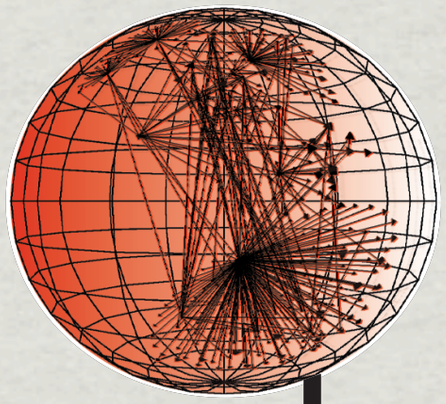

caida's Topology Updates and Analysis

Bradley Huffaker
CAIDA

2009 CAIDA/WIDE/CASFI Workshop (Korea)
April 3rd, 2009

Outline

- Archipelago
  - Monitor Deployment
  - Comparison of AS links (Ark, Dimes, Routeviews)
  - Future Work
- Alias Resolution
  - List of Technics Used
  - Evaluation of various Technics
  - Future work (RADAR GUN)

Archipelago

Measurement Infrastructure

Young Hyun

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

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

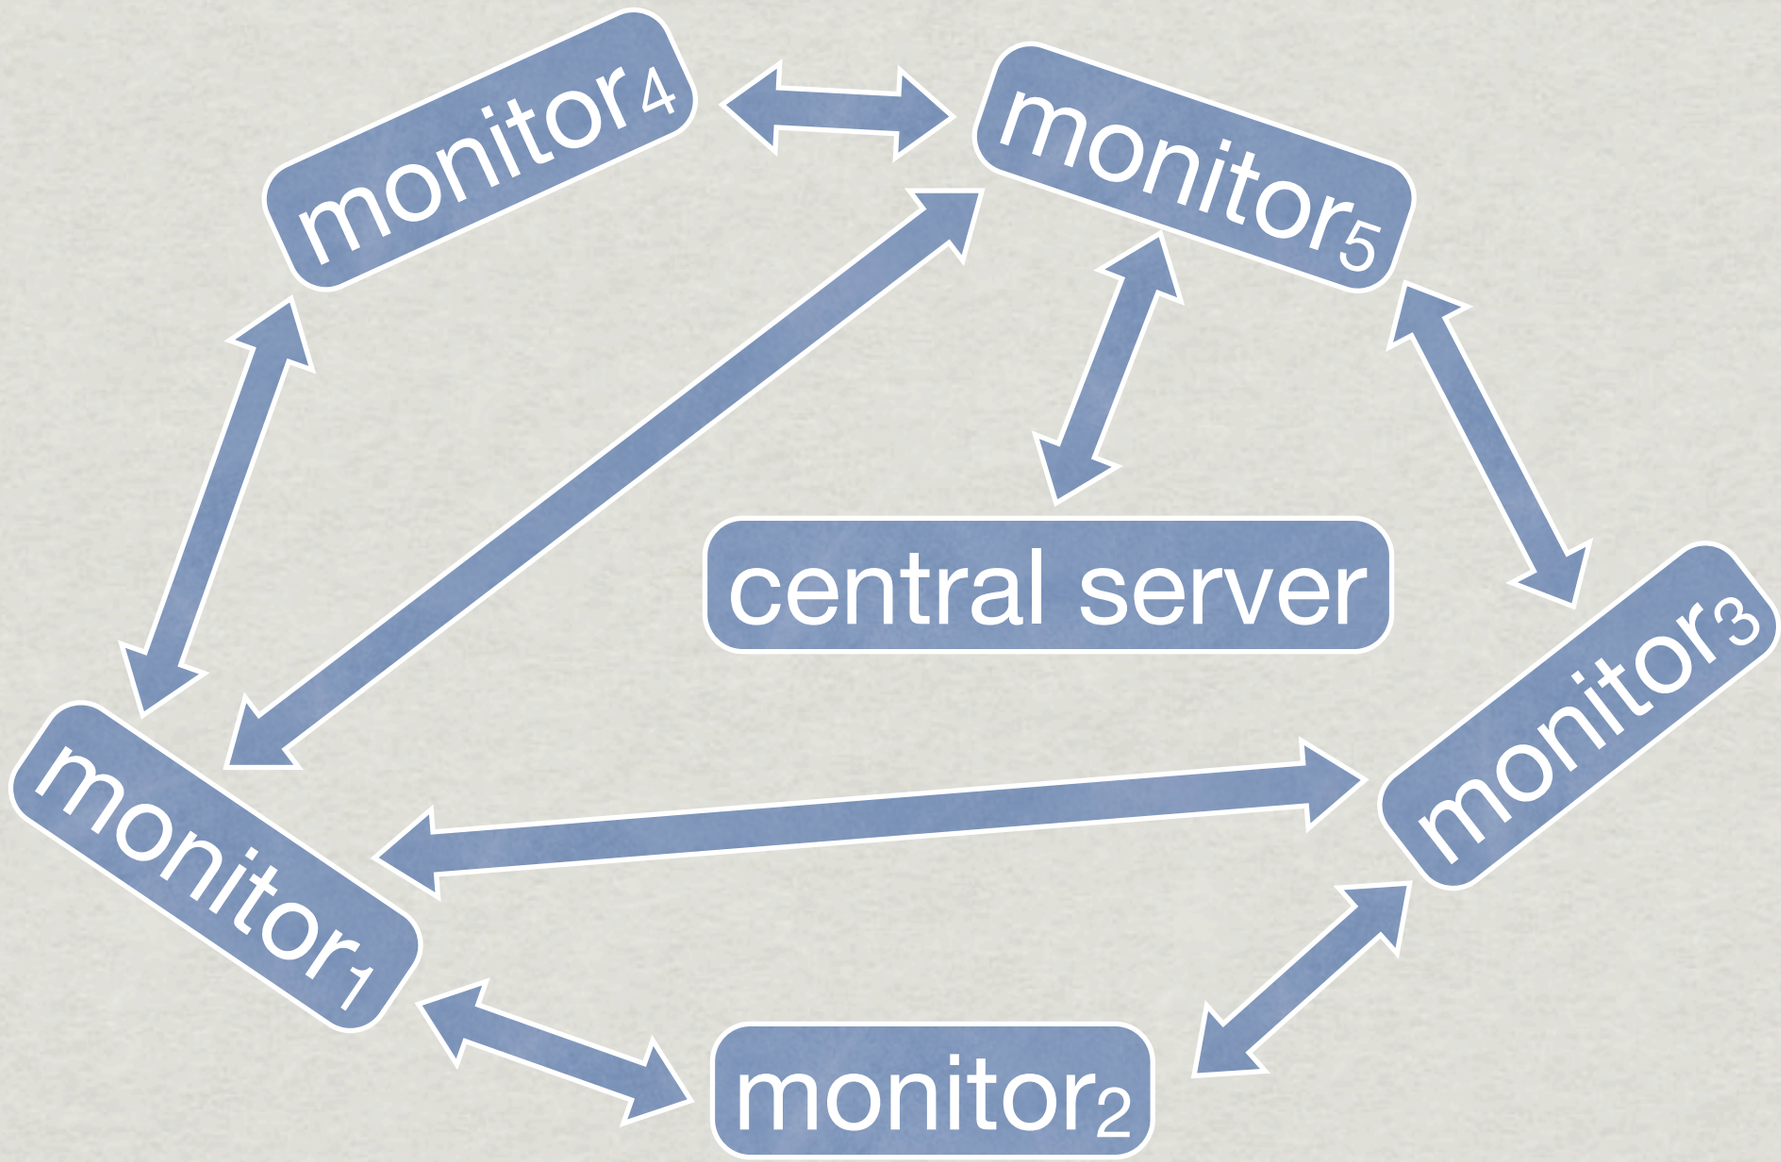

Monitor Deployment

* 33 monitors in 22 countries

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

	Organization
19	academic
9	research network
2	network infrastructure
1	commercial network
1	community network
1	military research

Measurements

- IPv4 Routed /24 Topology
- IPv4 Routed /24 AS Links
- DNS Names
- DNS Query/Response Traffic
- IPv6 Topology
- Spoofer Project Collaboration

IPv4 Routed /24 AS Links

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

	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*

Ark IPv**6** Topology

- ongoing “large-scale” IPv6 measurements since Dec 12, 2008
- 6 monitors: 3 in US, 3 in Europe
  - 2 IPv6 boxes down
  - 3 more IPv6 boxes 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

AS Core IPv4 vs IPv6

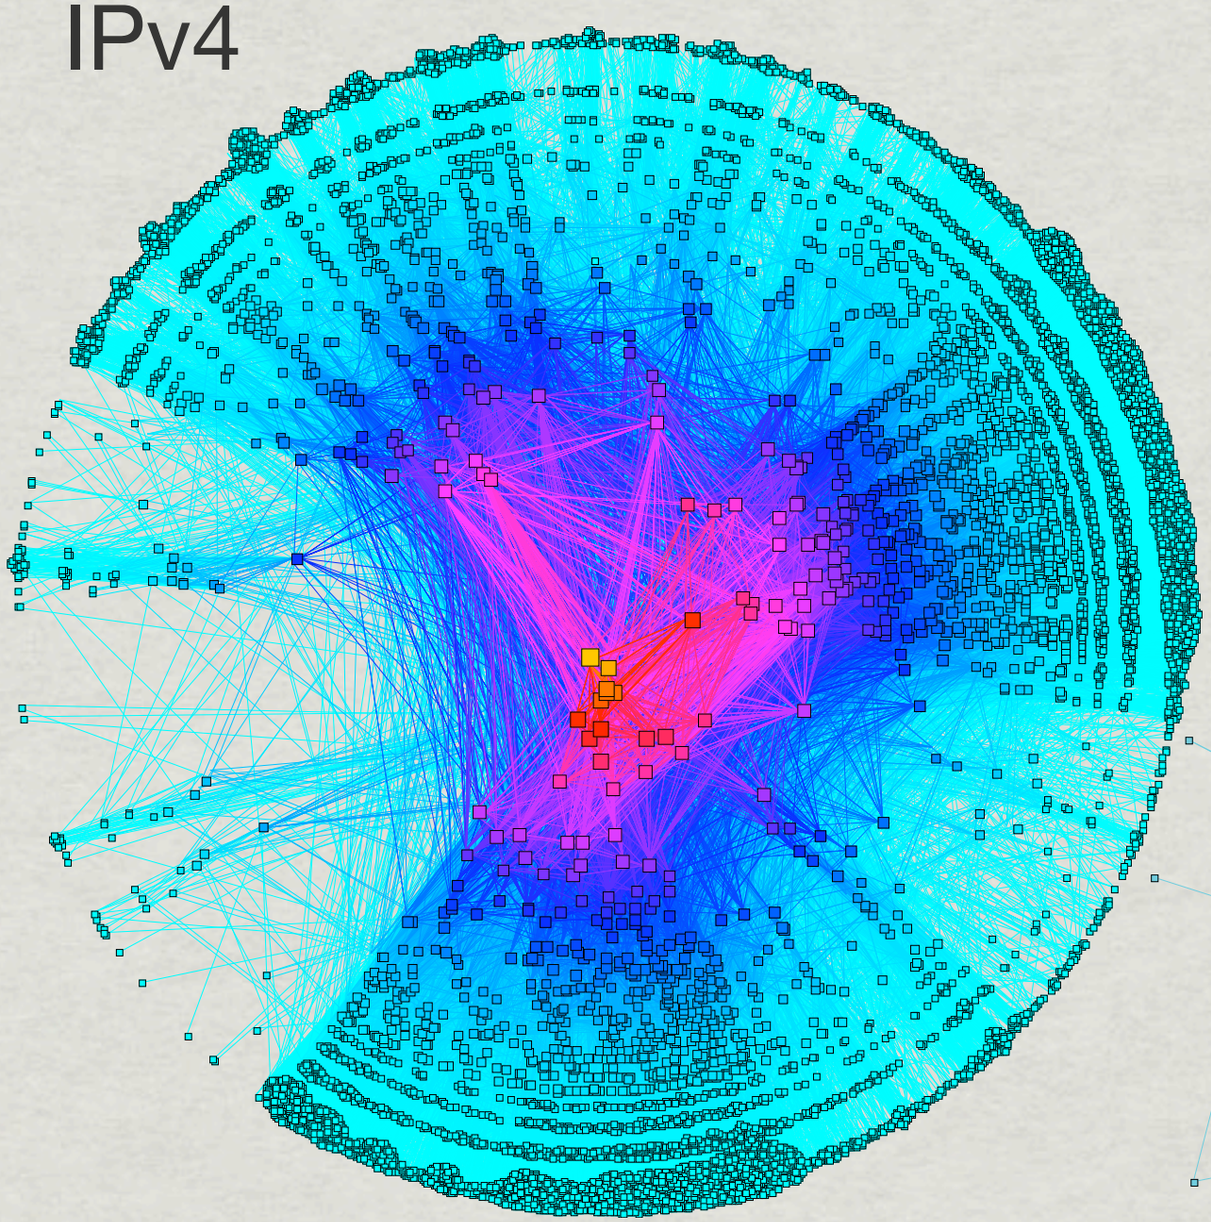

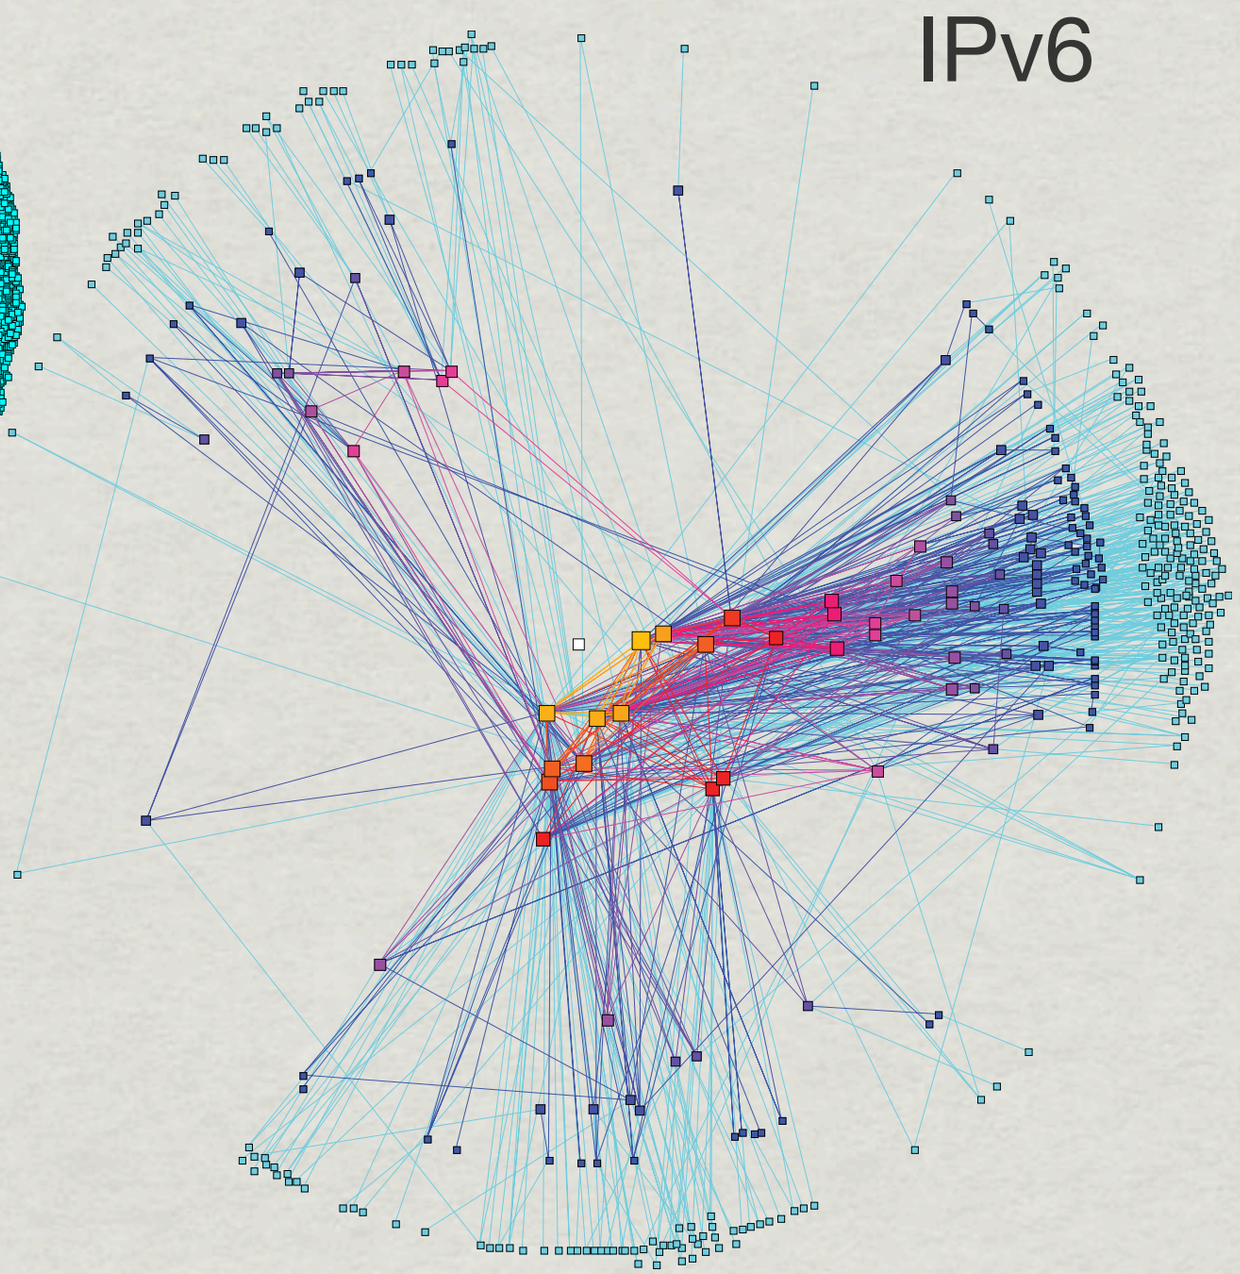

Spoofer Project

- 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

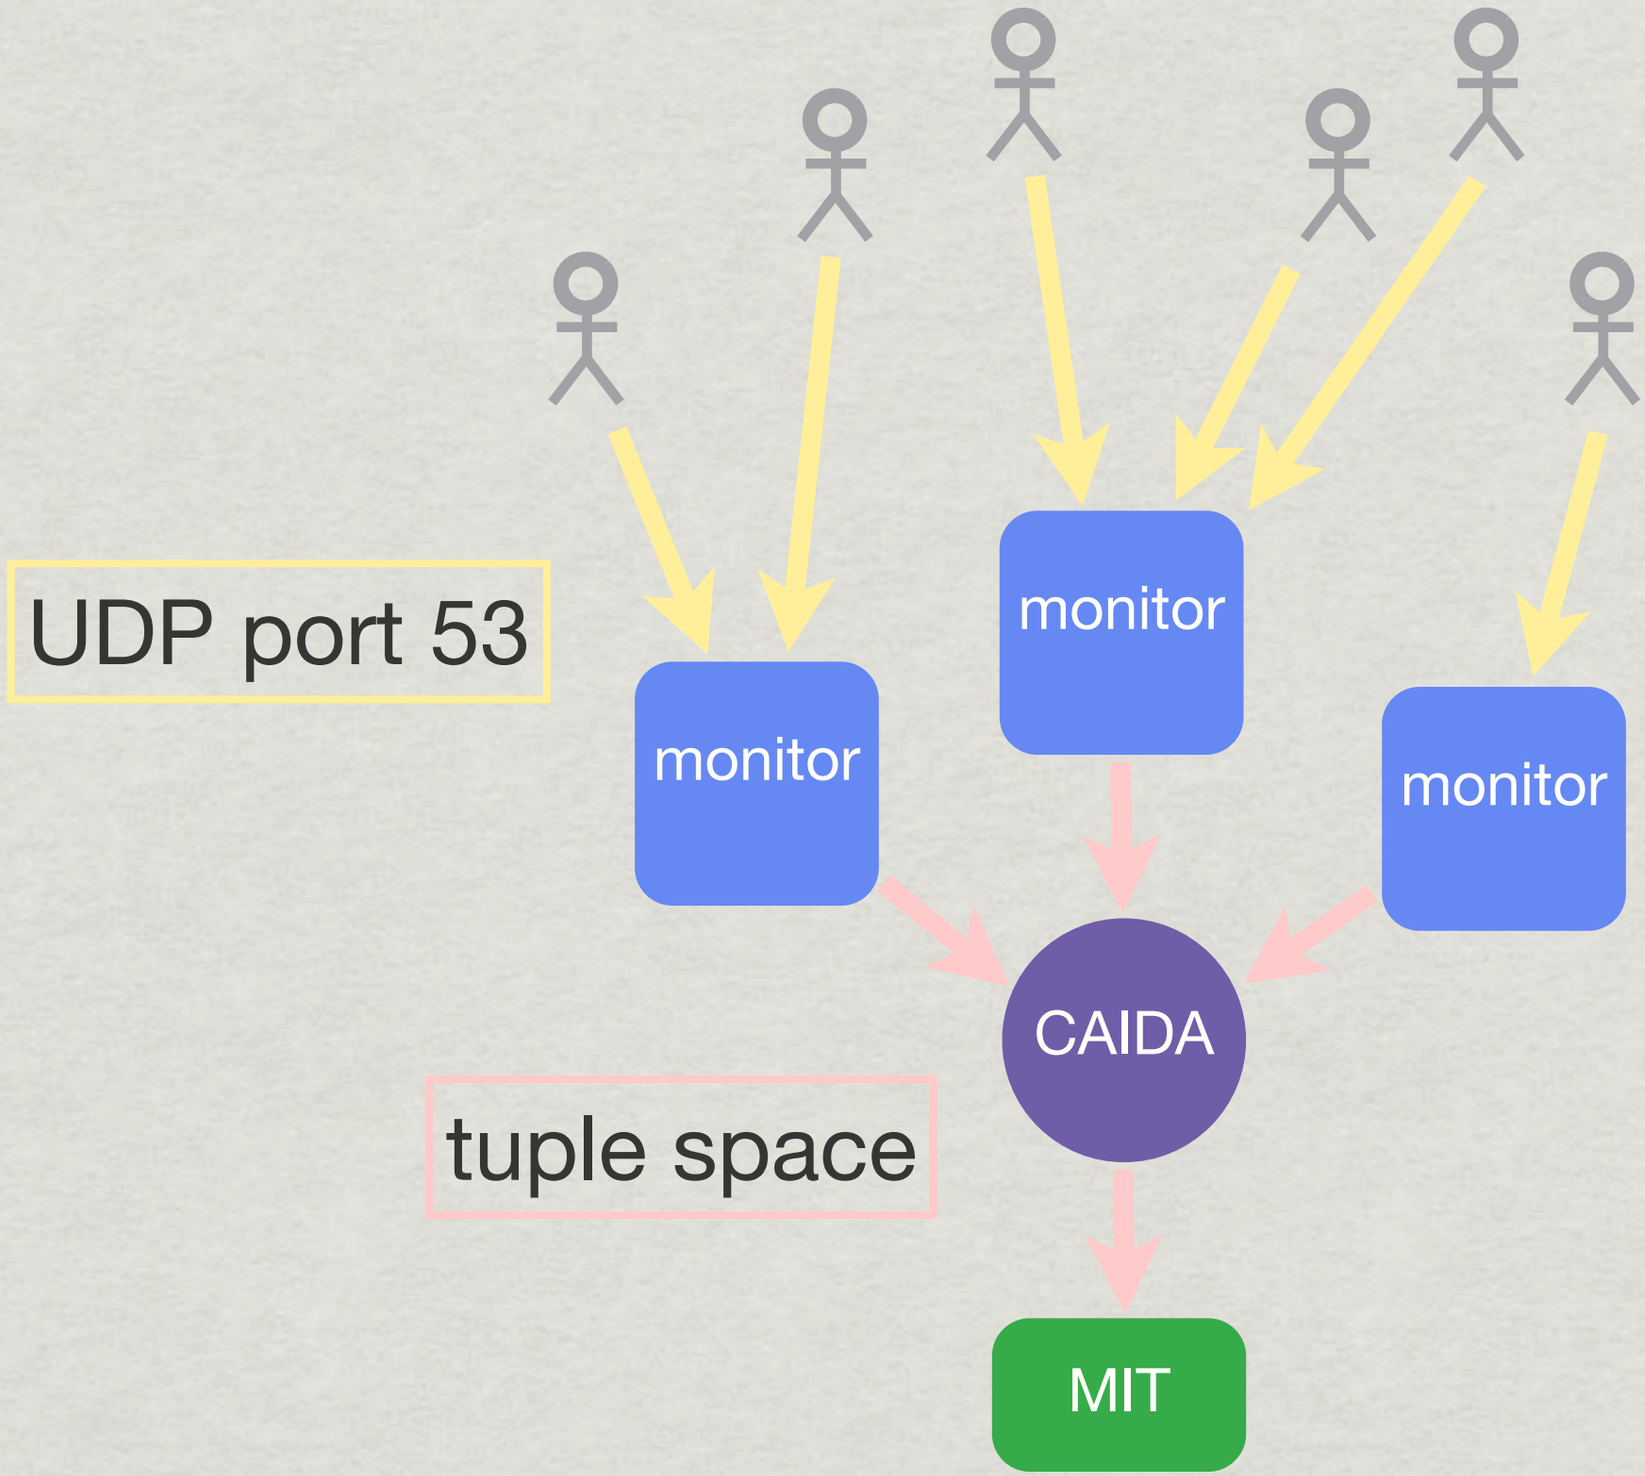

Ark Statistics Pages

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

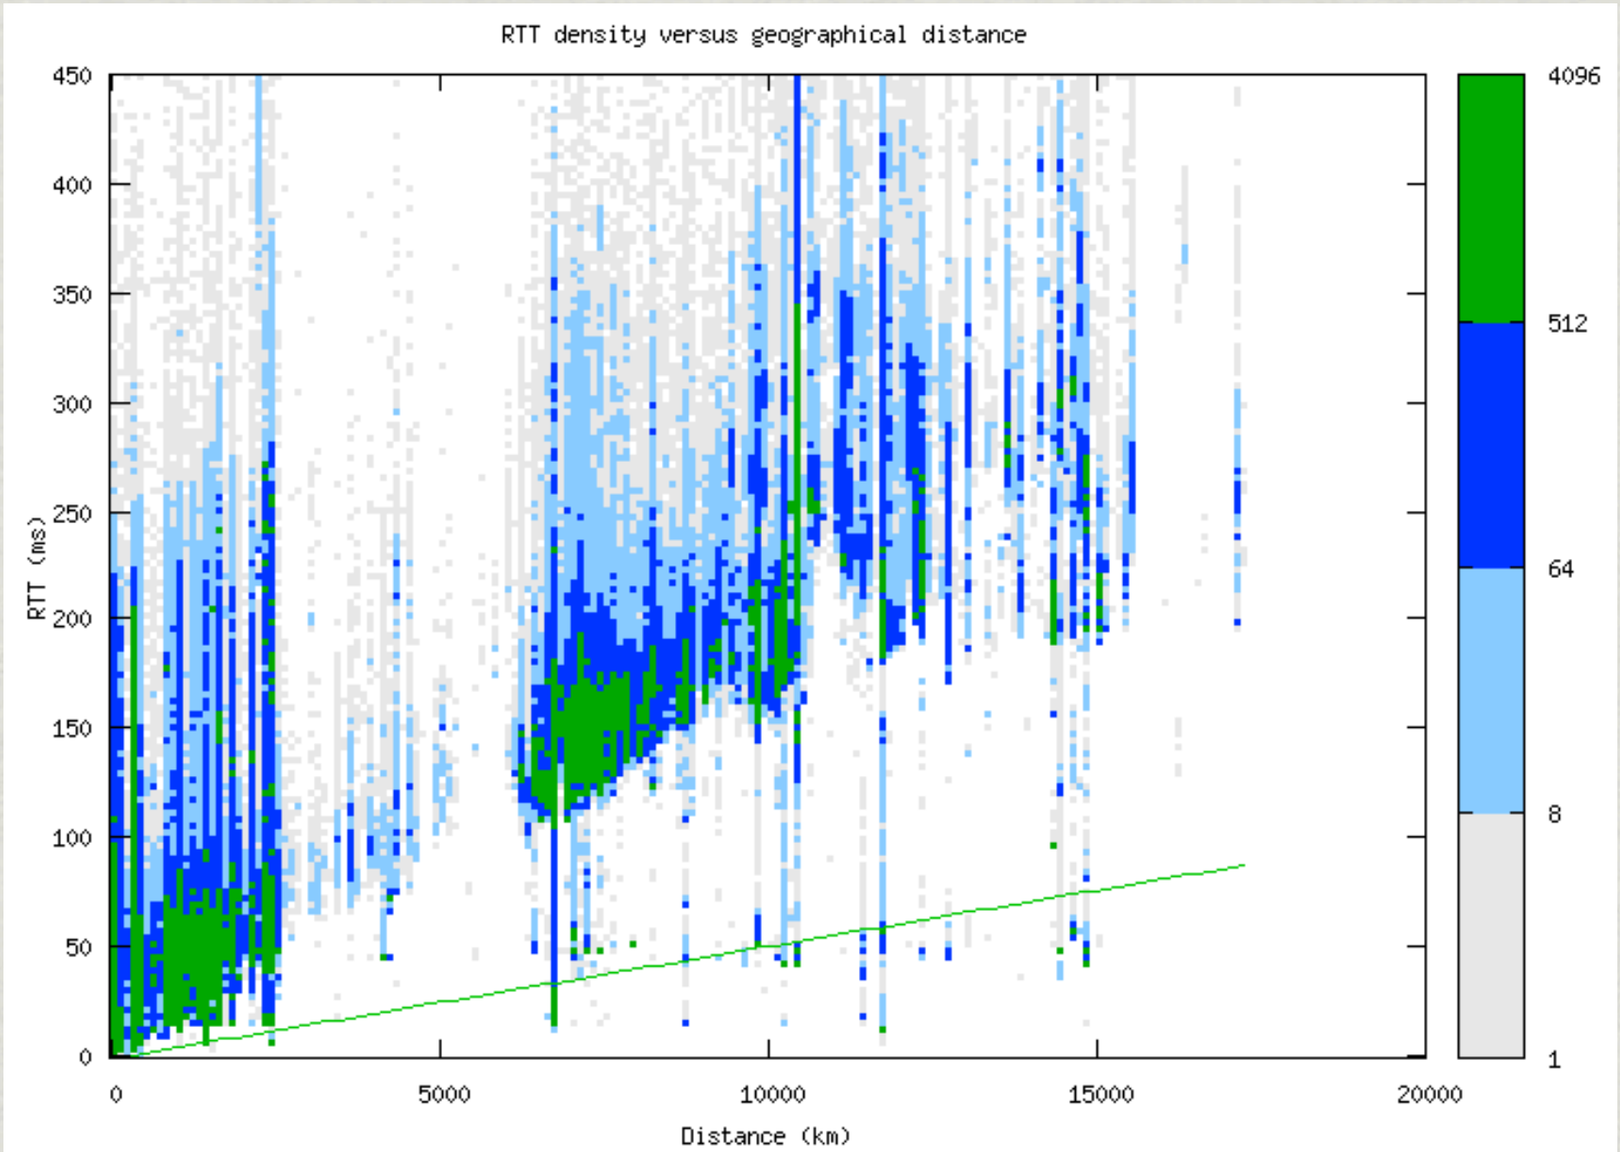

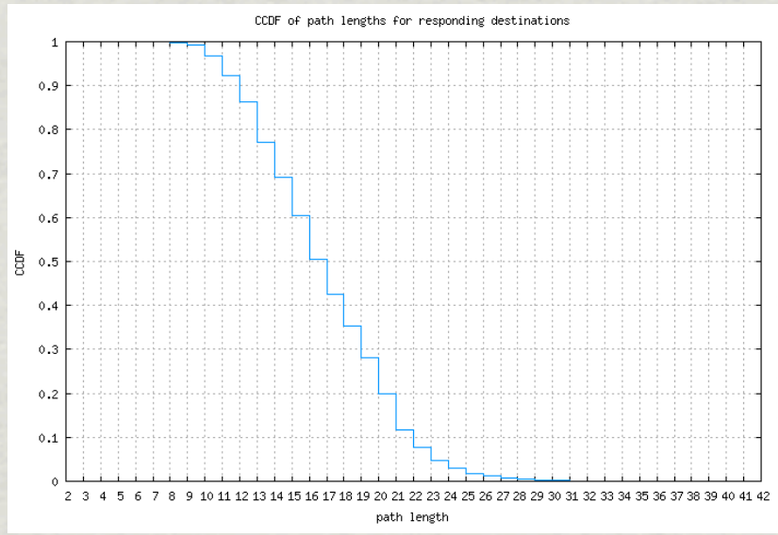

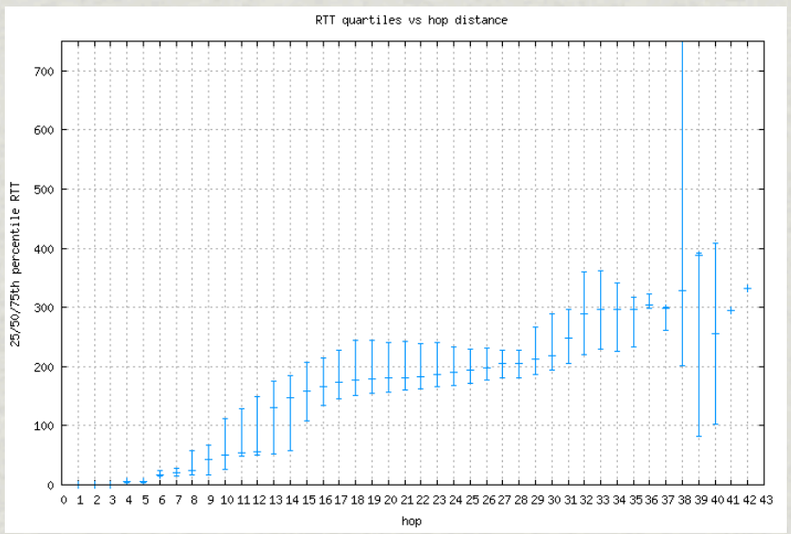

http://www.caida.org/projects/ark/statistics

Future Work

- Goals of Ark:
  - make it easy to develop and deploy measurements
    - ☑ easy to use communication and coordination facilities
      - Marinda tuple space
    - ☐ high-level packet generation, capture, and analysis API
      - inspiration from Scriptroute, Metasploit, Scapy, Racket
  - allow semi-trusted 3rd parties to conduct measurements
    - ☐ isolation between users and between measurements
    - ☐ enforce policies
      - bandwidth usage, destination selection, type of packets

Alias Resolution

Ken Keys

- Goal: collapse interfaces observed in traceroute paths into routers
- toward a router-level map of the Internet

The Alias Problem

- Traceroute reveals only one interface address on each router along a path.
- Given a set of IP paths, we can not tell which addresses belong to the same router.

Common Source Address: iffinder

- Send UDP or TCP packet to unused port at address A.
- If ICMP Port Unreachable response comes from address B, then A and B are aliases.

Graph Analysis: APAR

- Compare paths that cross the same subnets in opposite directions to infer aliases:

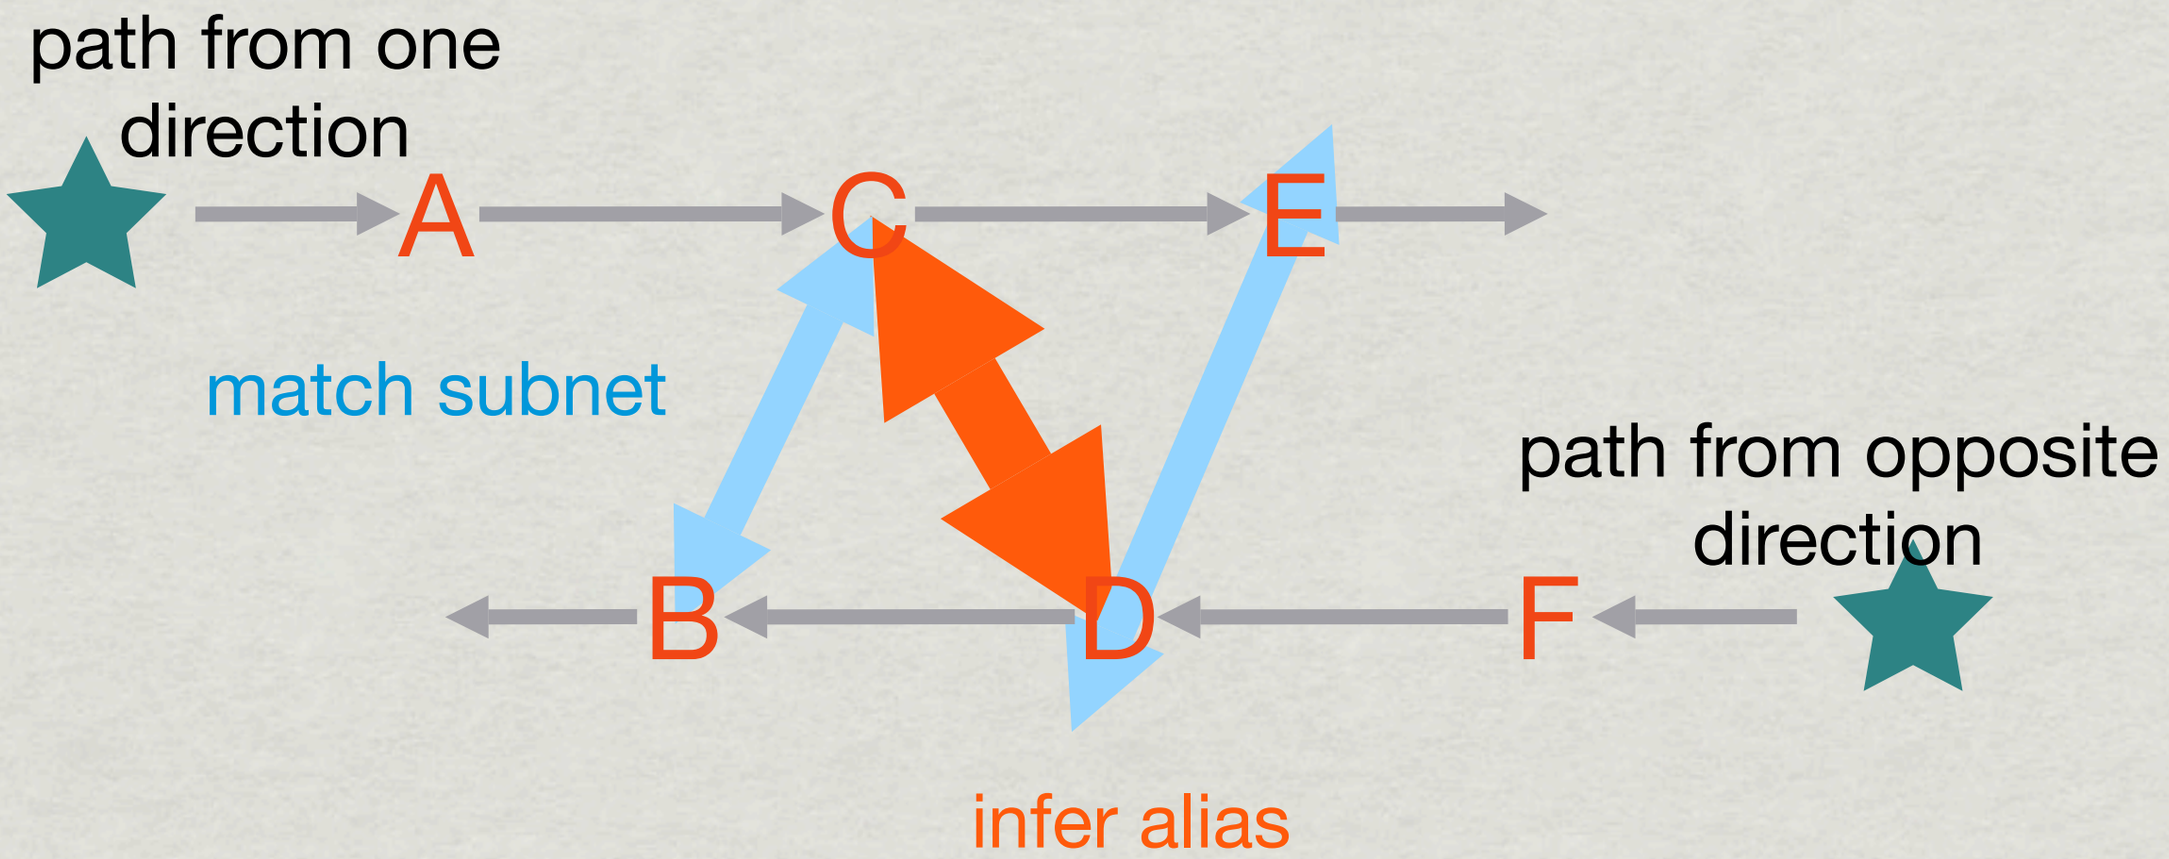

Graph Analysis: kapar

- CAIDA implementation of the APAR algorithm
  - Optimized
  - Additional heuristics
    - TTLs from *multiple* vantage points
    - Stricter subnet inference rules
    - Additional probes to broadcast addresses of potential subnets

Evaluation: data

- 373 M traceroutes from 26 Ark monitors
  - Found 2.4 M intermediate (router) addresses
  - Found 27 M total addresses
  - Ping each router address from all monitors, to collect TTLs
- Validated against known topology data from CANET, GÉANT, Internet2, NLR, and WIDE

Evaluation: results

	GEANT			Internet2			NLR		
	R	TP	FP	R	TP	FP	R	TP	FP
reality	18	540		9	713		7	231	
	0	0	0	0	0	0	6	100	0
kapar	14	75	6	15	193	26	9	61	7
kapar + TTL	11	80	6	12	163	6	8	67	6
iffinder + kapar + TTL	16	63	6	15	209	13	6	132	0
iffinder + kapar + TTL	11	84	6	14	167	4	7	127	0

R = routers with multiple interfaces
TP = true positive alias pairs
FP = false positive alias pairs

Evaluation: iffinder

- Ran on all 26 monitors to all router addresses
- Finds many aliases on networks where routers respond to direct probes, but finds no aliases on networks where routers do not respond
- Negligible false positive rate
- Using TTL constraints to check for false positives does more harm than good

Evaluation: APAR / kapar

- Works more evenly than iffinder across Internet
  - Finds 7 times as many alias pairs
- False positive rate is low, but significant
- Compared to APAR, kapar's stricter subnet rules and broadcast probes helped slightly
- TTL constraints reduce false positives (good), but also reduce true positives (bad); the net effect is a small benefit

Evaluation: iffinder + kapar

- Combines strengths of both methods
- In case of conflict, an iffinder alias is considered more reliable, because of iffinder's low false positive rate
- Even on parts of the Internet where iffinder does not find any aliases, results for iffinder+kapar are better than for kapar alone

Common IP ID counter: RadarGun

- Iterates over IP list multiple times, probing each address.
- Calculates “velocity”, or rate of change of IP ID counter over time, for each address.
- Any two addresses with similar velocity and predicted ID values are likely aliases.
- Improves upon Ally
  - Requires only *O(n)* probes
  - More tolerant of noise

RadarGun velocity example

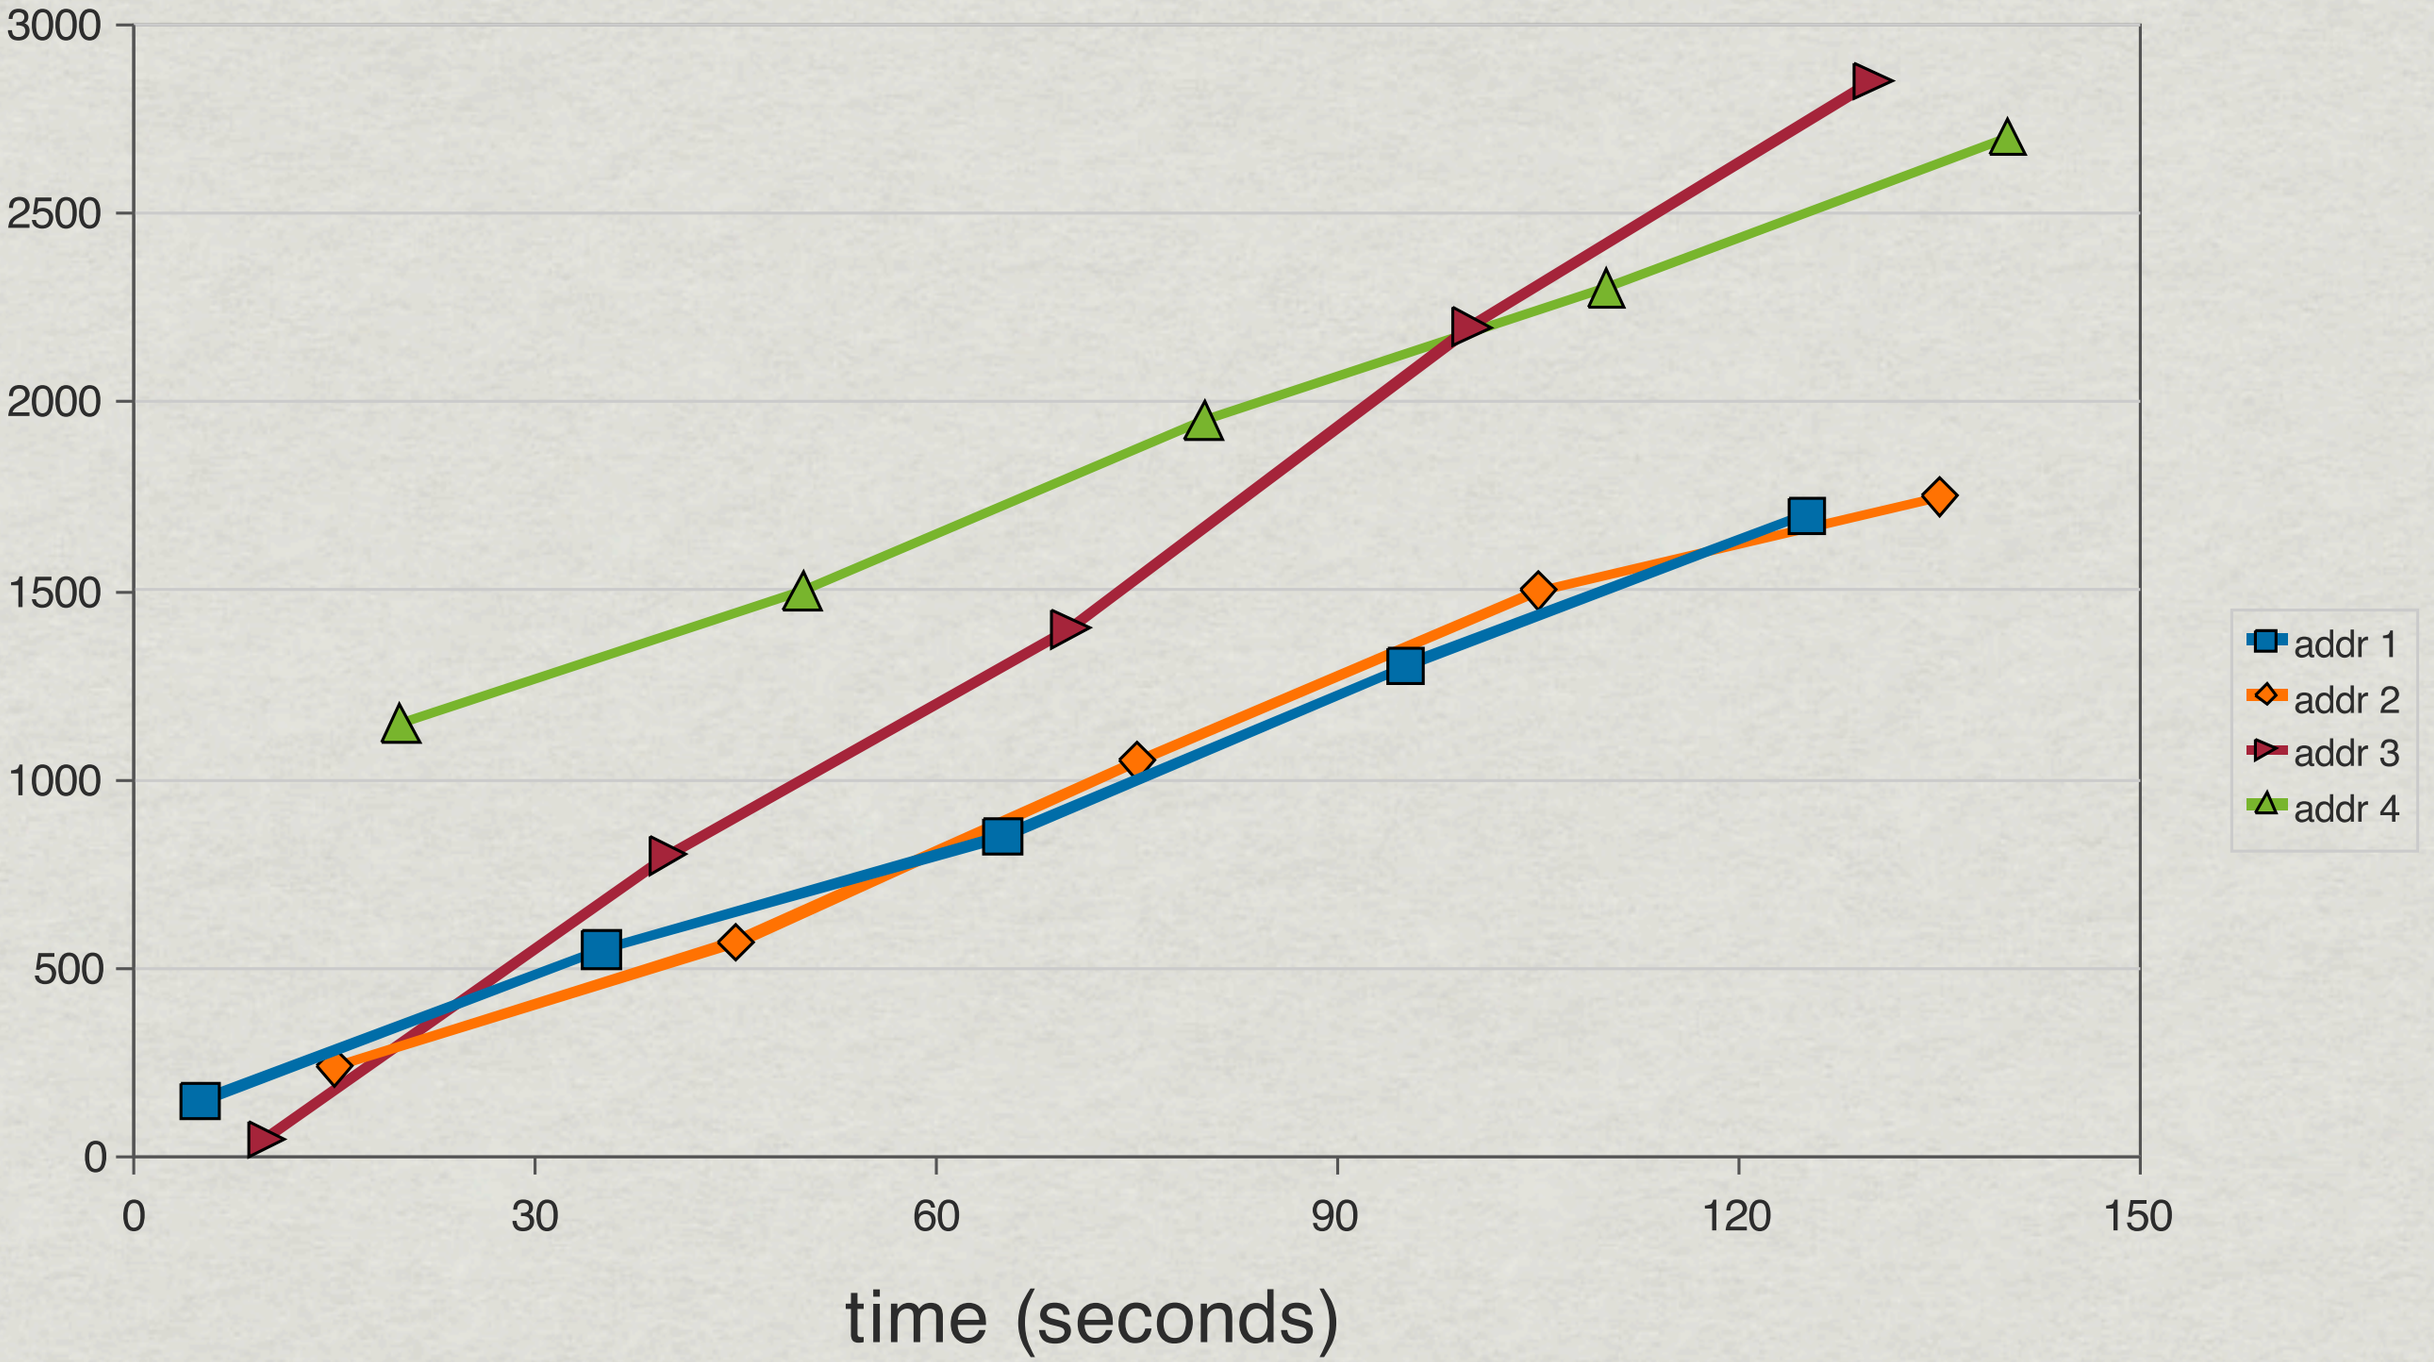

Interface vs Router graphs

* statistics for 1 month of 2009

	nodes	links	max degree	average degree	average neighbor degree	mean clustering
interface	23M	25M	47,658	2.19	1,170	0.001
router	23M	25M	47,661	2.68	1,264	0.077

Super Nodes

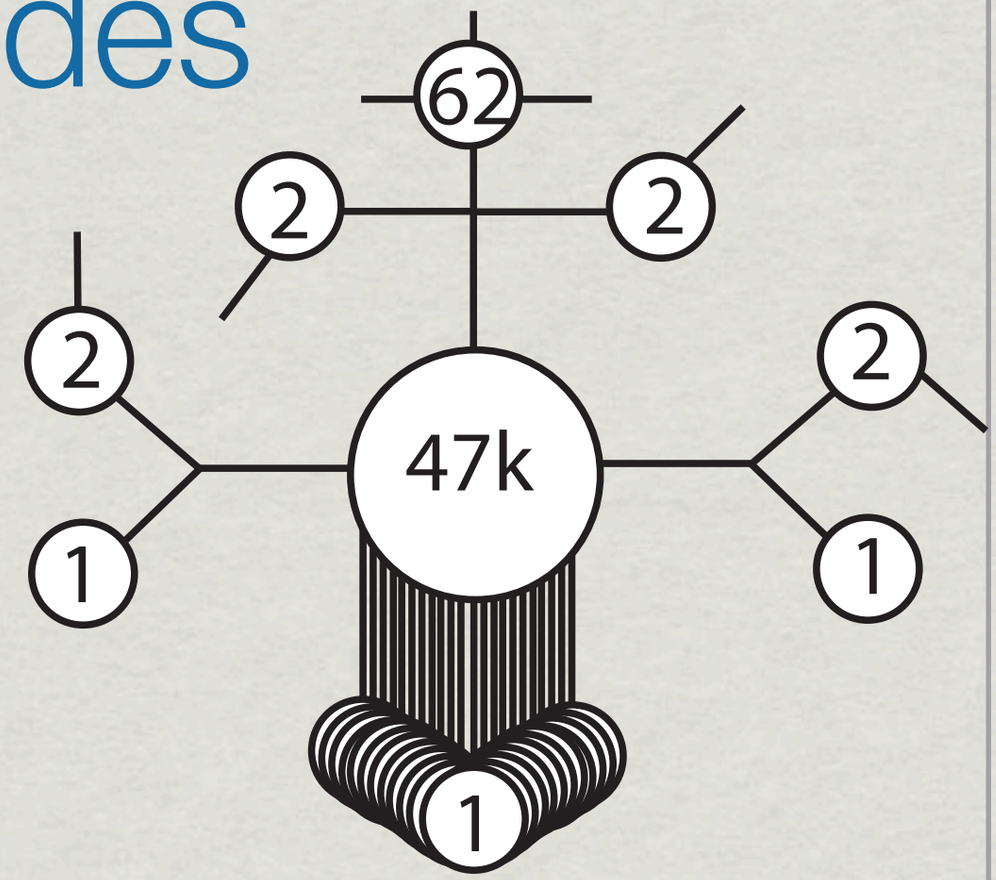

node id: N2899333
number of interfaces: 1
number of links: 47657
interfaces: 193.1.196.225

Max Degree	containing prefix	IP	AS	link id	node_id(**degree**):ip_on_link
62	193.1.196.225/32	193.1.196.225	1213	L21220817	N52248(**62**) N4720493(**2**) N4720494(**2**)
2	87.45.12.162/32		1213	L13644964	N2899335(**2**) N21887064(**1**):87.45.12.162
	87.46.171.65/32		1213	L13664444	N2899344(**2**) N21906544(**1**):87.46.171.65
(group) 1			1213		(number of) links:**47654** nodes:**95308**

Future work

- RadarGun
  - Still doesn’t scale to CAIDA’s IP graph
  - Using TTL-limited probes instead of direct probes should significantly improve response rate
  - Combine with iffinder and kapar
- TTLs
  - With multiple TTL probes, we hope to identify and discard inconsistent TTLs that hurt kapar’s results

Thank you for listening

- Archipelago
  - http://www.caida.org/projects/ark
  - http://www.caida.org/projects/ark/statistics
- Kapar Technical Report
  - http://www.caida.org/publications/papers/2008/alias_resolution_techreport/