425	56,760	2,509	4.85	467.3	0.354

Ark IPv6 AS Links

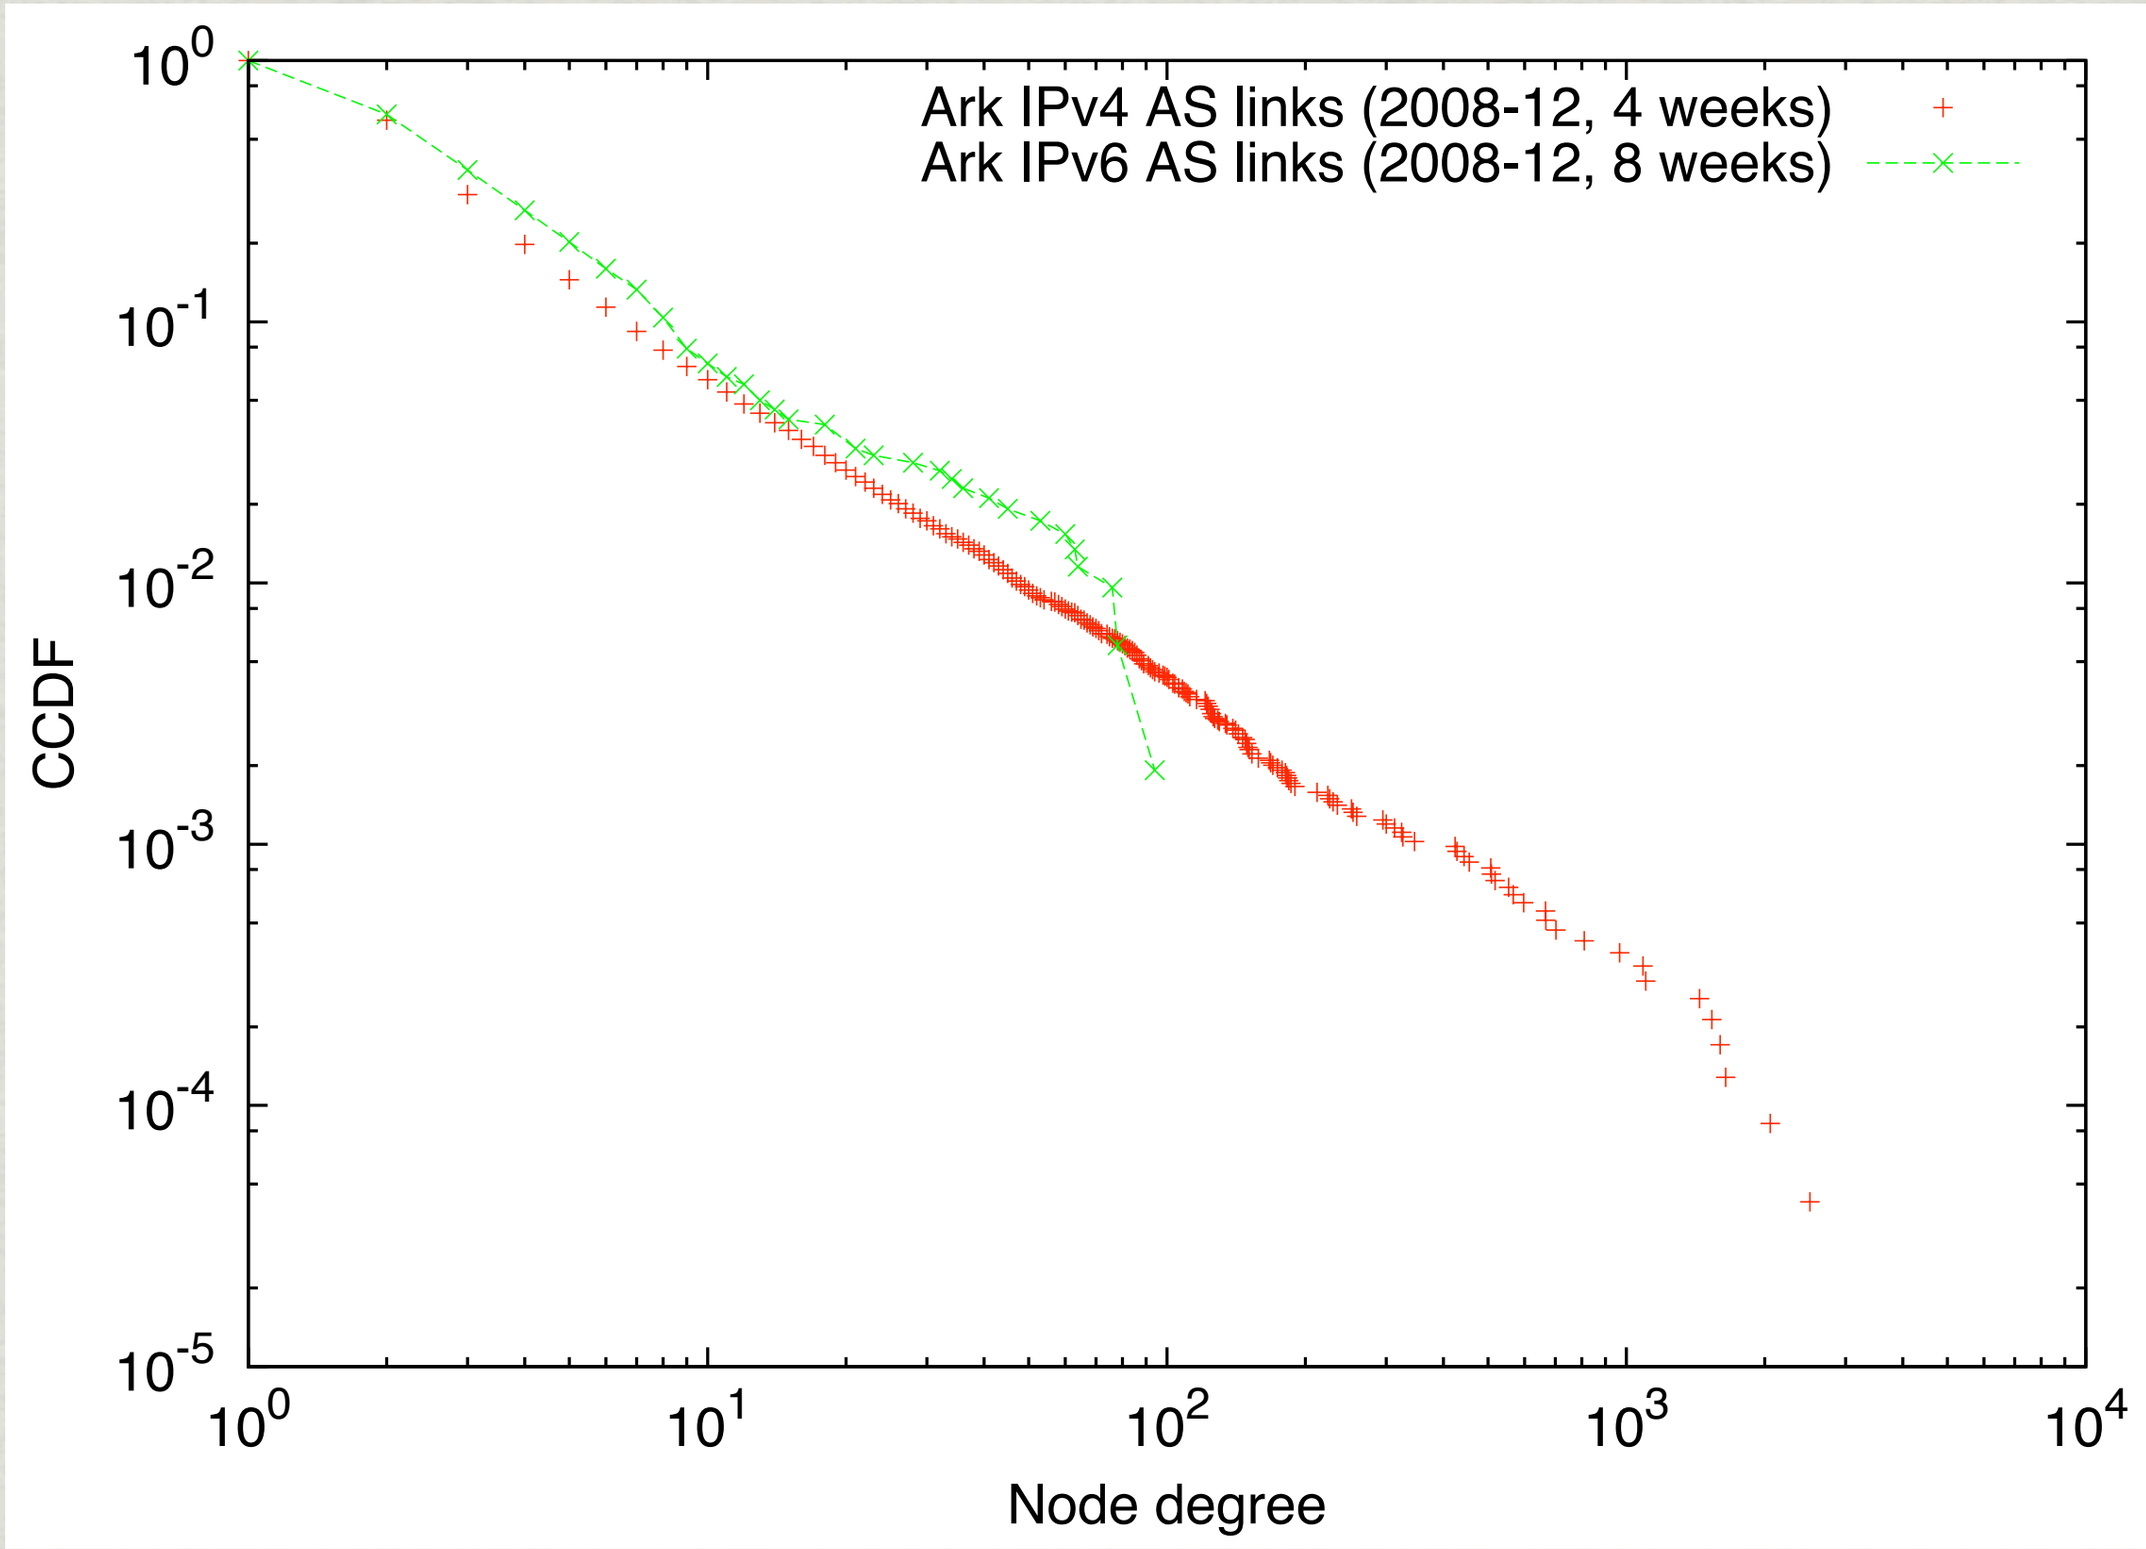

Ark IPv6 AS Links

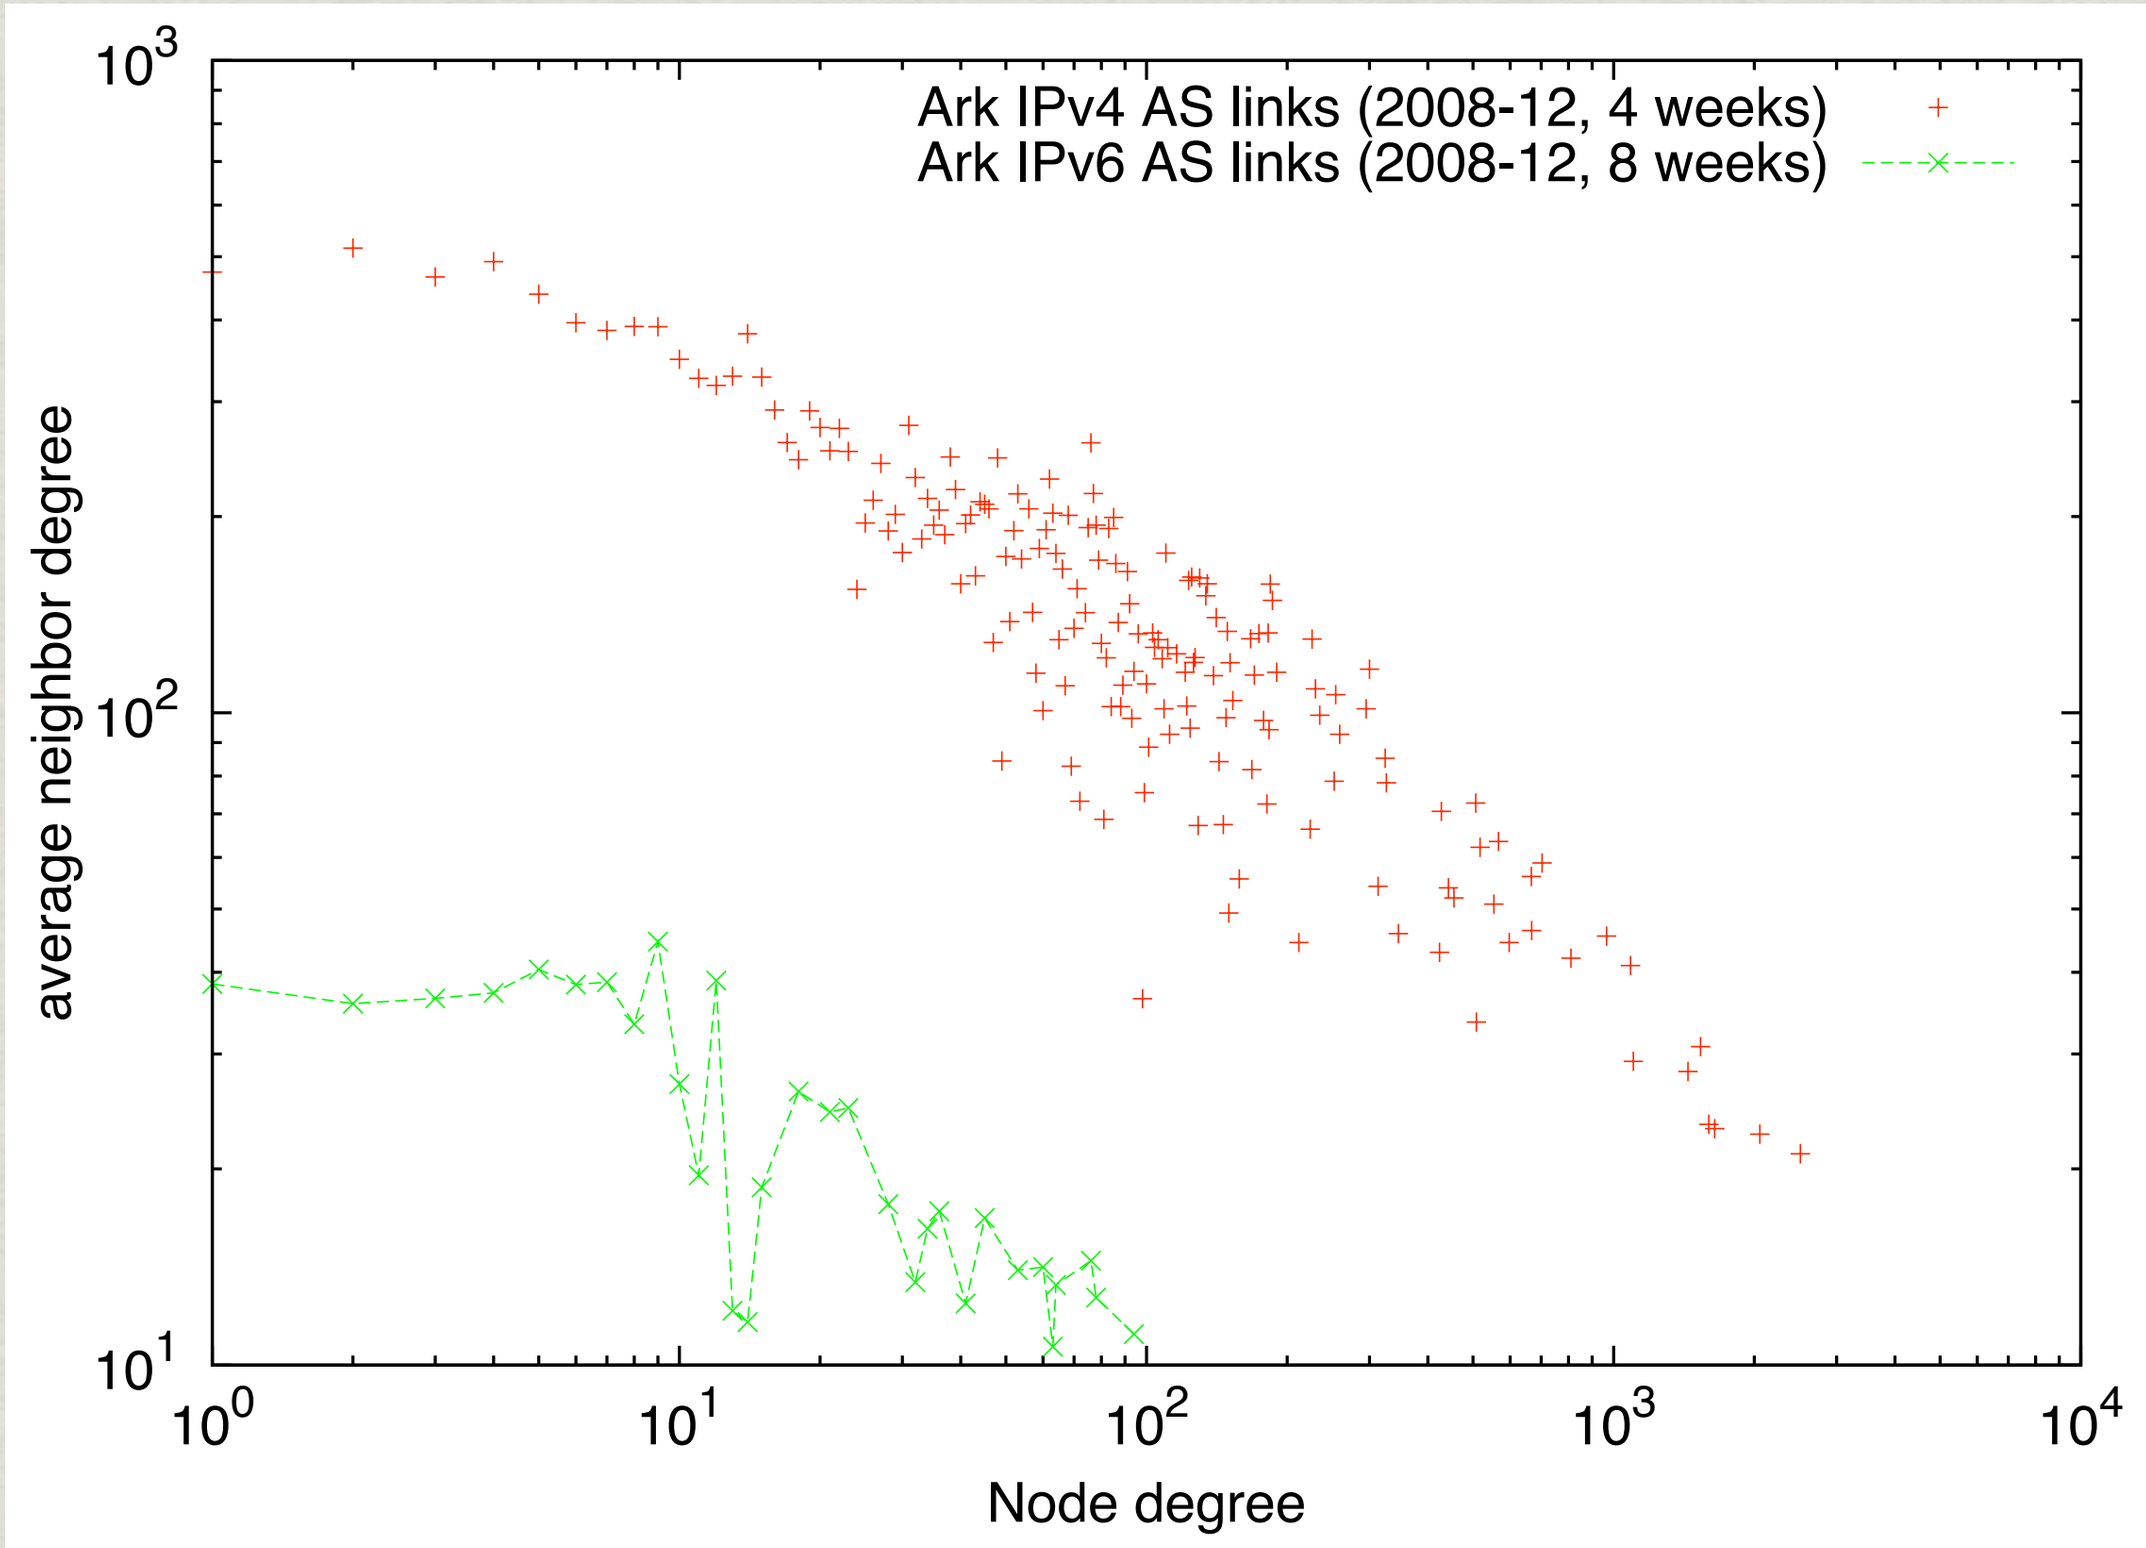

Ark IPv6 AS Links

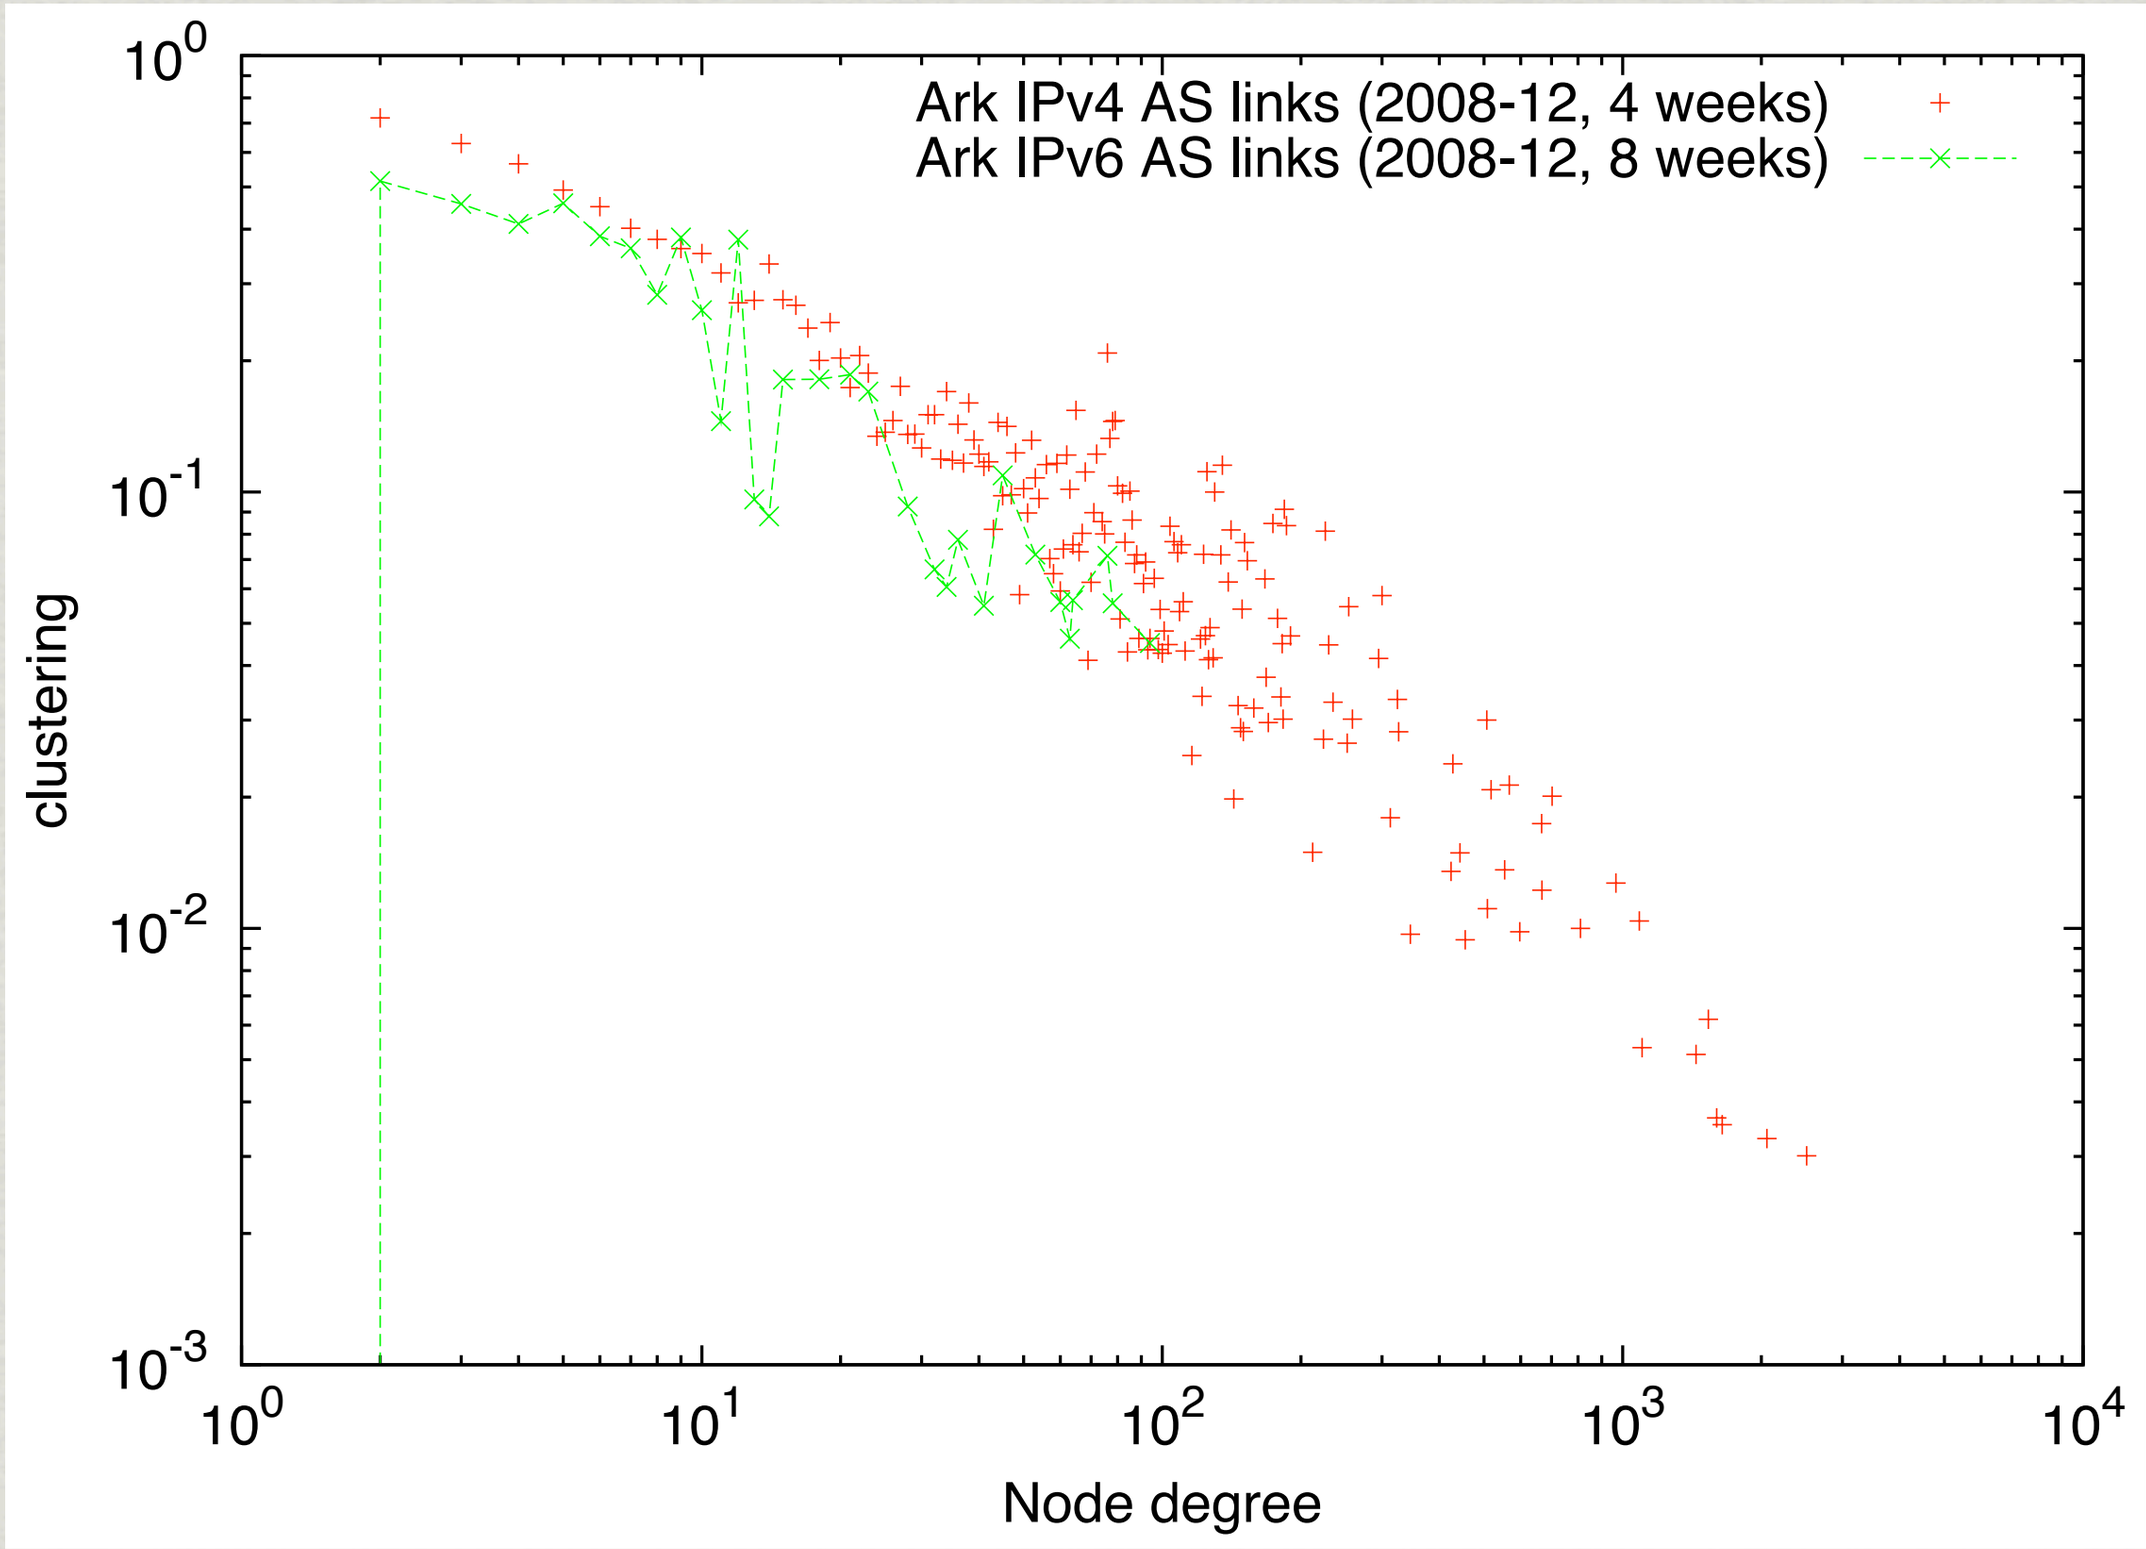

DNS Names

- automated ongoing DNS lookup of IP addresses seen in the Routed /24 Topology traces
  - all intermediate addresses and *responding* destinations
  - using our in-house bulk DNS lookup service (HostDB)
    - can look up millions of addresses per day
- 257M lookups since March 2008

DNS Traffic

- tcpdump capture of DNS query/response traffic
  - only for lookups of Routed /24 Topology addresses
  - continuous collection of 3-5M packets per day
  - can download most recent 30 days of pcap files
- a broad sampling of the nameservers on the Internet due to the broad coverage of the routed space in traces
- how many nameservers have IPv6 glue records? DNSSEC records? support EDNS? typical TTLs?

Alias Resolution

* Goal: collapse interfaces observed in traceroute paths into routers
  * toward a router-level map of the Internet
* alias resolution work led by Ken Keys

Spoofer Project

- collaboration with Rob Beverly on MIT Spoofer Project
  - how many networks allow packets with spoofed IP addresses to leave their network?
- Ark monitors act as targets for spoofed probes sent by willing participants
  - forwards received probe data to MIT server

Spoofer Project

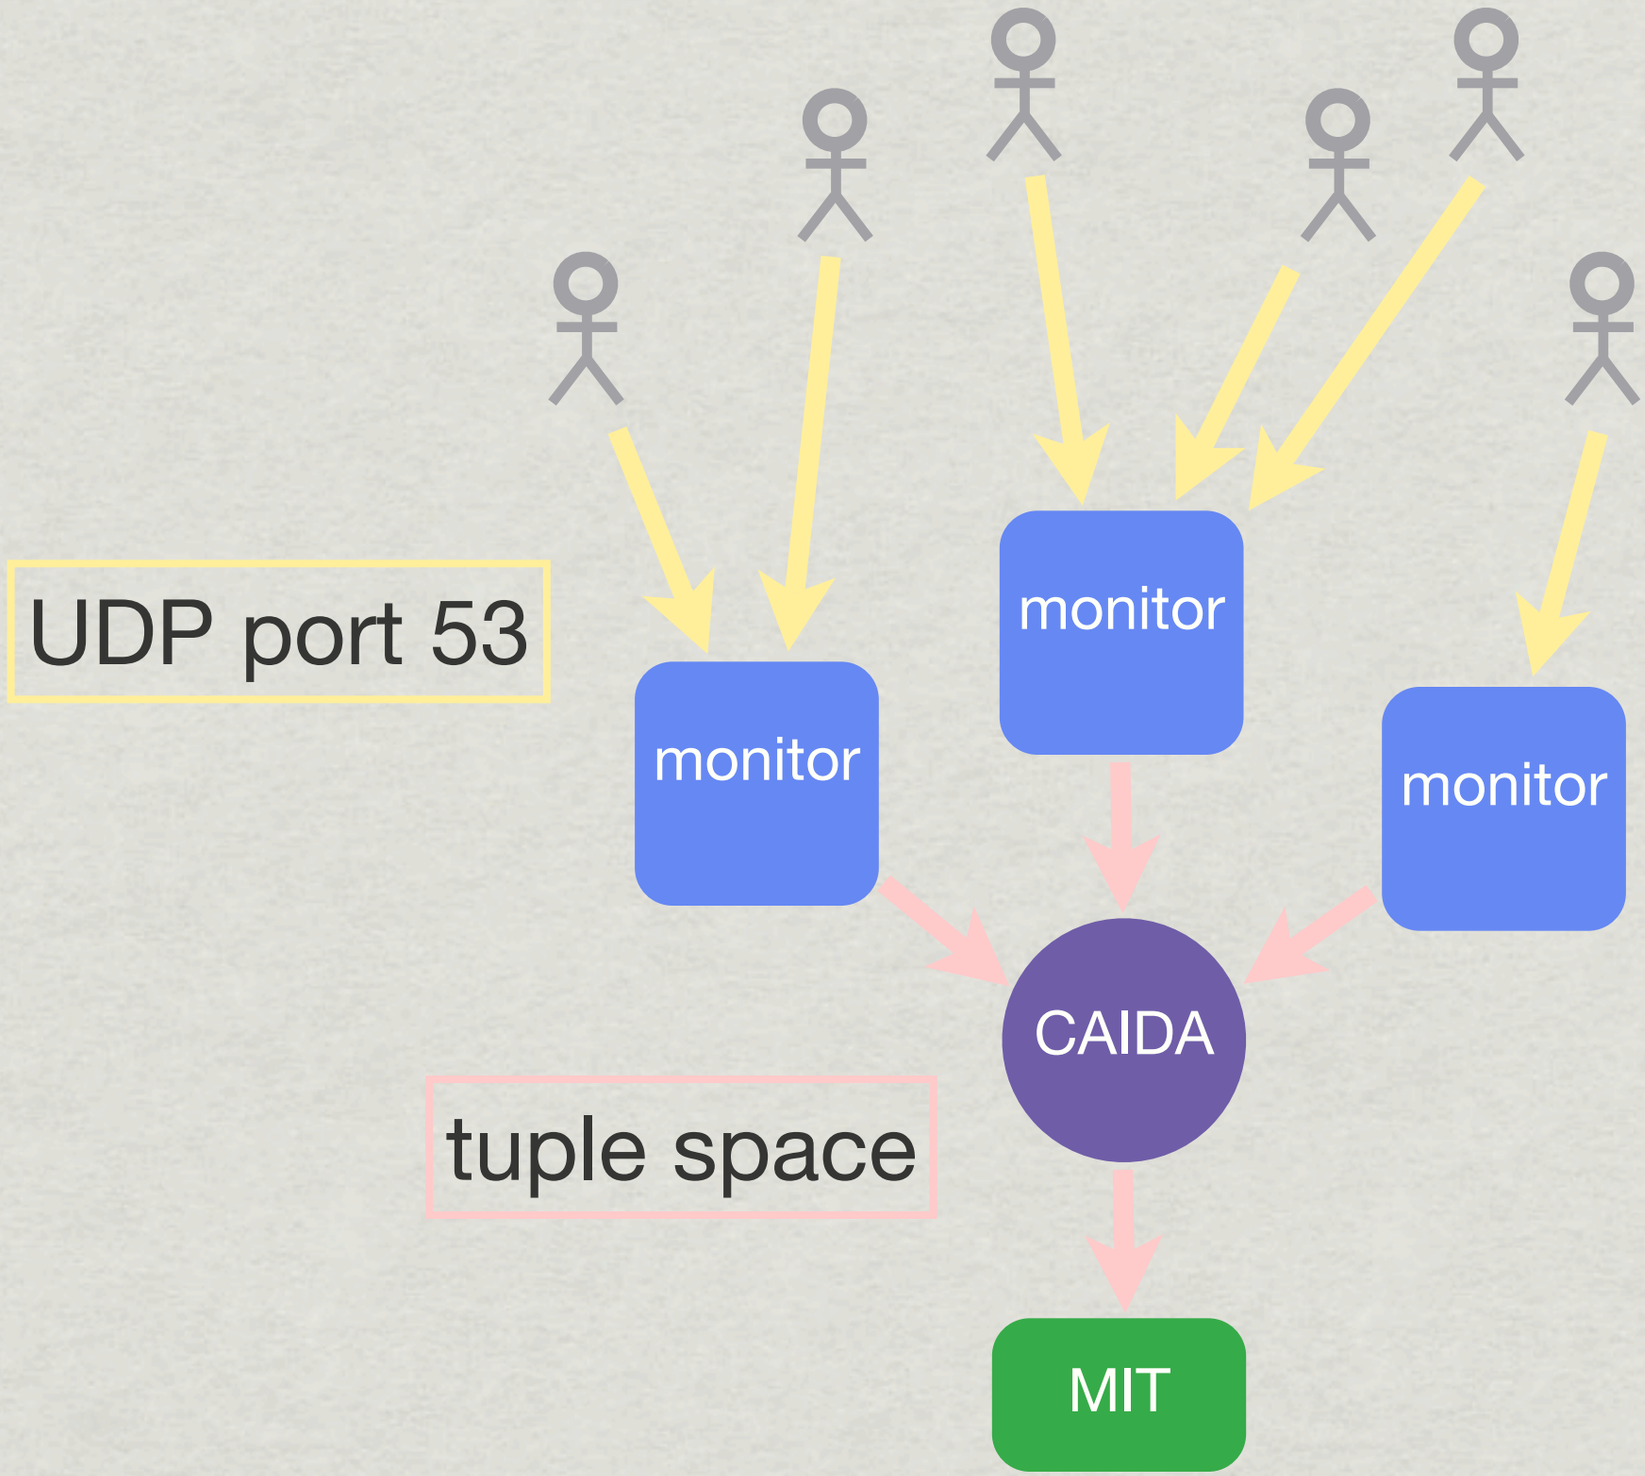

Ark Statistics Pages

- per-monitor analysis of IPv4 topology data
  - RTT, path length, RTT vs. distance

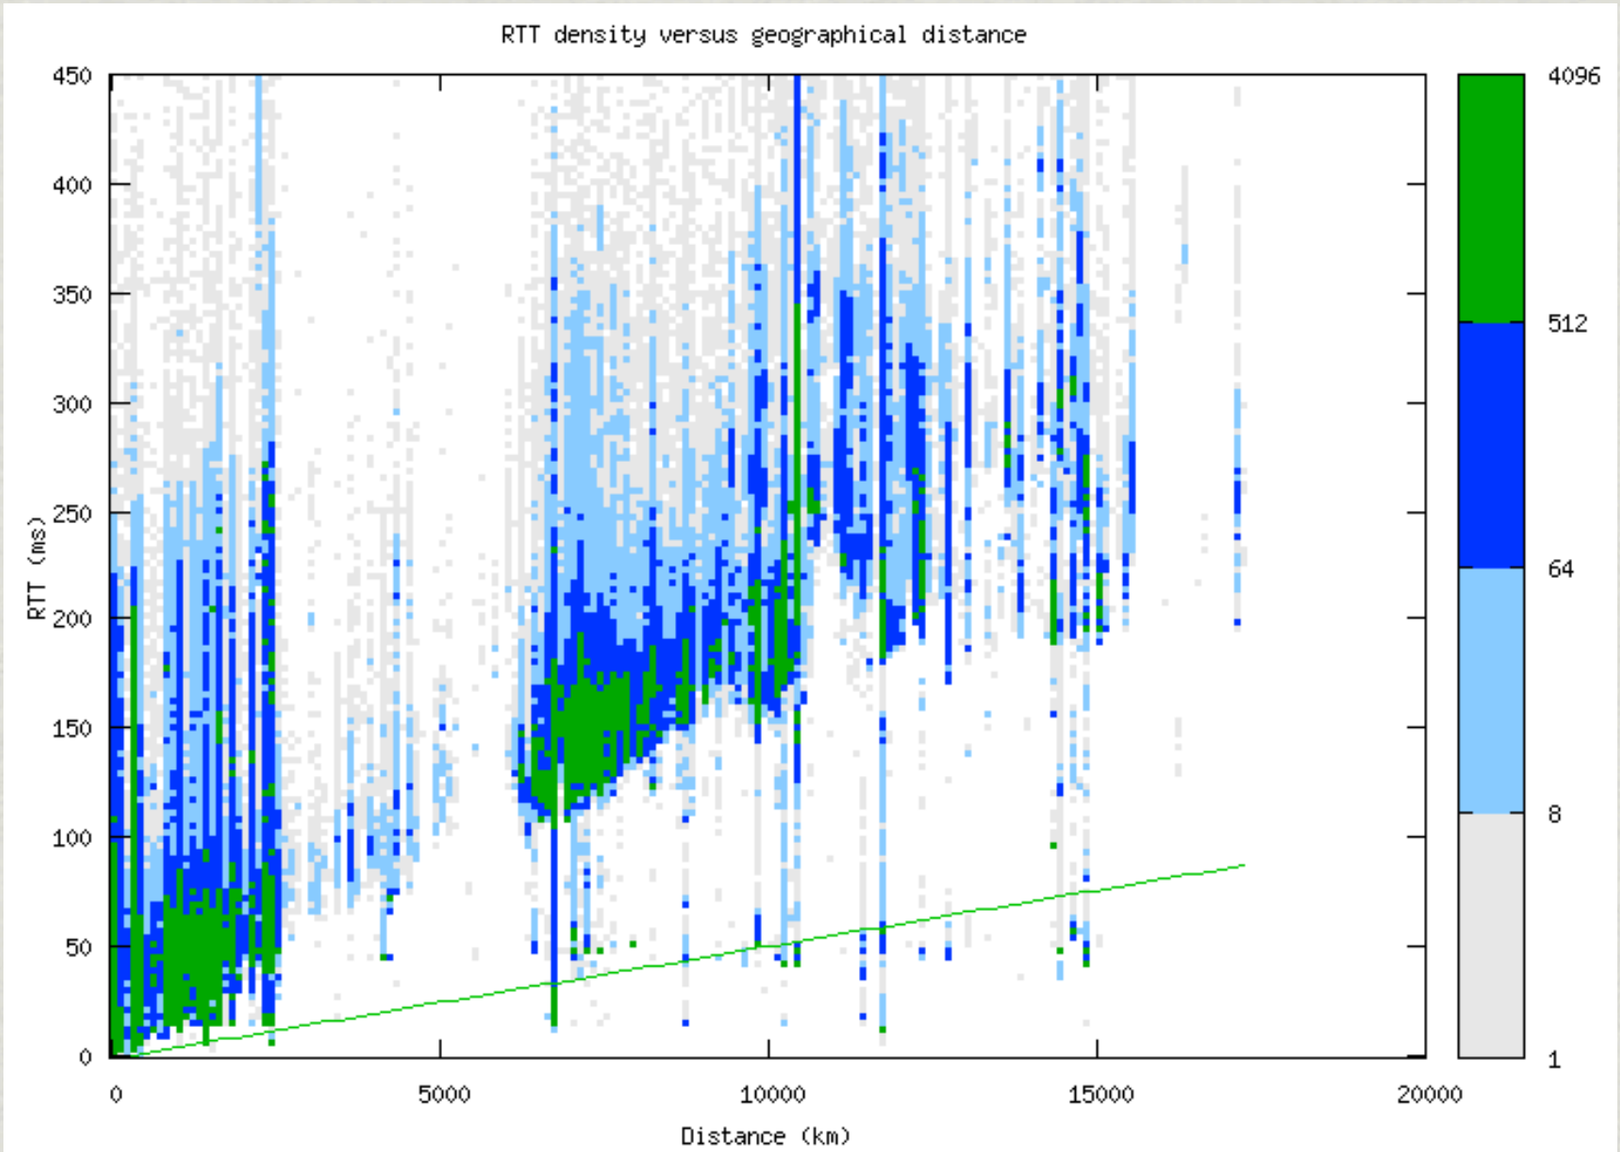

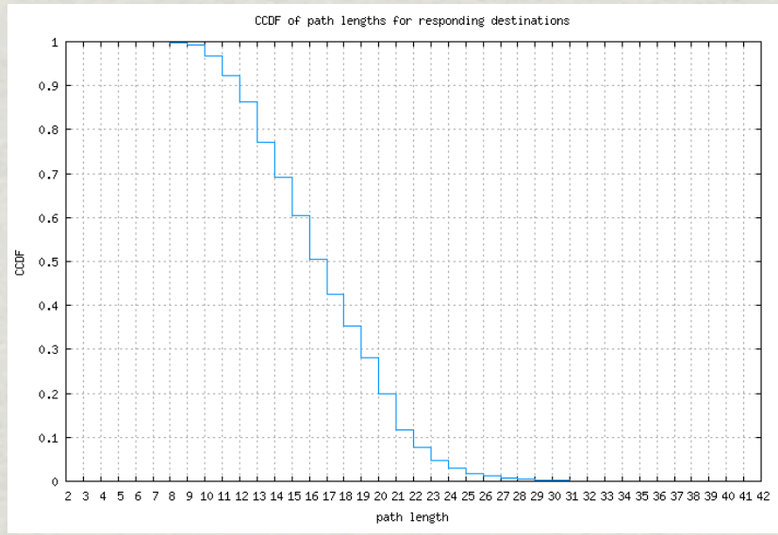

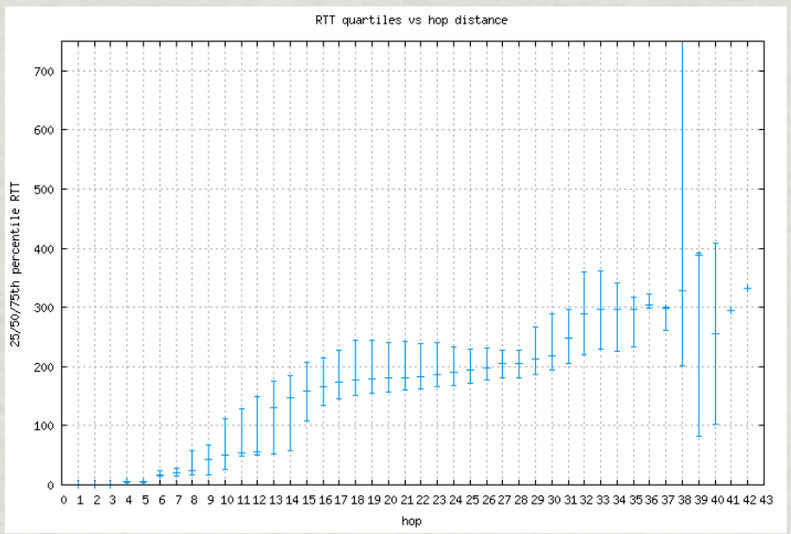

www.caida.org/projects/ark/statistics

Future Work

* release Marinda tuple space under GPL
* implement large-scale RadarGun measurements
* more in-depth analysis of data for stats pages
* investigate AS link densification
* DNS open resolver surveys?
* high-level packet generation, capture, and analysis API
* allow semi-trusted 3rd parties to conduct measurements

Thanks!

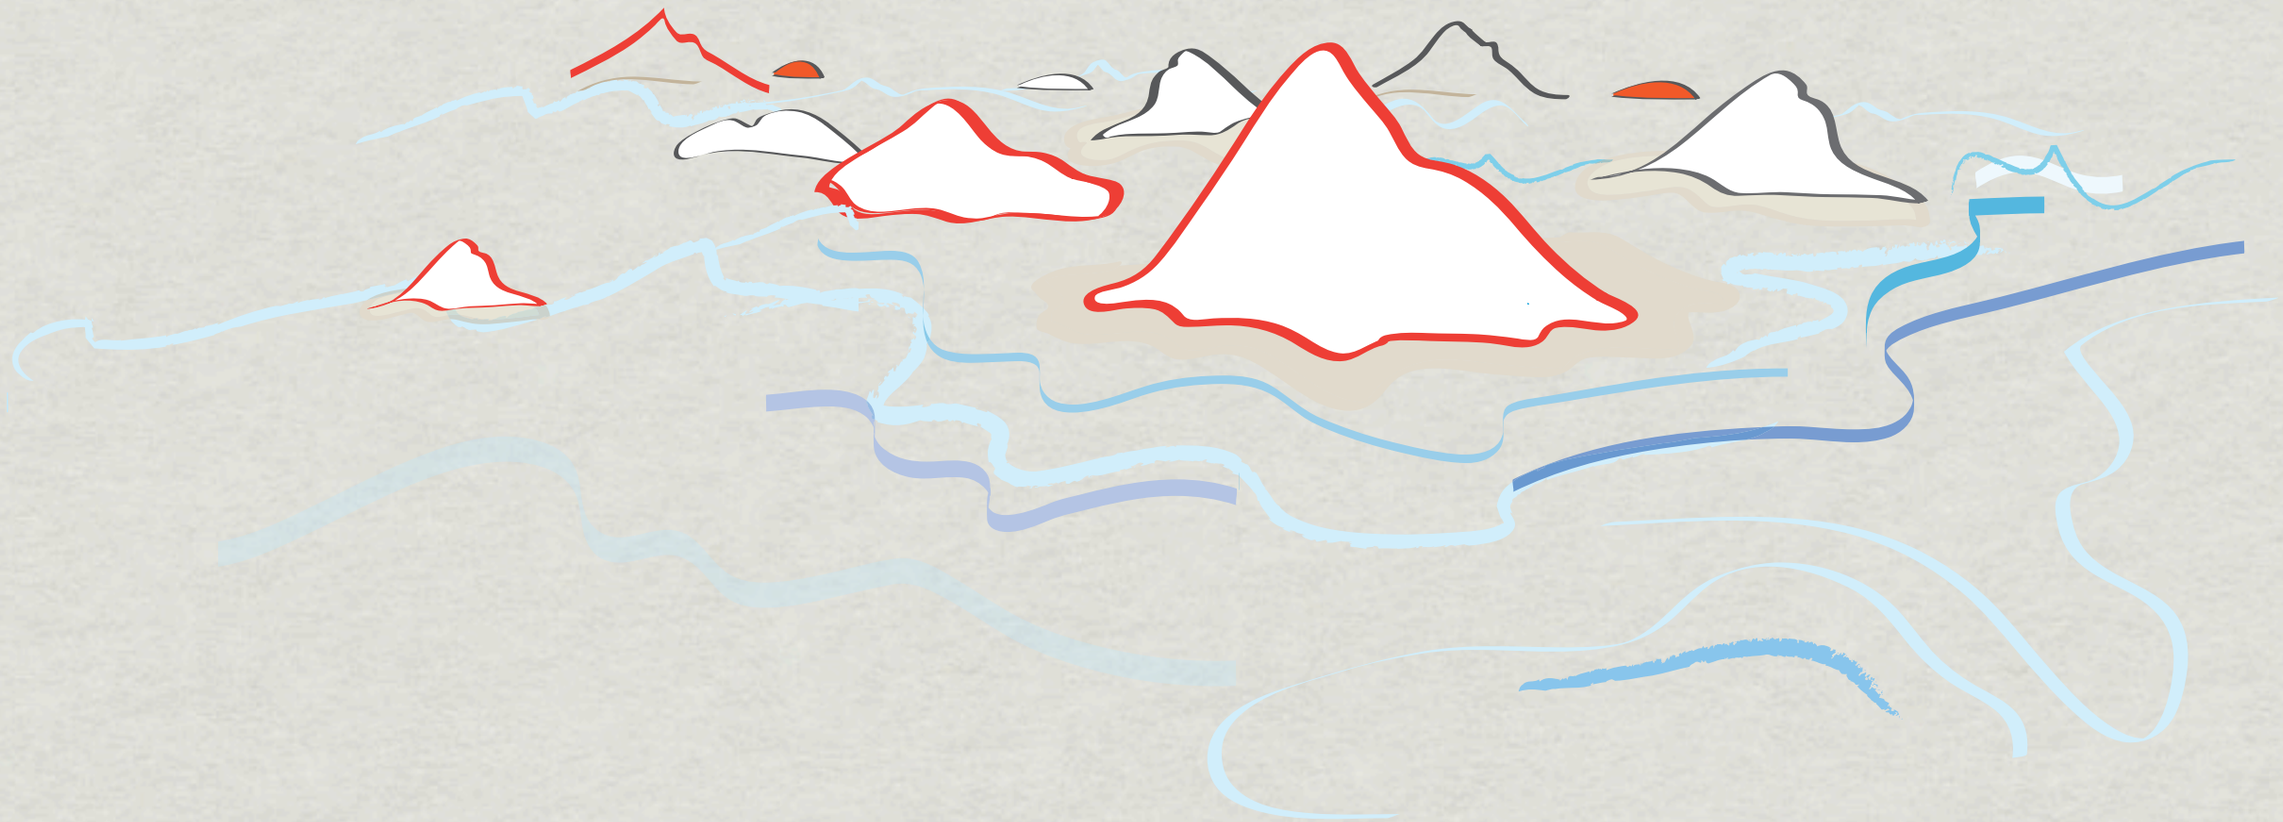

For more information and to request data:

www.caida.org/projects/ark