Segmentation of Internet Paths for Capacity Estimation

Matthew Luckie (PhD Student)
Tony McGregor (Supervisor)

WAND Group, University of Waikato
NLANR Measurement and Network Analysis

mjl@luckie.org.nz
tonym@cs.waikato.ac.nz

Introduction

- We are talking about capacity estimation of individual hops in an end-to-end path

- If we had a method to segment a path into hops, what problems might it solve, and what might remain?

Current Techniques

- pathchar - Jacobson, original TTL method
- clink and pchar
- Downey suggests statistical methods to detect convergence on a link capacity faster

What is hard?

- Capacity estimation of each hop in a path is hard to do with the TTL method
  - additive noise from prior links brings doubt
  - can only estimate the forward path
  - L2 store-and-forward devices interfere
  - each hop is probed with a separate series of packets

A goal

- A method to segment a path into links could be useful
  - remove additive error
  - measure capacity of each link in a long path
- Identified as future research by Jacobson in MSRI talk

Other goals

- Measure the capacity of hops on the reverse path
- Be kind to the network
- Measure capacities of hops correctly when Layer-2 Store and Forward devices are in the path

A method

- Use timestamps inserted at each hop [modify the forwarding path, IPMP]
- Use the packet-tailgating technique [nettimer] to encourage a packet pair to queue together through a network
- (extremely simplified) Estimate the capacity of each hop using the time difference between the first packet and the second packet

IPMP in a nutshell

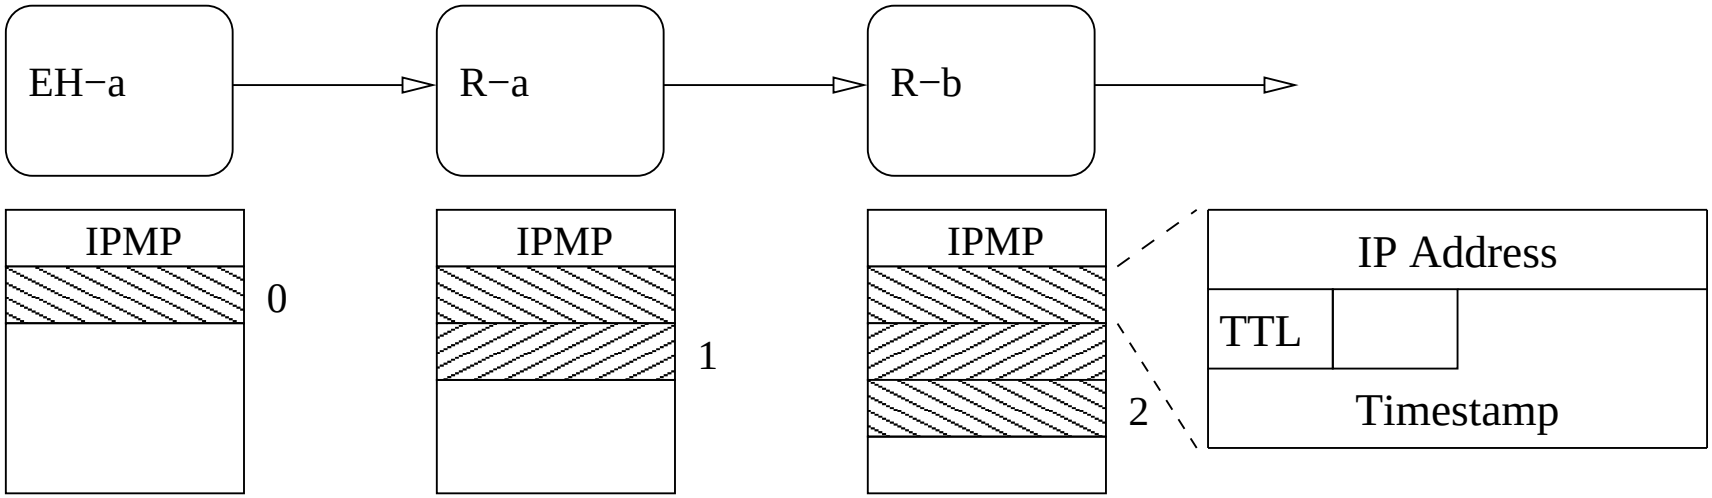

A specially-marked packet passes through a network,
collecting timestamps in "path records" as it goes

A method

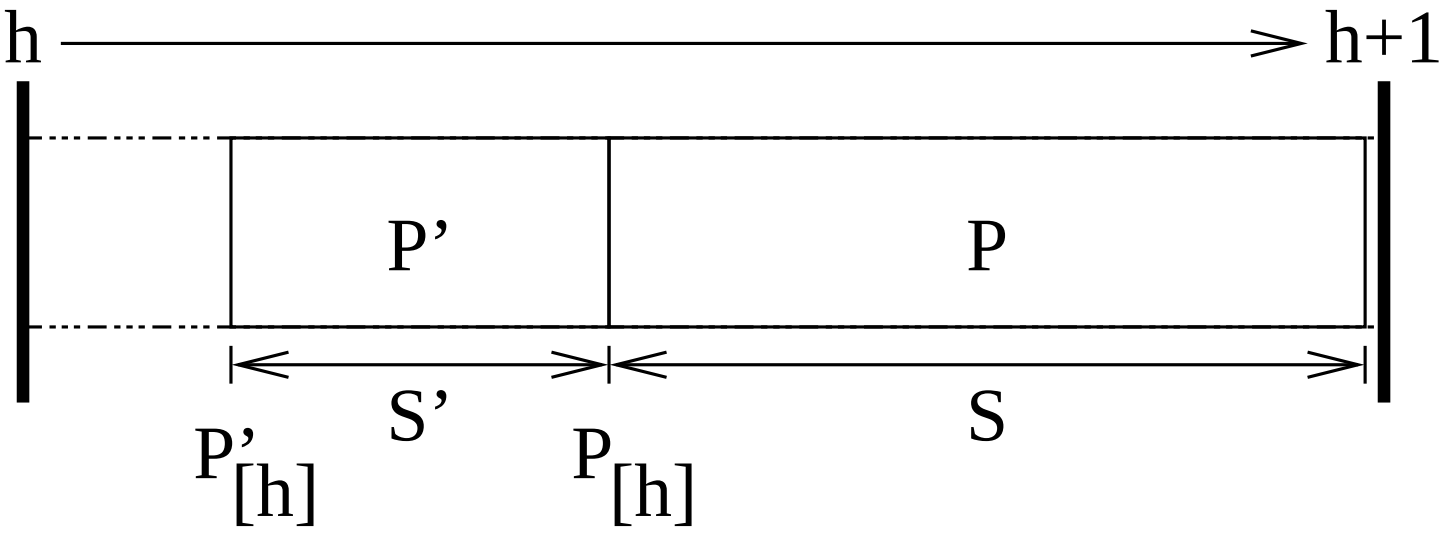

$$C = \min_{h=0\ldots H} \left\{ \frac{S'}{P_{[h]} - P'_{[h]}} \right\}$$

What does this buy us?

- the behaviour of each segment is confined to that segment unless the packets become separated [more on this shortly]
- measures the forward and reverse paths
- measures the capacity of hops with L2 store-and-forward devices correctly as it uses packet-pairs

Experiments

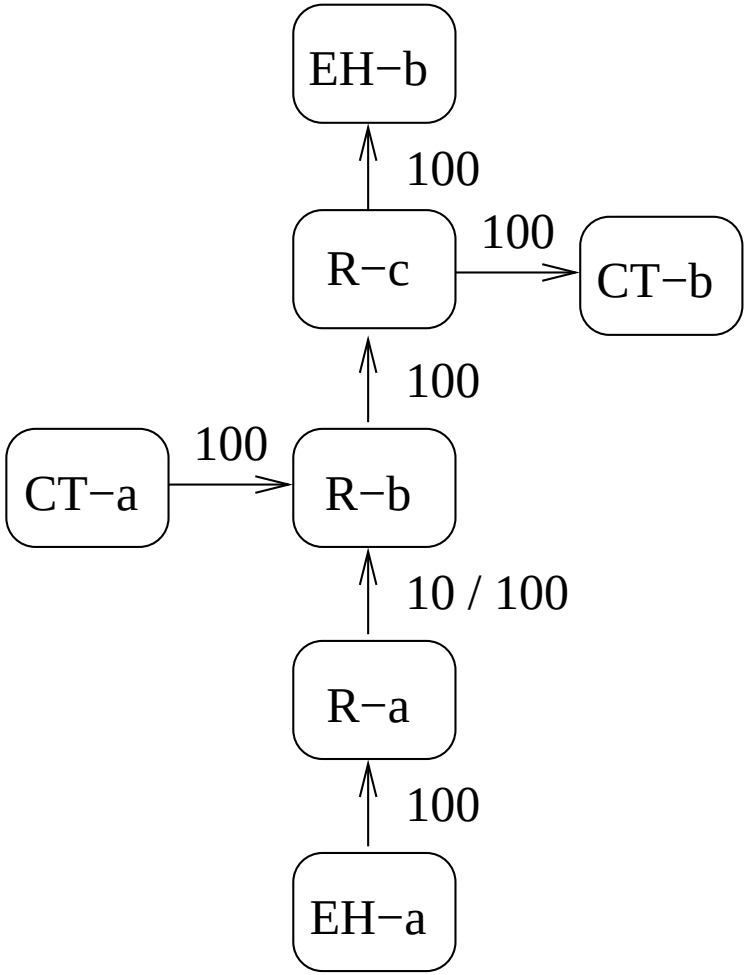

EH = Echo Host
R = Router
CT = Cross Traffic Host

Linux 2.4.20 + PPS Kit +
IPMP Kernel Implementation +
Crossover Cables

Arrows indicate direction
from source to destination

R−a to R−b is forced 10mbps
with a 10mbps 3COM hub
in some experiments

An aside: ctft

- Wrote cross-traffic-from-trace (ctft)
  - using an Auckland trace
  - given a probability of a packet of a given size
  - and the probability of a time-to-next packet
  - send combinations of size and delay randomly that fits the profile

Experiments

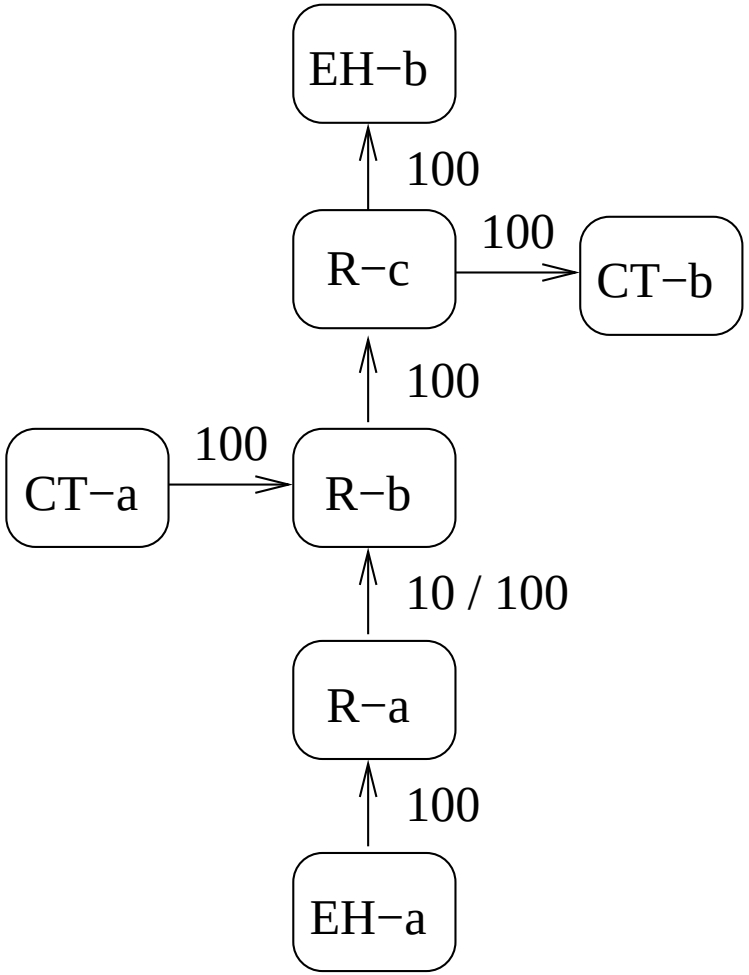

EH = Echo Host
R = Router
CT = Cross Traffic Host

Linux 2.4.20 + PPS Kit +
IPMP Kernel Implementation +
Crossover Cables

Arrows indicate direction
from source to destination

R−a to R−b is forced 10mbps
with a 10mbps 3COM hub
in some experiments

100mbps, no cross traffic, Hop 4

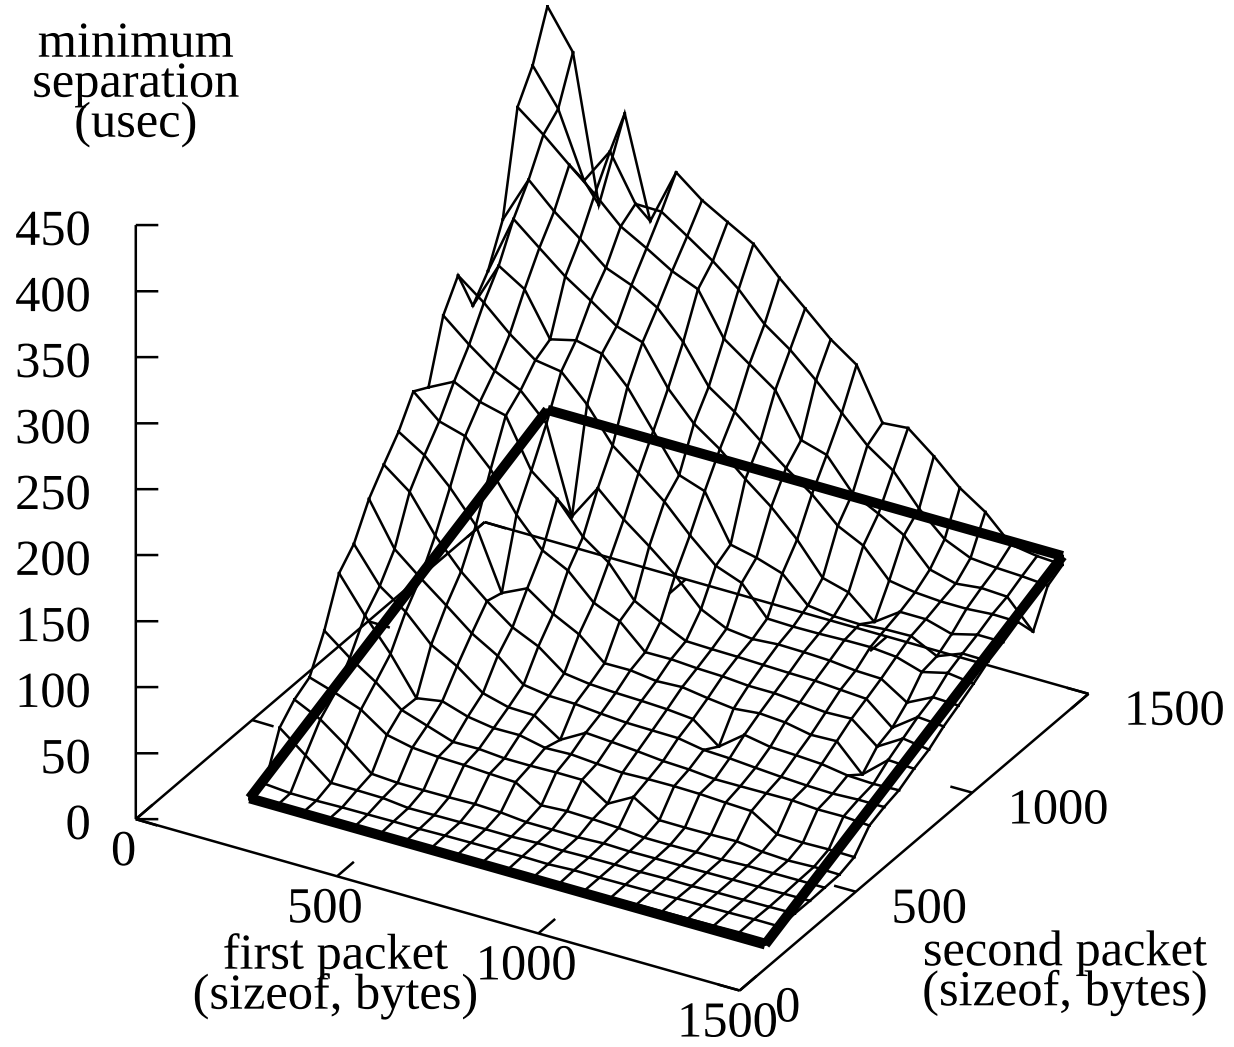

10mbps, cross traffic

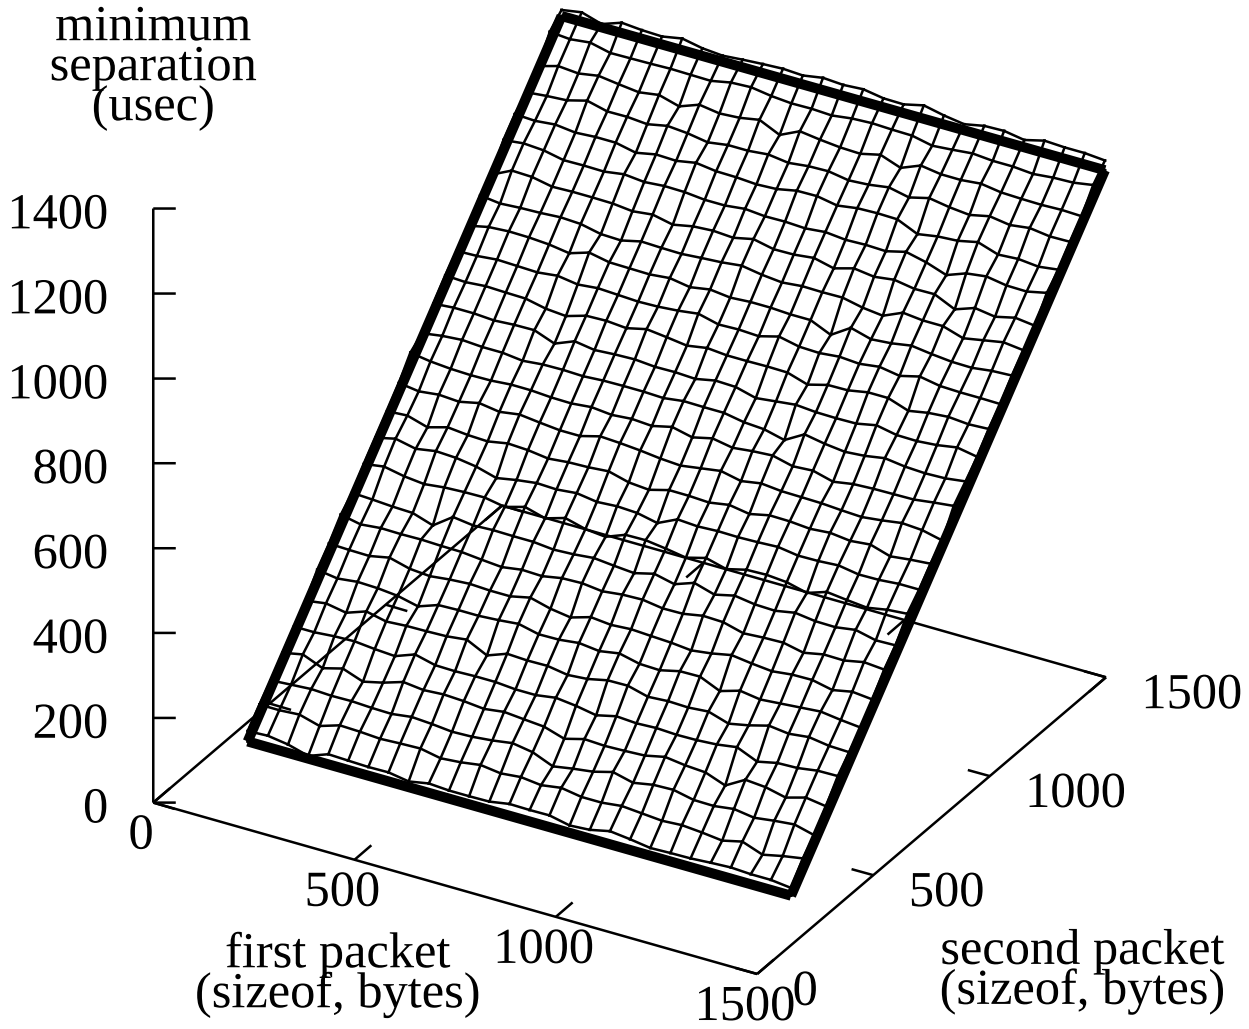

- What problems still remain?

- given an egress link n times faster than the ingress link to a router, then S’ must be no larger than S/n if the pair is to remain together

  - i.e. measuring a 100mbps hop immediately after a 10mbps hop is problematic; cross traffic can help here, though

  - limited number of path records in a small packet

10->100mbps, cross traffic

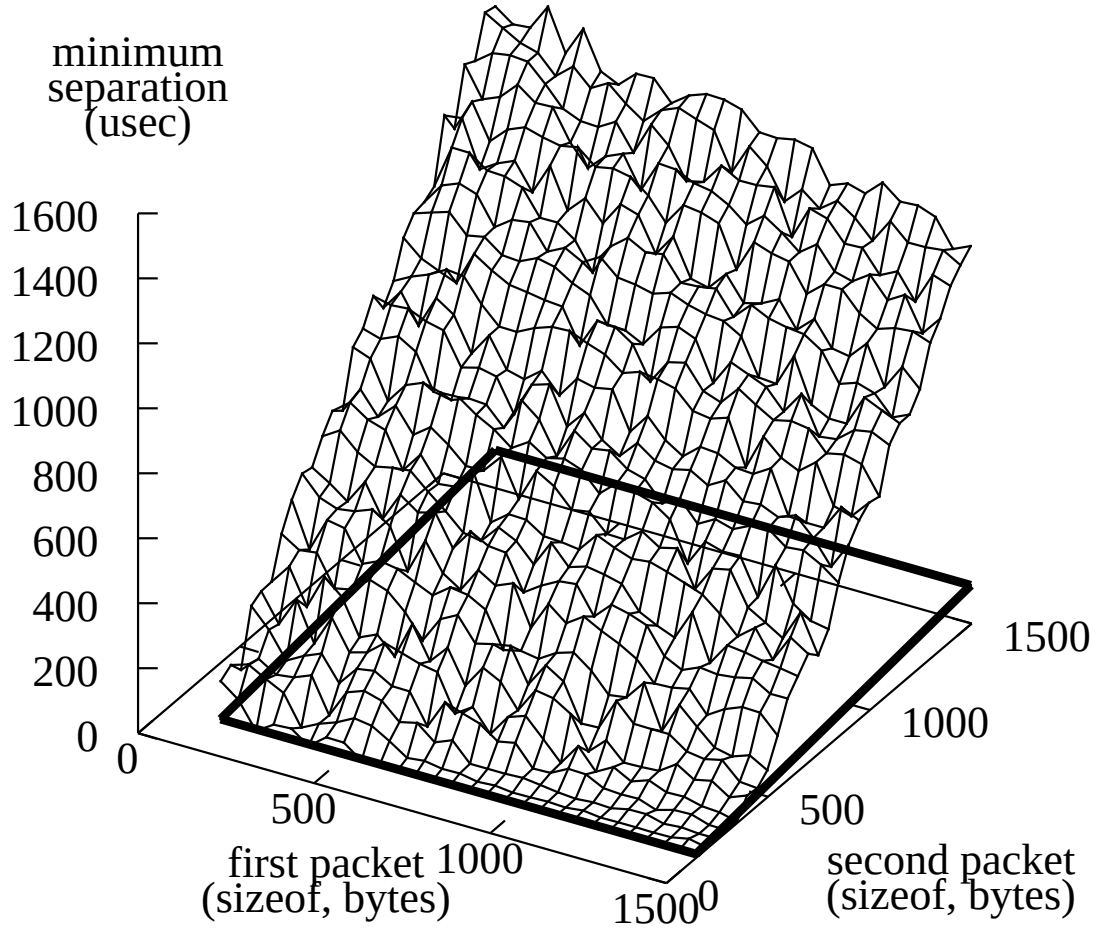

Disadvantages

- need the protocol deployed
- not everywhere before it becomes useful for capacity estimation, though
- can estimate the capacity of the path through two points by looking for the minimum time for a pair through the points and the minimum separation of said pair

Future work

- work at estimating congestion / queue length at intermediate nodes
- looking for Layer 2 capacities
  - cross traffic (CT) is likely to be either a 40 / 576 / 1500 byte packet
  - look for CT separating the pair (cross traffic) at various sizes and infer what the capacity of the link was that caused the separation

Conclusion

- Presented a method to segment a path into hops for capacity estimation
- Discussed the advantages ...
- ... and caveats

Segmentation of Internet Paths for Capacity Estimation

Matthew Luckie (PhD Student)
Tony McGregor (Supervisor)

WAND Group, University of Waikato
NLANR Measurement and Network Analysis

mjl@luckie.org.nz
tonym@cs.waikato.ac.nz

802.11b Network [an aside]