

Pchar: Child of Pathchar



Bruce A. Mah

bmah@ca.sandia.gov

Security and Networking Research Department (8910)

Sandia National Laboratories / California

21 July 1999

What can we learn about the network?

- Measure network characteristics
 - Bandwidth
 - Delay
 - Loss Rate
 - End-to-end or hop-by-hop
- Constraints
 - Algorithms execute on endhosts only
 - Avoid disrupting existing network traffic
 - Avoid depending on specific network protocols

Existing Work

- Ping (Mike Muuss, ARL)
 - Round-trip times only
- SNMP
 - Detailed information at intermediate network nodes
 - Requires administrative access on routers
- ttcp, netperf, et al.
 - TCP-specific
 - Large TCP bulk transfers can disrupt network


Existing Work, Continued

- Traceroute (Van Jacobson)
 - Like ping, but uses IP TTL to discover routing information, round-trip times
- Pathchar (Van Jacobson)
 - Try to characterize individual network hops
 - Delay
 - Available bandwidth
 - Queuing
 - Loss

pchar: Child of Pathchar

- An independent implementation of pathchar
- Based on pathchar algorithms
 - `ftp://ftp.ee.lbl.gov/pathchar/msri-talk.ps.gz`
- Designed as a framework for examining different measurement and analysis algorithms

pchar/pathchar measurements



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

pchar Output

Path to cs.berkeley.edu (169.229.60.56)

Packet size increments by 32 to 1500

46 test(s) per repetition

32 repetition(s) per hop

0:

Partial loss: 0 / 1472 (0%)

Partial char: rtt = 0.653832 ms, r2 = 0.992239

Hop char: rtt = 0.653832 ms, bw = 21678 Kbps, r2 = 0.992239

Partial queueing: avg = 0.000194 ms (526 bytes)

1: 146.246.243.254 (con243.ca.sandia.gov)

Partial loss: 0 / 1472 (0%)

Partial char: rtt = 0.871580 ms, r2 = 0.992828

Hop char: rtt = 0.217748 ms, bw = 89483 Kbps, r2 = 0.787175

Partial queueing: avg = 0.000257 ms (560 bytes)

2: 146.246.250.254 (snl-outnet.ca.sandia.gov)

Partial loss: 0 / 1472 (0%)

Partial char: rtt = 0.783626 ms, r2 = 0.982747


Hop char: rtt = -0.087954 ms, bw = 172435 Kbps, r2 = 0.296921

Partial queueing: avg = 0.019788 ms (39197 bytes)

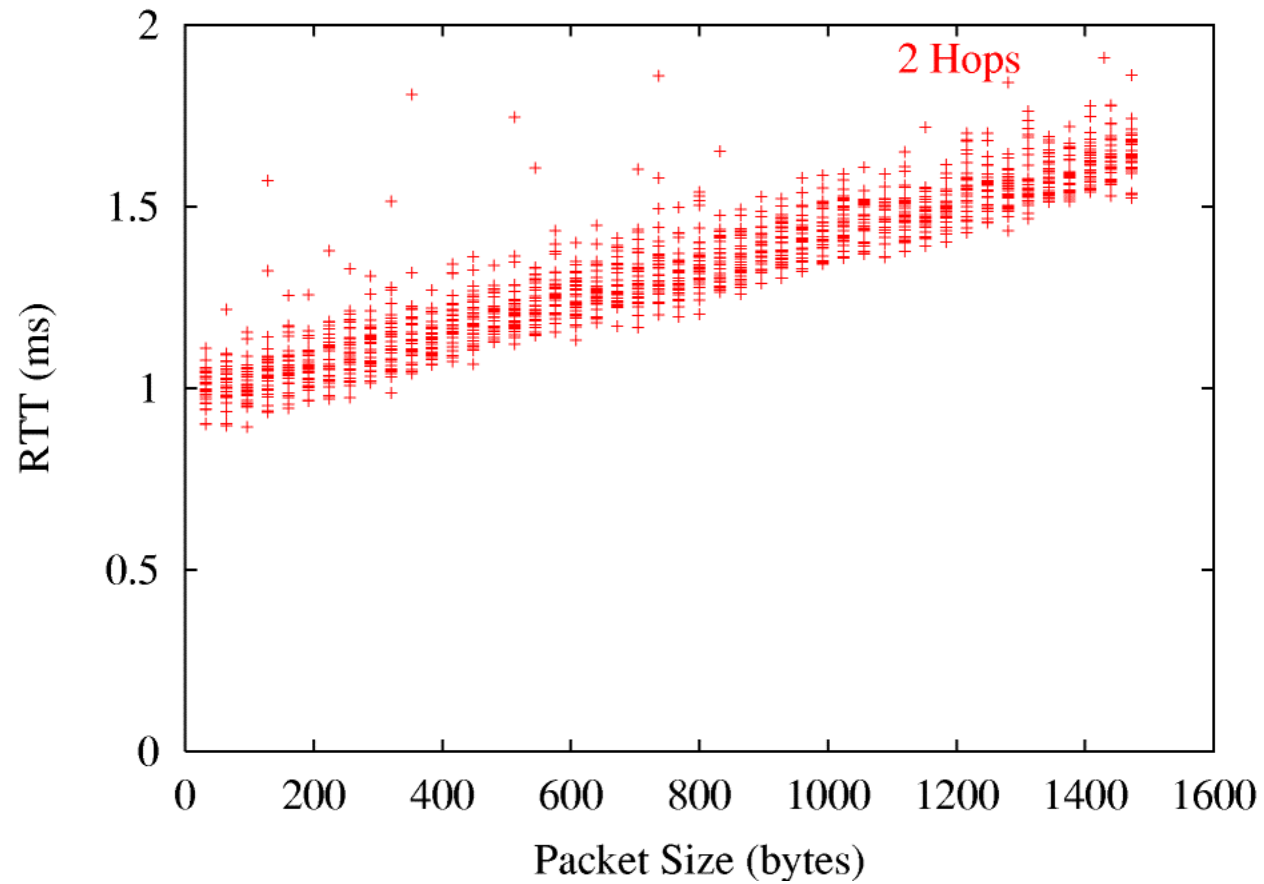
pchar Output, Continued

```
.  
.
10: 128.32.120.181 (f1-0.inr-181-soda.Berkeley.EDU)
    Partial loss: 0 / 1472 (0%)
    Partial char: rtt = 5.956725 ms, r2 = 0.949262
    Hop char: rtt = -0.071023 ms, bw = 108351 Kbps, r2 = 0.044718
    Partial queueing: avg = 0.001376 ms (1112 bytes)
11: 128.32.40.202 (128.32.40.202)
    Partial loss: 0 / 1472 (0%)
    Partial char: rtt = 6.673852 ms, r2 = 0.986403
    Hop char: rtt = 0.717128 ms, bw = 7960 Kbps, r2 = 0.892486
    Partial queueing: avg = 0.001291 ms (575 bytes)
12: 169.229.60.56 (cs2.CS.Berkeley.EDU)
    Path length:      12 hops
    Path char:        rtt = 6.673852 ms, r2 = 0.986403
    Path bottleneck: 7960.846285 Kbps
    Path pipe:        6641 bytes
    Path queueing:    average = 0.001291 ms (575 bytes)
```

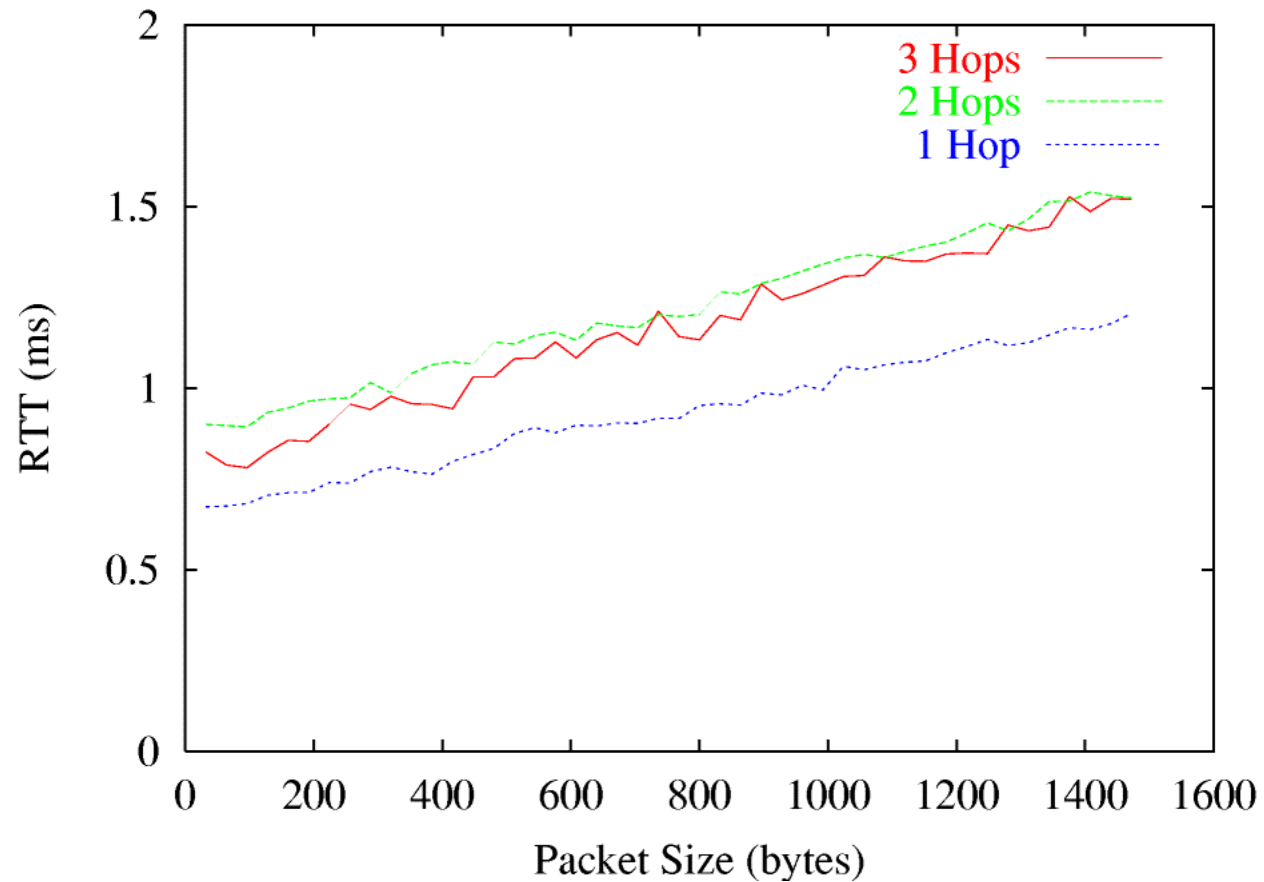

One-hop Packet Response Times



Two-Hop Packet Response Times



Partial Path and Per-Hop Characteristics



Future Work


- New or improved measurement algorithms
 - Reduce impact on the network
 - Reduce measurement time
 - Produce useful results over switched networks
- New or improved analysis algorithms
 - Reduce the effects of experimental errors
 - Adaptive analysis and measurement
- Programming interface (API) for applications

Where to get it?

- **Web:**
 - `http://www.ca.sandia.gov/~bmah/Software/pchar/`
- **Email:**
 - `bmah@ca.sandia.gov`
- **Platforms currently supported:**
 - FreeBSD, Solaris, Linux, IRIX

Finis

Diversity in Networks is Increasing



The Internet is becoming increasingly heterogeneous.

Who Cares?

- Applications
 - Network video: set coding parameters
 - Collaboration applications: set data resolution
 - Distributed resource managers: Find network resources
- Network managers, architects, and researchers
 - How well is the network working under current load?