Estimating Bandwidth and Other Network Properties

Bruce A. Mah
bmah@cisco.com

(With Allen Downey, Wellesley College)

ISMA Winter 2000 Workshop

8 December 2000

Outline

- The goal and our approach
- Implementations
- Fun with linear regression

- Adaptive probing
- Selected results
- Problems we've seen
- netchar

Characterizing Networks at a Distance

- Measure network/link characteristics along a path
  - Bandwidth
  - Delay
  - Loss Rate
  - Queueing
- Assumption: No explicit support in routers or hosts
- Three similar implementations
  - *pathchar* by Van Jacobson (Packet Design)
  - *clink* by Allen Downey (Wellesley College)
  - *pchar* by Bruce Mah (Cisco Systems)

Approach

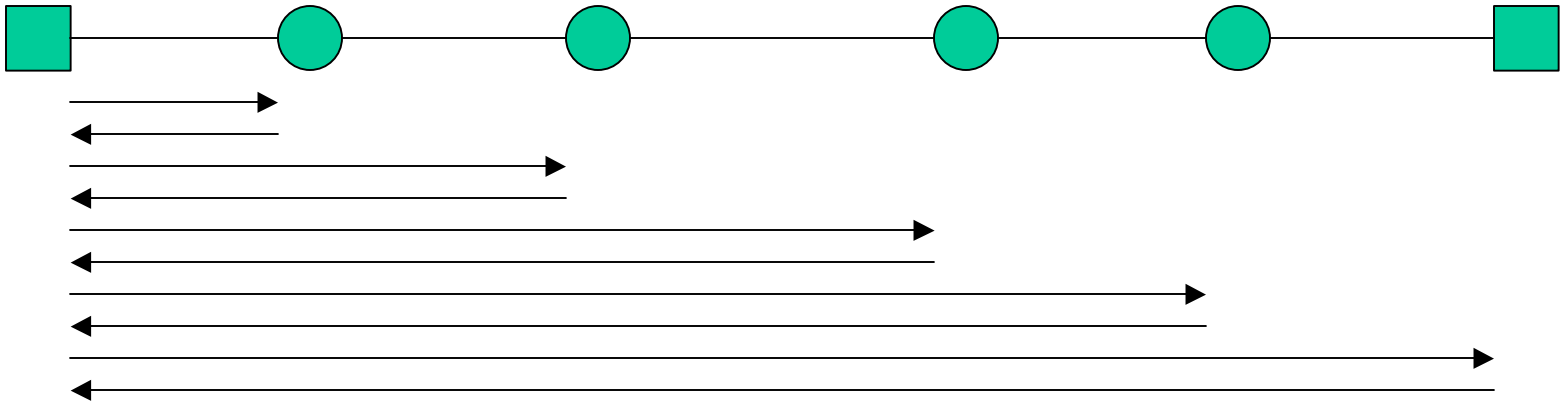

- Send packets along path, wait for responses (ICMP)
- Vary IP TTL to control how far into network packets can travel: gives links traversed (a la *traceroute*)
- Varying packet sizes gives bandwidth and latency
- Multiple repetitions give queuing and loss information

Model (One Hop, No Queueing)

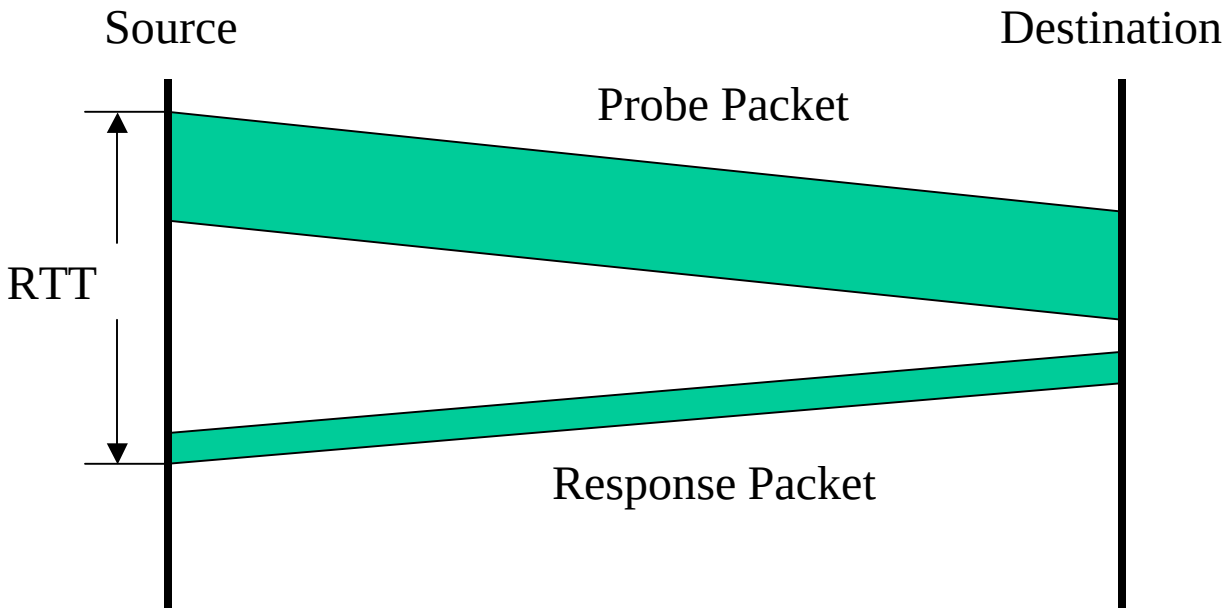

- Vary probe packet size, determine response time as a function of probe size

Response Time vs. Packet Size

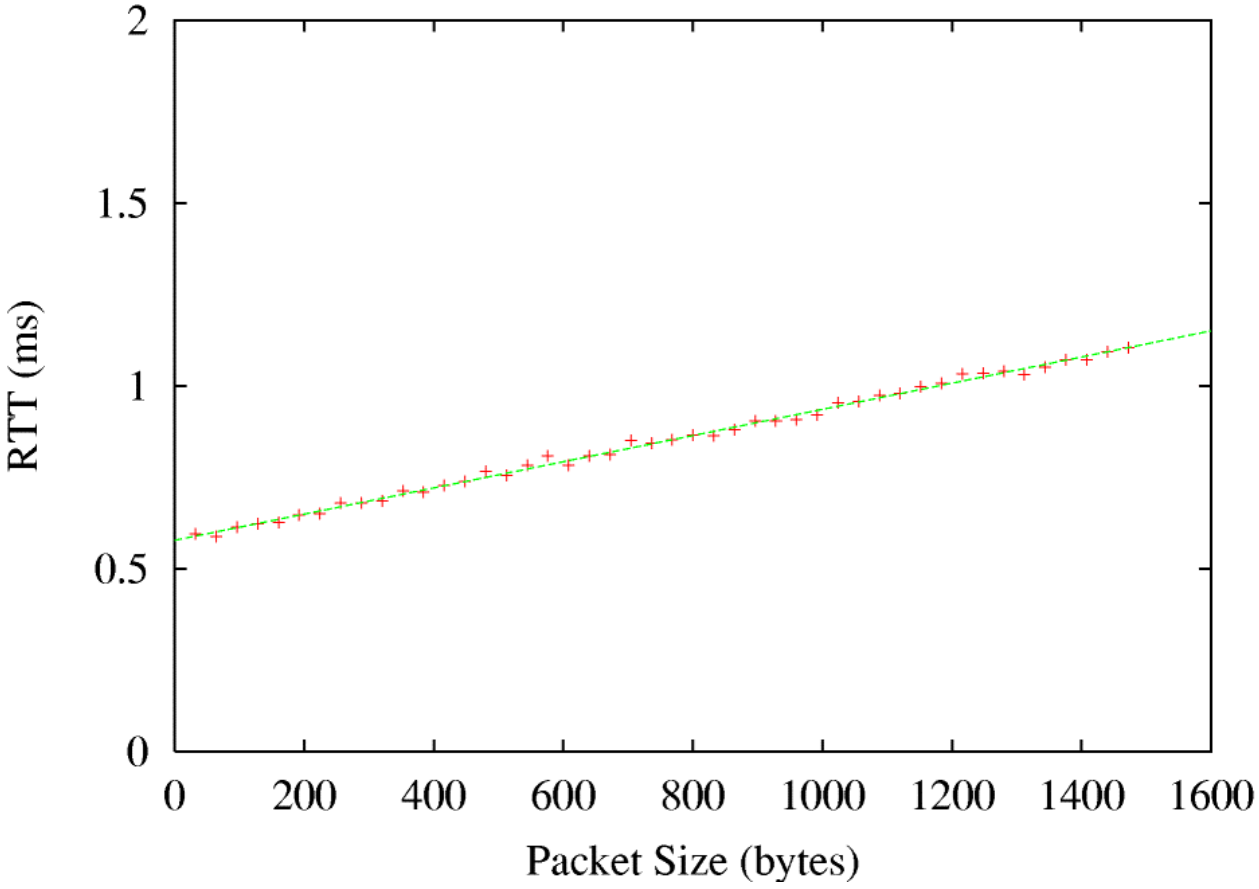

Fitting

- Slope of linear fit estimates “time per byte”
  - The inverse of bandwidth
- Intercept of linear fit estimates the round-trip time.
- What if there’s queueing?
  - Send lots of packets
  - Take only the minimum response time at each packet size
  - Hopefully that packet experienced no queueing

Minimum Filter and Linear Fit

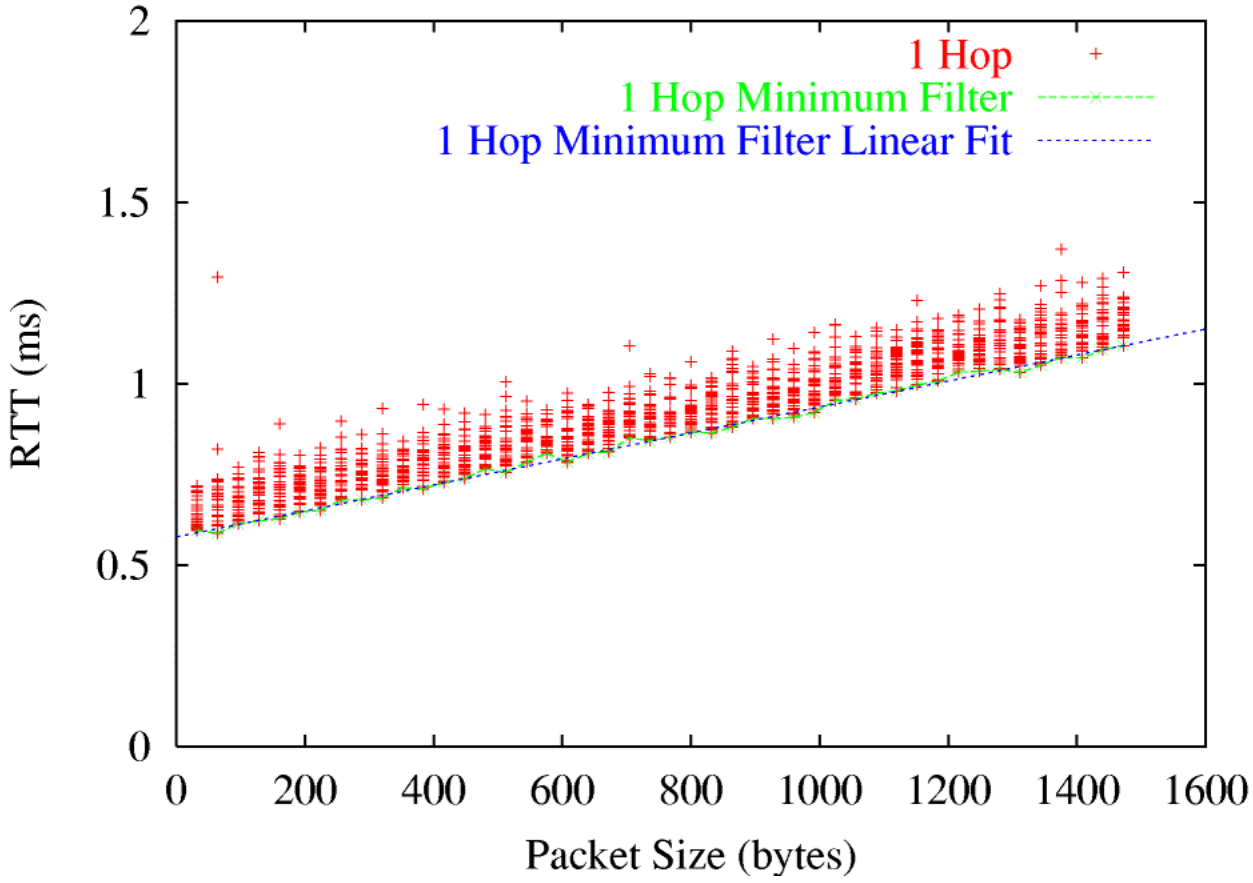

Multiple Queues in Series

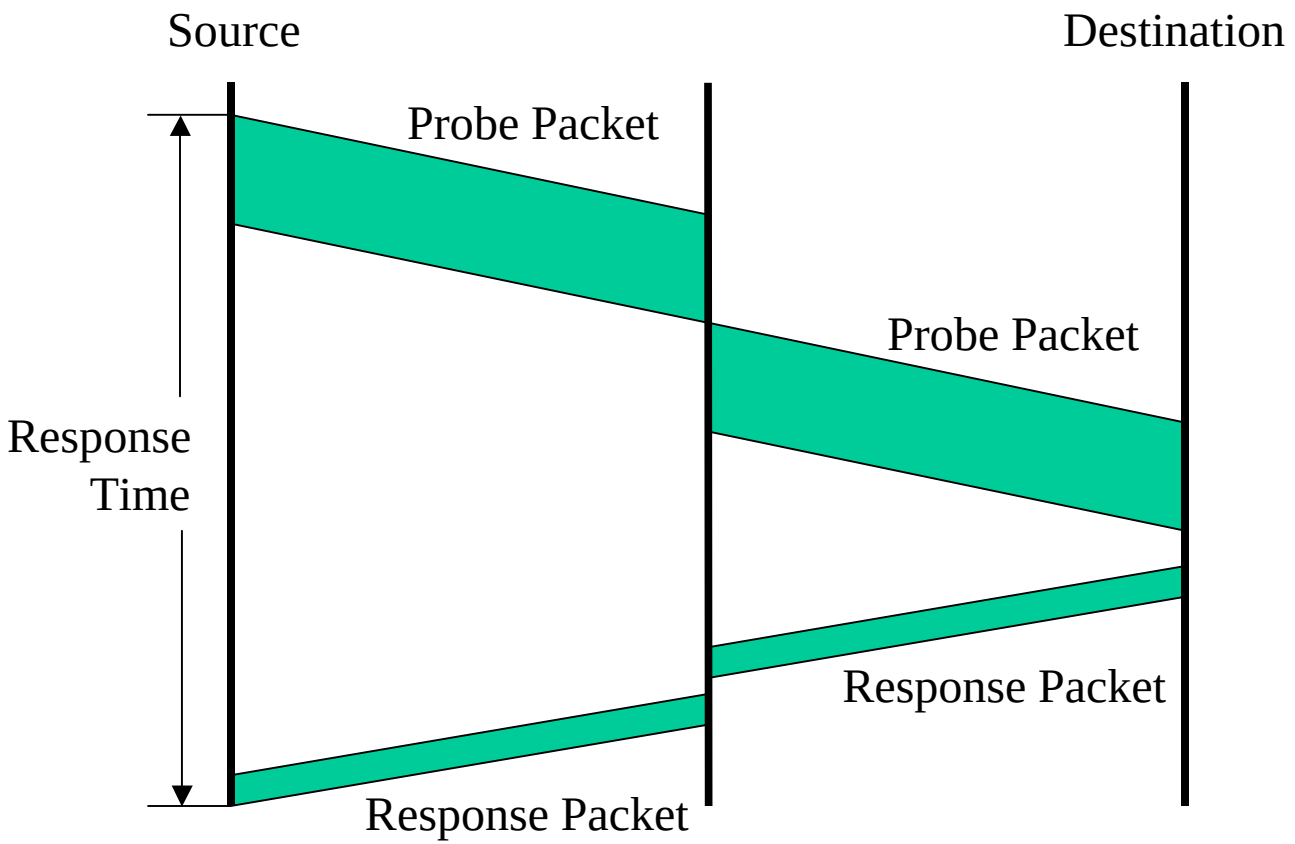

Multiple Queues in Series

- Works the same way as with one link, sort of
- Still send lots of packets, take minima, do linear fit
  - Vary packet size as before
  - Vary TTL to control how far into the network probes travel
- For each “partial path”:
  - Slope of linear fit estimates sum of time-per-byte
  - Intercept of linear fit estimates sum of round-trip times
- Do differencing to determine individual links

Minimum Filter on 11-Hop Path

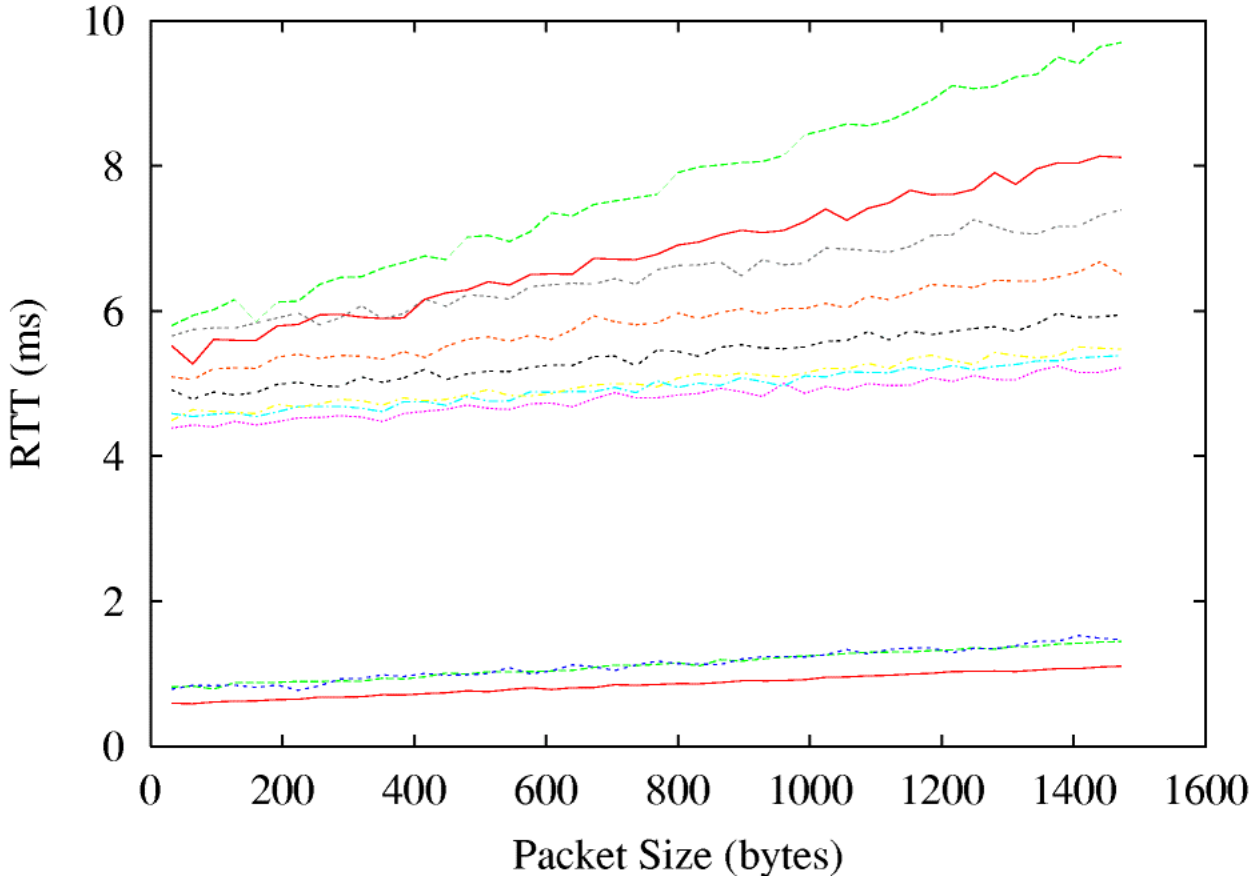

Linear Fit for 11-Hop Path

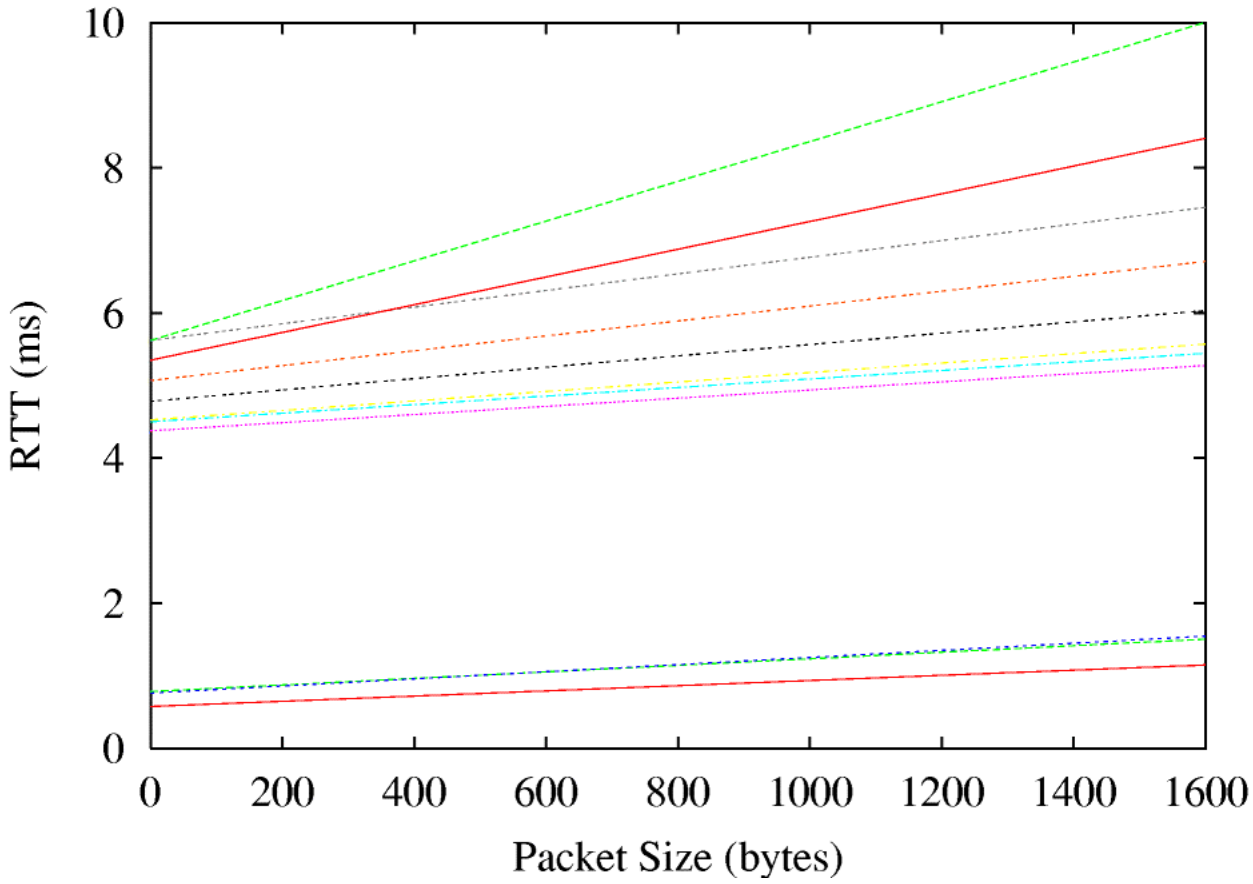

Summary of Approach

- Get (packet size, response time) points along partial path to hop *i*
- Apply min filter to response times
  - Remove variable queueing delay
  - Try to find packets that experienced no queueing
- Linear fit to response times
  - Yield “time per byte” and round-trip time to hop *i*
  - Possible because link bandwidths are constant
- Differencing
  - Subtract linear fit parameters from partial path to hop *i-1*
  - Result is time per byte and round-trip time for hop *i*

Implementations

- *pathchar* (Van Jacobson, 1997)
  - C on FreeBSD, Linux, NetBSD/alpha, OSF/1, Solaris
- *clink* (Allen Downey, 1999)
  - C on Linux
  - Kernel-level timing, adaptive probing
- *pchar* (Bruce Mah, 1999)
  - C++ on *BSD, Solaris, Linux, OSF/1, Solaris, IRIX
  - Working on IOS implementation
  - IPv6 support, alternate linear regression models

Implementation Lessons Learned from *pchar*

- Multi-platform support is hard
  - Tru64 timing resolution
  - Solaris `-lnsl`
  - Linux *send*(2) and `SOCK_DGRAM`
  - FreeBSD `<sys/param.h>`
  - IPv6 API changes
- Tweakability is good (but messy)
  - Multiple types of probe packets (UDP, ICMP, maybe TCP)
  - DiffServ CodePoint settable
  - Conservative defaults (people won't read manual anyways)

Alternate Linear Regression Algorithms

- Least Sum of Squares
  - Pick the linear fit that minimizes the sum of the errors
  - Simple and easy to understand
  - Estimates affected by outliers
  - Needs floating point operations

Alternate Linear Fit Algorithms

- Nonparametric method
  - Compute all slopes between all pairs of points, take median slope
  - Based on ranks, slightly resistant to outliers
  - Gives some improvement
- Least Median of Squares
  - Compute lines between all pairs of points, pick the one that minimizes the median error
  - Robust: Up to half of dataset can contain outliers
  - Computationally expensive

Comparing Linear Fit Algorithms

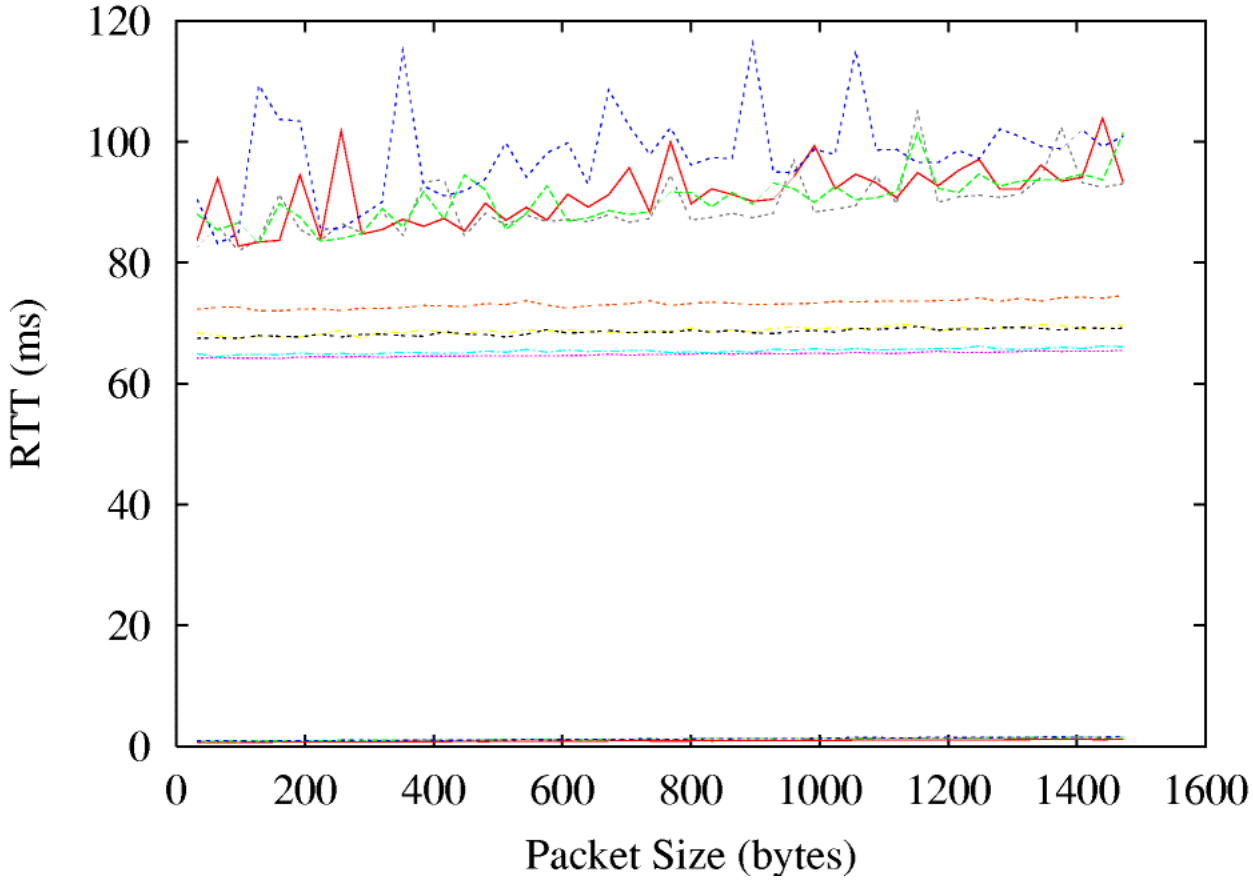

Linear Fits for One Partial Path

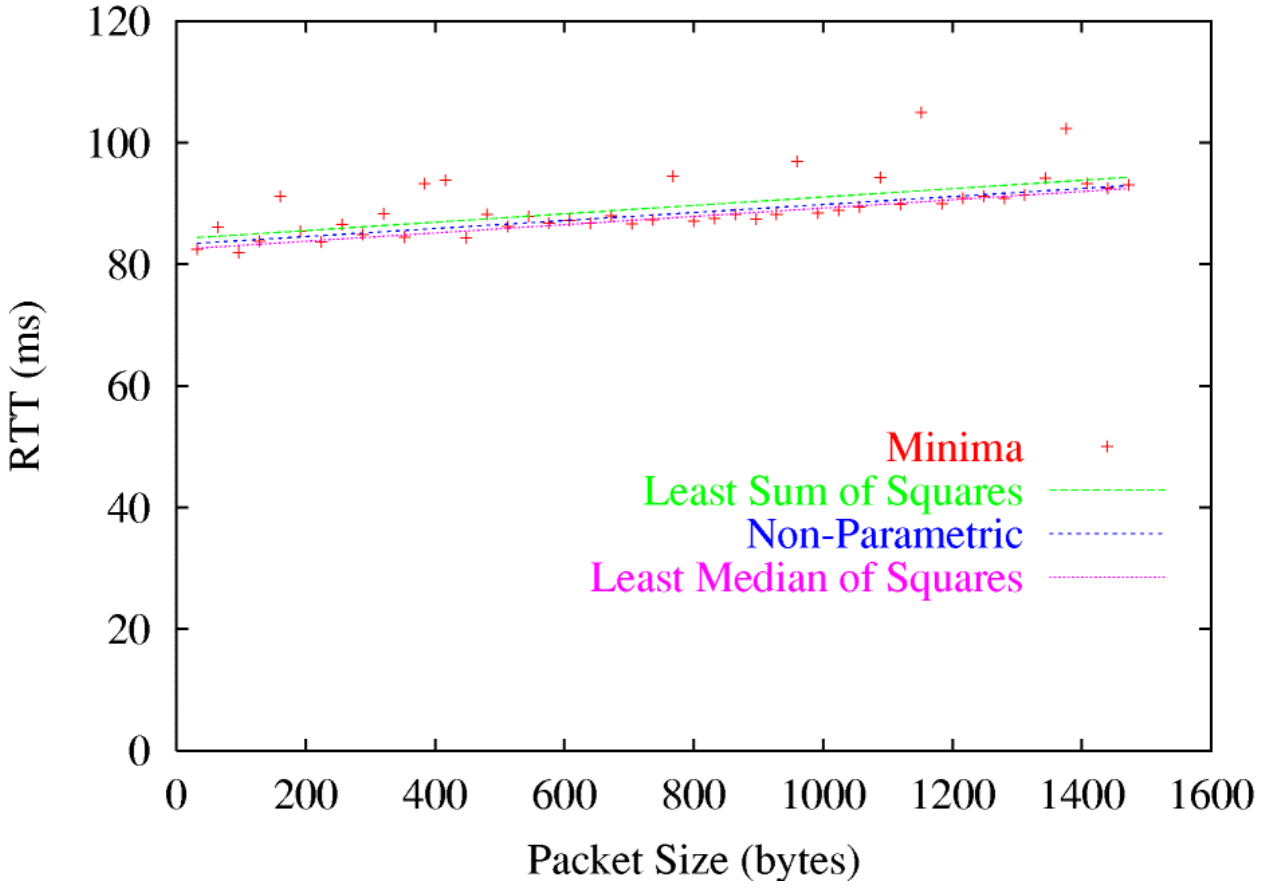

Linear Fits for One Partial Path

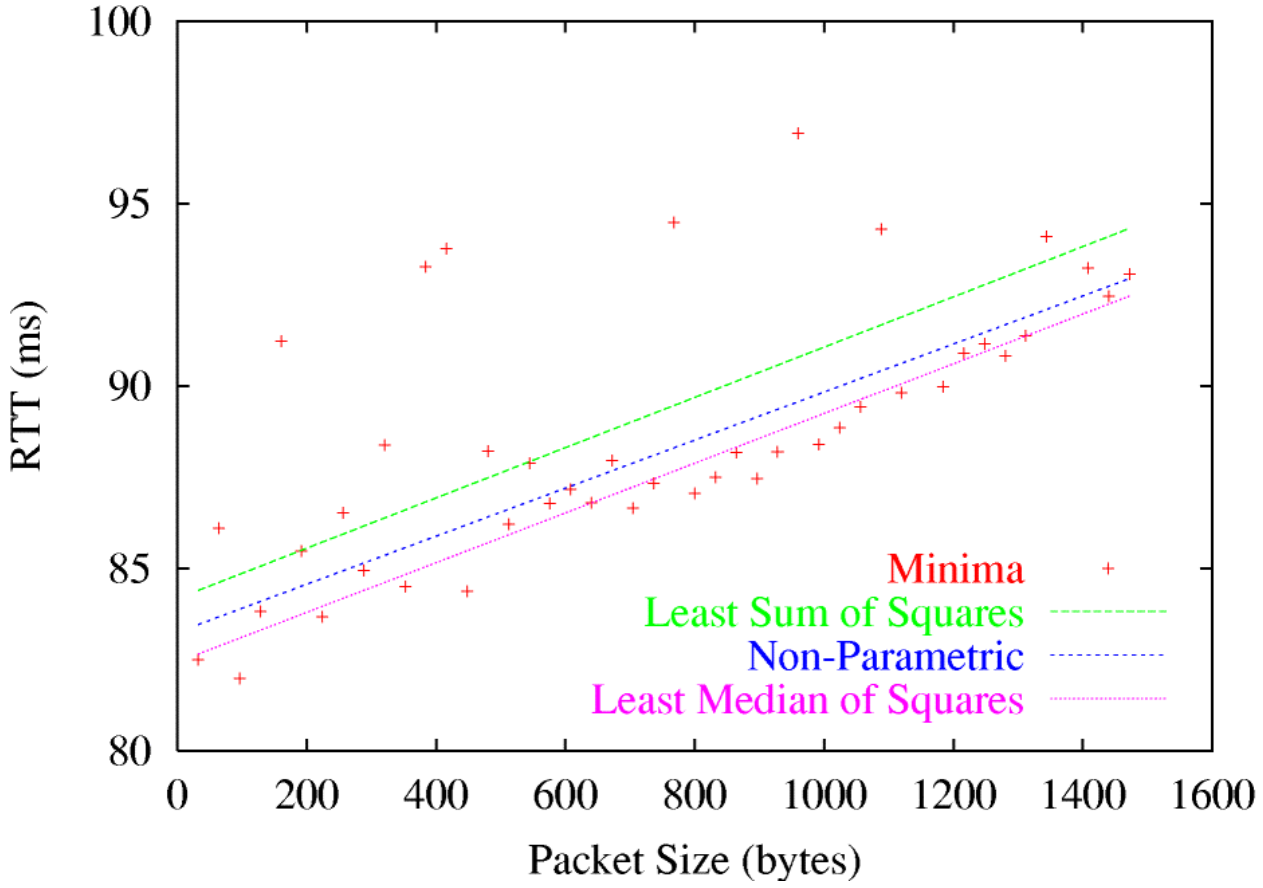

Effect on Bandwidth Estimation

	Hop 9	Hop 10	Hop 11
Least Sum of Squares	1451 Kbps	53997 Kbps	-47836 Kbps
Non-Parametric	1547 Kbps	5286 Kbps	-8562 Kbps
Least Median of Squares	1527 Kbps	12175 Kbps	13360 Kbps

For More Info

- *pathchar*
  - ftp://ftp.ee.lbl.gov/pathchar/
- *clink*
  - http://rocky.wellesley.edu/downey/clink/
- *pchar*
  - http://www.employees.org/~bmah/Software/pchar/
- Related information
  - http://www.caida.org/analysis/performance/bandwidth/