



Experience with the Surveyor IP Performance Measurement Infrastructure

Matt Zekauskas

`matt@advanced.org`

IMIC Workshop

30 August 1999



Outline

- Surveyor overview (+IPPM)
- System description
- Lessons and some results
- Futures & opportunities (+IPPM)





Surveyor Program Objective

- Create technology and infrastructure to allow users and service providers (at all levels) to have an accurate common understanding of the performance and reliability of paths through the Internet.





IP Performance Metrics

- IETF IPPM effort
 - Framework and Connectivity RFC
 - One-way delay, packet loss, round-trip delay to proposed standard RFCs
 - Others: bulk transfer, delay variation
 - Surveyor: implementation of one-way delay and packet loss metrics
-
-



Why Measure Delay?

- Minimum of delay:
transmission/propagation delay
 - Variation of delay: queuing delay
 - Large delay makes sustaining high-bandwidth flows harder
 - Erratic variation in delay makes real-time applications harder
-
-



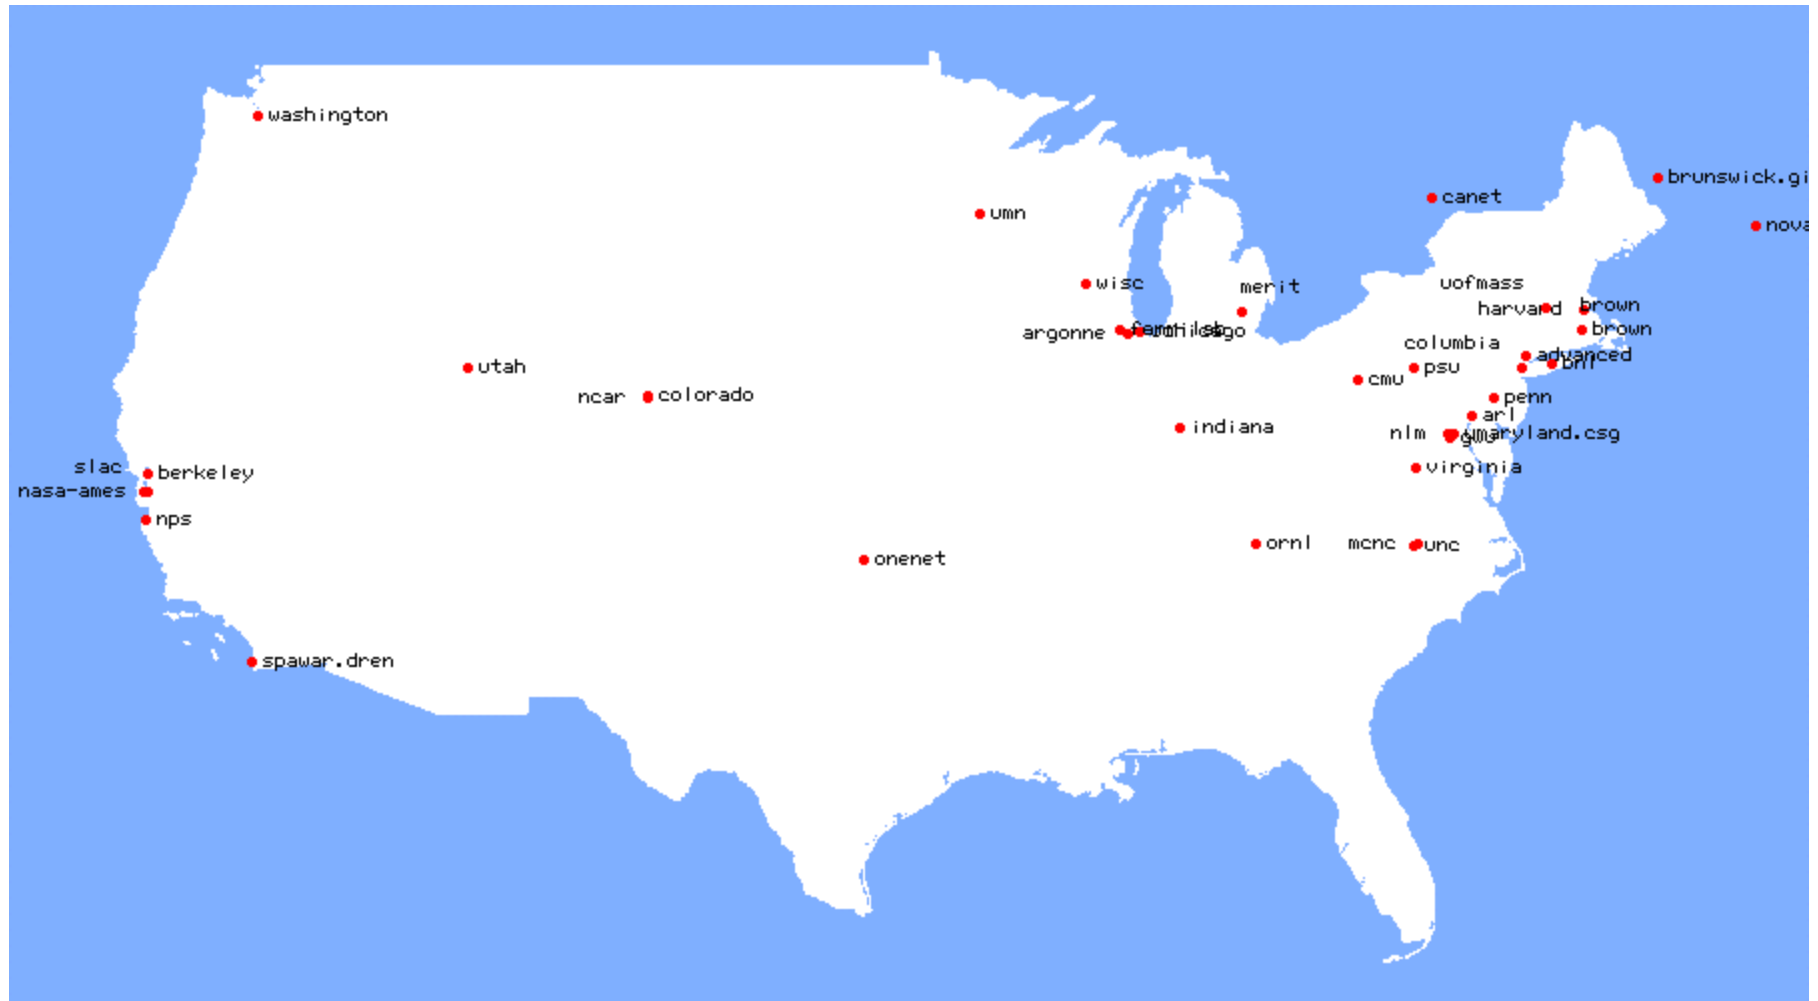
The Surveyor Infrastructure

- Measurement machines at campuses and at other interesting places along paths (e.g., gigaPoPs, interconnects)
 - GPS to synchronize clocks
 - Centralized database to store measurement data
 - Web based reporting and analysis tools
-
-



Current Surveyor Deployment

- 56 machines
 - Universities
 - Tele-Immersion Labs
 - National Labs
 - ...others
 - 2098 paths
 - NASA Ames XP
 - I2 gigaPoPs (some)
 - CA*net2 gigaPoPs
 - APAN sites
 - Abilene router nodes up with NTP, awaiting GPS
-
-





Measurement Machines

- Dell 400 MHz Pentium Pro
 - 128 MBytes RAM; 8 GBytes disk
 - BSDI Unix
 - TrueTime GPS card and antenna (coax)
 - Network Interface (10/100bT, FDDI, OC3 ATM in beta)
 - Special driver for the GPS card
-
-



Measurement Technology

- Active tests of one-way delay and loss
 - Test packets time-stamped with GPS time
 - Back-to-back calibration: 95% of measurements $\pm 100 \mu\text{s}$ \rightarrow $10 \mu\text{s}$ soon
 - Measurements centrally managed
- Concurrent routing measurements





Ongoing Tests - Delay

- Type-P
 - 12 byte UDP packets, 40 bytes total
 - Port “random” per session
 - QoS (via EF DS byte) ready
 - Scheduled using a Poisson Process
 - average rate: 1 or 2 per second
 - “Mostly” full mesh
-
-



Ongoing Tests - Routing

- Traceroute to same sites as One-Way delay
- Scheduled with Poisson process
 - average rate: one every 10 minutes





Reporting and analysis tools

- Web based tools
- Daily summary reports
- Integration with route measurements
- Java applet to dynamically graph
- Developing asynchronous notification of anomalies





Lessons so far

- Routing is asymmetric
 - Even when routing is symmetric, queuing is asymmetric
 - Can detect level 2 changes (SONET failover or ATM routing)
 - Have observed
 - low delay with loss
 - high delay without loss
-
-



Lessons

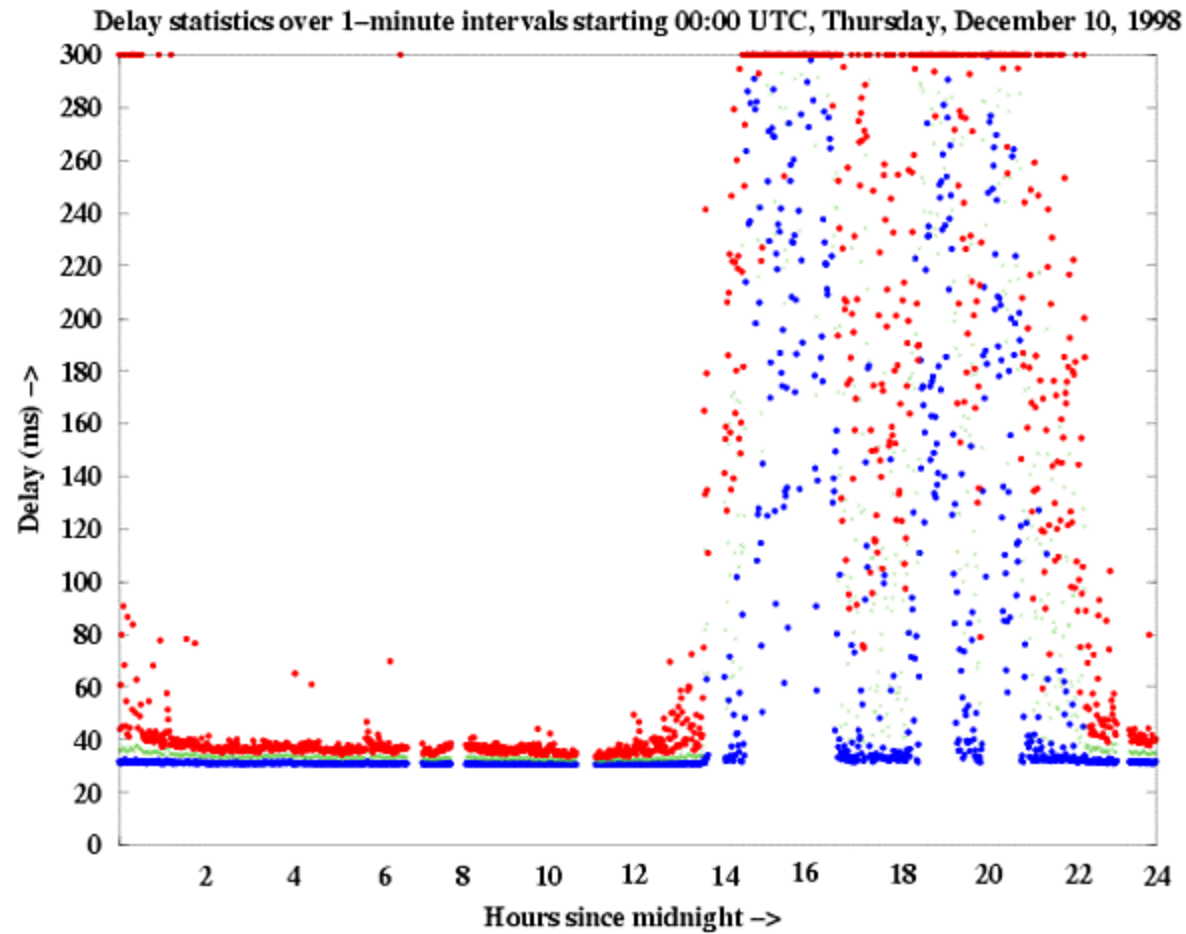
- Commodity path losses, even with this traffic, high (10% over 4 hour period!)
- HPC connections indeed provide low-latency low-loss paths
- HPC connections do fall back to commodity paths, sometimes frequently or for long periods






- minimum delay
- 50th percentile delay
- 90th percentile delay

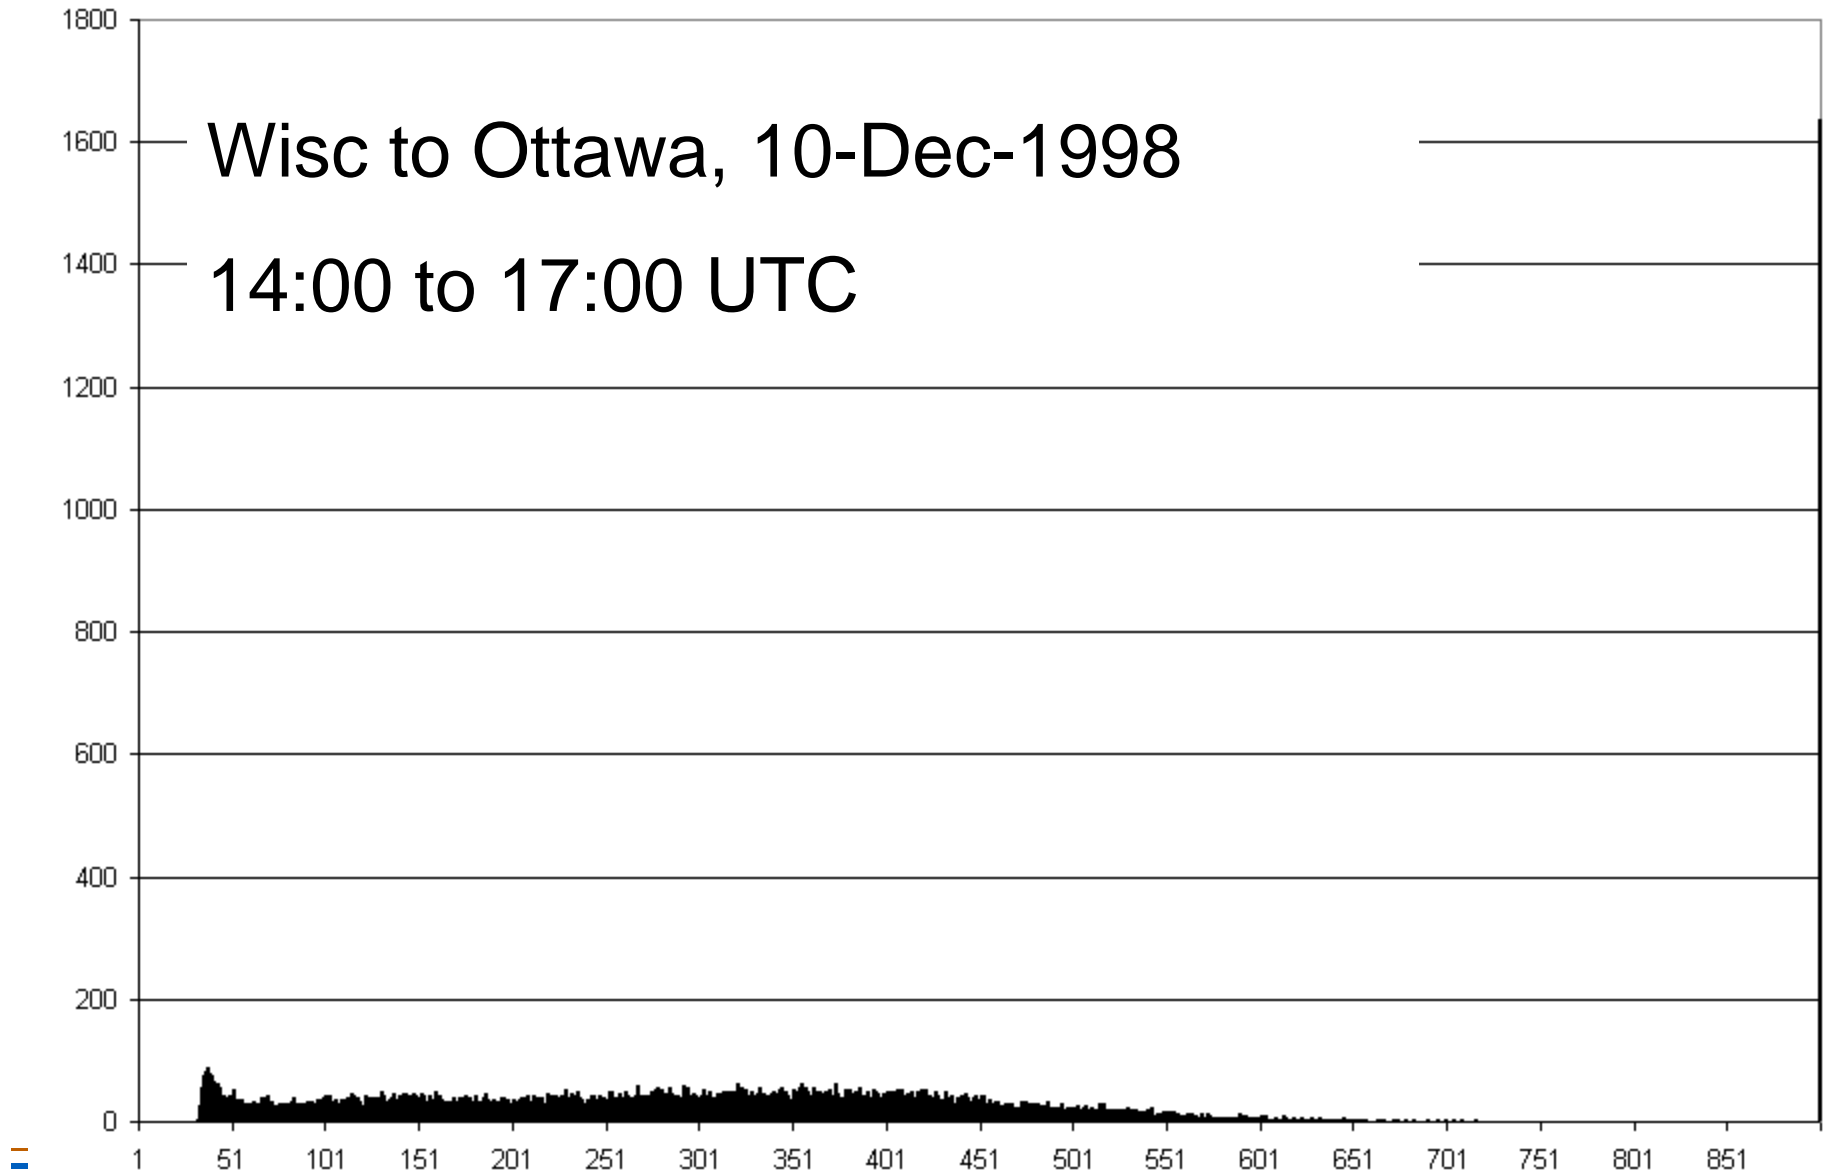
Univ Wisconsin to CANARIE-I2 Gigapop







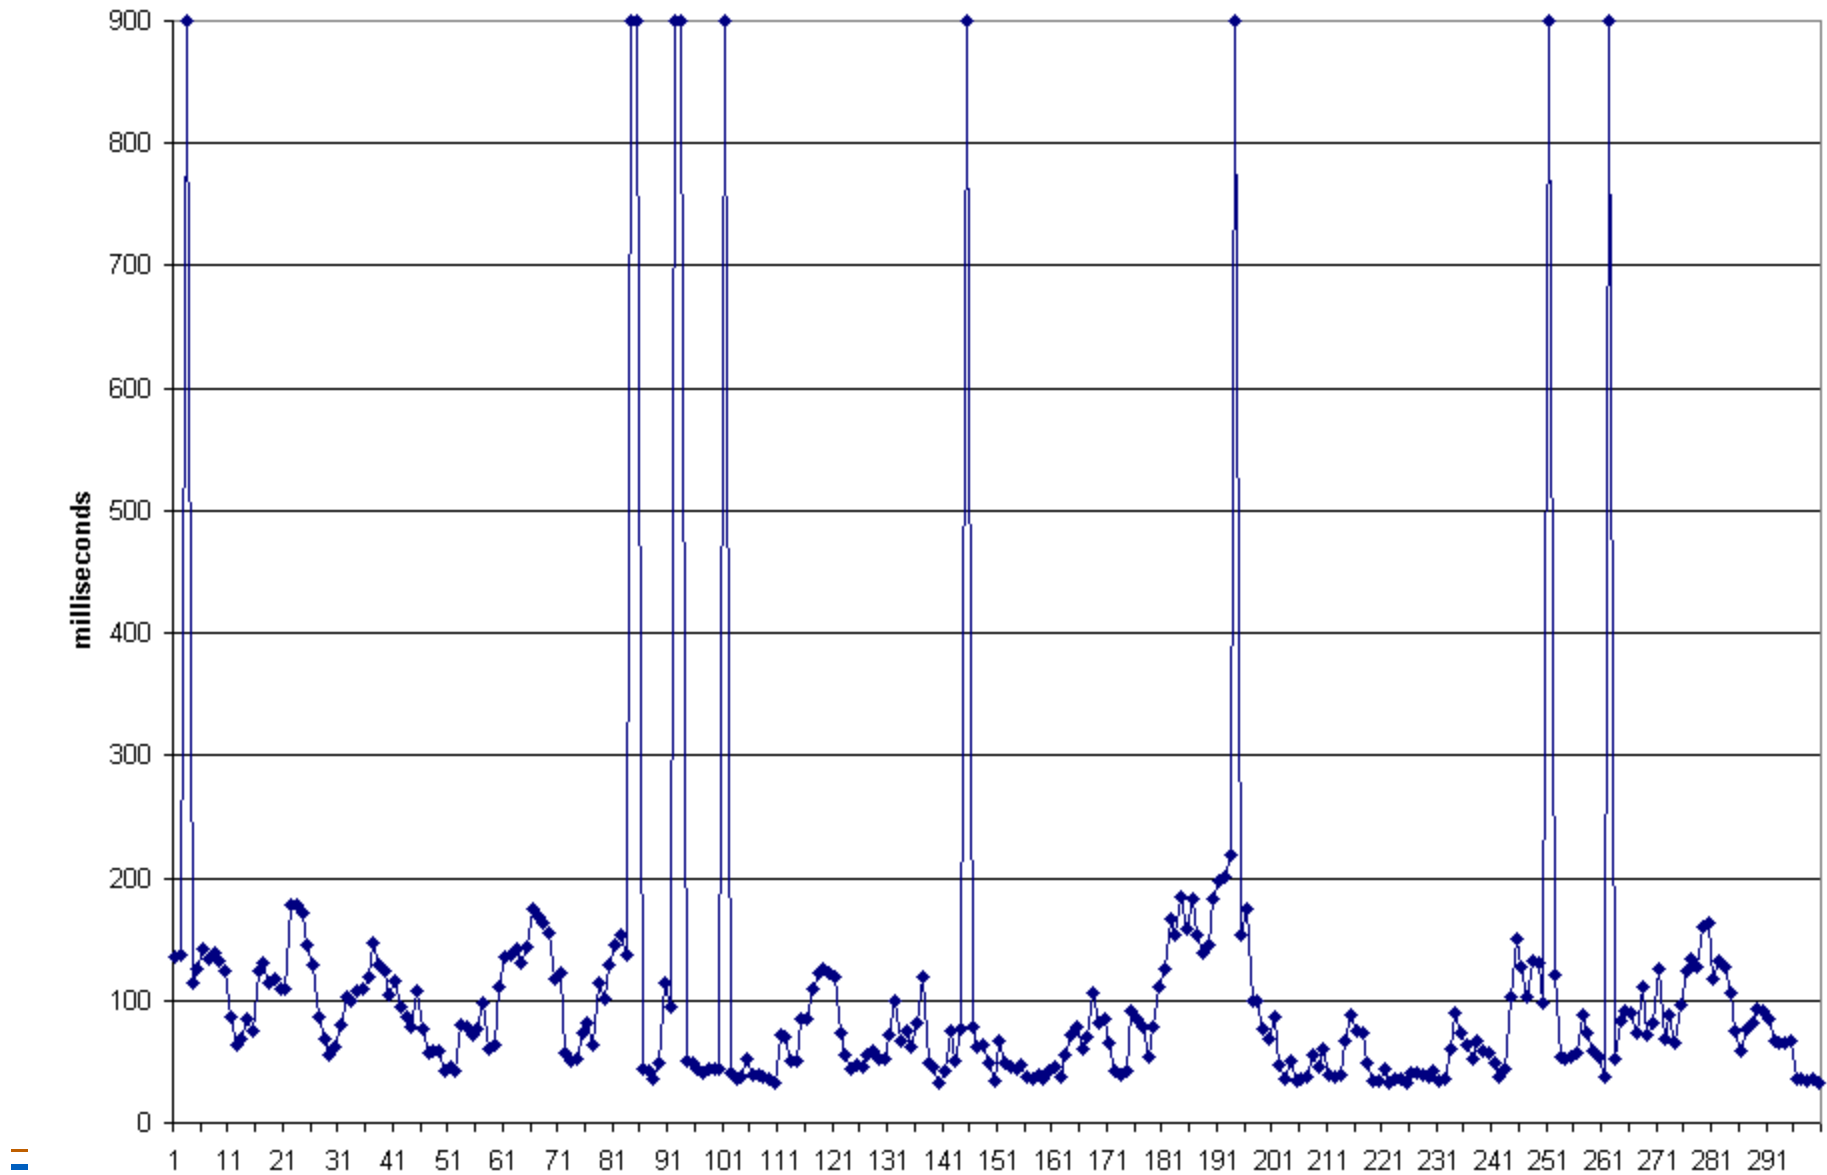

Wisc to Ottawa, 10-Dec-1998

14:00 to 17:00 UTC





Wisc to Ottawa, 10-Dec-1998, 14:00 to 17:00 UTC



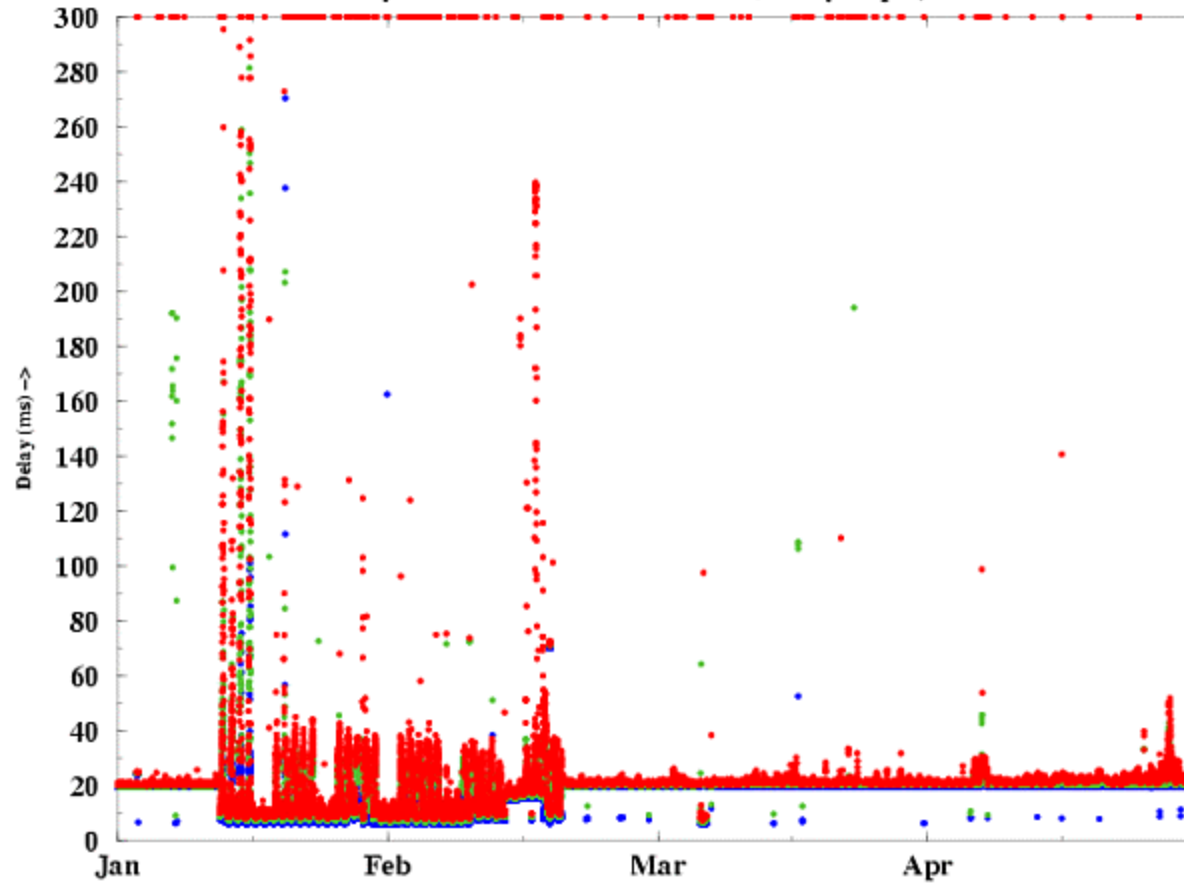
Wisc to Ottawa, 10-Dec-1998, 14:00:00 to 14:02:30 UTC



- min
- 50th percentile
- 90th percentile

Univ North Carolina-TI to Univ Pennsylvania-TI

Delay stats over 5 minute intervals from January - April, 1999





Problems

- As Universities obtain I2 connections, we measure fewer commodity paths (lose the “interesting” stuff).
 - Too much data, must present “interesting” material
 - n^2
 - Placing machines is difficult; GPS
-
-



Research Opportunities

- Correlation between this kind of network performance data and application performance
 - Correlation between active and passive
 - Is it possible to meaningfully summarize this data?
 - Comparing metrics? (Verify two implementations show same thing.)
-
-



Surveyor Futures

- General release of Java tools
 - New summary server
 - Add some kind of throughput measure (NSF, University admins)
 - Application focus (high-end/"I2"):
 - different active probes (back-to-back?)
 - correlation; feedback
-
-



IPPM Issues

- What is a valid throughput metric?
 - How many `auxilliary metrics' might you need to understand throughput?
 - As with Surveyor, comparing metrics
 - Multicast. Base metrics probably OK, what's the right sample and statistic?
 - QoS summaries
-
-



Wrap-up

- Surveyor: active tests of one-way delay, loss and routing
 - Surveyor infrastructure has good university coverage
 - Data available for research projects
 - We are going to focus on applications
 - ..could use some stats people
-
-



More info

- Surveyor project info
 - <http://www.advanced.org/surveyor/>
 - Access to plots
 - <http://ippm-db.advanced.org/plots/>
 - IETF IPPM WG
 - <http://www.advanced.org/IPPM/>
 - My email: matt@advanced.org
-
-