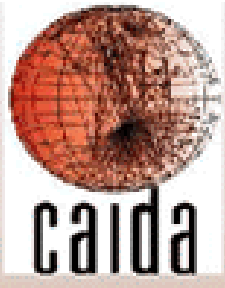

ISMA 2001:
routing & topology analysis

no aphorism is more frequently repeated...
than that we must ask Nature few questions,
or ideally, one question at a time.
...this view is wholly mistaken.
Nature will best respond to a logically
and artfully thought out questionnaire;
indeed if we ask her a single question, she will often refuse
to answer until some other topic has been dicussed.
-- perspectives in medicine and biology 1973 sir ronald a fisher

17 dec 01
www.caida.org

outline of my remarks

- motivation: background and questions
- sources of routing and topology data we use
  - topology
  - routing
- nanog panel on bgp (october 22)
- agenda for next three days
- goals and non–goals of workshop
- what i need from you

routing/topology

- motivation: macroscopic topology study
  - infrastructure-relevant research questions
  - led us to analyze BGP data on its own
  - realized that routing & topology inextricably related
- describe both kinds of data sets
- techniques for analysis of Internet 'core'
  - graph theoretic model
  - combinatorial core vs giant component
  - dual graph
  - subprefix connected components
- metrics for analysis
  - indegree/outdegree characteristics
  - transit versus origin role
- new contributions to the field of topology analysis
  - new granularities: BGP atoms, dual AS graph, ramified atoms
  - size distributions (weibull) with explanatory model (coalesence)
  - reachability functions
  - topological resilience

motivating questions (a sample)

- number of network prefixes in table
  - reduction beneficial to [routing system] infrastructure
- trends in routability of IP space given by registries
- which ASes are most highly:
  - connected, vulnerable, controlling connectivity
- is there an Internet 'core'
- does the core grow due to multihoming?
  - what is driving up the table size/churn?
- hop diameter of AS, AS diameter of IP paths
- optimizing active measurement architectures

can topology be inferred from BGP tables for the purposes of modeling, simulation, and infrastructure analysis?

motivating questions (more specific)

- how many prefixes does it take to cover the IPv4 address space?
- how many routing entries would be required if arbitrary intervals (as opposed to CIDR blocks) in IP space would be allowed as ranges of network addresses?
- what is the complexity of the system of AS paths associated with individual prefix?
- how many different types of networks are globally distinguishable with respect to routing policies?
- how many routing policies are applied by the Internet to addresses originated by one AS?

global AS topology (from skitter data)

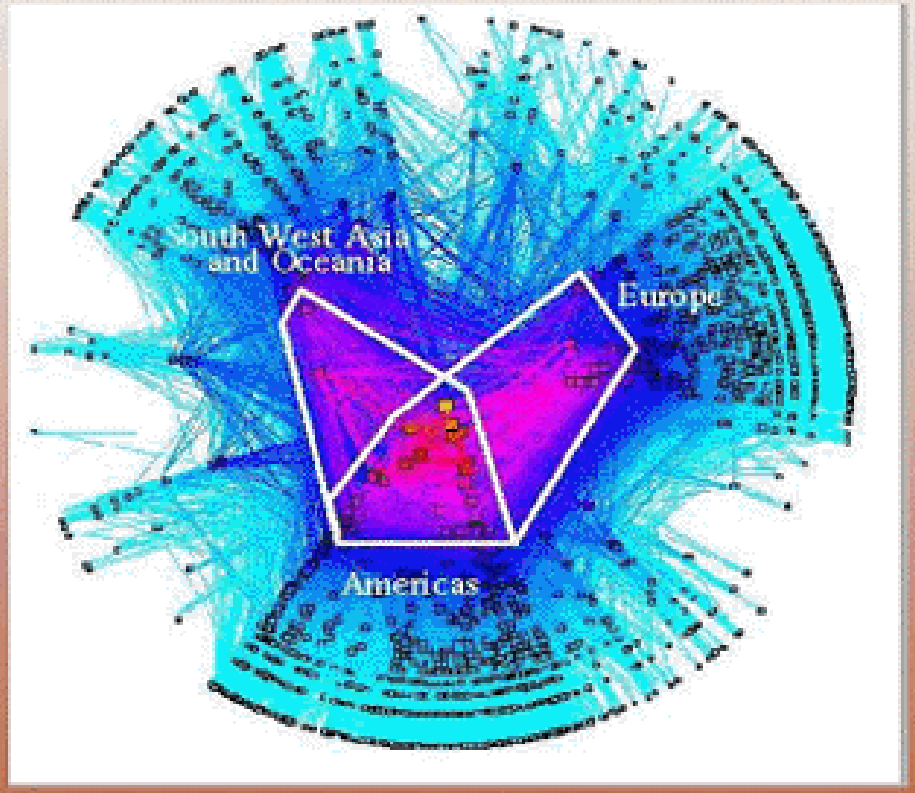

topology data: skitter

- about 26 sources around world
- 500k destinations (diff by box)
- over 60% of prefixes covered (still building lists)
- parallel ICMP echo request reply
- lightweight, coarse temporal granularity
- used for variety of macroscopic studies
- most comprehensive in world (low bar)
- data available to researchers

www.caida.org/tools/measurement/skitter/

topology data: dynamic properties

- all graphs change with time
- new equipment (nodes, links, firewalls) added
- new IP blocks alocated
- renumbering out of old IP blocks
- "death rate" of IP addresses
- paths oscillate (load balancing)
- paths fluctuate (routing instability)
- paths flip (manual config)
- outages, cuts, blackouts

okay so data is messy...

topology data: ambiguities

- non-responding hops (no IP returned)
- rate-limited response
- private addresses
- multicast addresses
- addresses in 0.-2./8 blocks
- no matching BGP prefix
- prefixes with multiple origin AS

okay so data is icky, too...

inter-domain routing (BGP) data

- data sources
  - UO's www.routeviews.org (Joel)
  - RIPE's www.ripe.net/ris (Hank)
  - Merit's IPMA /www.merit.edu/ipma (historic)
- really need real-time instrumentation
  - w/o impacting forwarding performance
  - ideal: route-lookups in real-time w/o kernel
- need data for realistic inter-domain routing models
  - real-time identification/vis of flaps, outages, critical pts
  - correlation of perf. w/some measure of path 'length'
  - comparison of forward path with
    - BGP path
    - shortest path
    - reverse path

other uses for inter-domain BGP data

- mapping topology data
  - aggregating IP to network prefixes
  - aggregating prefixes to origin AS
- inferring contractual relations
- "bird's eye view" of the net – (AScore)
- predicting AS path taken by a packet

important question: can you get essentially the same information from either dataset? (hint: no)

indeed, even though often covering fewer ASes than a full BGP table, skitter data shows bidirectional and transit connectivity for a significantly more ASes than BGP data of the best available quality and sampling. (totally non-obvious...)

routing analysis: nanog–relevant perspective

- entire day devoted at last nanog (22 oct 01)
- routing research increasingly recognized as critical
  - at least bgp
- http://www.maoz.com/nanog23/
- http://www.nanog.org/
- essential for researchers to pay attention to this community
  - bridge the gap

somebody has to do something, and it's just incredibly pathetic that it has to be us....
-- jerry garcia of the grateful dead

routing analysis: available data (on the rise!)

- RIPE RIS (www.ripe.net/ris)
  - zebra bgpd
    - www.zebra.org
  - MRT format RIB and UPDATES
    - www.merit.edu/mrt
- route-views.oregon-ix.net (archive.routeviews.org)
  - 'sh ip bgp' format at 2 hour intervals
  - collection script from Sean McCreary
  - hwb archive on moat.nlanr.net
- PCH (www.pch.net)
- looking glasses
  - good list at http://neptune.dti.ad.jp
- private listeners/archives

highlights of BGP nanog track (last nanog)

- routing table growth (ietf talk)
- dark Internet address space (not announced)
- global routing instabilities during code red II and nimda worm propagation (ato/jim)
- BGP MOAS conflicts (multiple origin AS)
  - prefixes announced twice
  - operational attempt at engineering vs misconfiguration
- CIDR at work (cengiz)

rough agenda for next 3 days

monday
- sources of routing data (bill's talk moved to tue)
- routing instabilities and anomalies
- routing performance
- macroscopic statistics
- (...reception in lobby/discussions here)

tuesday
- topology analysis
- statistical methods (donoho couldn't make it)
- leave early for wild animal park

wednesday
- routing management
- open discussion: future needs for archiving, analysis

goals and non-goals for next 3 days

model after andy's leiden workshop

- more academic than nanog
- less academic than sigcomm
- more like IRTF, working group workshop
- lots of discussion, relaxed agenda/timekeeping
- encourage working together
- no proceedings
- NSF sponsorship (means we can bring quality people here that might not otherwise be able to attend)
- publications/citations is an explicit non-goal
  - @unpublished{}

what i need from you

- your slides to marg@caida.org
- any minutes you take
  - send to kc@caida.org to help me write final report
- reviewing final report before i put on web
  - make sure you're represented accurately
  - feel free to preempt me with exec summary as you see it
- use microphone because we're being webcast
- thank nsf (chip should be here)
  - good for you to meet him anyway

summary

- a definite increase in receptiveness to research and analysis of routing data and topological behavior
- a lot of good work starting to happen
  - field in very early stages
- data sources increasingly high integrity
  - researchers can get as good data as operators (!) (mostly)
  - still not enough
- operators more inclined to work with researchers
  - if they're asking the right questions
  - they need our help
- excellent opportunity for researchers
  - www.ietf.org (next one minneapolis, march)
  - www.nanog.org (next one miami, 10 feb 2002)
  - isma mailing list isma0112@caida.org