Network Delay Tomography

Calculating estimations for network internal delays from end-to-end TTM measurements

Jan-Pascal van Best (janpascal@vanbest.org)
Willem Vree (w.g.vree@tbm.tudelft.nl)

Delft University of Technology

TU Delft
Technische Universiteit Delft

Outline

- Idea
- Challenges
- Approach
- Results
- Reliability
- Tool
- Conclusions

Idea

- TTM gathers a LOT of traceroute and delay measurement data between the test boxes...

- Wouldn't it be nice if we could also determine (some of) the delays between internal hosts?
  - ... more information on the state of the network
  - ... locating problem areas in the network
  - ... also information on the networks of *non-participating* ISPs

Challenges

- Lots of data
- Raw traceroute data often not very reliable
- Theoretical limits: typical snapshot contains 6,000 links, with only 4100 end-to-end measurements

Approach

- Find traceroutes for more-or-less stable paths
- Take *minimum* measured delay over this path during one day
- Identify router interfaces (using UDP probes)
- Segmentation (eliminate strings of links)
- Use stable least-squares algorithm to calculate link delays from end-to-end delays

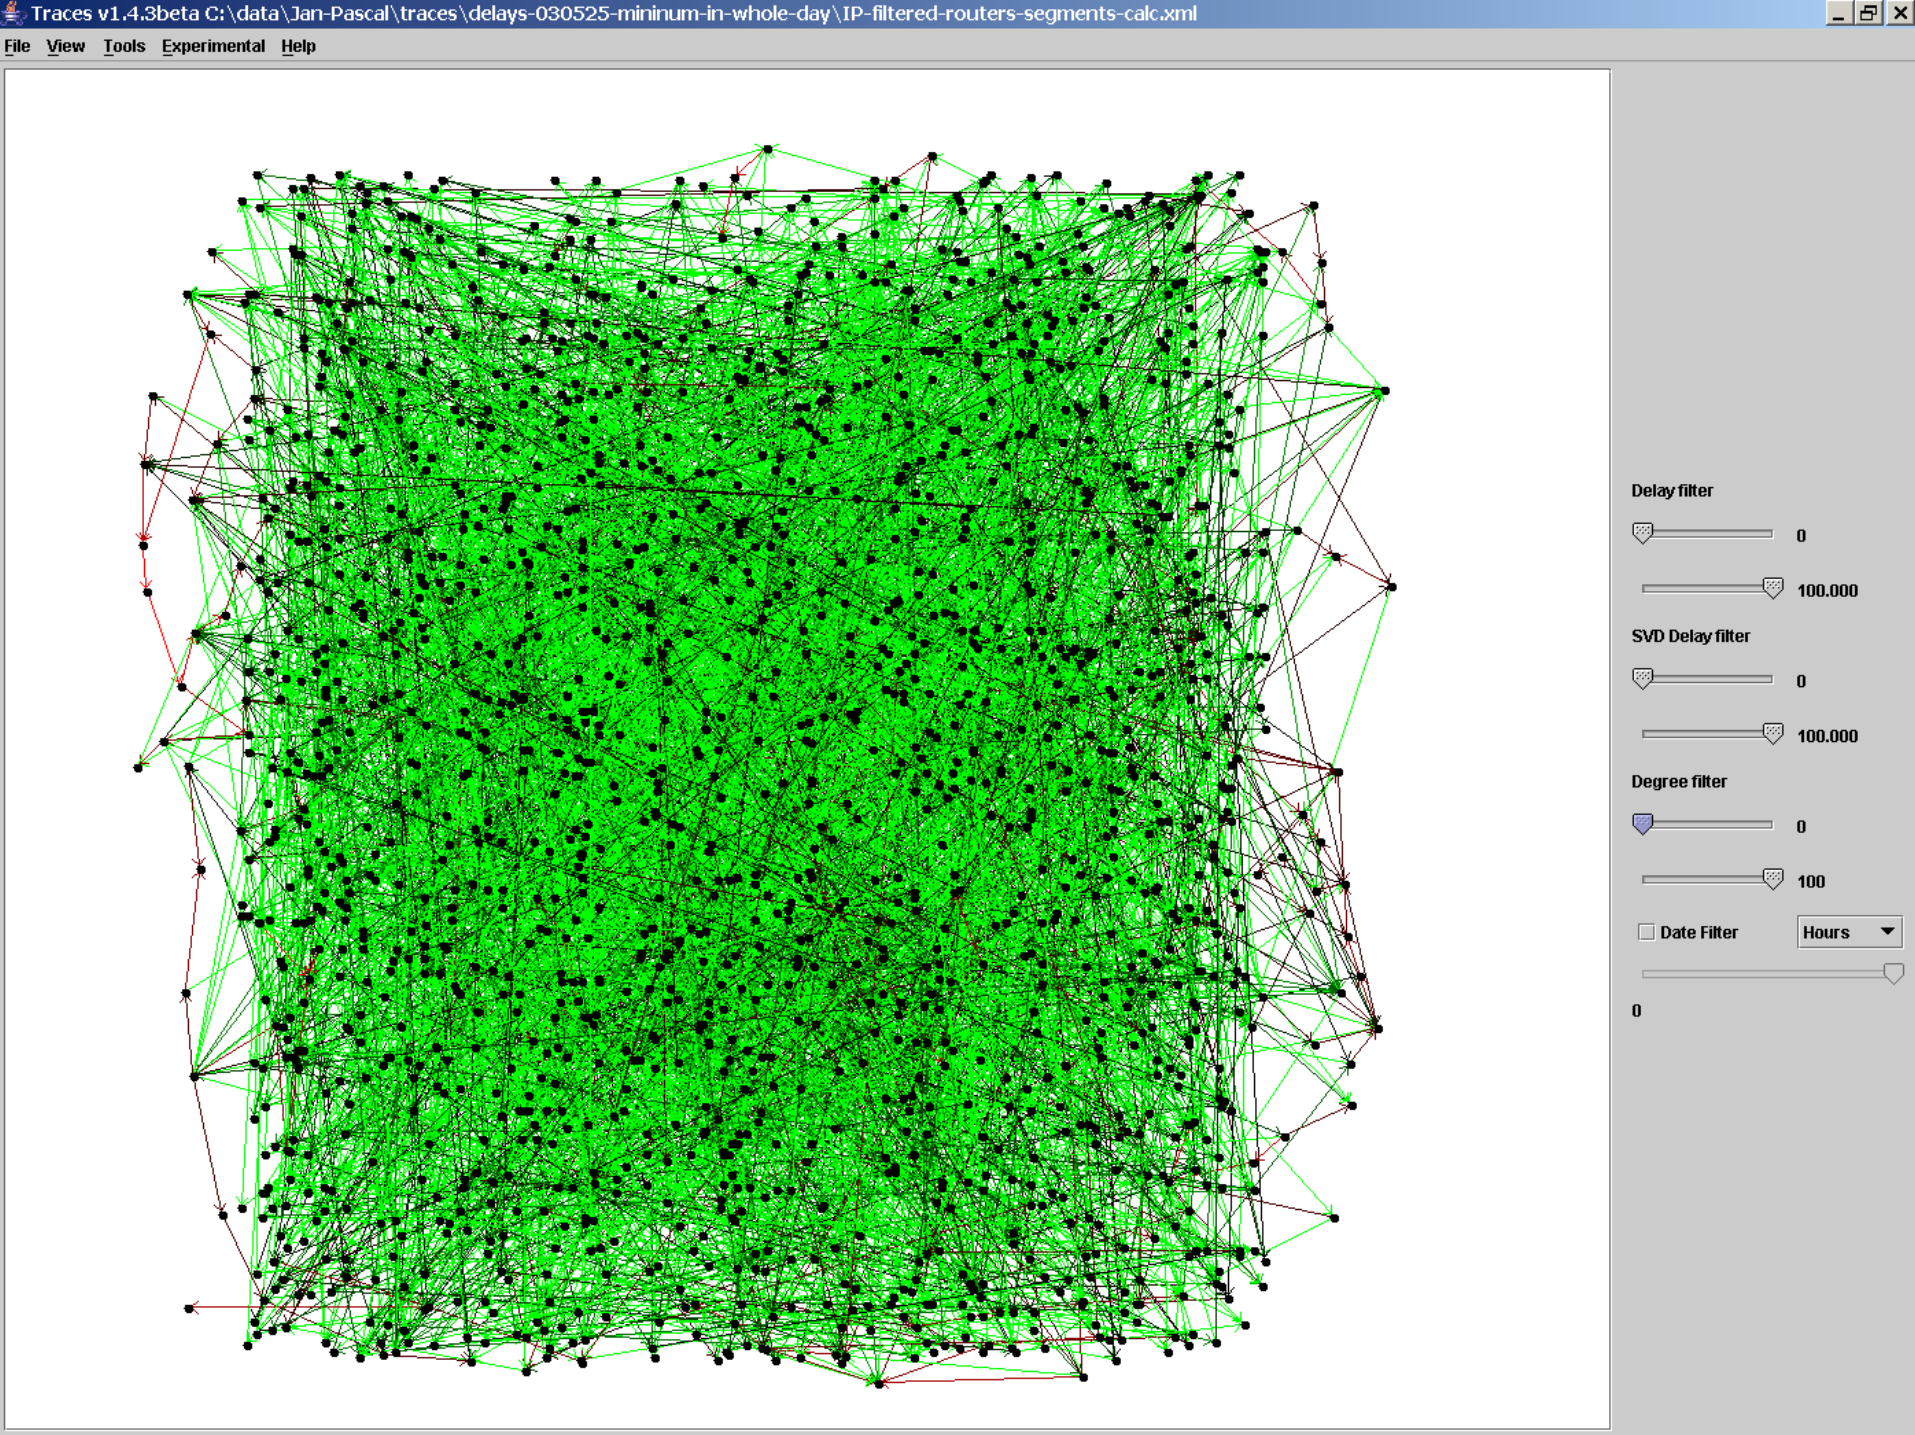
Traces v1.4.3beta C:\data\Jan-Pascal\traces\delays-030525-mininum-in-whole-day\IP-filtered-routers-segments-calc.xml
File View Tools Experimental Help
Delay filter
0
100.000
SVD Delay filter
0
100.000
Degree filter
0
100
Date Filter
Hours
0

Traces v1.4.3beta C:\data\Jan-Pascal\traces\delays-030525-mininum-in-whole-day\IP-filtered-routers-segments-calc.xml

File View Tools Experimental Help

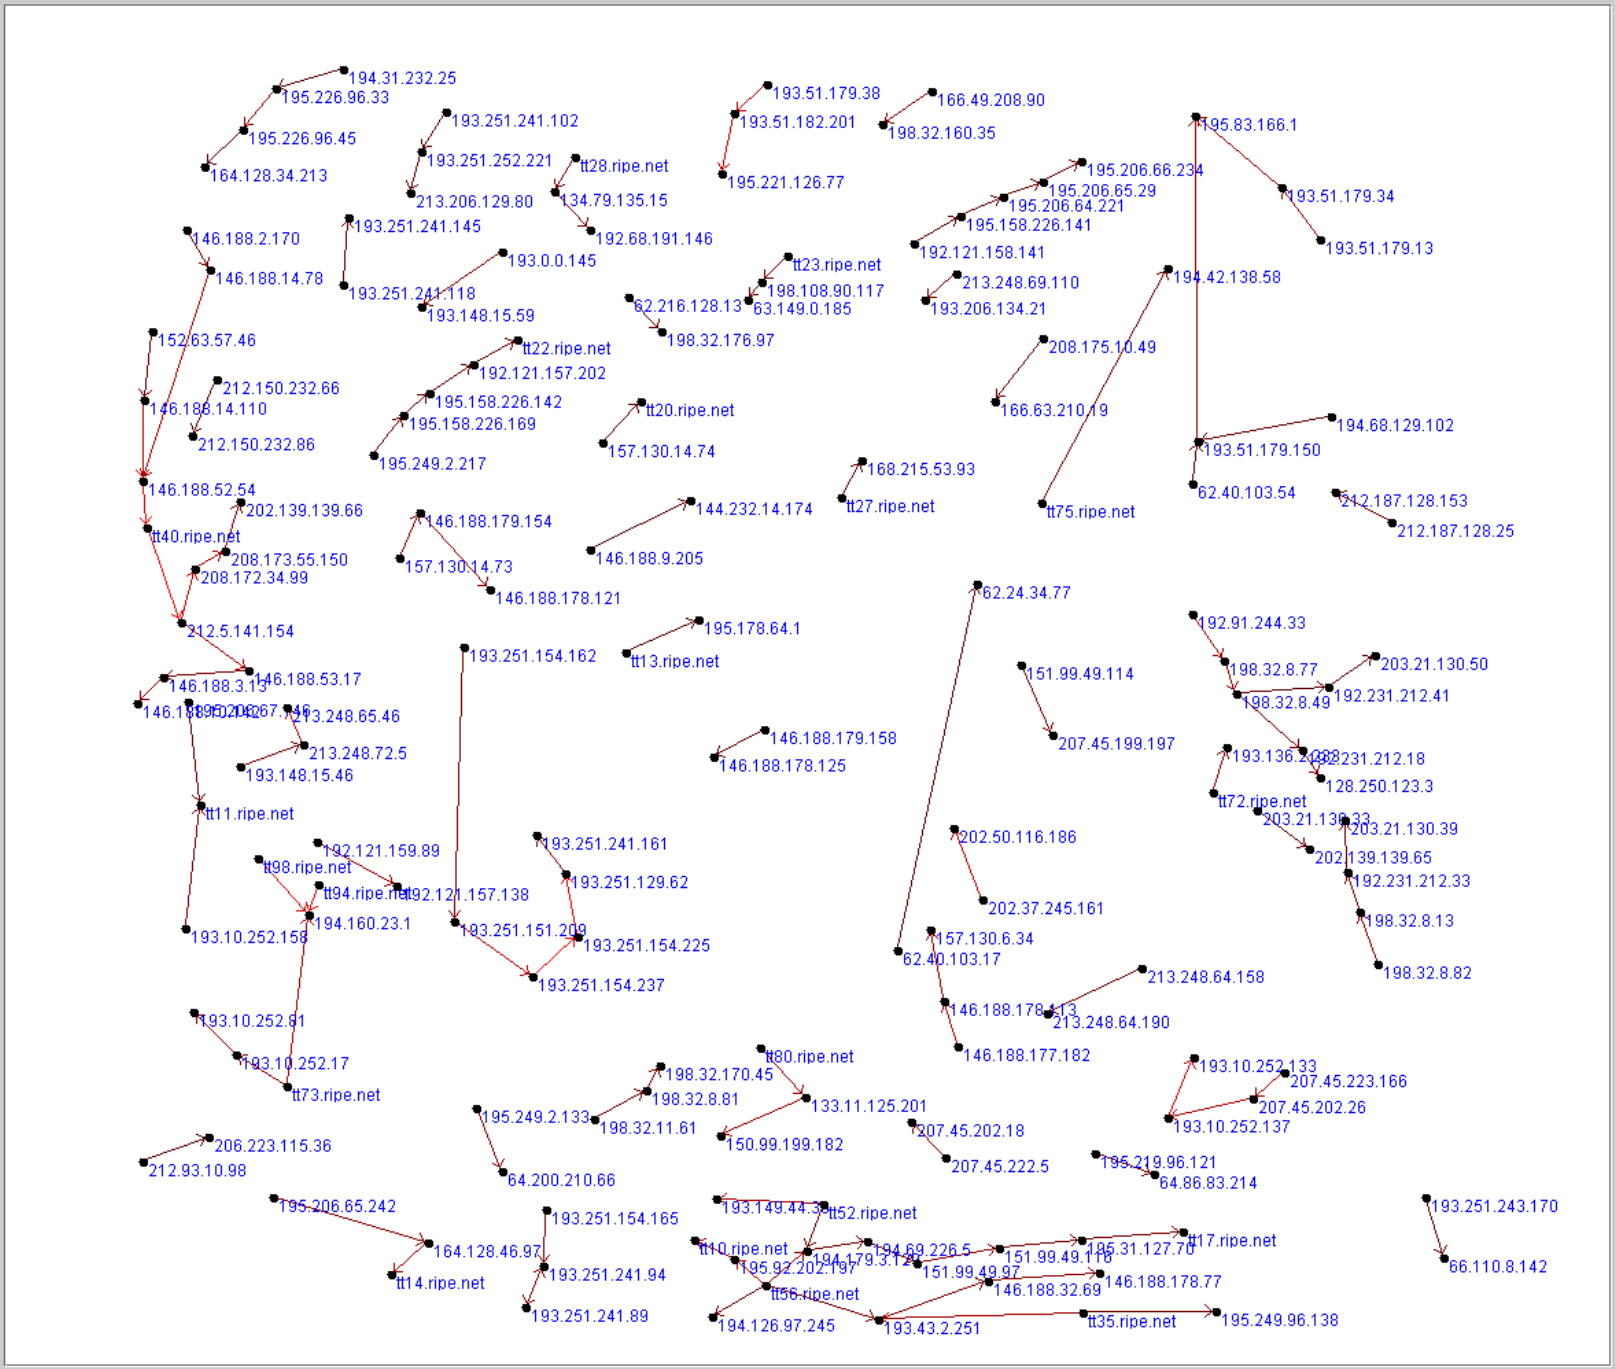

Delay filter

0

100.000

SVD Delay filter

31.741

100.000

Degree filter

0

100

Date Filter Hours

0

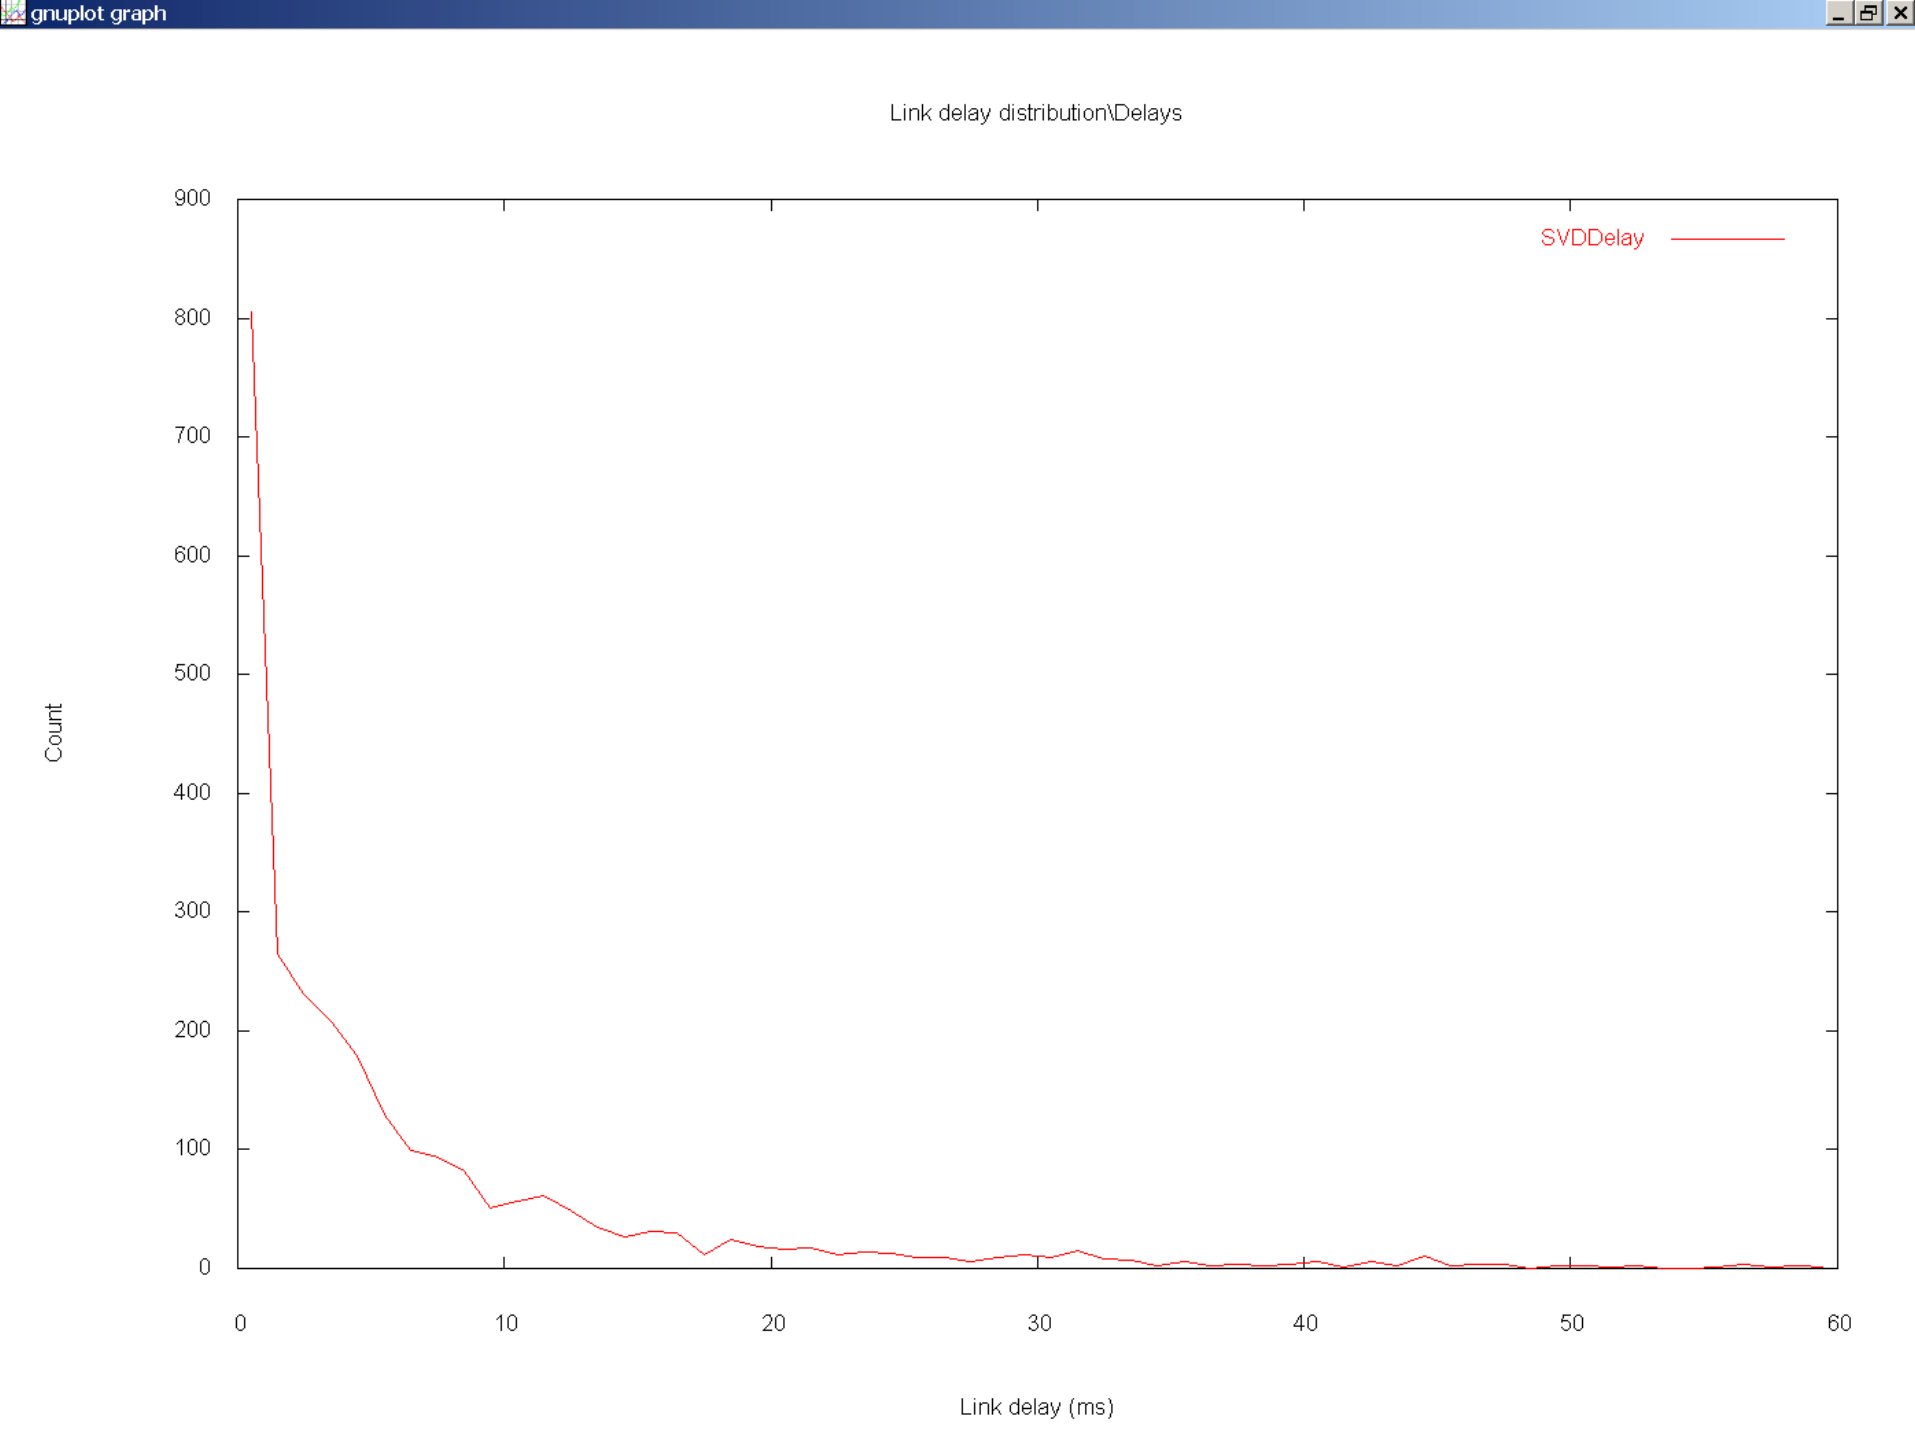
gnuplot graph
Link delay distribution\Delays
SVDDelay
Count
Link delay (ms)
900
800
700
600
500
400
300
200
100
0
0
10
20
30
40
50
60

Verification

- Are these estimates any good?

- Putting synthetic link delays into a real network
- Calculate end-to-end delays
- Then use our method to find estimates for the link delays
- Compare (input) synthesized delays to estimates

Distribution of error from synthesized delay

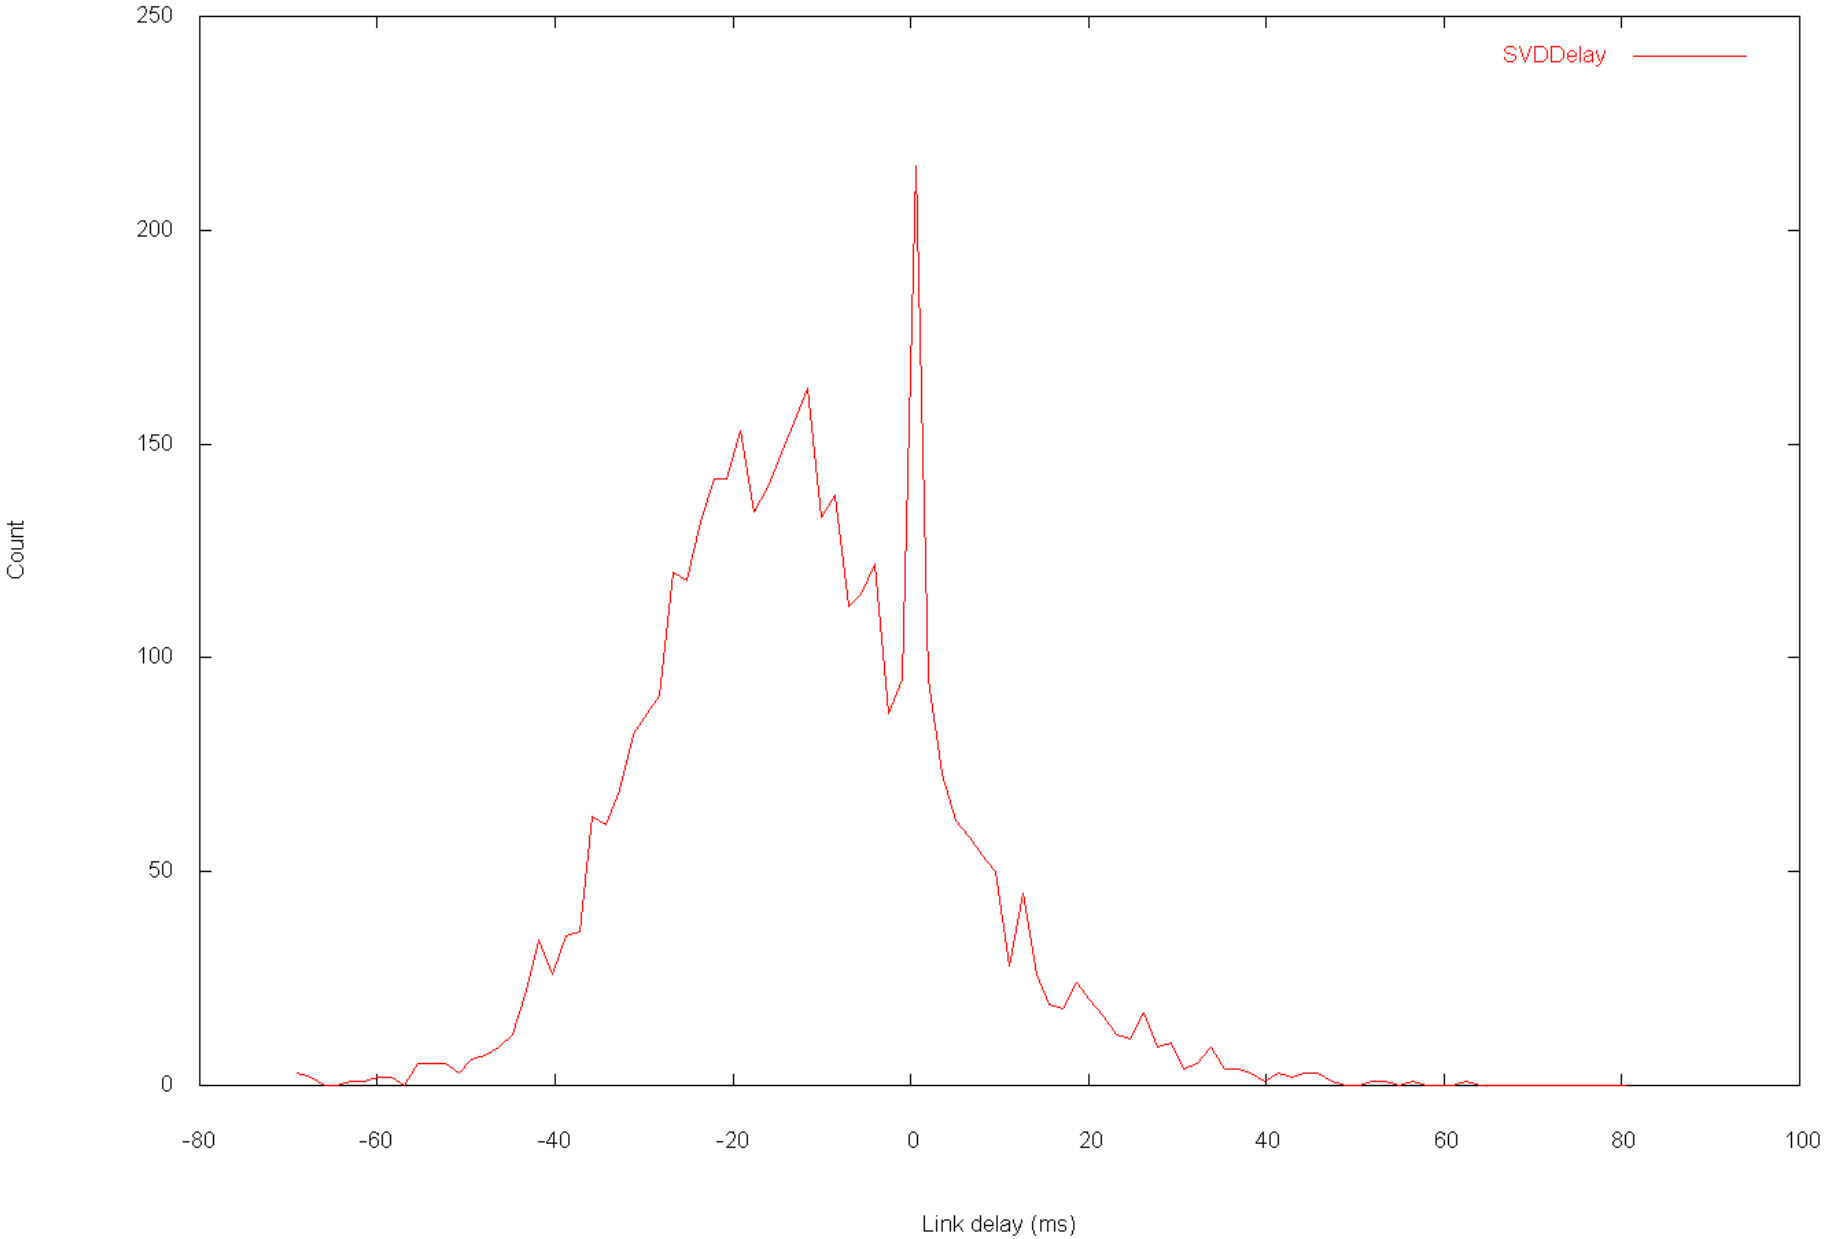

gnuplot graph

Distribution of error from synthesized delay, for delays>40ms

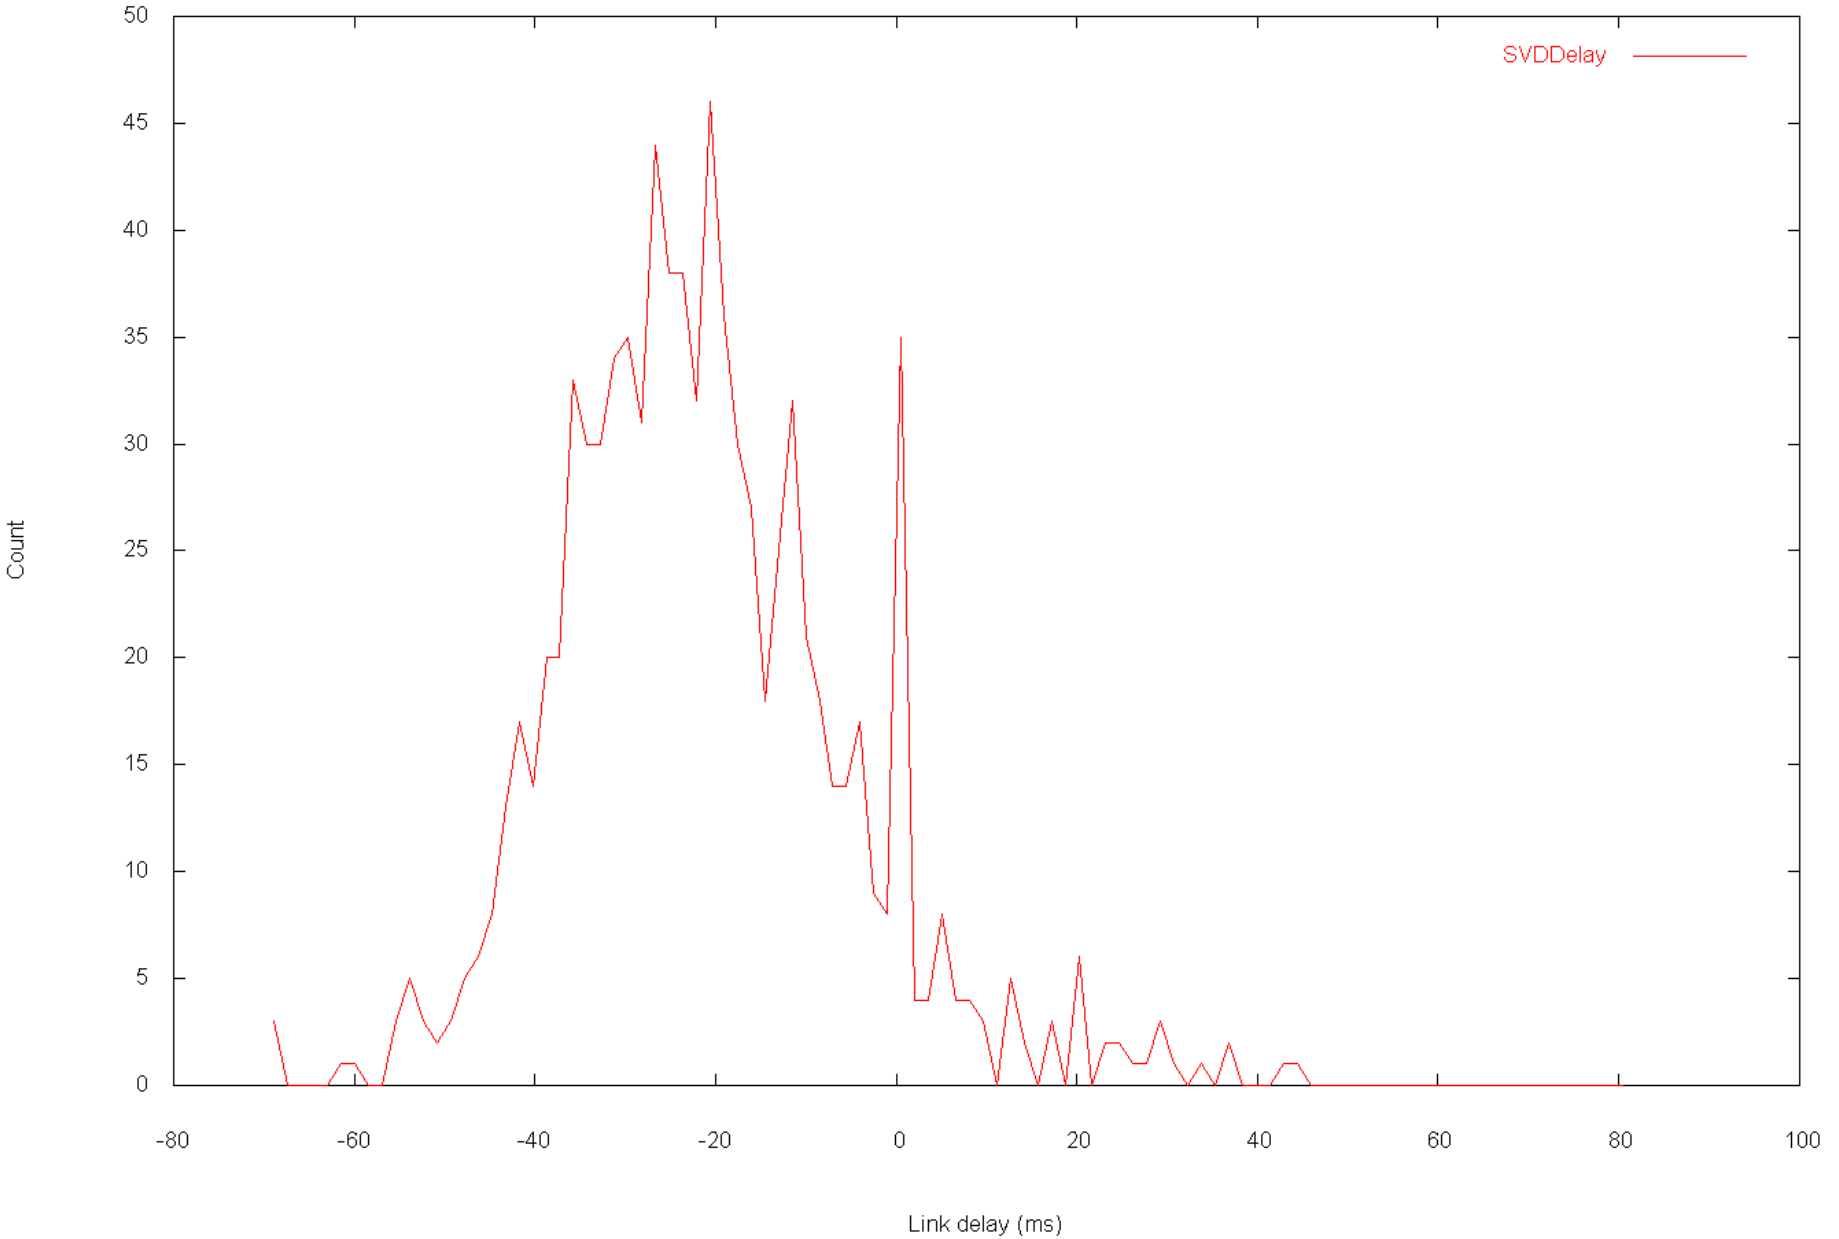

Verification (4)

- Are these estimates any good?

- Answer:
  - Error margins larger than most delays
  - But if a link has a large estimated link delay, it really *is* large!

Tools

- `traces` Java-based analysis tool
- Imports Test Traffic and RIS data
- Graph display & layout
- Trace filtering
- Router mapping (multiple IPs per router)
- Segmentation
- AS mapping
- Delay calculations
- Synthetic delays
- Histograms

Conclusions

- Method useful for identifying slow or congested links
- Not (yet) reliable for general link delay estimation
- Limitations, like MPLS

- Future work:
  - Try to find link-level data and check results
  - Even larger data sets
  - Other methods

- Demo (in the break)