# A Measurement Study of Available Bandwidth Estimation Tools

Dina Katabi

with Jacob Strauss & Frans Kaashoek

MIT - CSAIL

## What is available bandwidth?

- ABW is the left-over capacity along an Internet path

## Why is it useful?

1. Route Selection in Overlays
2. Traffic Engineering
3. QoS Verification

# Why is Measuring ABW Hard?

- Tools must limit probing rate to not interfere with other traffic
- Available bandwidth changes over time → measurements must be quick
- Tool validation is hard because we don't know the true available bandwidth

# Our Contributions

- First large scale Internet measurement study of available bandwidth tools
  - 400 paths
  - PlanetLab & Ron Testbeds
  - A variety of RTT, capacity, link technology, load
- Spruce, a new light-weight tool for measuring available bandwidth
- Reveal characteristics of Pathload, IGI, and Spruce

## Three Tools

1. Pathload
2. IGI
3. Spruce

# Pathload

Idea:

1. Send probes with increasing rate
2. Look for the transition where you begin to observe queuing delay
3. ABW is the probe rate at transition

Caveat:

Finding the transition is difficult because of transient queues, which though not a sign of congestion do increase delay

## Three Tools

1. Pathload
2. IGI
3. Spruce

# IGI

Idea

1. Measures cross traffic rate
2. Send a train of probes with constant spacing, and use output gaps to measure cross traffic

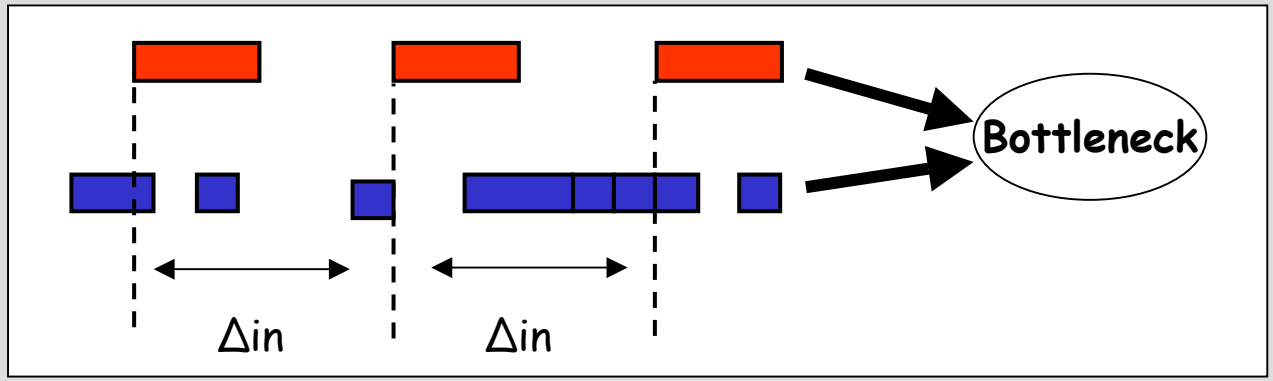


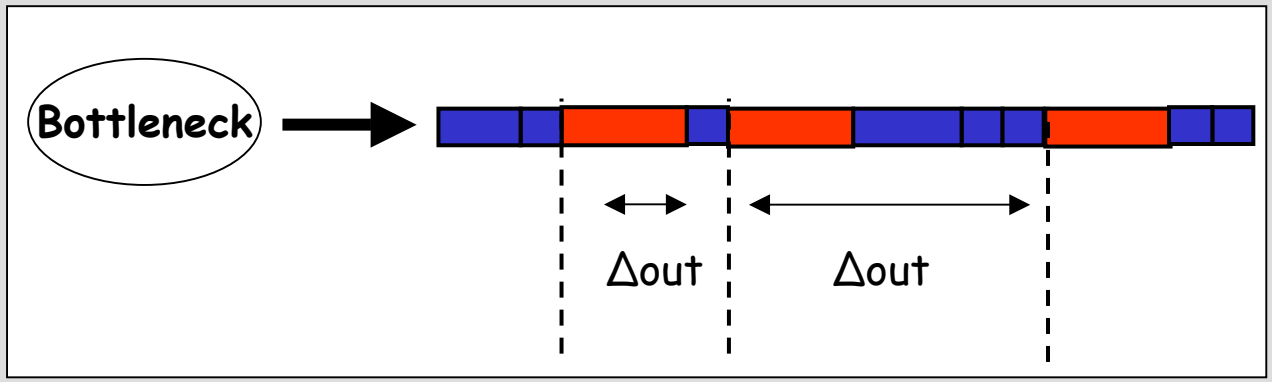

## Three Tools

1. Pathload
2. IGI
3. Spruce

# IGI

Idea

1. Measures cross traffic rate
2. Send a train of probes with constant spacing, and use output gaps to measure cross traffic

Caveat:

1. Gaps between pairs of packets in the train are correlated
2. IGI ignores the deceasing gaps causing over-estimation of ABW

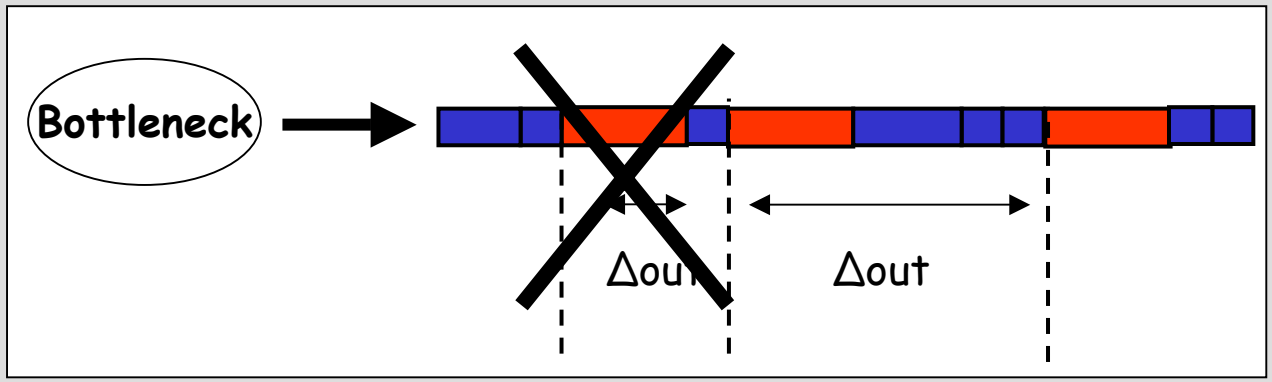

# Three Tools

1. Pathload
2. IGI
3. Spruce

# Our Tool: Spruce

Idea

- Also uses a gap model but doesn't send a train of packets → Rather send a Poisson process of packet pairs
  - For correctness, set intra-pair gap Δin to probe size / capacity
  - To prevent intrusiveness, choose a large inter-pair gap

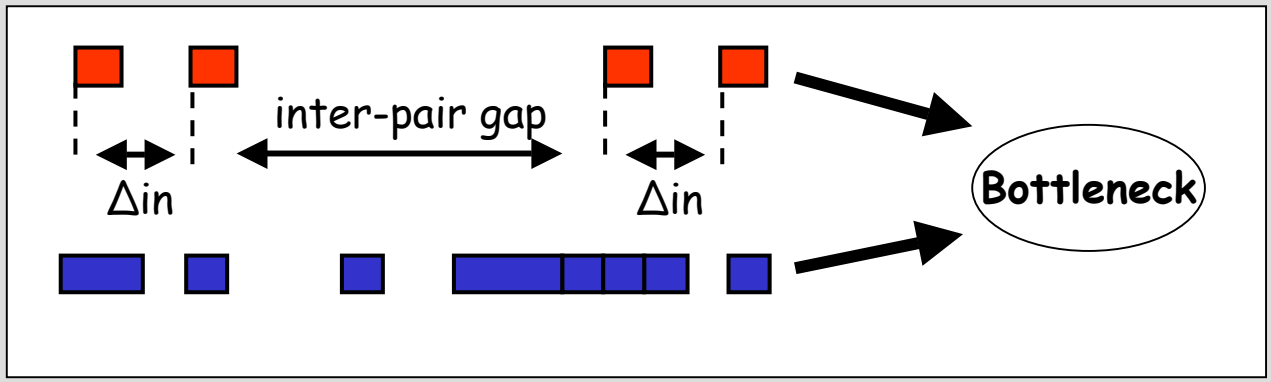

## Three Tools

1. Pathload
2. IGI
3. Spruce

# Our Tool: Spruce

Idea

- Also uses a gap model but doesn't send a train of packets → Rather send a Poisson process of packet pairs
  - For correctness, set intra-pair gap Δin to probe size / capacity
  - To prevent intrusiveness, choose a large inter-pair gap
- Compute ABW from

$$Cross\ \ Traffic = C \times \frac{\Delta_{out} - \Delta_{in}}{\Delta_{in}}$$

- Use one of the capacity estimation tools and amortize the cost

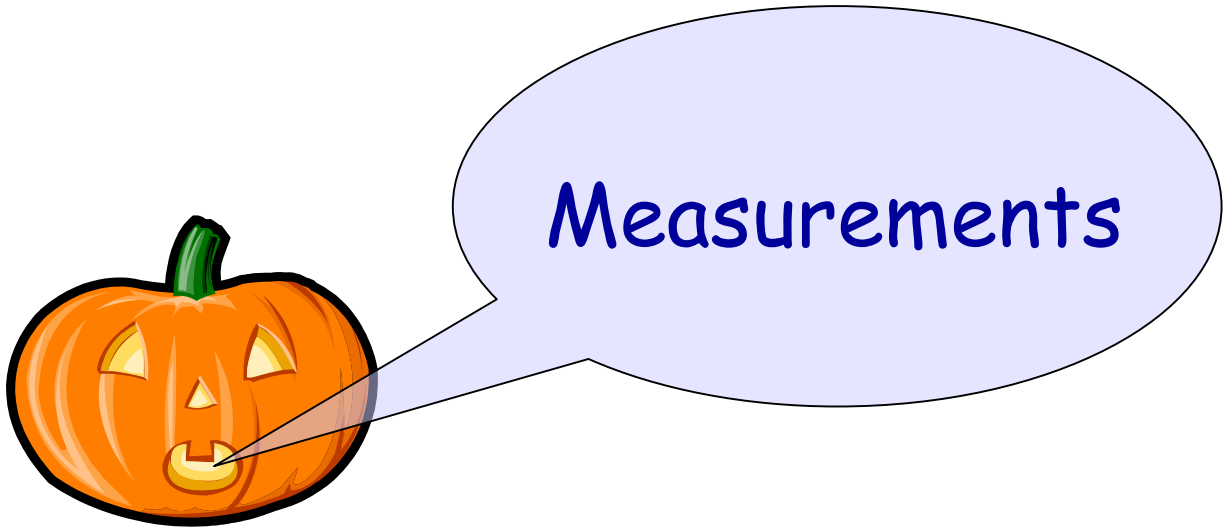
Measurements

# MRTG Test: Absolute Accuracy

- Use router logs of traffic rates
- Can only measure paths for which we have full access to all routers

# Pathload is Inaccurate

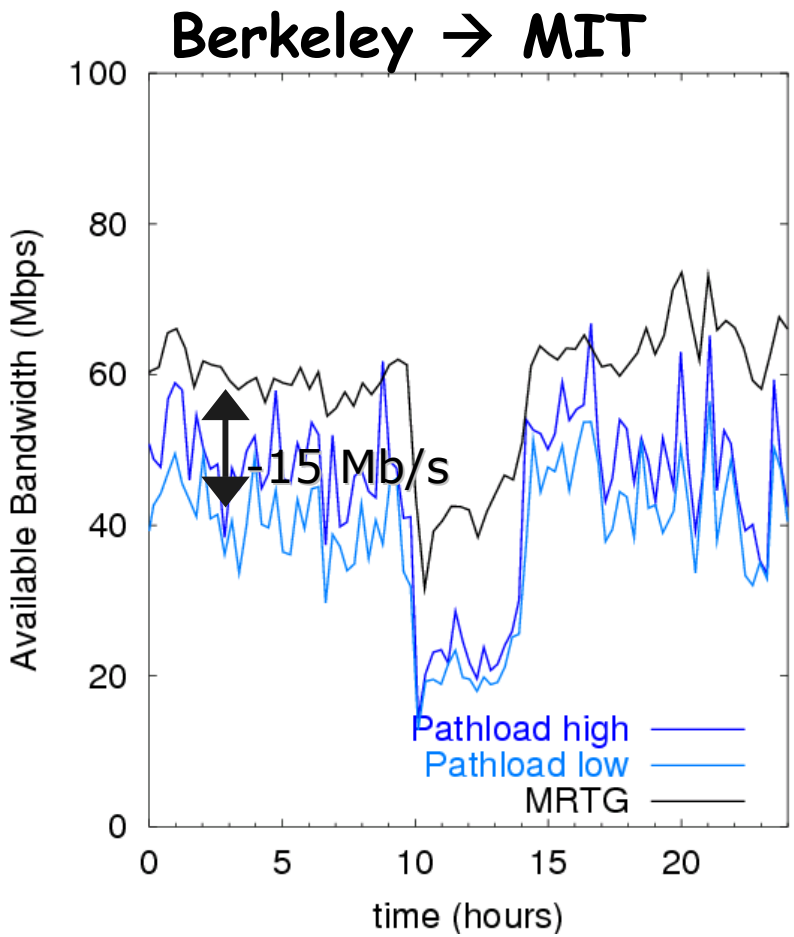


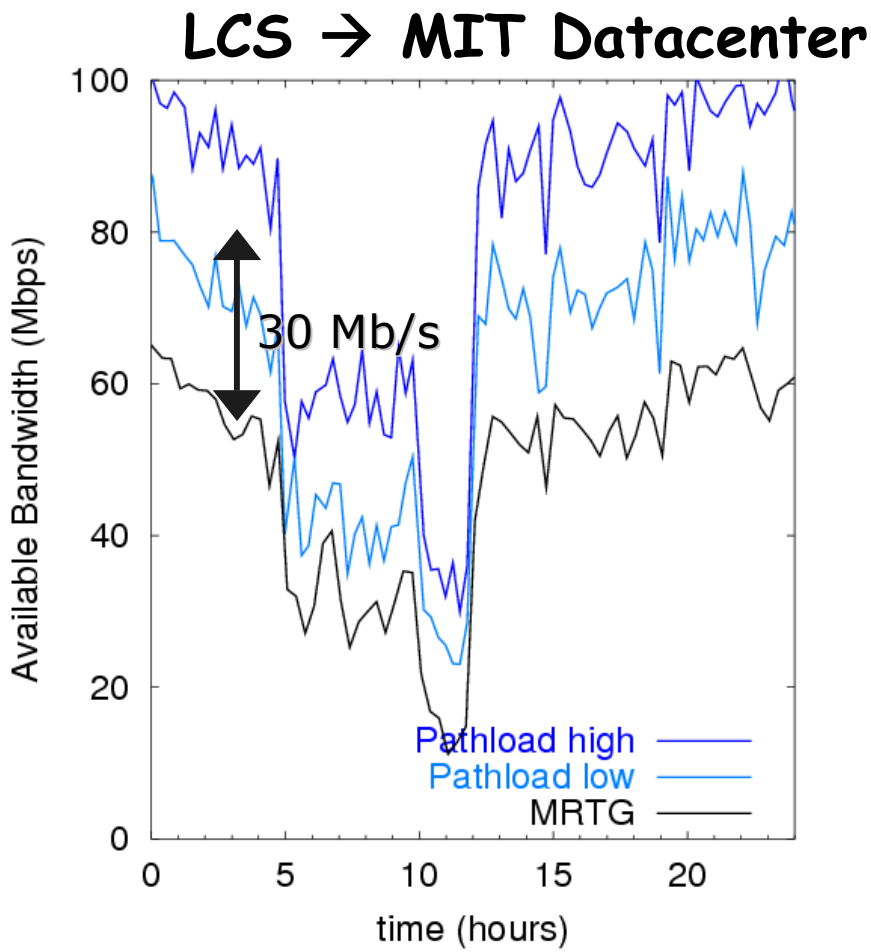


**Problem with parameter tuning:**
**Threshold parameters are not adequate for both paths**

# IGI is non-reactive to ABW decrease

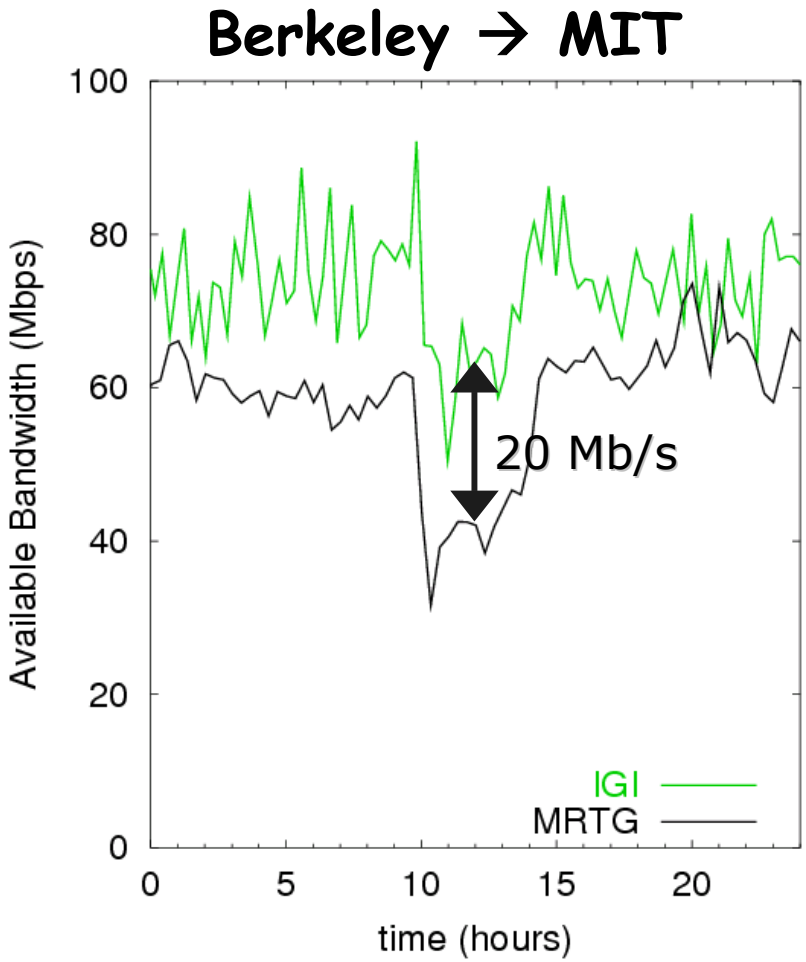


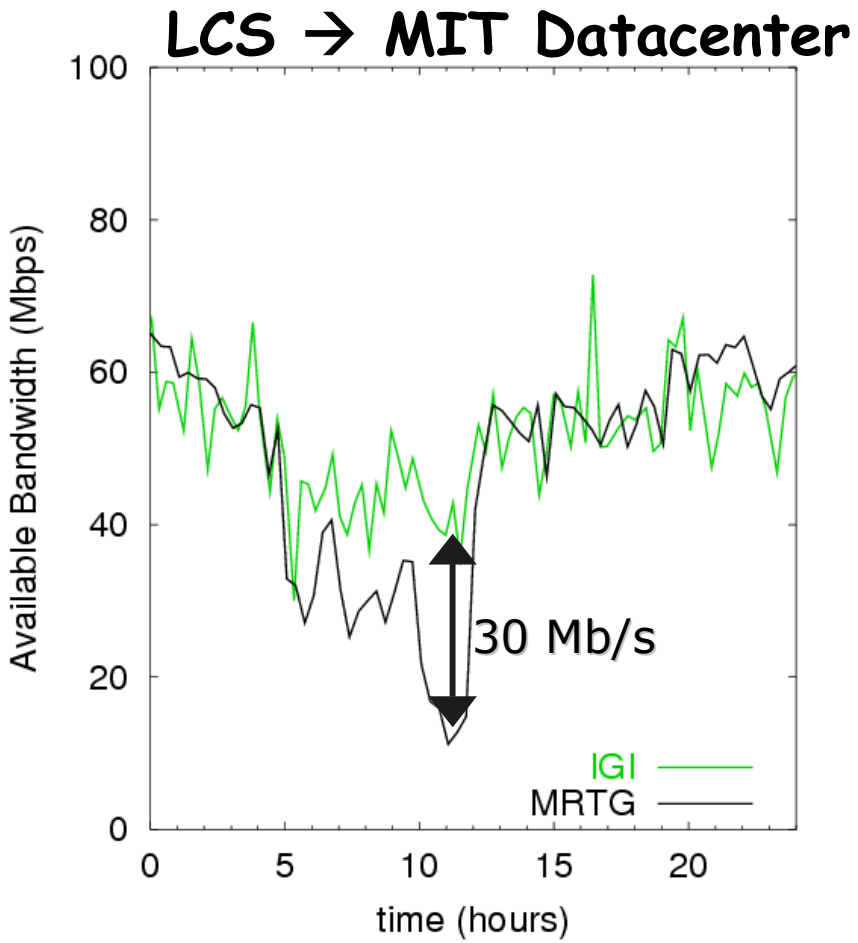


**Does not account for all cross traffic when load is high**

# Spruce is Accurate

**Berkeley → MIT**

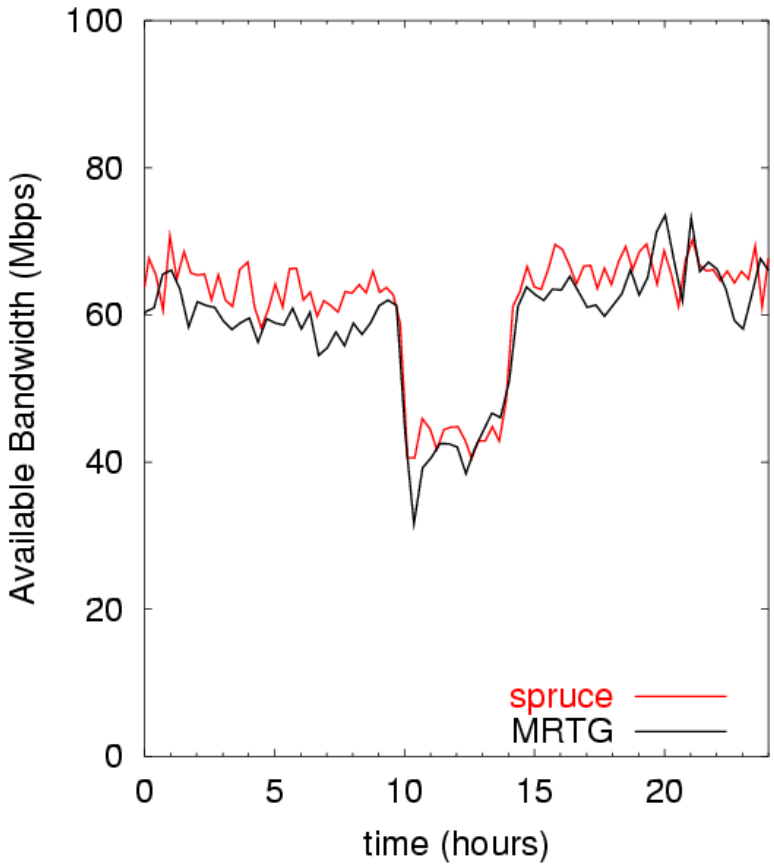


**LCS → MIT Datacenter**

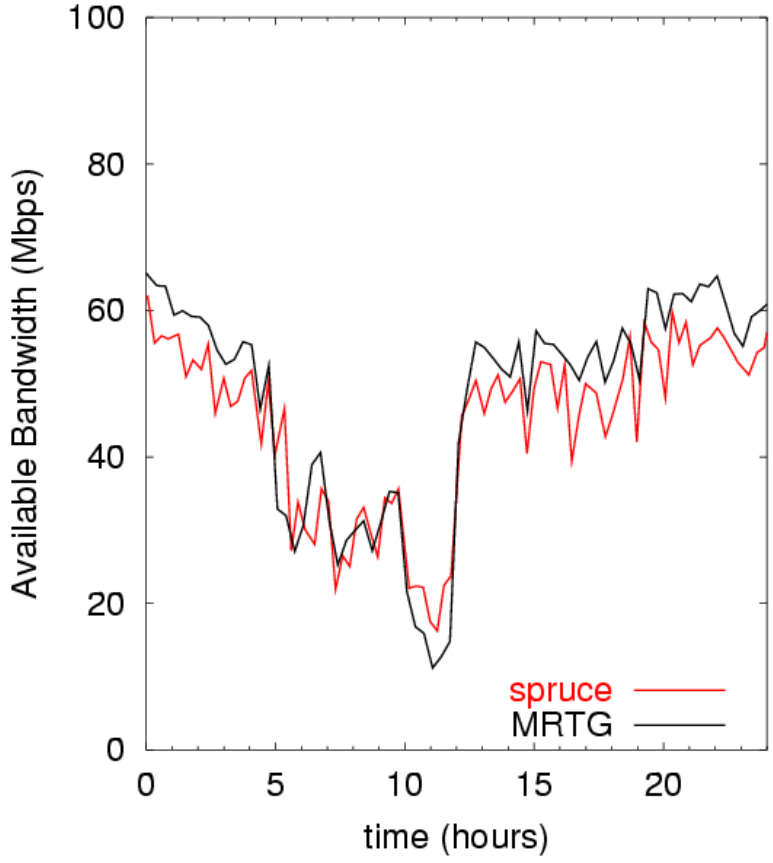


**No tunable parameters**

**Sampling with a Poisson process of pkt pairs**

# Larger Study

- 400 different Internet paths
  - PlanetLab nodes
  - RON testbed
  - T1, DSL, 10 Mb/s Ethernet, 100 Mb/s Ethernet, fractional T3
  - Variety of Loads, capacities, hop-counts, etc.

But, how to verify without MRTG data?

# D-Test: Relative Accuracy

1. Run tool and read its estimate of the available bandwidth, M1
2. Inject cross-traffic at a rate S=0.5*M1
3. Run tool again and read estimate, M2

$$\text{Relative Error} = \frac{S - (M_1 - M_2)}{S}$$

Assumption: Network state doesn't change between M1 and M2

Reasonable:

- Intrinsic limitation
- Time between M1 and M2 is < 1 minute
- Eliminated outliers

# D-Test Results

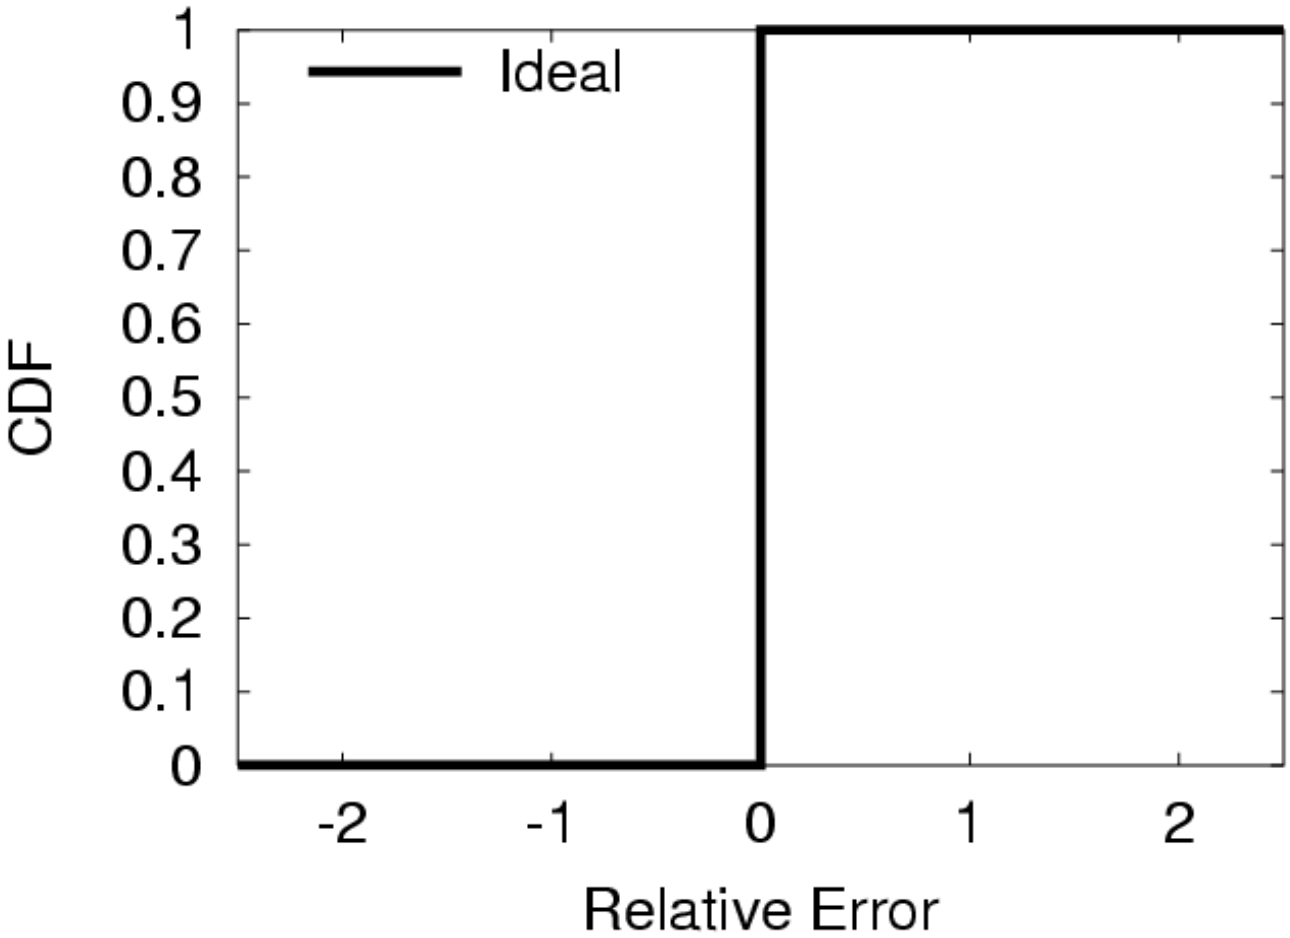


Ideally the CDF should be a step function at 0

# D-Test Results

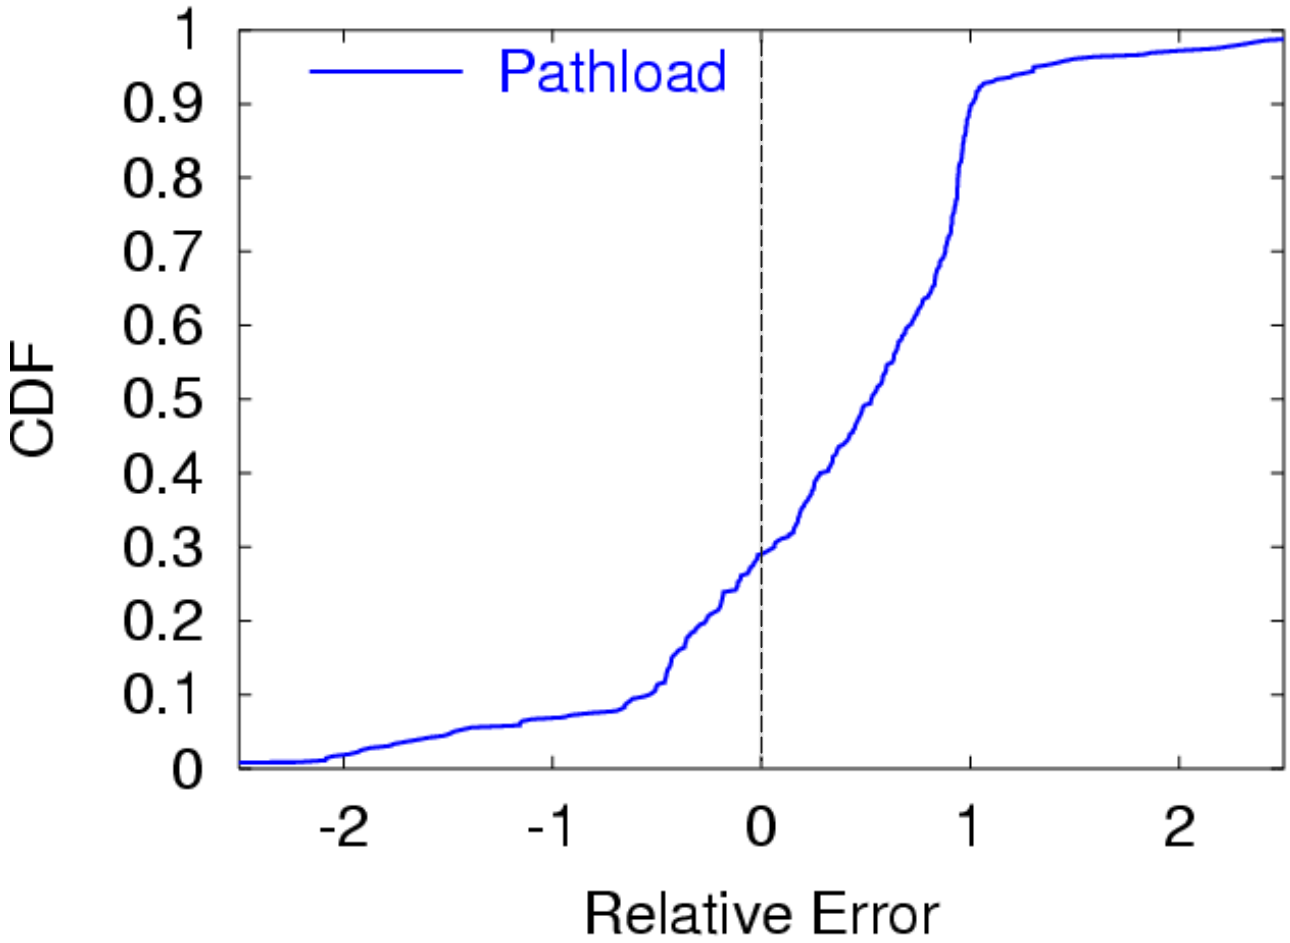


Pathload is Inaccurate

# D-Test Results

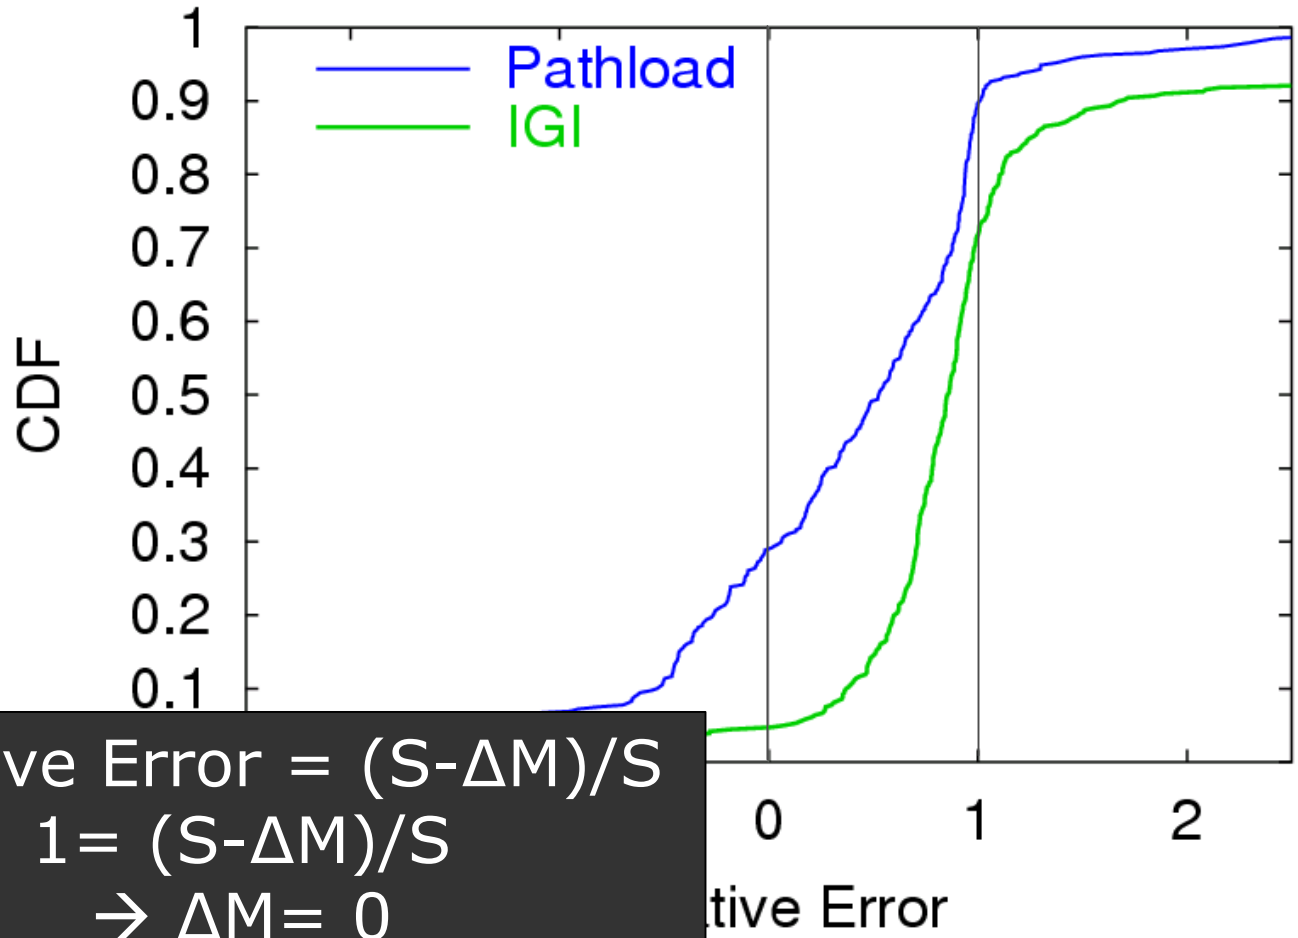


Relative Error = (S-ΔM)/S
1= (S-ΔM)/S
→ ΔM= 0

IGI is almost a step function at 1 →
IGI is non-reactive

# D-Test Results

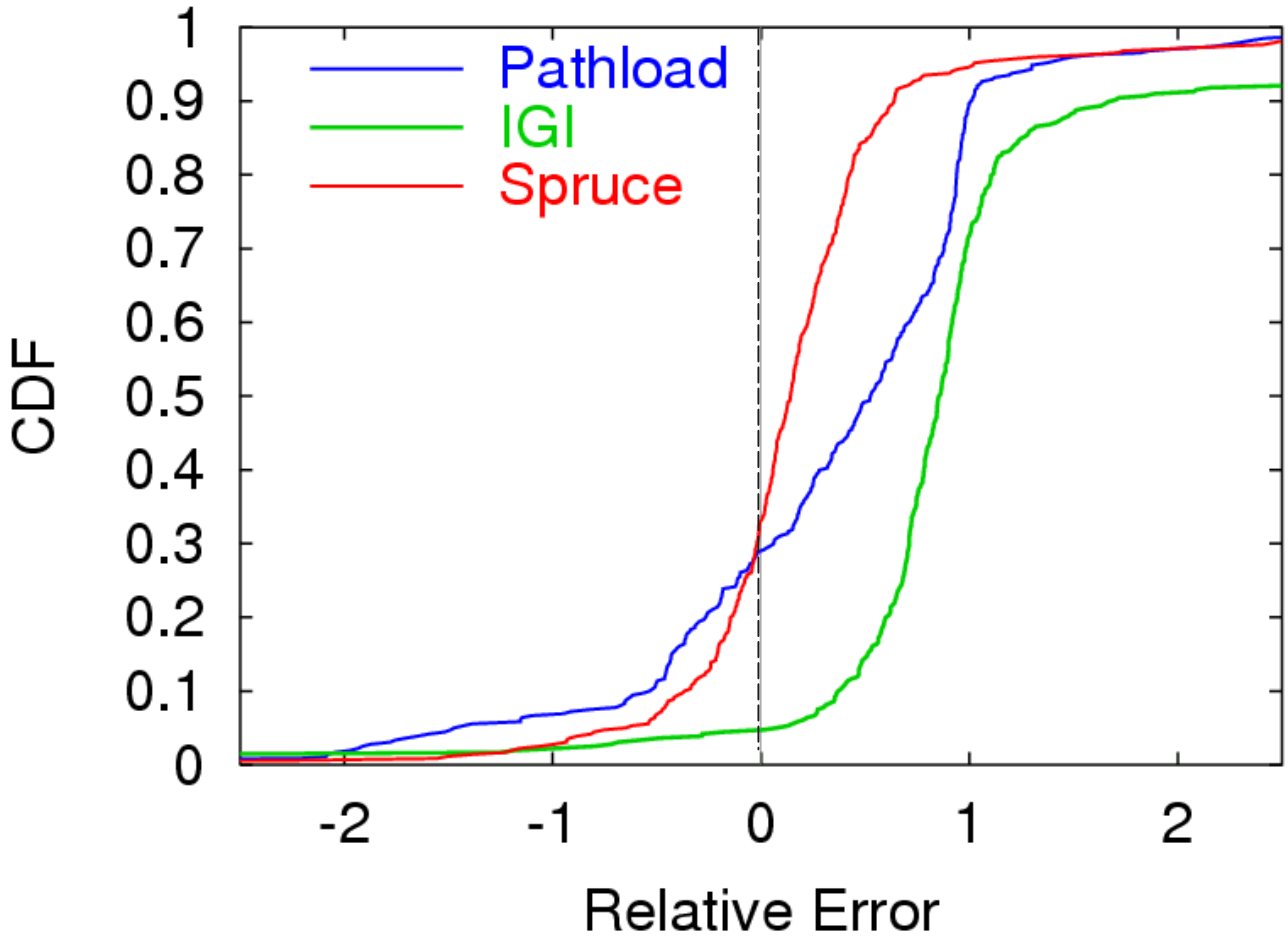


Spruce is more accurate
70% of measurements are within 30% of correct values

# Probe Traffic Overhead

- Average Per Measurement
  - Pathload: Between 2.5 MB and 10 MB
  - IGI: Around 130 KB
  - Spruce: About 300 KB

# Conclusion

- Wide-scale Internet test of ABW tools
- Accuracy
  - Pathload is relatively inaccurate
  - IGI is non-reactive to decreases in available bandwidth
  - Spruce is more accurate
- Overhead
  - Pathload sends order of magnitude more probe traffic than IGI and Spruce
- Spruce is available at
  `http://project-iris.net/spruce/`