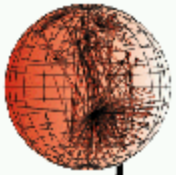

preparing for the future in the midst of a typhoon:

measurement and analysis of the Internet

scientific apparatus offers a window to knowledge,
but as they grow more elaborate,
scientists spend ever more time washing the windows.
-- Isaac Asimov

kc claffy, UCSD/SDSC/CAIDA
kc@caida.org
www.caida.org

CAIDA's evolution

- created late 1997 at ucsd/sdsc
  - seed grant from NSF
  - sponsorship from Cisco
- goals
  - foster engineering-level cooperation among industry, research, government
  - tools/analysis for robust, scalable Internet
  - topics beyond purview of single networks/vendor

CAIDA's evolution

- 1998
  - hardware (OC3/12 monitors with MCI)
  - viz, analysis of commercial infrastructure
    - (Nature of the Beast – iMCI)
  - cflowd; active measurement (skitter)
  - education (IEC), outreach (ISMAs), training
- 1999
  - measurement, analysis and viz tools
  - analysis results (active measurement & multicast)
    - skitter, mantra, cflowd, coral
  - education, outreach & training (EOT)
    - ITL

outline of today's talk

- strengthening academic Internet engineering education
- priorities for Internet measurement and analysis
- synergies in the Cisco / CAIDA relationship

[EOT]: Internet Eng. Curriculum

initiated Jan'98 w NSF & Cisco support

- lead: Prof. Evi Nemeth
- http://iec.caida.org

- problem: need to improve quality of graduating Internet engineers

- approach:
  - broad university–industry collaboration
  - dynamic repository of state–of–the–art Internet engineering curriculum materials for university coursework & continuing education

[EOT]: Internet Eng. Curriculum

activities:

- advisory group
  - (aiken, bellovin, bradner, bush, connolly, crowcroft, kurose, wilder)
- workshops for faculty
  - august '99, faculty (Cisco-supported – routing, tcp, ns, traffic analysis)
  - june '00, faculty
- traffic analysis training CD
  - lectures, exercises, data sets
  - beta released January 00
  - web version: http://traffic.caida.org
  - Cisco-supported

[EOT]: Internet teaching labs (ITL)

- problem: univ.'s lack hands–on facilities for teaching Internet engineering

approach

- jan 99: Cisco (Estrin) commits ~100 trade–in routers
- aug 99: RFP announced
  - round 1: twelve universities selected (38 applied)
  - round 2: pending
- dec 99: NSF commits funds for glue (training, collaboration, support)
- today: blocked on 7000s' removal from C&W POPs

[EOT]: other workshops

goal: bring together key players in communities

- ISMAs: passive meas't, viz, analysis
- SDNAP: BGP/MBGP for ISPs (SDNAP)
- ISP tools: cflowd, RRDtool

future workshops:

- correlation among datasets
  - routing, topology, performance, workload
- active measurement (SLAs)

priorities in measurement & analysis

- topology (mapping)
- workload characterization (passive)
- performance evaluation (active)
- routing (dynamics)

will show examples, priorities, obstacles

topology: skitter

- macroscopic, infrastructure–wide
- dynamically discover/depict topology (& b/w)
- correlate path perf. w events, e.g. BGP
- identify critical pieces of infrastructure

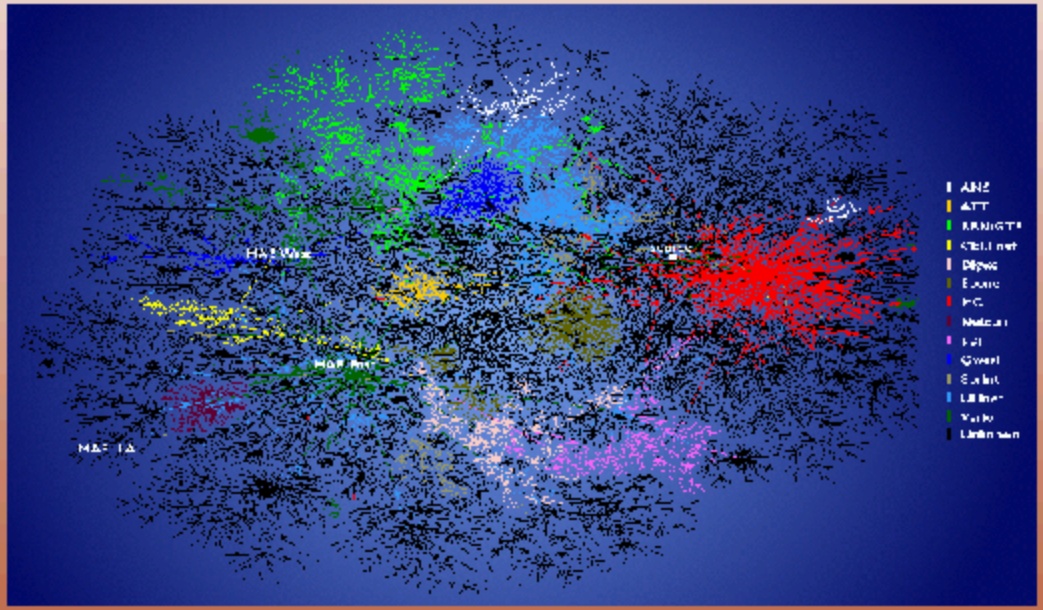

skitter: infrastructure–wide measurements

- 17 monitors (inc. root name servers)
- multiple dst lists (29k servers, 36k dns)
- architecture:
  - parallel ICMP probes
  - 52–byte packets
  - kernel time stamping
  - ssh / Kerberos

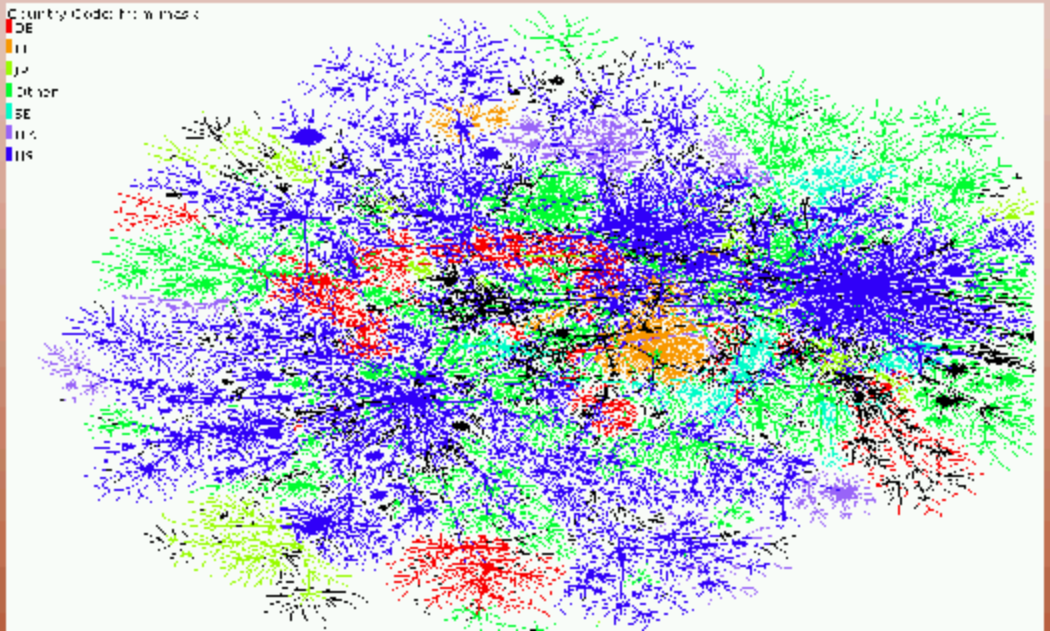

skitter: AS interconnectivity

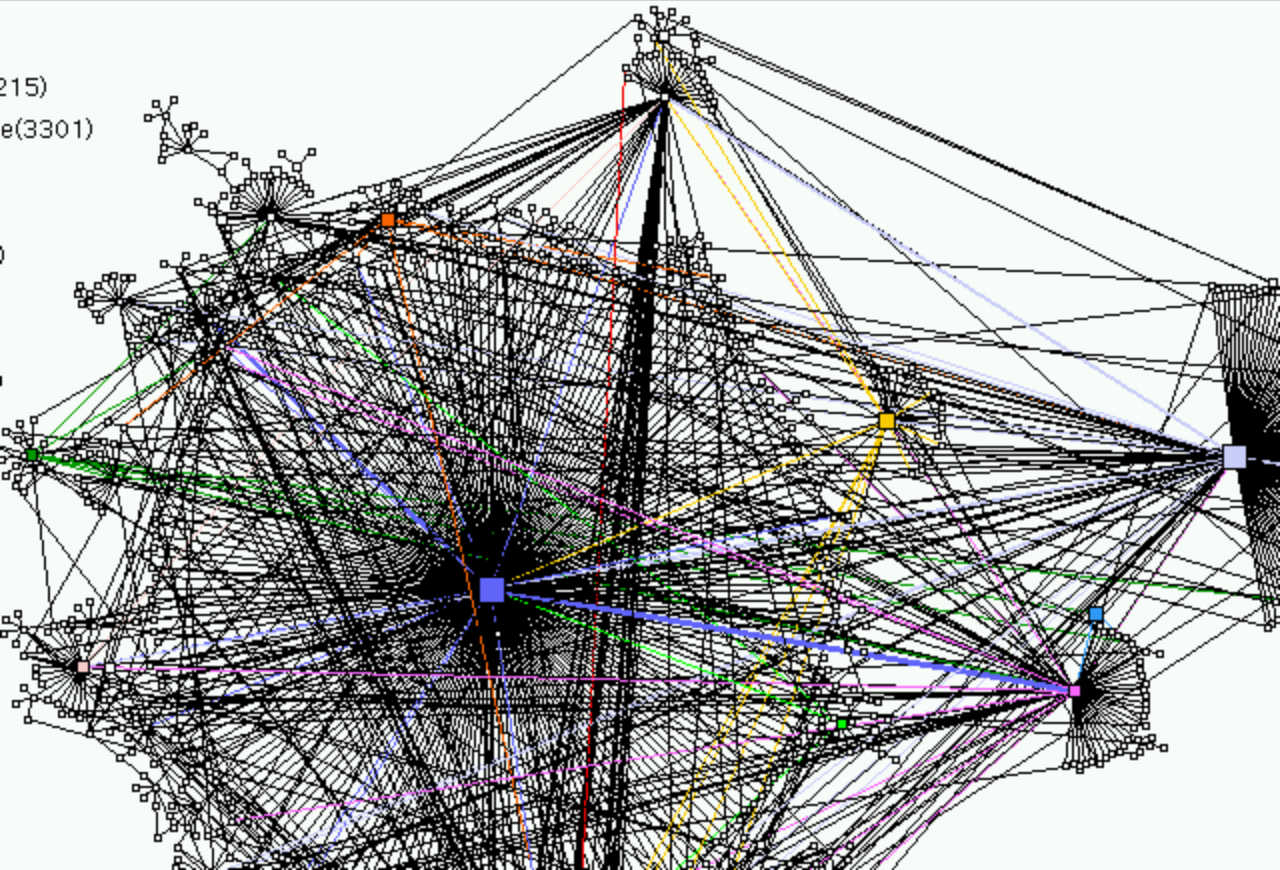

GTrace: geographic traceroute

www.caida.org/Tools/GTrace/

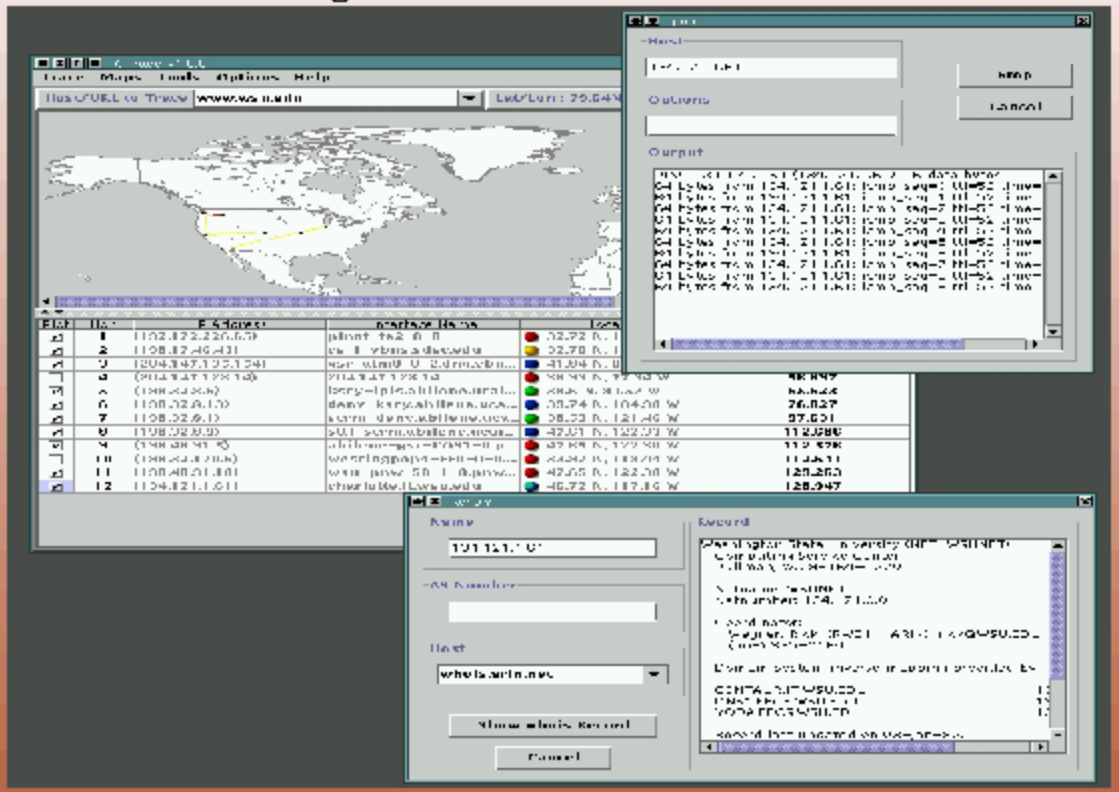

skitter: 3D hyperbolic

BGP routing table data (connectivity among ASes)

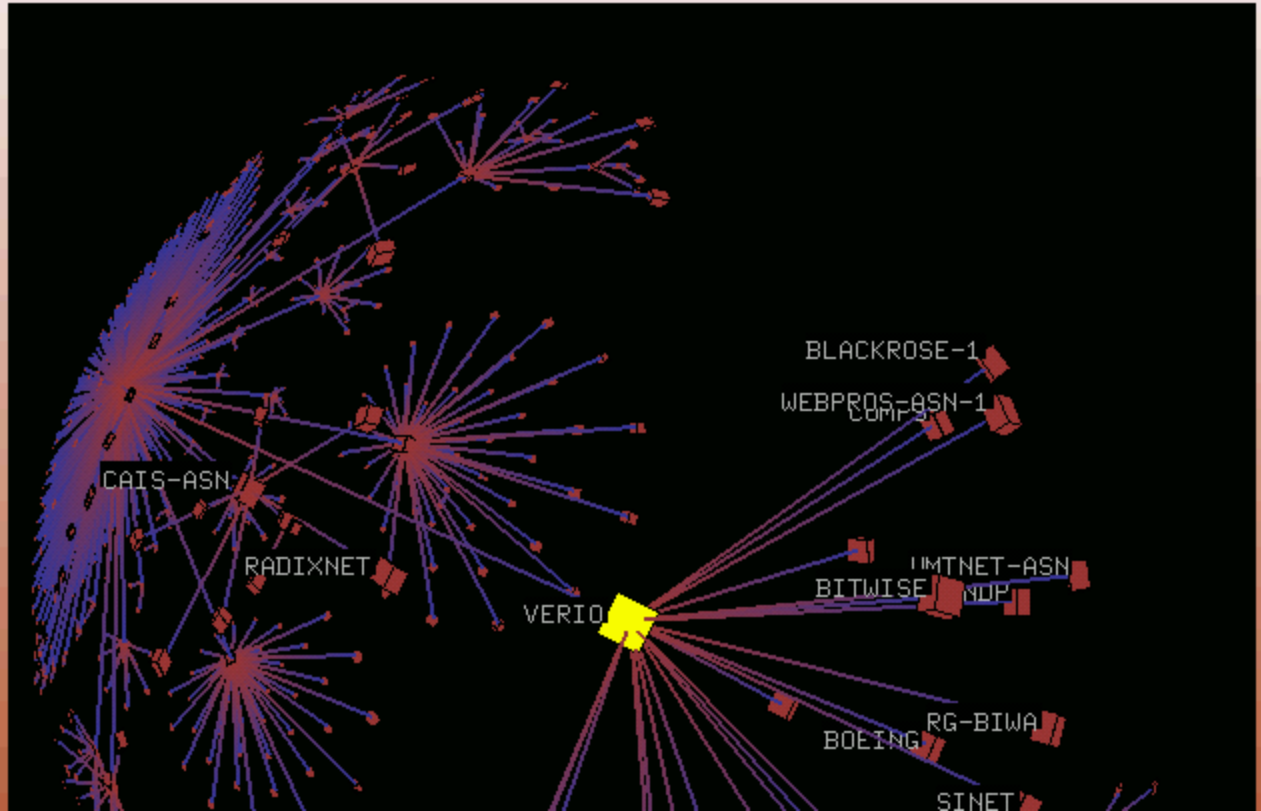

mantra: analyzing multicast

www.caida.org/Tools/Mantra

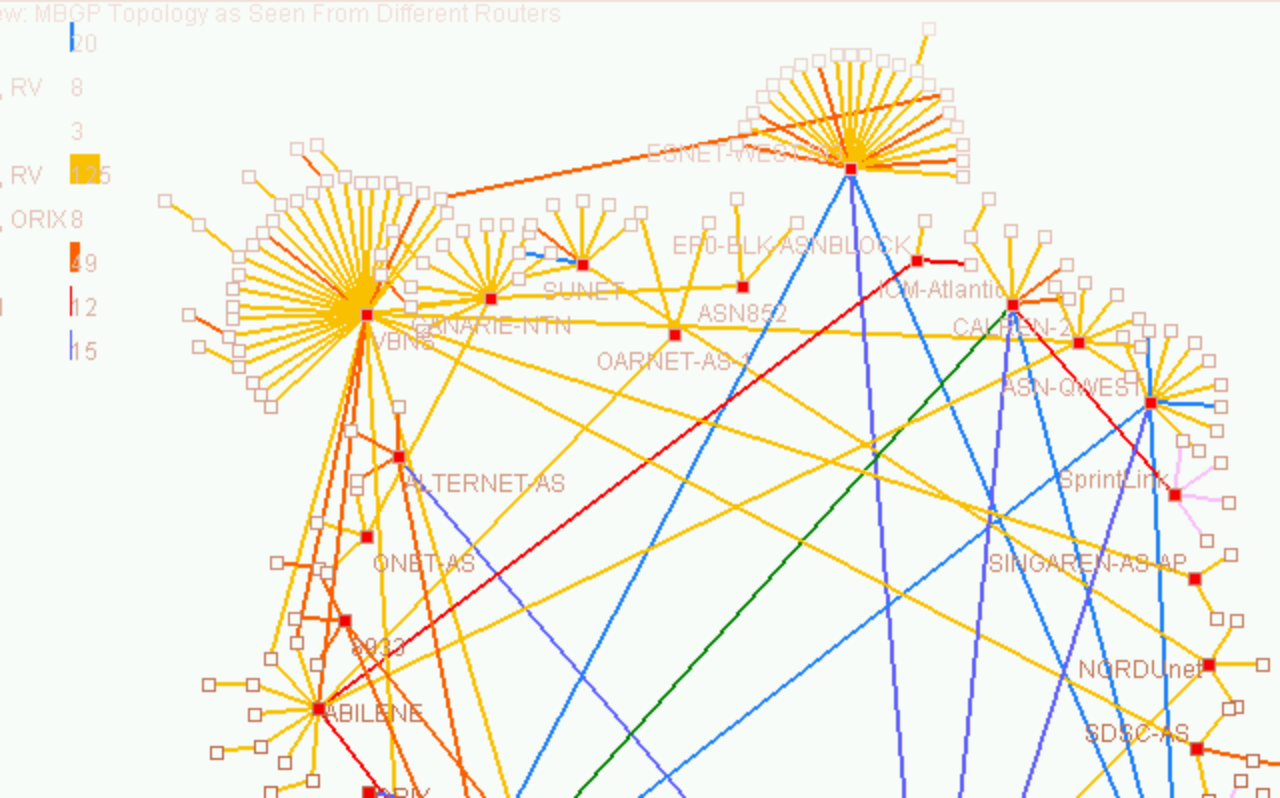

logical vs geographic topology

2–dimensional, hierarchical

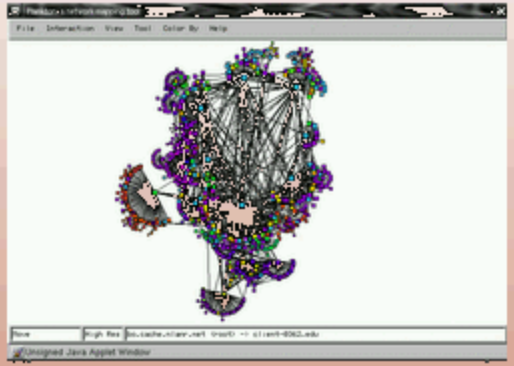

geographic

semi–geographical (otter)

www.caida.org/Tools/Otter

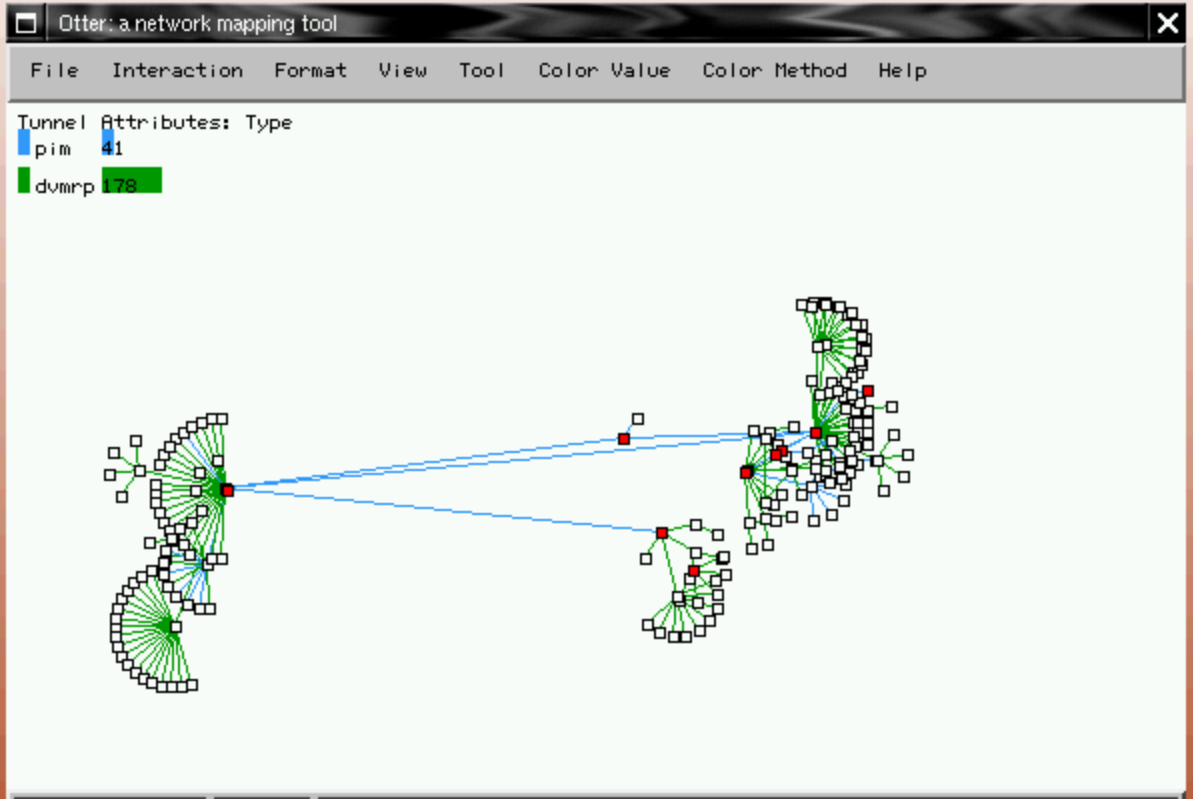

topology: research priorities

- accurate measurement & visualization
  - latency
  - key routers/networks
  - AS granularity
  - geographic
  - integration w mgt tools

- obstacles:

  mapping IP addresses to
  - router
  - geography
  - AS
  - service provider
  - country
  - anything...

  route changes faster than can measure

workload characterization

insights for

- usage profiling
  - h/w, protocol, application design
  - architecture optimizing
- capacity and peering planning
- network control and management
- security
- performance analysis
  - delay, loss, jitter?
- QOS assurance across ISPs
- accounting and billing

- tools: netramet, netflow, cflowd, coral
  - some suck less? ...evolution requires use

workload char: working w/vendors

cflowd

- www.caida.org/Tools/Cflowd
- primarily for capacity planning and trend analysis
- Cisco's netflow export

- AS-to-jS matrices
- net-to-net matrices
- port and protocol tables
- forward IP path

==> line rate

measurement specifications to vendors

workload char.: protocol

19 aug 99, ucsd–cerfnet
https://anala.caida.org/CoralReef/Demos/

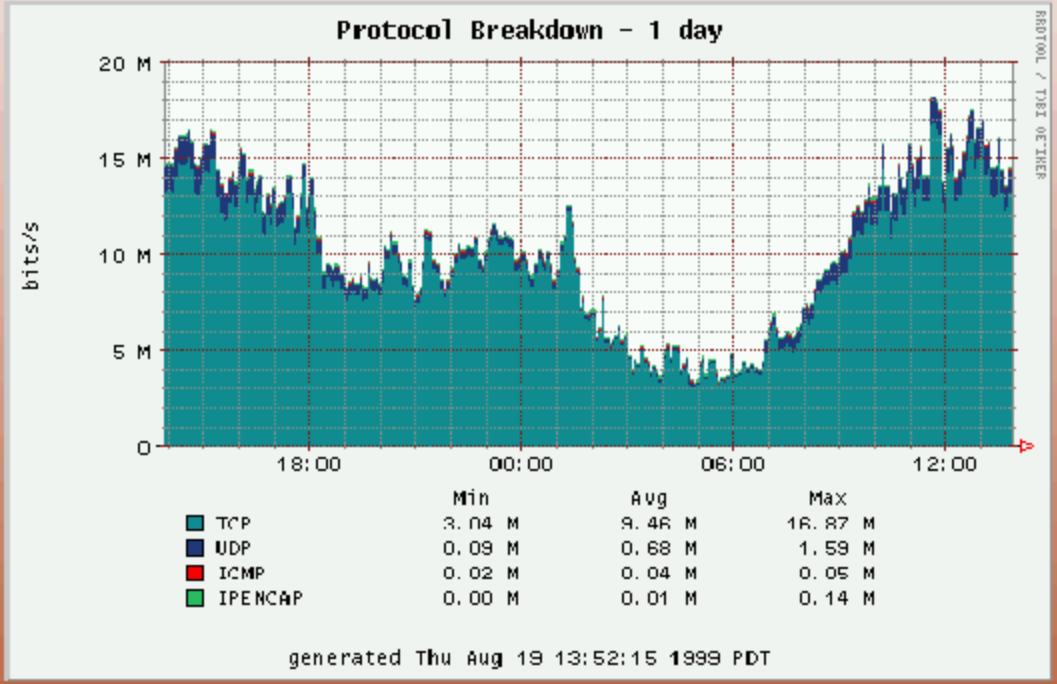

workload char.: protocol (proportion)

19 aug 99, ucsd–cerfnet

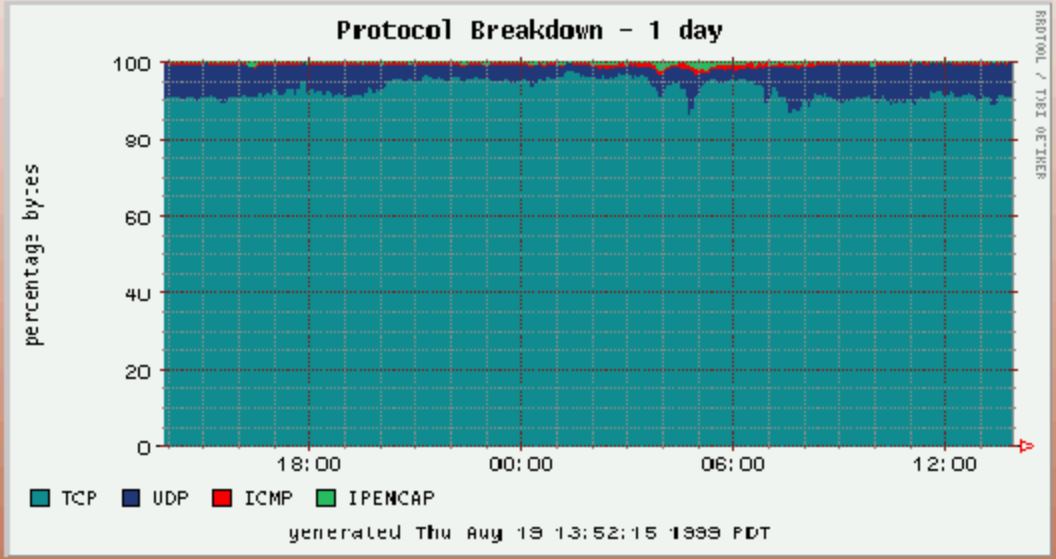

workload char: applications (ucsd–cerfnet)

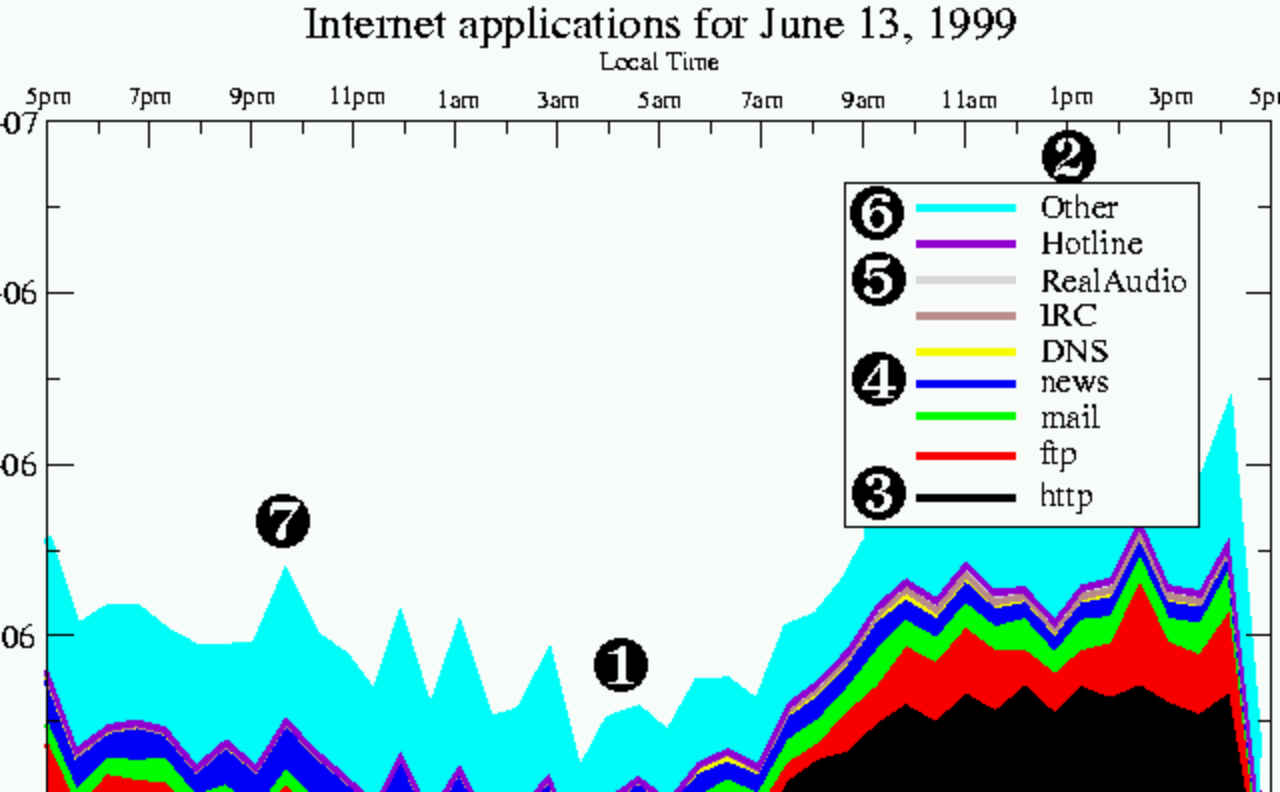

workload char.: mantra

18 oct 99, fix-west.mbone.nasa.gov
http://www.caida.org/Tools/Mantra/

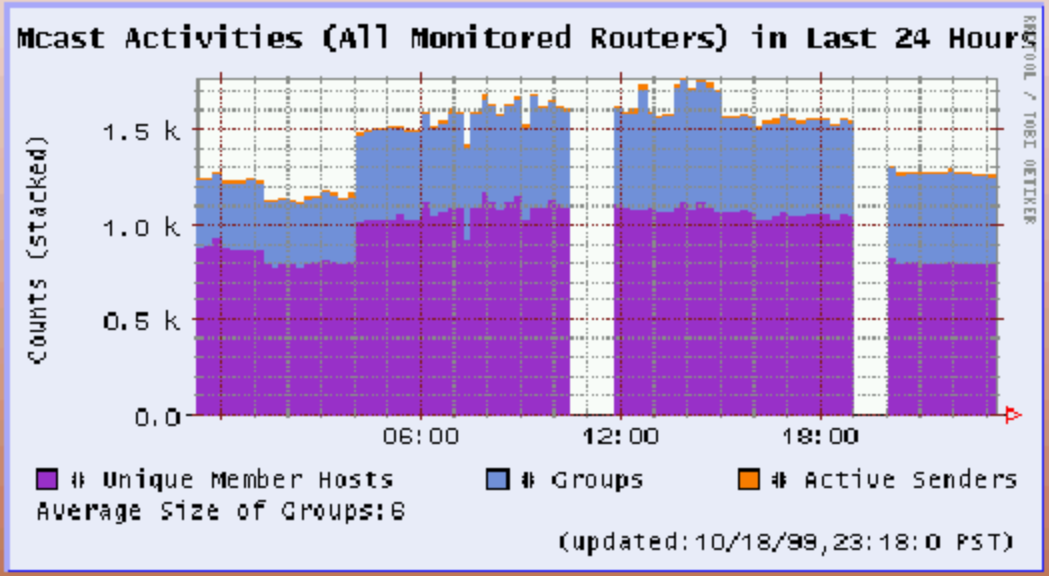

packet sizes, 4/98.mci vs 7/99 AIX

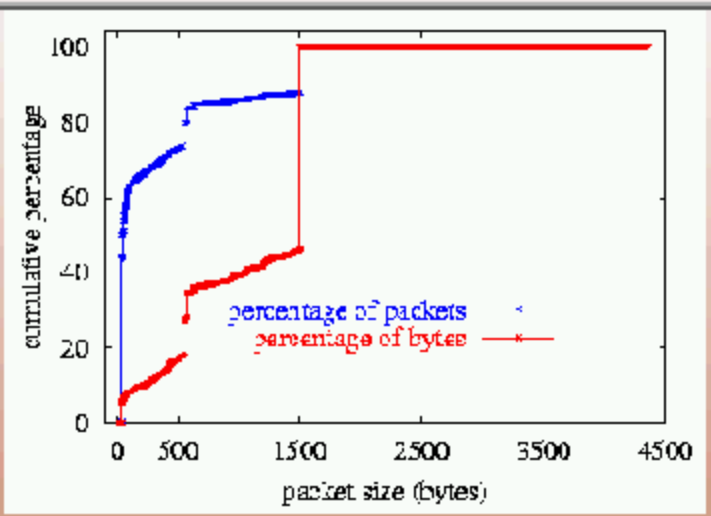

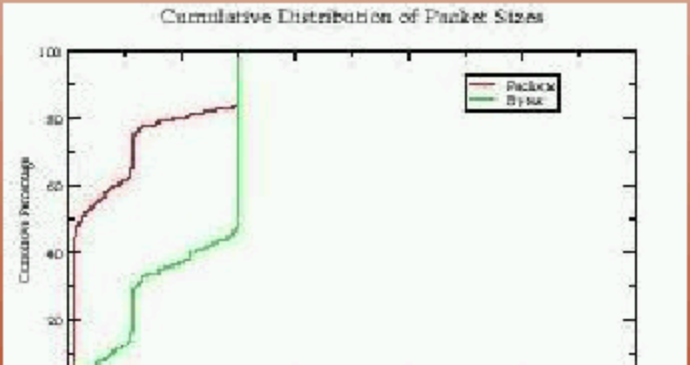

packet sizes by prefix length

larger packets from shorter prefixes (why?)

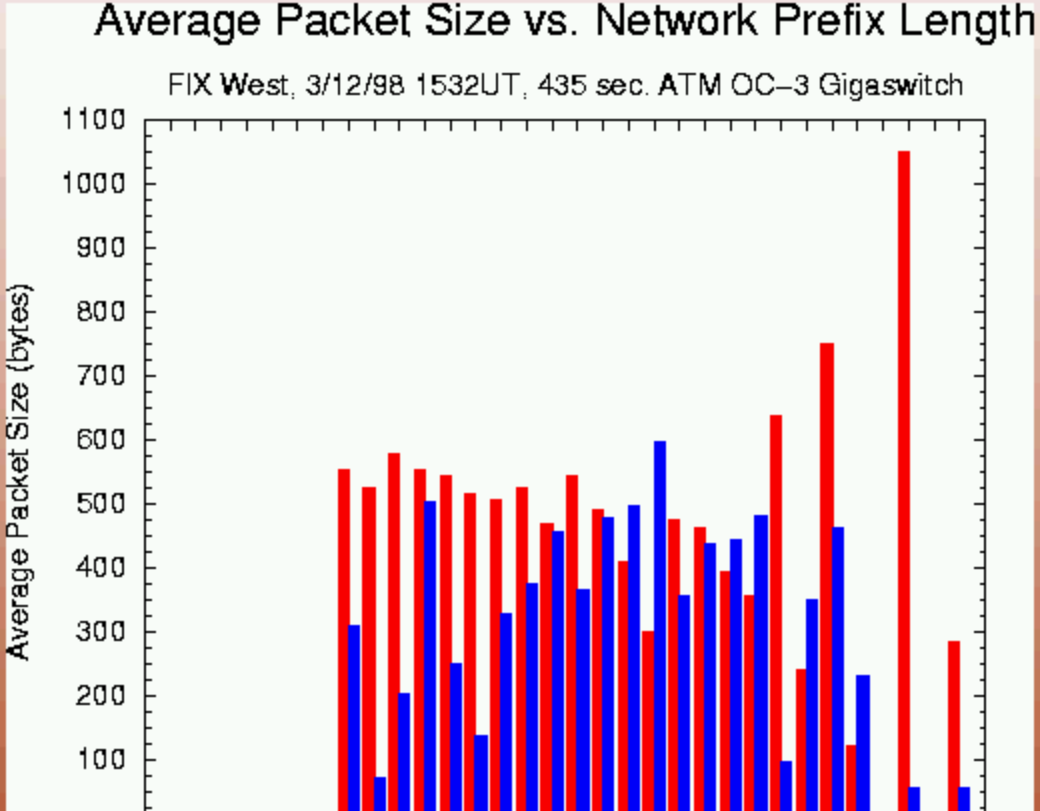

favoritism/locality by AS

80% of traffic from < 5% of ASes
60% of reachability from < 7% of ASes

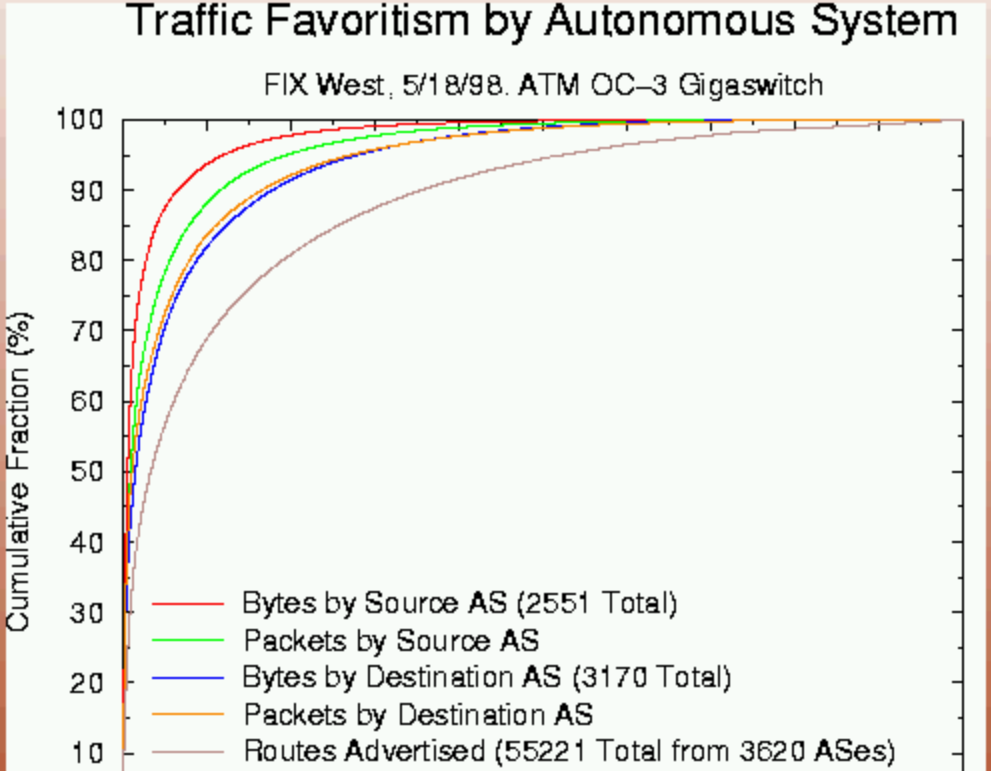

flow length distribution, 4/13/98 mci backbone

heavy tail (quite truncated)
if you only learn about one distribution this quarter...
flow defn creates artifacts (100 bytes)

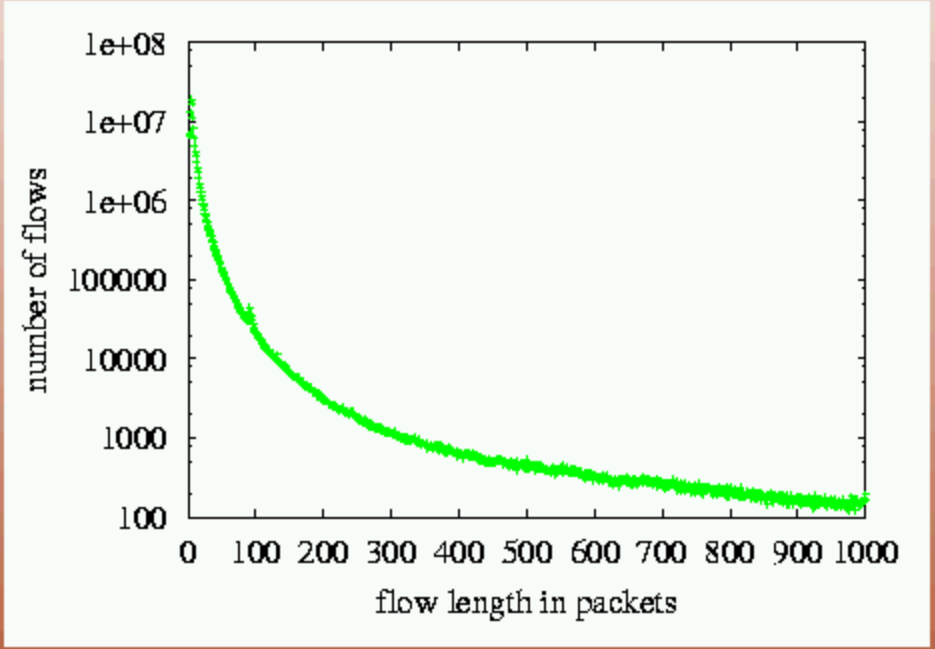

workload char of high perf. networks

- coral/ocXmon testing (oc3,12,48,gE)
- persistent real-time full frame collection
- integration w coralreef analysis s/w
- dag4.0 testing planned 1/00

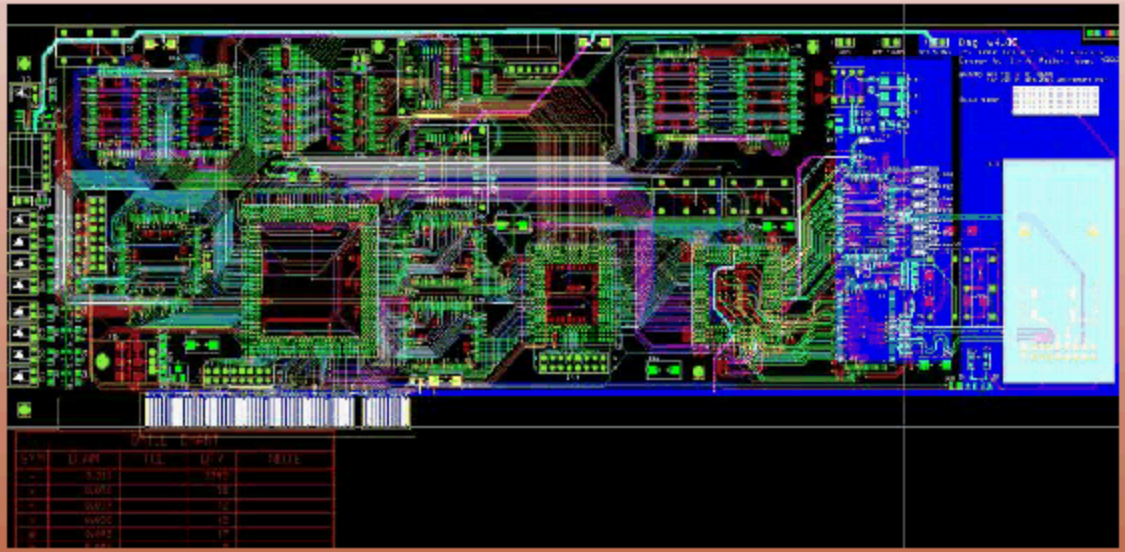

security of high perf. networks: coralreef

- CoralReef 3.2 release 12/16/99
- ATM: Applied Telecom and Fore OC3 (OC12 App. Tel. only)
- ATM/POS:OC3/12 (DAG3.2 testing 1/00)

- crl_portmap s/w module
  - listens for RPC portmap access
  - adds suspicious probing host to list
  - records all send-rec'd packets
  - full payload capture (default)
  - tcpdump output format

workload characterization: coralreef

- crl_filter
  - ATM reassembly
  - tcpdump output format
  - pipes to security tools
- future
  - OC12 (full line rates), OC48 need card support
    [obviously way (& getting further) behind switching...]
  - active enforcement module (yr 2000)
  - getting ISPs to use them

www.caida.org/Tools/CoralReef

workload characterization: priorities

- id and present 'useful' workload metrics, particularly given persistence of fire–fighting environment
- id significant patterns, timeframes, correlations
  - vary by user need
  - change as technologies and 'net change
- obstacles:
  - limited access to commercial networks
  - network performance impact
  - faster speeds and changing transport technologies complicate data acquisition and processing

workload characterization: priorities

- coral/ocXmons (OC3,12,48, gigE)
- persistent, realtime, full-frame collection
- security policy
  - compliance auditing (passive)
  - enforcement (active)
  - dynamic packet filtering triggered by attack precursors
- SLA support
- obstacles
  - hardware expensive
  - privacy issues
  - IPsec

performance evaluation (active)

- network engineers to diagnose problems
- ISPs & users to verify SLAs
- designers of real–time apps to predict software HCI
- Internet weather reports

perf.eval: skping (RTT, loss, analysis)

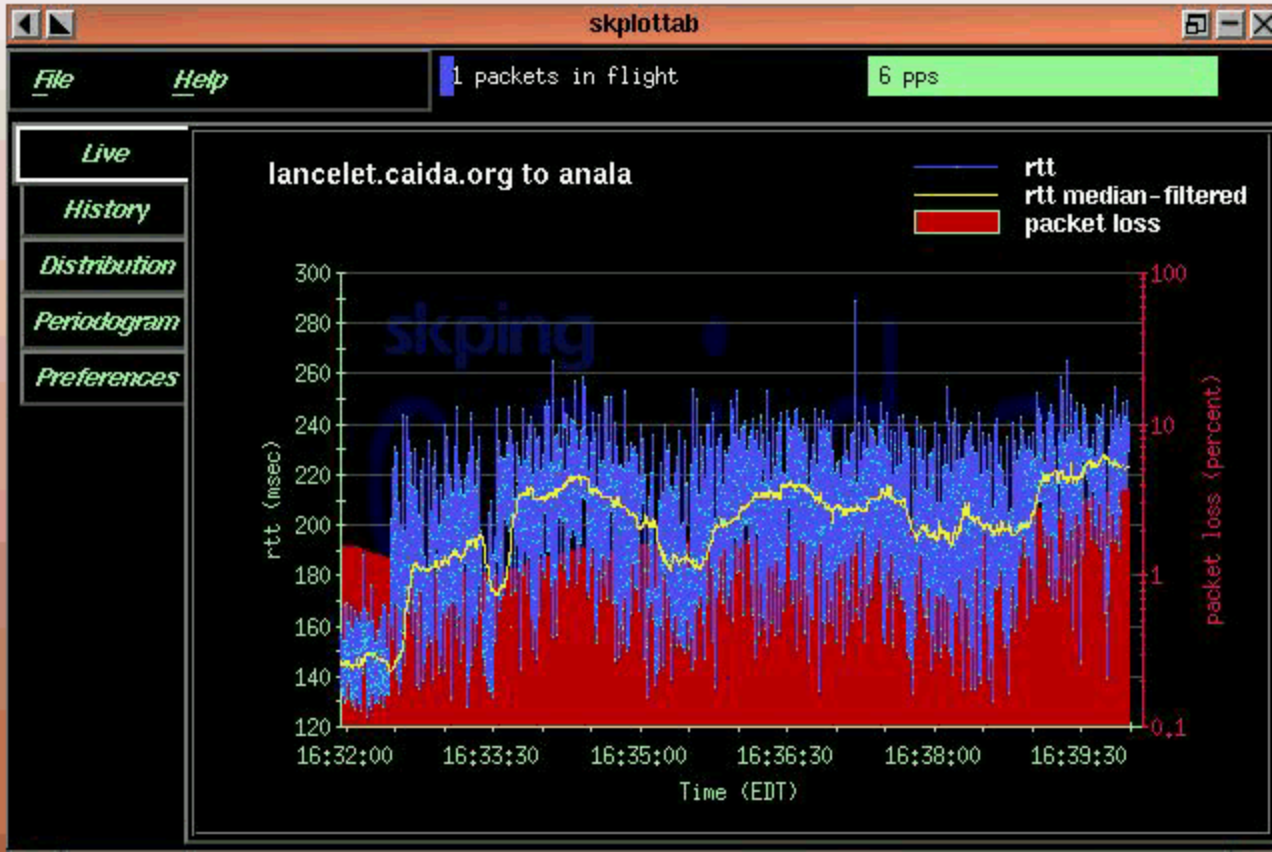

perf.eval: routing (path change)

perf.eval: sktrace (www.cnet.com)

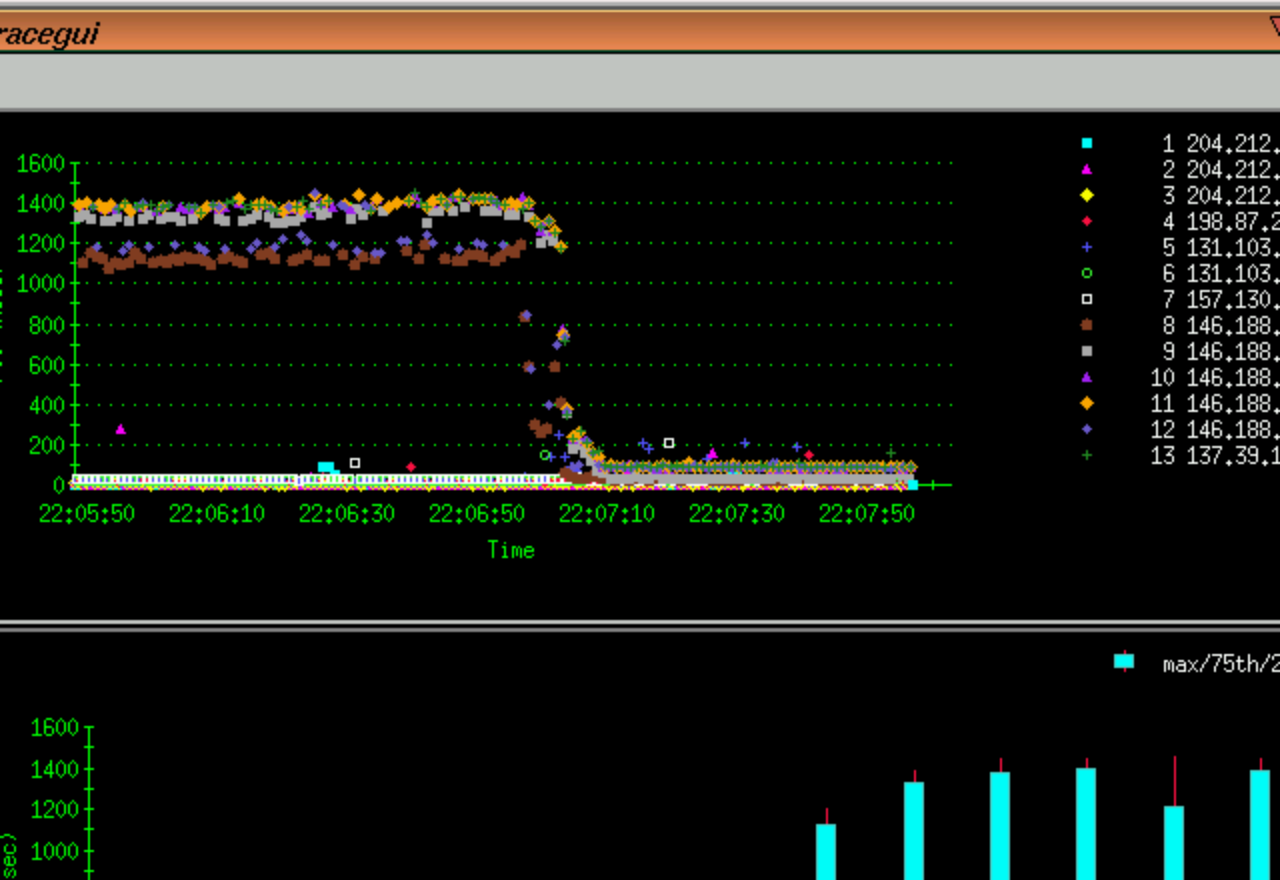

skitter: rtt distribution: tri-modal

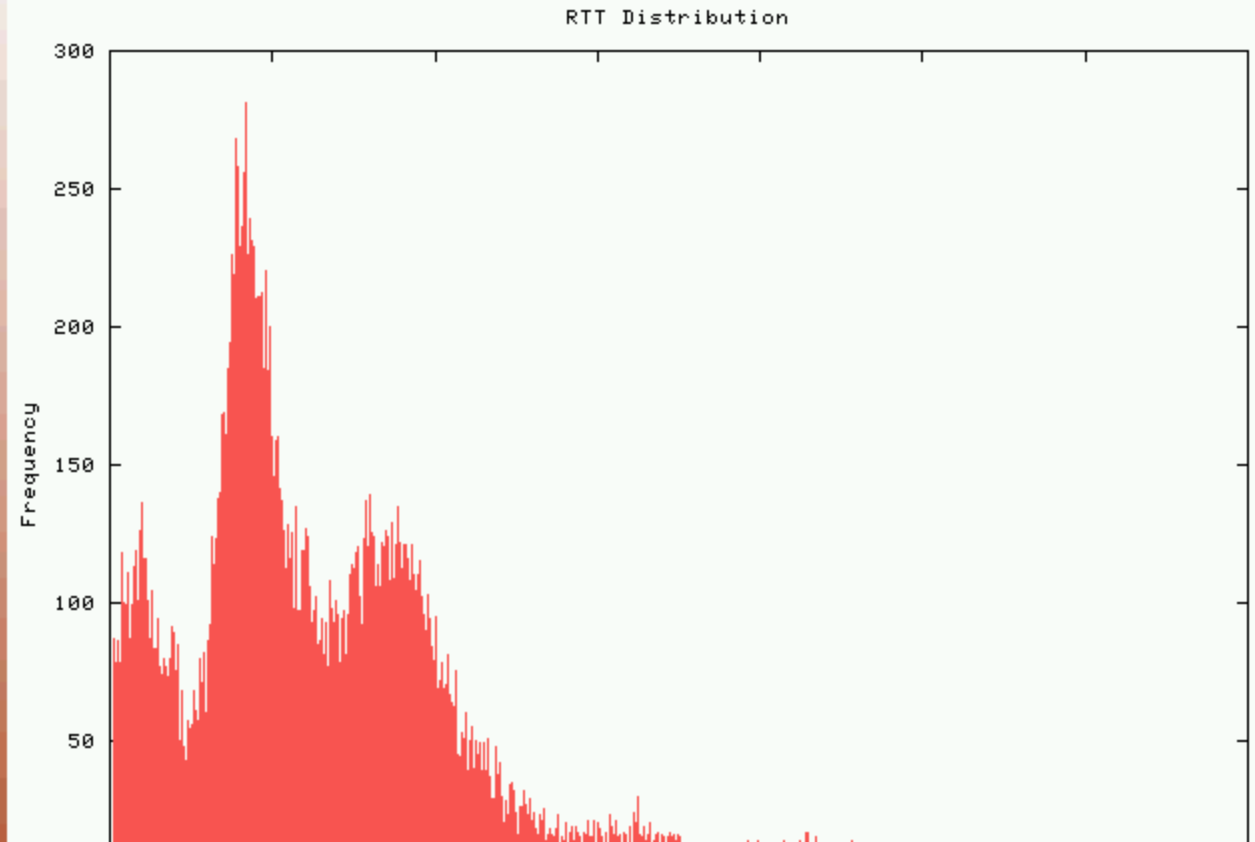

skitter: rtt vs longitude (light cone)

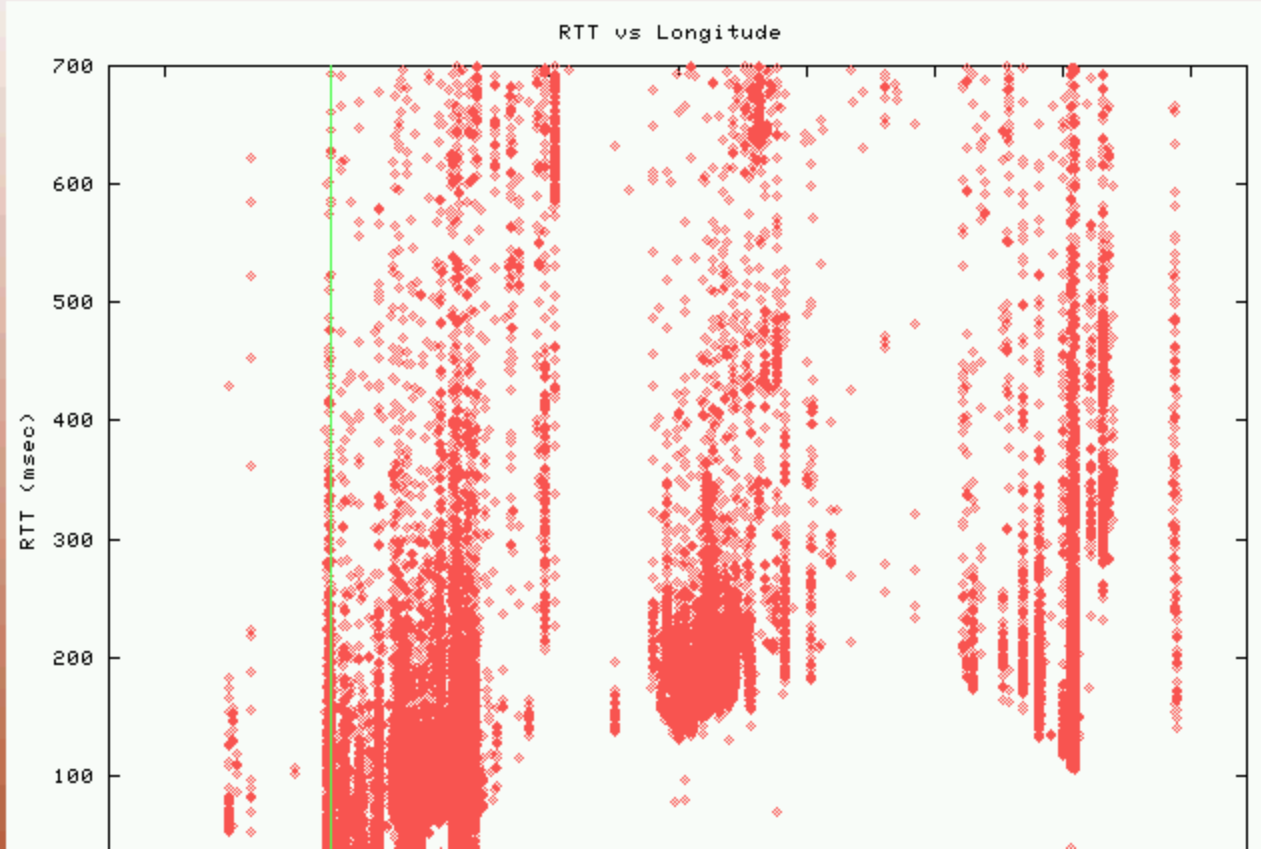

dispersion among ASes across paths

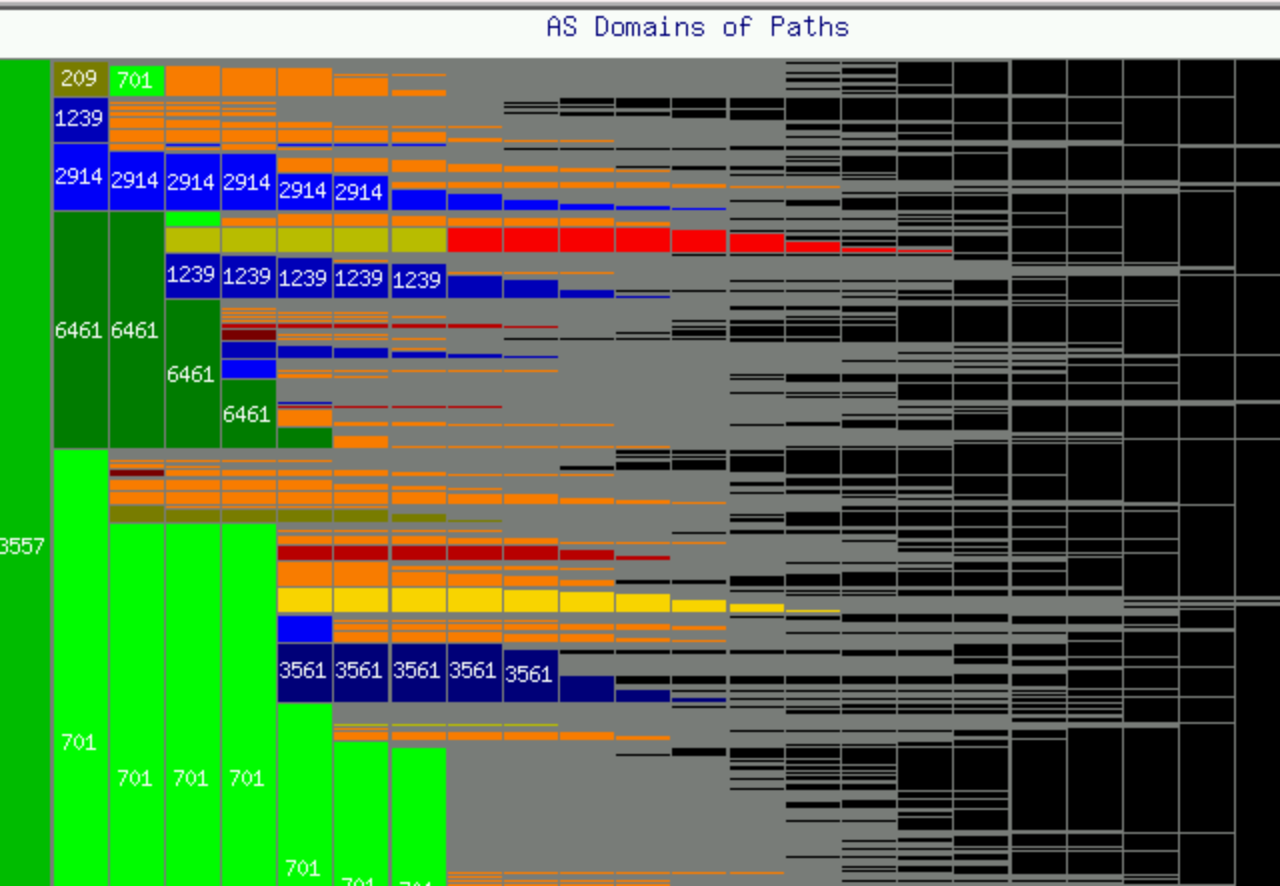

skitter: correlation (perf., topology, routing)

preliminary findings

- ~1% IP destinations disappearing monthly (re-addressing, firewalls)
- route announced path not matching forward path
- indication of potential routing configuration errors
  (by no means automatic)
- persistence of paths
- methods to identify critical infrastructure
- is there an Internet "core"?

datasets available to researchers

performance eval.: priorities

- faster collection, processing, rendering
- bandwidth assessment techniques
- need intuitive graphic presentations correlating:
  - performance across sources
  - comparisons w/topology, workload, routing analyses
- obstacles
  - poorly defined user requirements/interfaces
  - negative perceptions regarding quality and worth driven by explosive growth

routing dynamics

[nothing this room doesn't know...]

- 15-year-old technology
- well, it works...
- not much instrumentation/diagnostics
  - really need real-time
  - without interfering w performance
  - (or other engineering priorities)
- makes analysis hard

routing: example (instability)

- RTT data changes color if path changes
- 10 unique paths over 24 hour period
- lots of jitter in data
  - unlikely to be intentional
- heavy tails predominate

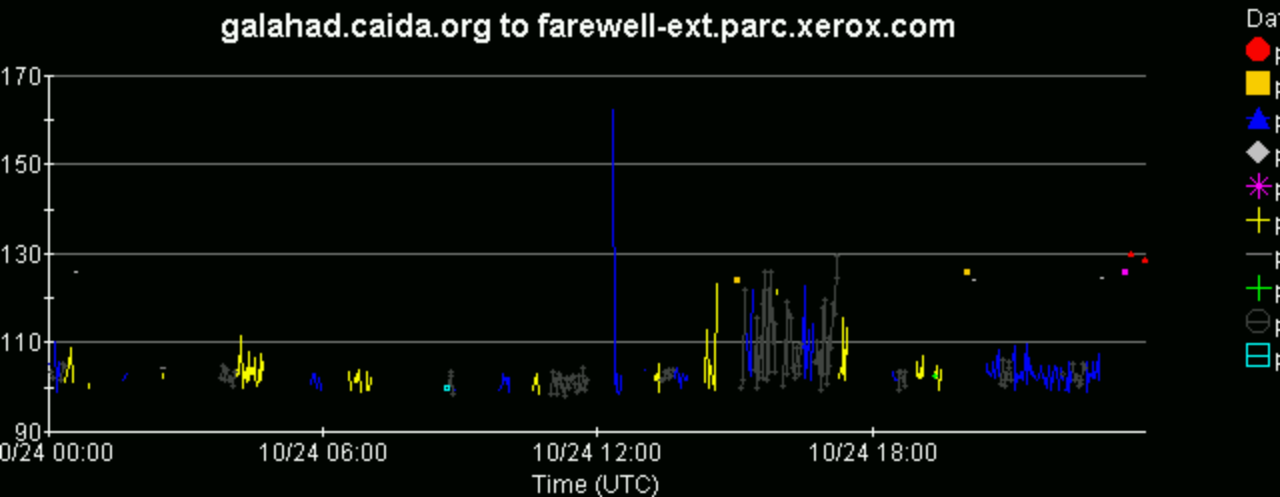

routing: example (load balancing)

- RTT similar over predominantly two paths
- likely intentional load balancing

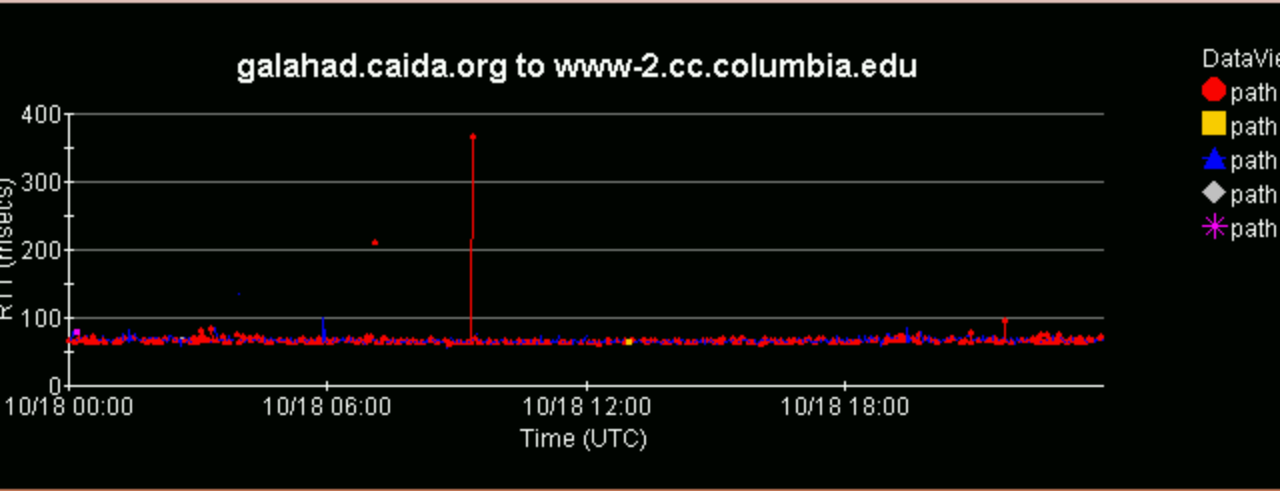

routing: RouteViews table analysis

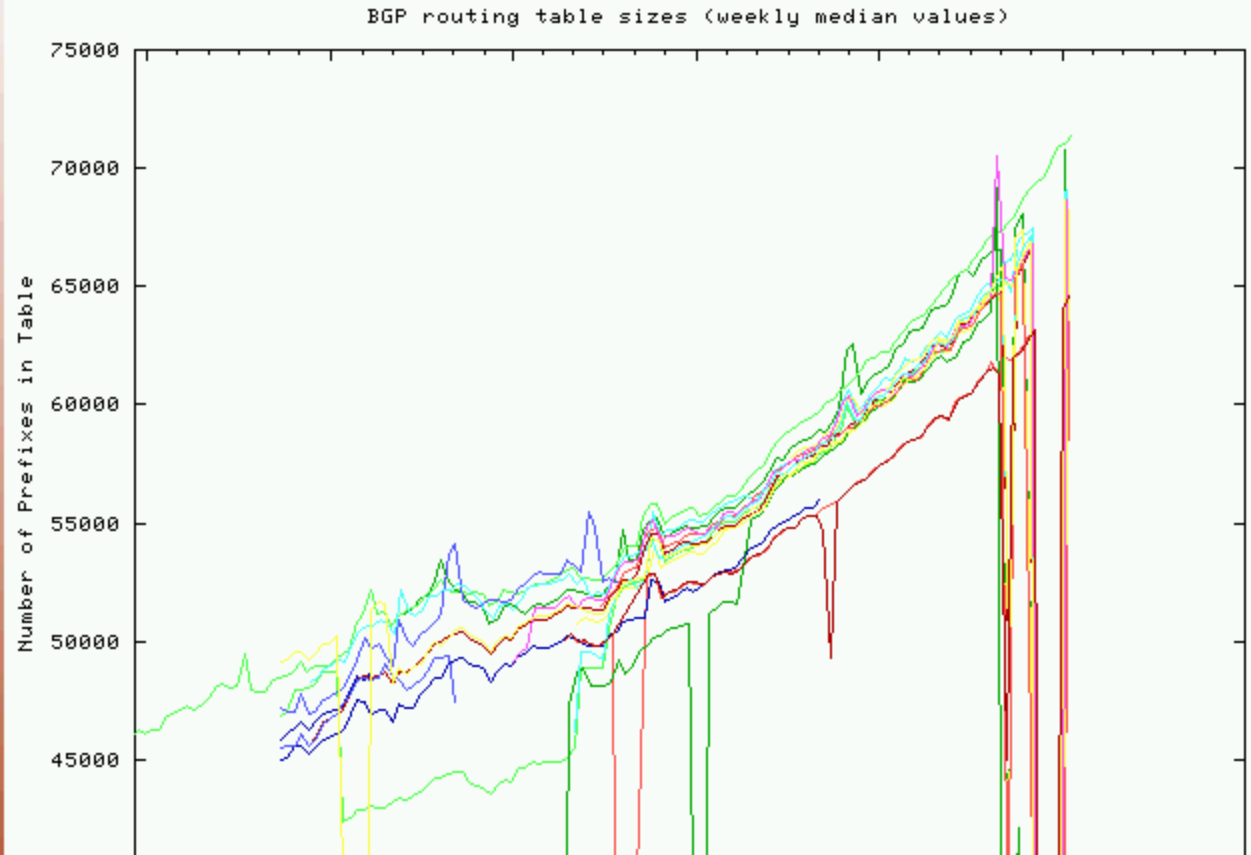

routing analysis: research priorities

- real-time identification & vis of flaps, outages, critical paths
- unintended consequences of new policies, topology
- propagation of change across ISPs
- realistic inter-domain routing models

==> requires better instrumentation
(w/o interfering with forwarding..)
ideal: route-lookups in real-time w/o kernel

routing: research obstacles

- canonical BGP (route table) data (not so much anymore)
- routes may change faster than ability to measure or analyze
- mapping IP addr to anything (deja vu)
- prudent security dictates making research difficult

overall: meas't & analysis challenges

- new methods for data collection, reduction, aggregation, mining, viz
- large, complex datasets (~Pbyte)
  - geographically and logically distributed
  - dynamically changing
  - enable inter- and intra-ISP analysis and feature detection
- correlation among data sources/types
- user-friendly integration with network utilities and control systems
- proactive participation
  - top-down (app devel's scope constr.)
  - bottom-up (ISP cooperation)
  - vendors in middle (to right of research)

summary: CAIDA/Cisco relationship

- education, outreach and training
  - IEC, ITL, ISMAs
- complementary tools/products
  - OC3/48mons
- router–based statistics
  - netflow, cflowd
  - route–views (Cisco/UO [Meyer])
- SNMP management tools
- infrastructure insights (multicast)
- traffic analysis (patterns, trends)

www.caida.org/Presentations/

kc claffy
UCSD/SDSC/CAIDA
kc@caida.org
www.caida.org

skitter: macroscopic study

DNS f root server (pv's): path wingspans
www.caida.org/Tools/Skitter

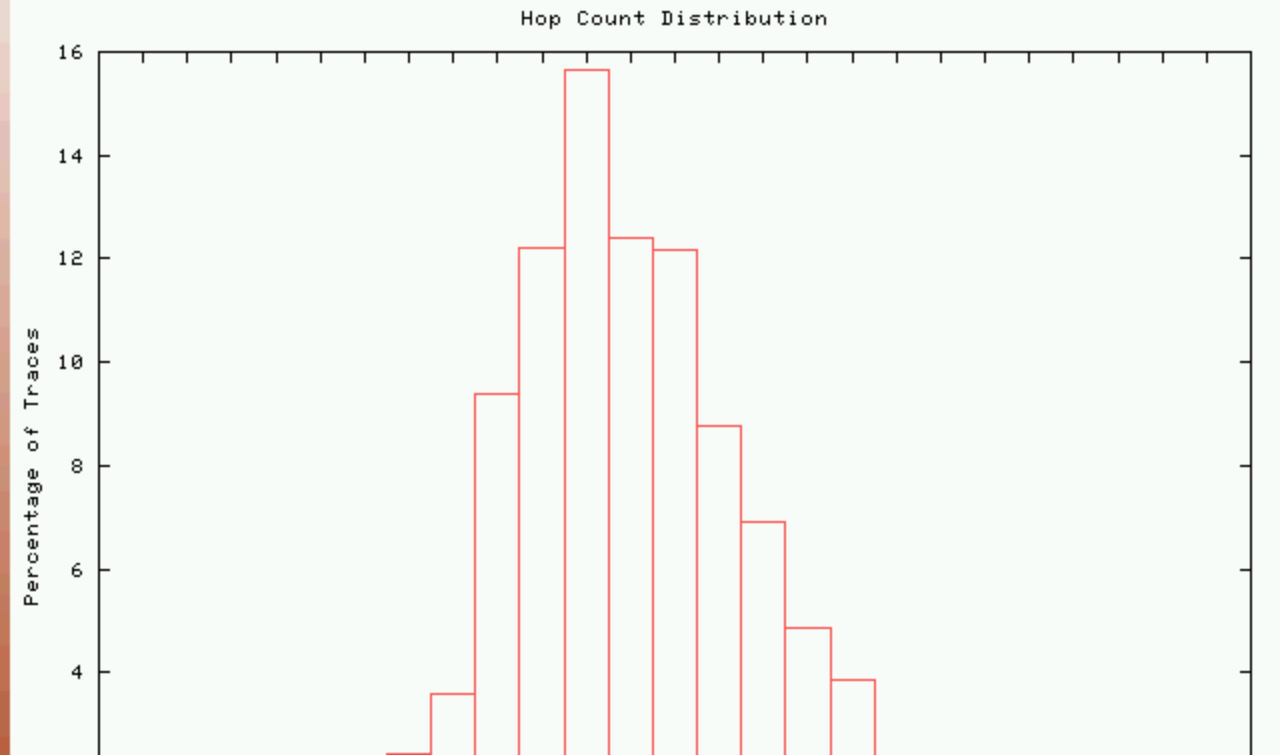

skitter: rtt vs hopcount (correlation?)

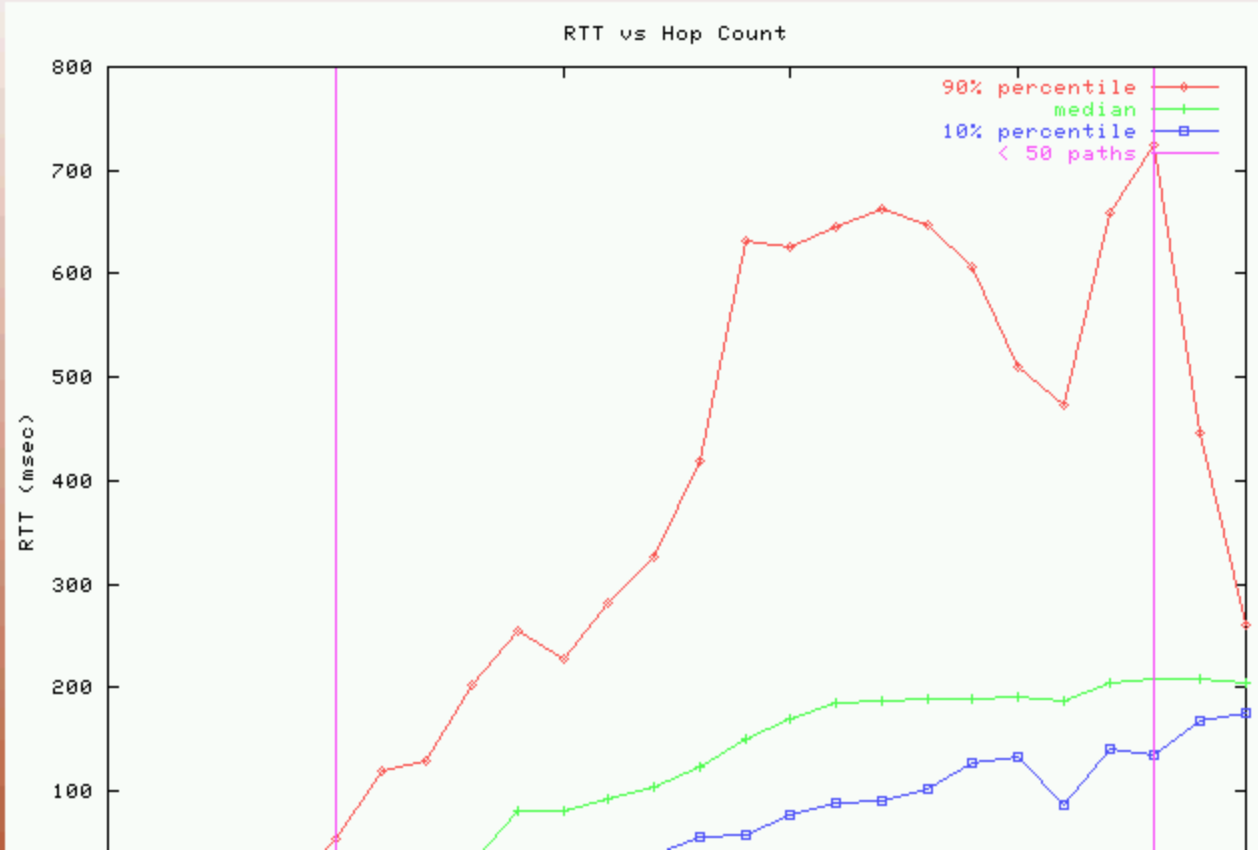

skitter: rtt distribution: tri-modal

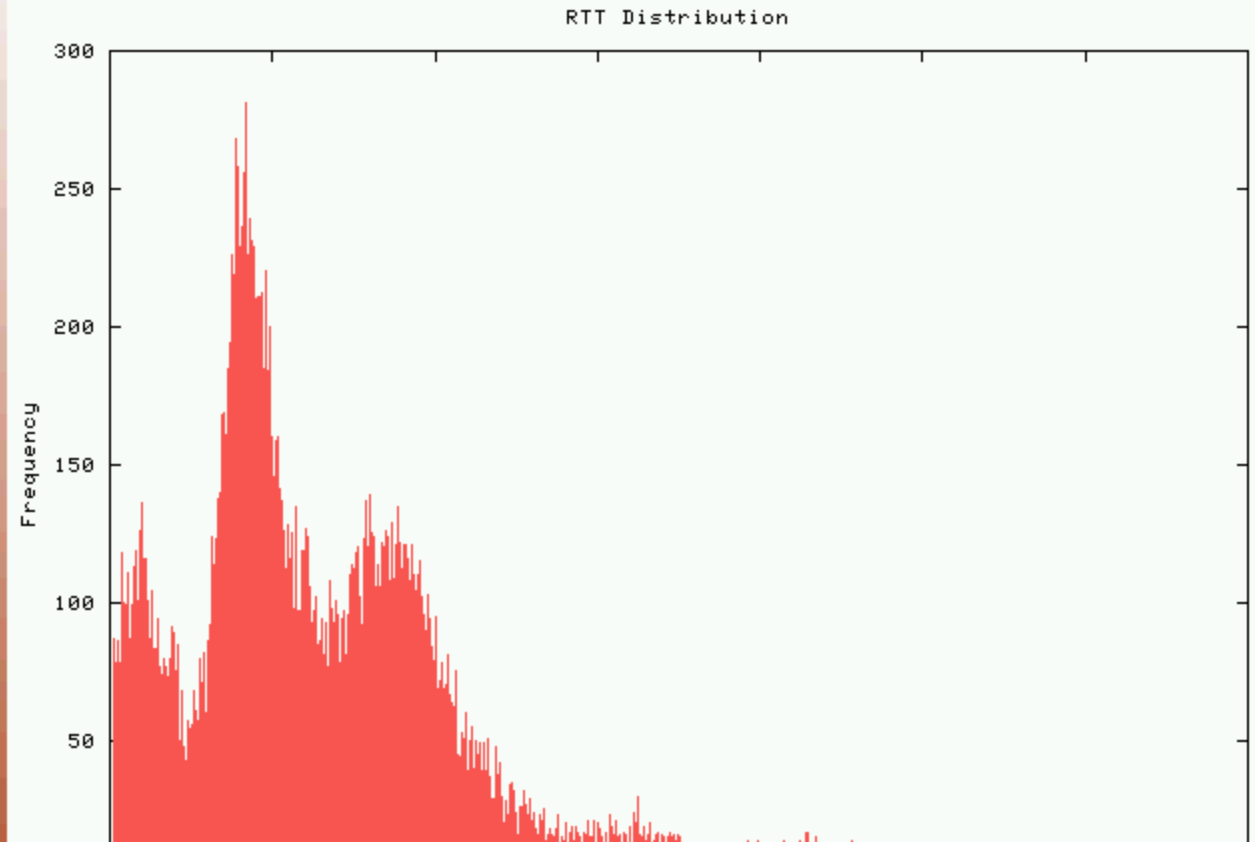

skitter: rtt vs longitude (light cone)

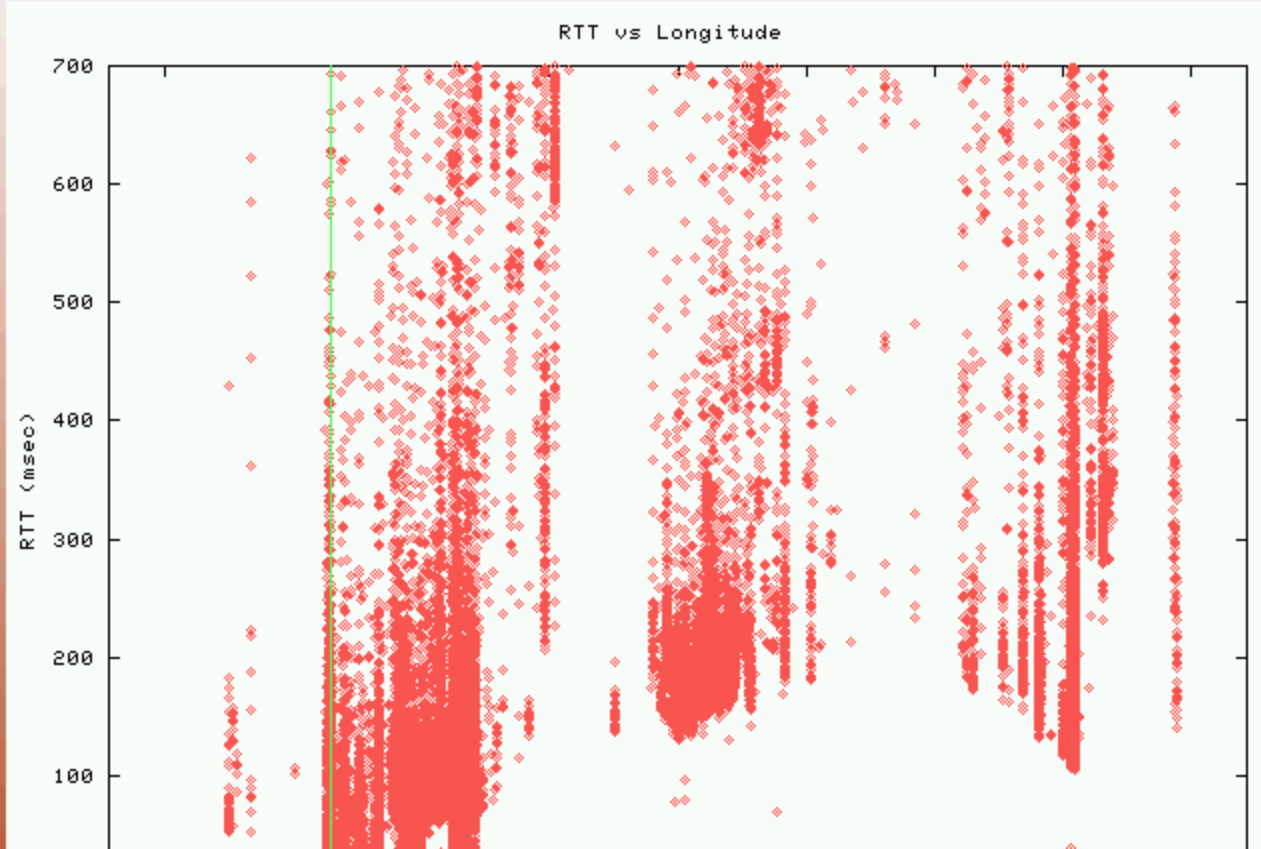

dispersion among ASes across paths

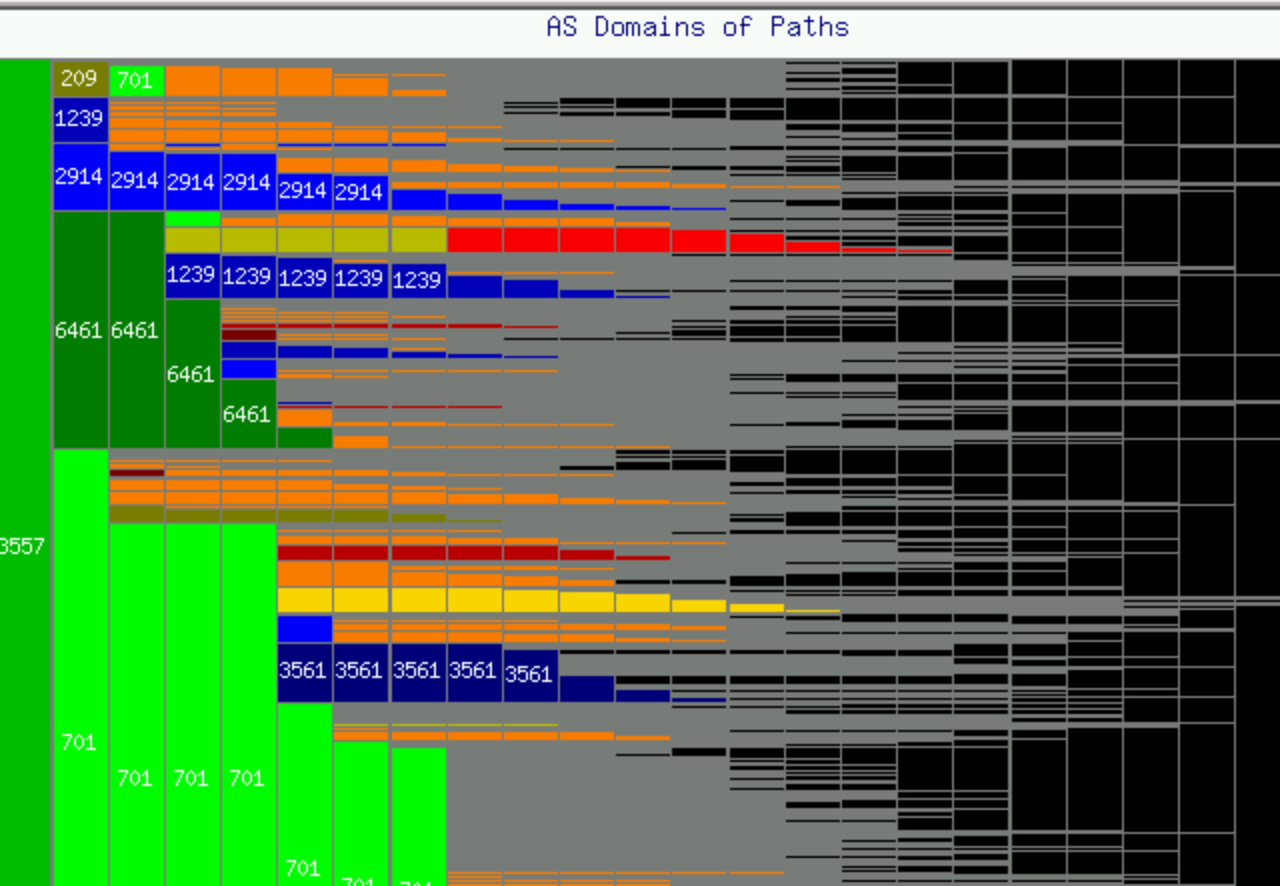

dispersion among ASes across paths (sdsc)

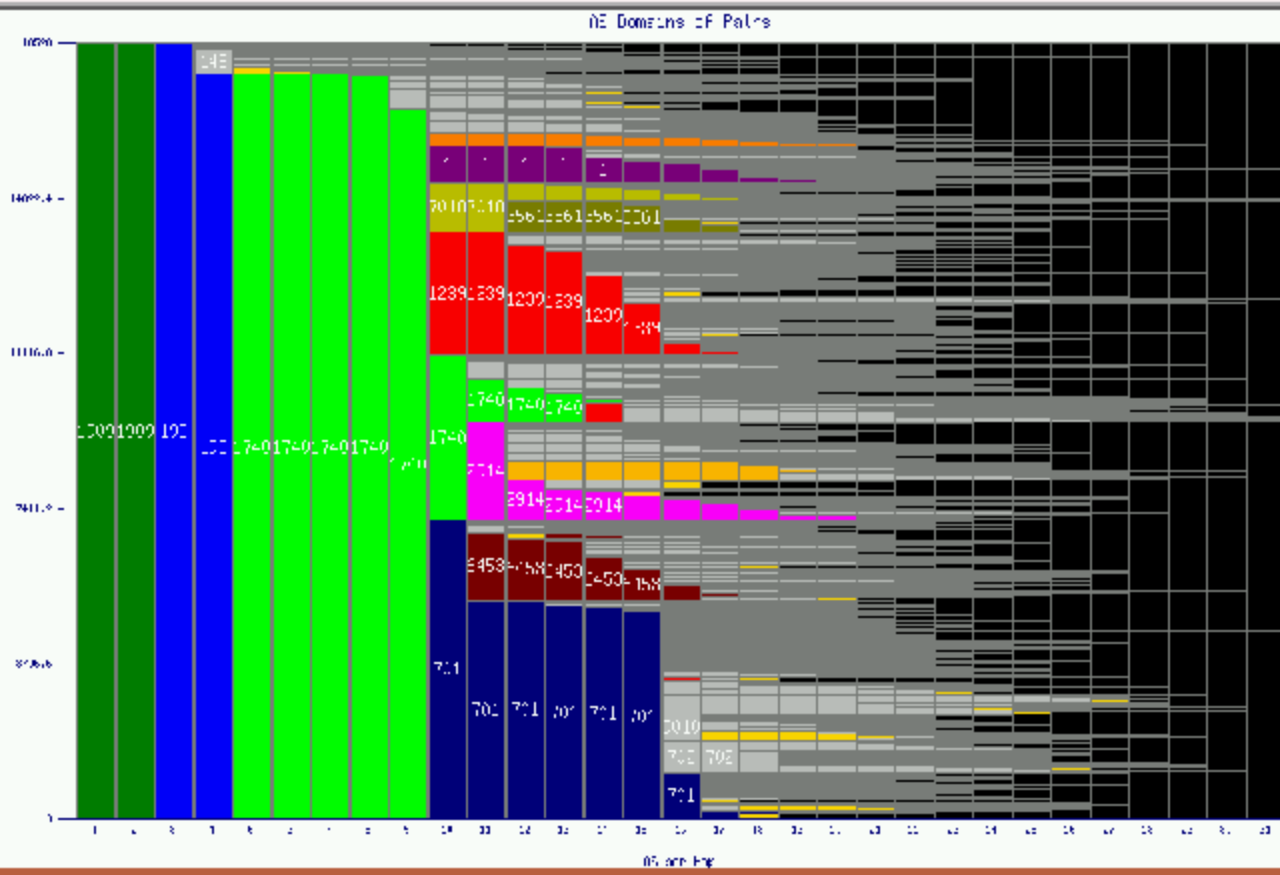

dispersion among countries across paths

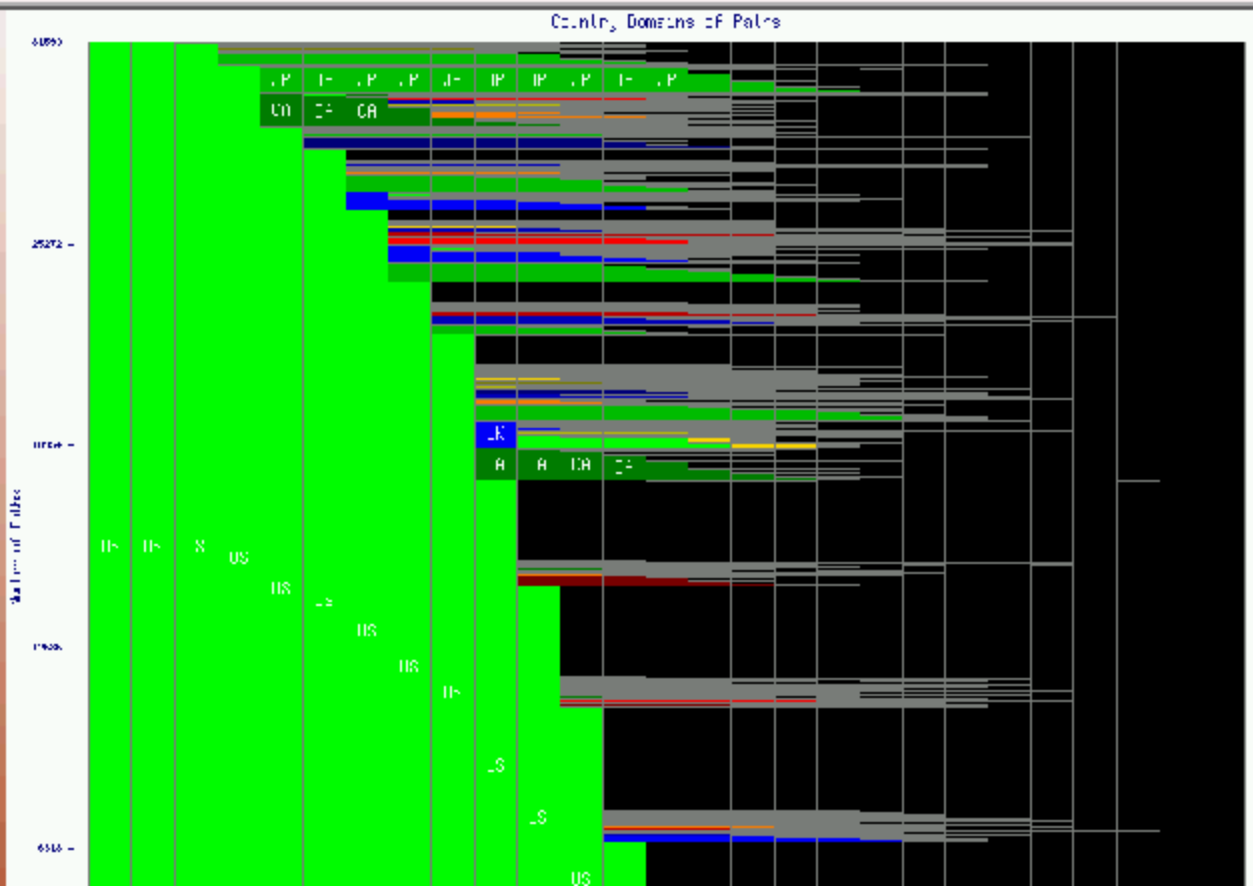

skitter analyses – preliminary findings

- ~1% IP destinations disappearing monthly (re-addressing, firewalls)
- route announced path not matching forward path
- indication of potential routing configuration errors
  (by no means automatic)
- persistence of paths
- methods to identify critical infrastructure
- is there an Internet "core"?

datasets available to researchers

CAIDA workscope summary

research pushing boundaries

- analysis of complex conditions
- management of large datasets
- correlation among different datasets
- development of timely, insightful visualization

recognized needs

- tool integration
- user interface improvements
- networked collaborative environments

summary: CAIDA/Cisco relationship

- education, outreach and training
  - IEC, ITL, ISMAs
- complementary tools/products
  - OC3/48mons
- router–based statistics
  - netflow, cflowd
  - route–views (UO/Meyer)
- SNMP management tools (new – dwm)
- infrastructure insights (multicast)
- traffic analysis (patterns, trends)

www.caida.org/Presentations/

kc claffy
UCSD/SDSC/CAIDA
kc@caida.org
www.caida.org