Next Generation Network Measurement

Mark Crovella

Internet Measurement: The Golden Years

[Kleinrock and Naylor, 1974]

- Original ARPANET had built-in abilities to: Trace a single packet's passage through the network
  - Obtain instantaneous traffic matrix
  - Obtain instantaneous queue lengths in IMPs
  - Obtain per-IMP traffic summaries and histograms
  - Obtain any IMP's routing table
- Measurement-based research on packet switching was a explicit goal of the ARPANET

What Happened?

- The focus of network research became concentrated on connectivity
- Detailed analysis of the network itself fell by the wayside
  - c. 1975 DCA took over operation of the ARPANET, Network Measurement Center shut down
- Pendulum has swung very far in this direction
  - Many Internet design decisions actually have impaired network measurement

Design Impediments

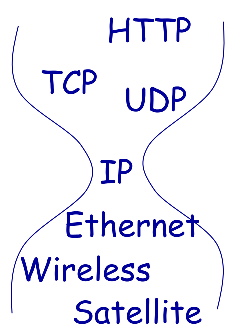

- The IP Hourglass
  - IP over everything
- Core simplicity
  - Stateless switching, stupid network, end-to-end argument
- Distributed Internetworking
  - Shutdown of NSFNET backbone in 1995 removed single measurement point

Practical Issues

- Data Volume
  - Hundreds of megabytes / sec on backbone links
- Data Sharing
  - Sensitive information on many levels
    - Privacy sensitive
    - Competition sensitive

Statistical Difficulties

- Long tails
  - Instability of metrics
    - Debate over stationary models
  - Difficulty in modeling (powerlaws? Lognormals?)
- Autocorrelation
  - Strong autocorrelations increase measurement needs
- High dimensionality
  - Address space, port space: huge

What can we do?

- Ad hoc, near term:
  - Continued research on methods to overcome limitations of current situation
- Fundamental, long term:
  - New Internet measurement paradigms
  - Construction of Internet measurement facilities operated and shared by the research community (not by providers)

Near Term / Ad hoc

- Too much data: Sample (intelligently)
  - Packet, trajectory, flow sampling
  - Topology capture
- Unavailable data: Infer (intelligently)
  - Network-internals (tomography)
  - Link types (wired/wireless)
- Lack of Visibility: Ingenious workarounds
  - Traceroute (many variants)
  - Rocketfuel

Some Successes

- Router & AS topology characterization
- Characterization of interdomain system
- Inference of hidden properties
- Traffic modeling (short and long timescales)
- Statistical invariants (mice & elephants, Zipf laws)
- Characterization of Web graph
- Models of worm propagation
- Science driven engineering (AT&T, Sprint,…)

Big challenges ahead

- Engineering
  - Performance evaluation
  - Capacity planning
  - Security
- Science
  - Interaction of network and people / society
  - Growth laws
  - Statistical properties

Fundamental, longer term approach

- New Internet measurement paradigms
- Construction of Internet measurement facilities operated and shared by the research community (not by providers)

How is Internet Measurement Done?

- Three models
  - Internet Measurement Organizations
    - CAIDA, NLANR, RIPE, …
  - PI driven projects
    - Local measurement infrastructures
    - Built by effort of a single PI / small group
  - Planetlab
    - Community-shared resources
    - But *very* limited measurement capability

Time Ripe for a Community Approach?

- Community Approach =
  well defined measurement community +
  well defined measurement scope +
  variety of research agendas +
  need for expensive measurement equipment +
  community self-organization

Active Community Exists

- IMW/IMC submissions
  2001: 53
  2002: 93
  2003: 109
  2004: 157
- PAM experienced similar growth
  2004: 184 submissions
- Books in area
  "Evolution and Structure of the Internet,"
  "Internet Measurement: Infrastructure, Traffic and Applications"

Internet Science

- Measurement Scope: Understanding the Internet at all layers, as it evolves in time
- Does this correspond to any other sciences?
- Can we learn from how other sciences organize their measurement infrastructures?

Astronomy

- Large collection of discrete objects (stars, galaxies, planets, etc)
- Interested in their emissions and reflections
- Can measure these objects, but can't really do much to affect the objects being measured

Biology

- Interested in describing systems (cells, populations) that are
  - Complex
  - Comprised of many interacting mechanisms with
  - Many feedback loops
- Can affect systems in some ways
  - Can "poke" a cell or organism to see what happens
- Can't usefully take apart a functioning system

Earth Science

- Scale of the system studied is global
- Many important effects concern interaction of human society with the system
- Many important effects depend on geography and physical distance

Example Community Approaches

- Astronomy: building and operating large telescopes
- Oceanography: building and operating research vessels

Telescopes

- Range of options (smaller -> more informal)
  - Owned/operated by small groups
    - BU/Lowell 2m telescope
      - BU supports at $150K/year (1/2 time)
  - National Facility
    - Keck
  - Space Based
    - Hubble

Astronomy

- Example: Keck Observatory
  - Governing board for telescope
    - One member per institution (Dean or Scientist)
  - Director appointed by Board
  - Time Allocation Committee
    - Not insiders – peers from across discipline
    - Serve on committee 2-4 years
    - Accepts short (2-page) proposals 1x or 2x / year
    - Ranks and forms a consensus list
    - 20 proposals / semester (one day's reviewing)

Telescope proposal process

- Two parts
  - Science proposed
  - Amount of time being requested
- TAC:
  - Ranks science 1-10
  - Ranks time, makes recommendation
    - Can say "try 10% of time, if it works, come back for more" or "We think you can do this in 1/3 the time"
- Director makes final call if telescope is oversubscribed

Telescope Data

- Most national facilities make data available after some proprietary period
  - 6 months to a year
  - To allow PI to get data analyzed and out
  - Data will become available even if not used by PI
- Smaller facilities may not do this
  - Due to archiving costs
- Sometimes the Director will arrange a "shotgun marriage" if two projects propose to collect similar data

How do you build a new telescope?

- There is something called a "decadal review" – what astronomy needs to be done in the next 10 years
  - The next one is 5 years out, there is already a lot of jockeying going on ☺
- Clearly needs to have community behind it
  - If you can get on the decadal review, you are in good shape
- Usually:
  - Donor + Institutions + NSF/NASA

Oceanography – Research Ships

- All research ships are handled by a single organization – UNOLS (61 institutions)
  - 27 research vessels in 20 home locations
  - All schedules publicly available
- Ships are owned/operated by home institutions
  - under contract to NSF
- Chair, Council, and Committees
  - Ex: Ship Scheduling Committee

UNOLS oversees, Funding agency allocates

- $50,000 / day ship time
- Ship time request submitted as part of proposal
  - PI specifies how much ship time is needed
  - About a year in advance
- NSF, ONR, NOAA panel reviews and approves ship time
- UNOLS Scheduling Committee
  - Implements NSF panel recommendations

Oceanography Data

- Ocean Core Drilling Program
  - 15 years $150M
  - All cores are kept forever (3 locations)
  - Professors send their students to sample cores
  - All data must be made available 1 year after collection
- UNOLS
  - All data must be made available 2 years after collection
  - Researchers on same cruise share data
  - UNOLS matches experiments

Time Ripe for a Community Approach?

- Community Approach =
  well defined measurement community +
  well defined measurement scope +
  variety of research agendas +
  need for expensive measurement equipment +
  community self-organization

What model makes sense for a CONMI?

- Not single-threaded like a telescope
  - Many experiments should be able to run simultaneously
  - We can exploit virtualization
- Should have some sense of "global" coverage like ocean science
- Data archival
  - Notion of "embargo" or "proprietary period" seems to work in other sciences