Spectroscopy Methods for Network Inference

Andre Broido
C A I D A

CAIDA / SDSC / UCSD

http://www.caida.org

WISP

Workshop on Internet Signal Processing

San Diego

2004-11-12

“It shall be, when I bring the cloud over the earth, that the rainbow shall be in the cloud;

“And I will remember My covenant which is between Me and you [...] the waters shall never again become a flood to destroy all flesh.”

Gen.9

Plan

Perspective
Definition
Others’ work
ATM, DSL, Cable
DNS updates
ICMP delay
Conclusion

Integers

5 25 1/2 37 30

Fundamentals

- Questions inspired by Kolmogorov:
- How much do we owe to measure theory?
- Can we call our measures probabilities?
- Are complexity and randomness synonyms?
- Should we treat unknown as random?
- How can we reduce descriptions?
- Relative to what knowledge base?

Descriptions

- Maxwell: dF=0, d*F=0
- gauge theories vs. fiber modes
- Which notation/concepts should we use?
- Is structured risk minimization the way to go?
- Should we reduce dimensions or bit counts?

Experiment design

- Which parameters affect data variation?
- How (in)dependent they are?
- How do we scan parameter space?
- (Exhaustively? Consecutively?)

Definition

Spec-tros-co-py, the science that deals with the use of the spectroscope and with spectrum analysis

Claim to fame: discovery of quantum mechanics

Features

- Spectroscopy = study of quantization
- Binary, discrete, qualitative inferences
- from contuniuous/numeric data
- Typical method: a clever transform
- to focus relevant data
- followed by thresholding

Distinctions

- Find network properties from spectra
- Periods, frequencies, delays
- Inverse problem
- Classification vs. estimation
- Narrow spikes vs. continuous density
- Integers vs. reals
- Numerology vs. numeric analysis

Methods

- Autocorrelation
- Fourier transform
- Lomb periodograms
- Radon transform
- EM
- Eyeballing
- Hand-picking
- 500 page specs (DOCSIS, 802.11)

Timescales

- Months/days: Traffic per yearl, week
- Minutes: BGP timers and keepalives
- Seconds: TCP timeouts
- (Milli)seconds: RTT, TCP states
- Milliseconds: Interrupt latency

Related work

- Timestamping & Timekeeping
- Single-hop and point-to-point delay
- Cross-traffic interpretation
- Capacity and rate estimates
- Tomographic inference
- OS/TCP stack fingerprinting (RING)
- Router tests

Contributors

- Sue Moon - skew estimation
- Dina Katabi - cross-traffic
- Stephen Donnelly - timestamping
- Alefiya Hussain - identifying attacks
- Vinay Ribeiro - bitrate estimation
- Rajesh Krishnan - hidden flow detection
- Dina Papagiannaki - router delays
- Attila Pasztor - packet probing design
- Yolanda Tsang - tomography
- Rui Castro - topology inference
- Jorma Kilpi - wireless
- and their advisors...

Timescales vs. applications

- Hour: DNS updates
- (Sub)second: TCP dynamics
- Millisecond: Bitrate estimation
- Microsecond: SONET clock accuracy
- Nanosecond: Packet timestamp quality

How can delay be quantized?

- Bit, byte, word grids
- Finite timestamp resolution
- Fixed cell/slot time
- Layer 2 technologies:
- Time-division multiplexing
- Combined with frequency/code division
- Router switching fabrics
- Frame hierarchies in GSM/GPRS
- ATM, DSL, Wireless, Cable

Our work

- Radon tranform for ATM rate evaluation
- DSL rates
- Cable modems' rates
- DNS update analysis
- papers - see www
- more in the pipeline

ATM (2000)

- Stepwise size-delay dependence
- A jump every 48 bytes
- min delay = d. + ceil(L/48)/R
- What is the cell rate/time?

Algorithm

- Idea: substract a step sequence
- find the marginal with min spread

- Scan all possible cell times
- Compute residual inter-packet delays for each tested cell time
- Choose one with the sharpest spike (min entropy)
- A simple solution to an inverse problem

Answer

- The entropy minimum is at 18.48 usec
- OC-3 allows 2.7 usec/cell
- Rate is limited 7.5-fold
- Slightly below contract (19.3 Mbps)

DSL (2002)

- Send batches of same-size packets
- Scan all sizes, 40-1500 bytes
- Find size-delay dependence

Answer

- DSL is ATM based
- PPP over Ethernet over ATM
- Typical cell times:
  - 3.31 ms (128 Kbps)
  - 2.65 ms (160 Kbps)
  - location-dependent

Cable data

- Delay quanta for cable are mostly 2,3,6 ms
- 3 and 6 ms can arise via aliasing
- Spurious spikes for rational fractions
- 2 ms = providers' choice of 500 ”maps”/sec
- See DOCSIS for details

ICMP takes a break,

or

Nonlinear ICMP delays (2004)

Motivation

1. Test axioms
"Ground truth" for delay analysis

2. Solve a forward problem
to enable inversion

3. Use traceroute RTT to find:
link capacities
link latencies
same-router IPs
network geography
pop-level maps (plm)

Why not previous work?

Light Reading 2001 (Newman e.a):
Stress testing routers
Full line rate loads
Sonet only

Sprint 2002, 2004 (Dina e.a.)
Operational routers
No control of traffic
Single device

Axioms

- delay increases with packet size
- delay is linear in size, d = d. + L/C
- delay over minimum = cross-traffic
- delay is payload-independent

serious people use these facts
serious work is based on them
They must be correct

Sample problem

Packet-over-Sonet uses HDLC framing.
Every flag (frame delimiter) char is escaped
All flags' payload doubles packet size
Can we discover Sonet by delay increment?
Could solve backbone capacity inference
OC48: sensing 5 usec delta over mult hops
Aside: HDLC stuffing not logged
Utilization can be twice the byte count

Experiment

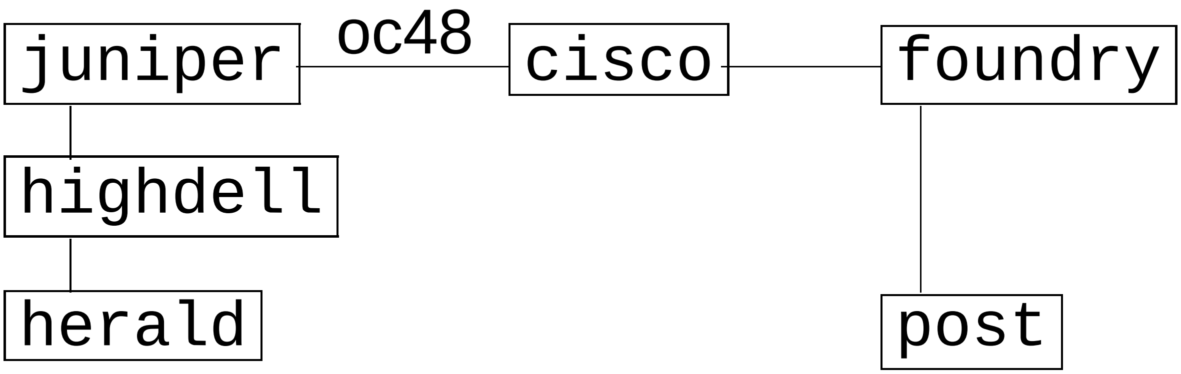

Equipment (clockwise):
IBM eServer `herald`
Dell PowerConnect 5212 switch
Juniper M20 router
Cisco 12008 router
Foundry BigIron 8000 router/switch
IBM eServer `post`
Links: oc48 (Juniper to Cisco)
GigabitEthernet (all other links)
more FreeBSD and Linux boxes

Factors of design space

- Medium to high-end routers
- Three router vendors
- Two switch vendors
- Gigabit capacities
- Sonet and Ethernet
- 9000 byte MTUs
- DAG4 OC48 and GigE monitors
- Several host vendors
- Two host OSes

ICMP tests

- TimeExceeded, PortUnreachable, EchoReply
- 40 to 9000 bytes
- unloaded routers (no other traffic)
- one packet at a time
- packet spacing of 200 usec-20 ms

Parameter scan

- hopping over product space:
- (40-9000 bytes) x 2 hops x 10 ToS x 4 pkt...
- hopping avoids damage from
  - burst errors
  - edge effects
  - time dependence
- hopping by powers of a primitive root
- in mixed-radix expansion

Observed

- Size-delay growth rate changes at 1500 bt
- Flipping (high-low) rate (piecewise linearity)
- Convex/concave bends (curvature)
- Jumps or drops (discontinuity)
- Stepwise growth (64 byte cells)
- Negative (decreasing) slope

ICMP gen.rate != input link capacity

More issues with ICMP

- Type-dependent drop and bit rates
- Uniform-like size-independent delay spread
- "bands" of preferred size-independent delays
- "Simple" sizes (32n bytes) served faster
- Occasional extra delay on empty router
- Cache warm-up causes extra latency
- Close packets postponed by 9-10 ms
- Confirmed some for forwarding delay

Conclusions

- Delay quantization is ubiquitous
- Spectroscopy can be used for
  - Layer 2 identification
  - bitrate estimation
  - SLA verification
  - source recognition
- ICMP delay is nonlinear for 40-9000 bytes
- Same for forwarding delay (under study)

The raw DNS and OC-48 data
is available on-site

Acknowledgements:

- kc claffy
- Young Hyun
- UCLA IPAM
- Ryan King
- Yoshi Kohno
- Margaret Murray
- Evi Nemeth
- Robert Nowak