ICE^3 : Ongoing Development and Results Around Netsniff(Presentation II)

Urs Keller

CAIA Lab

urs.keller@epfl.ch

Talk abstract

- This is my last presentation here at CAIA. I guess you want to know, what I have been up to ...

Talk abstract
Talk outline
Review
Paper 2: ACE2005
Paper 1: RTT
Conclusion
Questions
Thanks

Talk outline

- Where we stopped last time (Review)
- Paper 2: Caching paper
- Paper 1: RTT
- Conclusions
- Questions
- Thanks

- Talk abstract
- Talk outline

Review

Paper 2: ACE2005

Paper 1: RTT

Conclusion

Questions

Thanks

Review

- Talk abstract
- Talk outline

Review
- Netsniff deployment
- Traffic analysed
- Logfile parser
- Database

Paper 2: ACE2005

Paper 1: RTT

Conclusion

Questions

Thanks

Netsniff deployment

- Talk abstract
- Talk outline

Review
- Netsniff deployment
- Traffic analysed
- Logfile parser
- Database

Paper 2: ACE2005

Paper 1: RTT

Conclusion

Questions

Thanks

Netsniff is currently deployed at Irena's, Jason's and Sebastian's home.

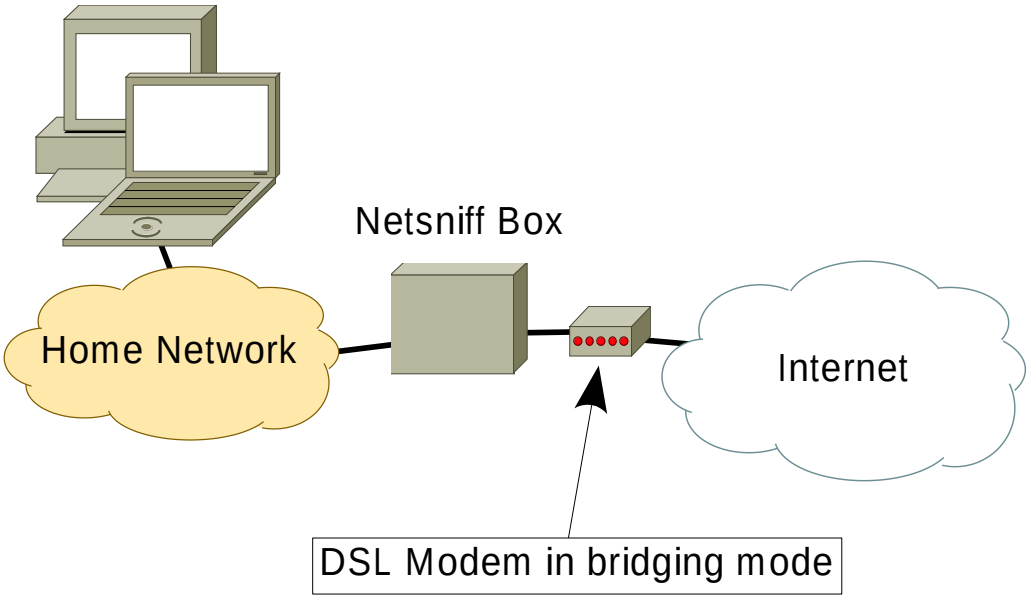

Traffic analysed

- Talk abstract
- Talk outline

Review
- Netsniff deployment
- Traffic analysed
- Logfile parser
- Database

Paper 2: ACE2005

Paper 1: RTT

Conclusion

Questions

Thanks

Traffic of the following protocols is currently analysed.

ARP	ICMP	DNS
HTTP	TLS	FTP
IMAP	POP3	SMTP

Traffic characteristics are logged into text files for later processing.

Logfile parser

- Talk abstract
- Talk outline

Review
- Netsniff deployment
- Traffic analysed
- Logfile parser
- Database

Paper 2: ACE2005

Paper 1: RTT

Conclusion

Questions

Thanks

Output from Netsniff needs to be parsed and stored in a manner, which makes data analysis easy and efficient.

→ Relational database

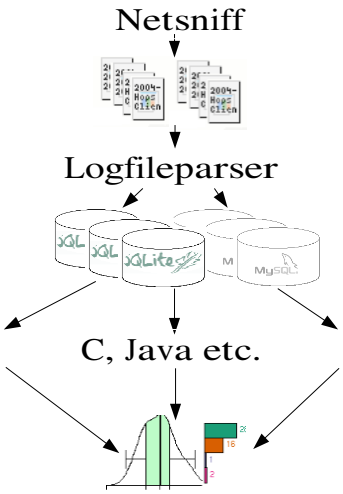

Currently the back-ends SQLite, SQLite3 and MySQL are supported. Where the later performs a lot better for large and complex queries.

Database

- Talk abstract
- Talk outline

Review
- Netsniff deployment
- Traffic analysed
- Logfile parser
- Database

Paper 2: ACE2005

Paper 1: RTT

Conclusion

Questions

Thanks

- Current schema has around 40 tables
- Some of them are still unused

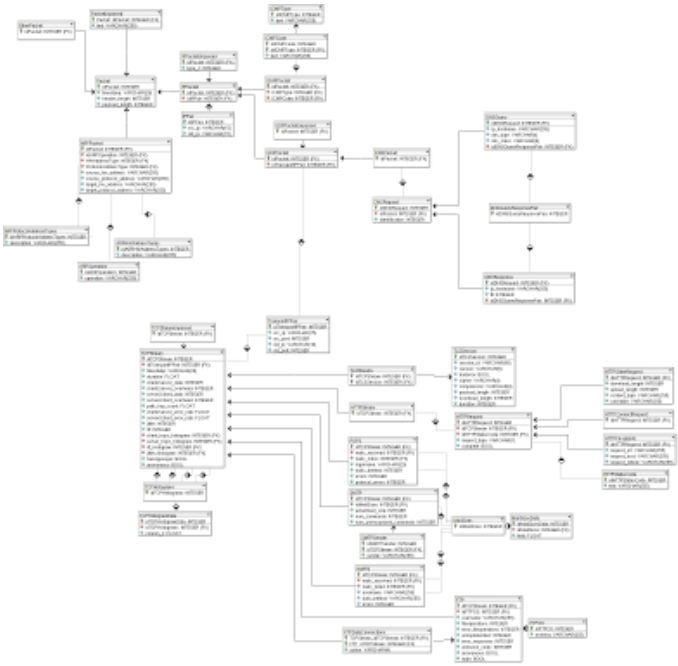

- Talk abstract
- Talk outline

Review

Paper 2: ACE2005
- Overview
- “Final results”
- The Steps
- Match HTTP and DNS
- RTT distribution
- Object sizes
- Some other results

Paper 1: RTT

Conclusion

Questions

Thanks

Paper 2: ACE 2005

(ACM SIGCHI International Conference on Advances in Computer Entertainment Technology)

SWINBURNE CENTRE FOR ADVANCED INTERNET ARCHITECTURES

Overview

- Talk abstract
- Talk outline

Review

Paper 2: ACE2005
- Overview
- “Final results”
- The Steps
- Match HTTP and DNS
- RTT distribution
- Object sizes
- Some other results

Paper 1: RTT

Conclusion

Questions

Thanks

Idea: Derive possible improvements on user experience from the data we have gathered with Netsniff over the last three months, in the context of ICE^3 .

→ **“rationale” for web caching**

- *“A rationale for web caching in consumer ISPs: The impact of DNS lookup times and HTTP session characteristics”* - TR
- *“A rationale for web caching in consumer ISPs: Evaluating the impact of DNS lookup times and HTTP session characteristics on consumer ISP web traffic”* - Paper

“Final results”

- Talk abstract
- Talk outline

Review

Paper 2: ACE2005
- Overview
- “Final results”
- The Steps
- Match HTTP and DNS
- RTT distribution
- Object sizes
- Some other results

Paper 1: RTT

Conclusion

Questions

Thanks

Let’s start at the end of the paper ...

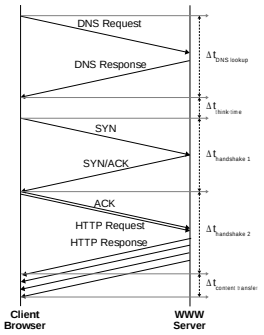

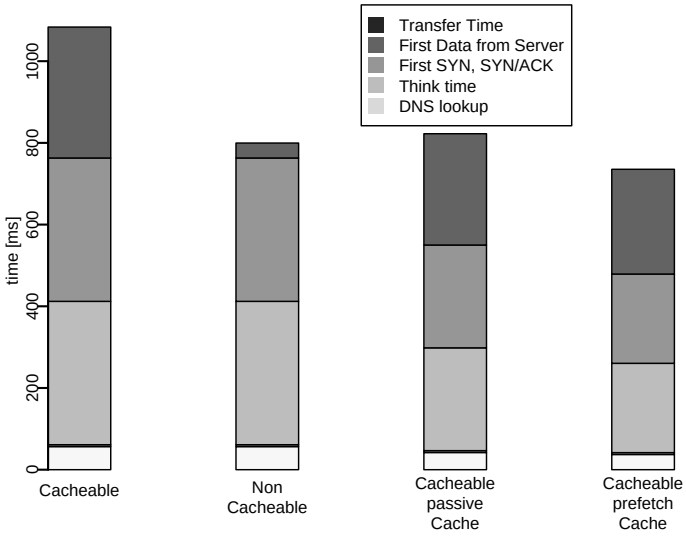

The Steps

- Talk abstract
- Talk outline

Review

Paper 2: ACE2005
- Overview
- “Final results”
- The Steps
- Match HTTP and DNS
- RTT distribution
- Object sizes
- Some other results

Paper 1: RTT

Conclusion

Questions

Thanks

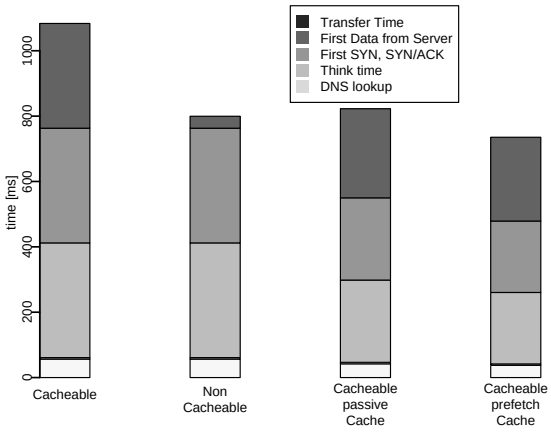

- Match HTTP and DNS (bottom two)
- RTT distribution (second and third)
- Object sizes (top part)

Match HTTP and DNS

- Talk abstract
- Talk outline

Review

Paper 2: ACE2005
- Overview
- "Final results"
- The Steps
- Match HTTP and DNS
- RTT distribution
- Object sizes
- Some other results

Paper 1: RTT

Conclusion

Questions

Thanks

- $DNS_Q = \pi_{qtime,id}(\delta(...))$
- $DNS_R = \pi_{rtime,id,ip\_res}(\delta(...))$
- $HTTP = \pi_{htime,dst\_ip}(\delta(...))$
- $DNS_{QR} = DNS_Q \bowtie_{DNS_Q.id=DNS_R.id} DNS_R$
- $HTTP_{DNS0} = DNS_{QR} \bowtie_{DNS_{QR}.ip\_res=HTTP.dst\_ip} HTTP$
- $HTTP_{DNS1} = \sigma_{htime \geq rtime}(HTTP_{DNS0})$
- $HTTP_{DNS} = \gamma_{(qtime,rtime),min(htime)}(HTTP_{DNS1})$

The actual implementation is a bit more complex. This is due to indexing and the deletion of temporary tables, once they are unused.

Match HTTP and DNS (2)

- Talk abstract
- Talk outline

Review

Paper 2: ACE2005
- Overview
- "Final results"
- The Steps
- Match HTTP and DNS
- RTT distribution
- Object sizes
- Some other results

Paper 1: RTT

Conclusion

Questions

Thanks

From the query in the previous slide (and some others), we get the following result:

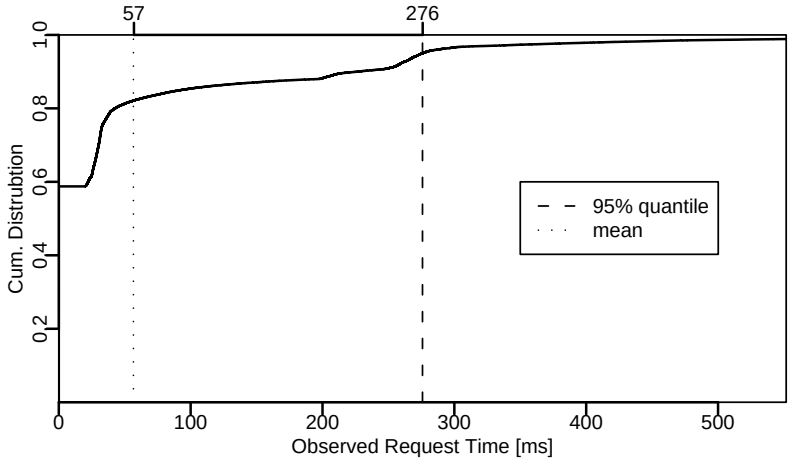

These show that around 60% are cached in either a browser cache or the operating system. Further 18% are locally cached and for another 22% the DNS server most likely had to recurse to retrieve the data.

RTT distribution

- Talk abstract
- Talk outline

Review

Paper 2: ACE2005
- Overview
- “Final results”
- The Steps
- Match HTTP and DNS
- RTT distribution
- Object sizes
- Some other results

Paper 1: RTT

Conclusion

Questions

Thanks

We analyzed the RTT of the TCP streams related to HTTP. The results are shown in the cumulative distribution below.

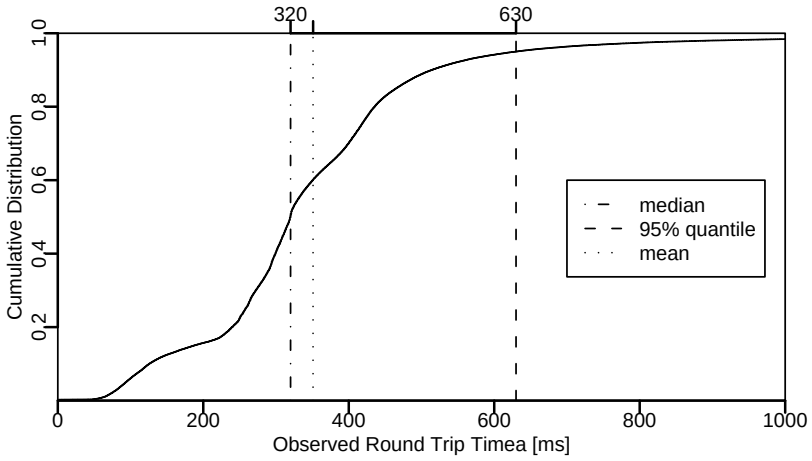

95% are below 630ms, 50% below 320ms and the average is 348 ms.

Object sizes

- Talk abstract
- Talk outline

Review

Paper 2: ACE2005
- Overview
- “Final results”
- The Steps
- Match HTTP and DNS
- RTT distribution
- Object sizes
- Some other results

Paper 1: RTT

Conclusion

Questions

Thanks

The last part is the load time of the objects. We assumed a typical DSL line speed and calculated an estimated load time.

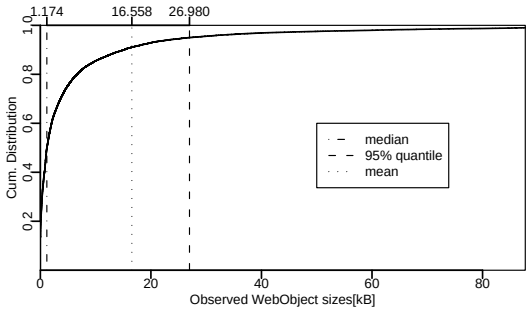

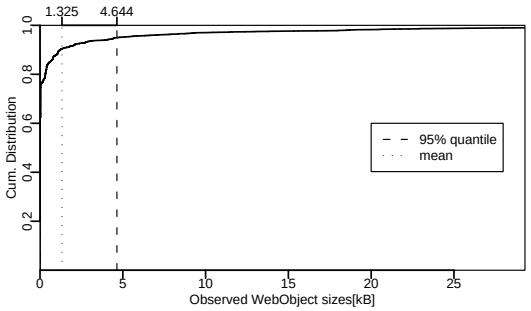

The results are very similar to what Sebastian obtained two years ago

Some other results

- Talk abstract
- Talk outline

Review

Paper 2: ACE2005
- Overview
- “Final results”
- The Steps
- Match HTTP and DNS
- RTT distribution
- Object sizes
- Some other results

Paper 1: RTT

Conclusion

Questions

Thanks

Hop count histogram

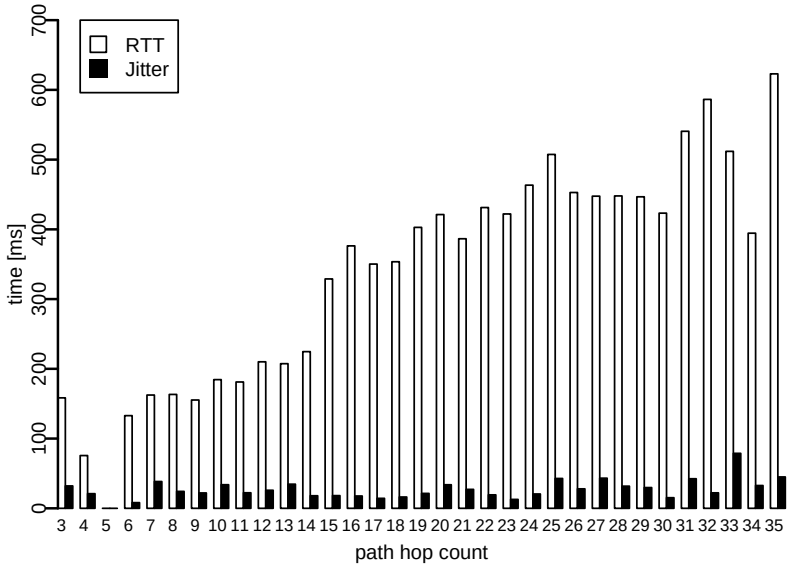

SWINBURNE CENTRE FOR ADVANCED INTERNET ARCHITECTURES

Some other results

- Talk abstract
- Talk outline

Review

Paper 2: ACE2005
- Overview
- “Final results”
- The Steps
- Match HTTP and DNS
- RTT distribution
- Object sizes
- Some other results

Paper 1: RTT

Conclusion

Questions

Thanks

Content types by count

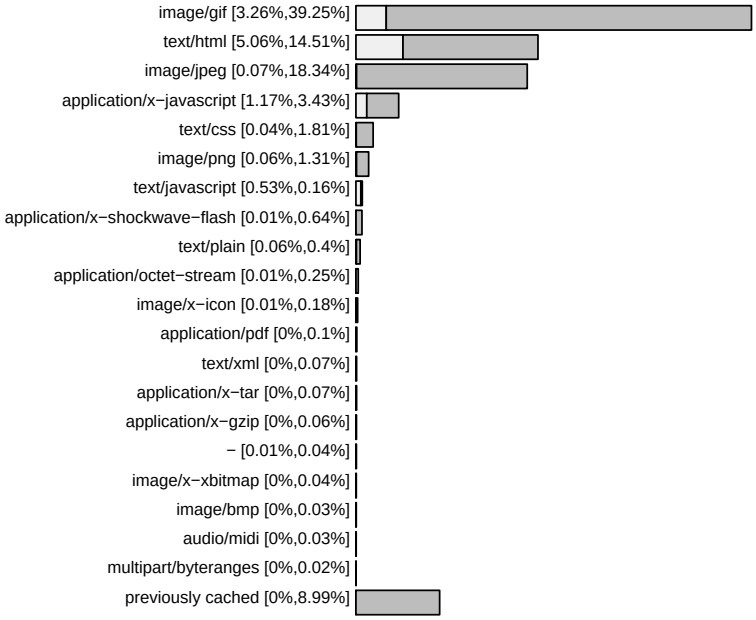

Some other results

- Talk abstract
- Talk outline

Review

Paper 2: ACE2005
- Overview
- “Final results”
- The Steps
- Match HTTP and DNS
- RTT distribution
- Object sizes
- Some other results

Paper 1: RTT

Conclusion

Questions

Thanks

Content types by size

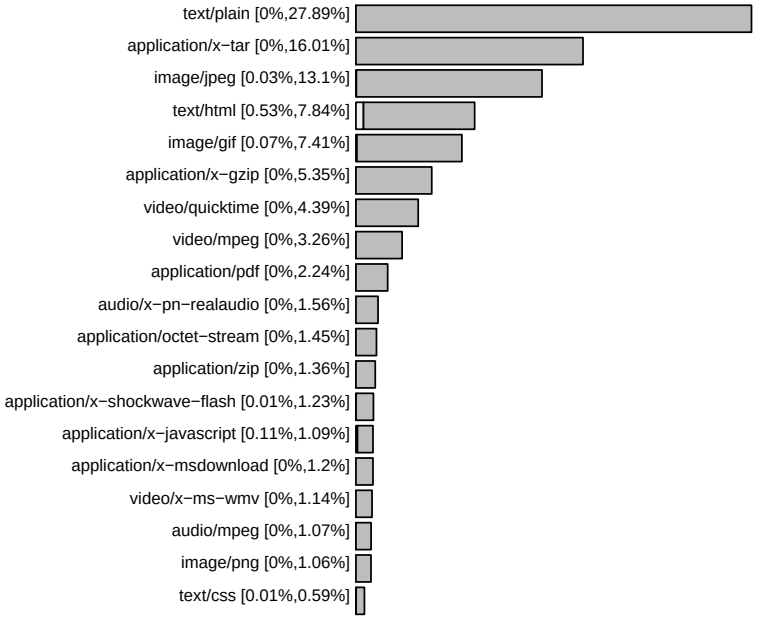

Some other results

etc.

- Talk abstract
- Talk outline

Review

Paper 2: ACE2005
- Overview
- “Final results”
- The Steps
- Match HTTP and DNS
- RTT distribution
- Object sizes
- Some other results

Paper 1: RTT

Conclusion

Questions

Thanks

- Talk abstract
- Talk outline

Review

Paper 2: ACE2005

Paper 1: RTT
- Setting
- TCP options
- TCP kernel modifications
- Test runs

Conclusion

Questions

Thanks

Paper 1: RTT

- Talk abstract
- Talk outline

Review

Paper 2: ACE2005

Paper 1: RTT
- Setting
- TCP options
- TCP kernel modifications
- Test runs

Conclusion

Questions

Thanks

Setting

Jason submitted a paper called “Passive TCP Stream Estimation of RTT and Jitter Parameters”, rejected due to “lack of experimental verification”.

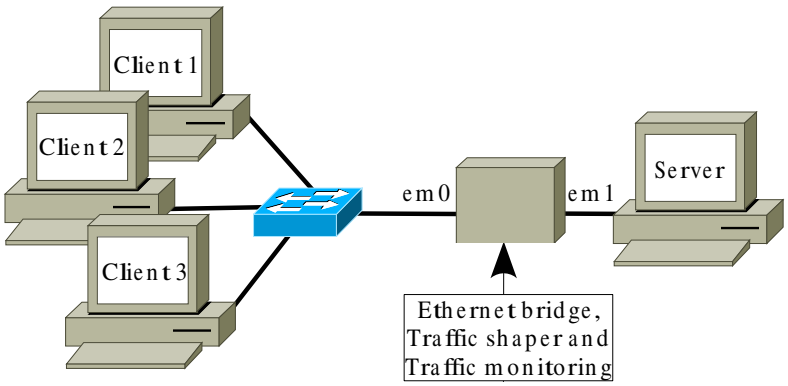

- Clients upload (via scp) data to the server. Various configurations were used.
- TCP stack and Netsniff estimated RTTs are observed

TCP options

- Talk abstract
- Talk outline

Review

Paper 2: ACE2005

Paper 1: RTT
- Setting
- TCP options
- TCP kernel modifications
- Test runs

Conclusion

Questions

Thanks

We want to compare the TCP stack estimated RTT and Netsniff's estimate.

- Put data to observe in TCP header on clients
  → modifications in the FreeBSD kernel

- Read data to observe with Netsniff
  → modifications in Netsniff

TCP kernel modifications

- Talk abstract
- Talk outline

Review

Paper 2: ACE2005

Paper 1: RTT
- Setting
- TCP options
- TCP kernel modifications
- Test runs

Conclusion

Questions

Thanks

- FreeBSD
  - netinet/tcp.h
    → magic numbers for new TCP options (2 lines)
  - netinet/tcp_output.c
    → output of options in TCP header (16 lines)
  - netinet/tcp_hostcache.c
    → flush hostcache and RTTSCALE (50 lines)
  - Compile, Reboot, done
- Netsniff
  - A couple of changes here and there
  - Easier, since our code, we know where goes what.

Test runs

- Talk abstract
- Talk outline

Review

Paper 2: ACE2005

Paper 1: RTT
- Setting
- TCP options
- TCP kernel modifications
- Test runs

Conclusion

Questions

Thanks

Netsniff →

We run around 132 test configurations. Of which three are shown in the paper.

FreeBSD →

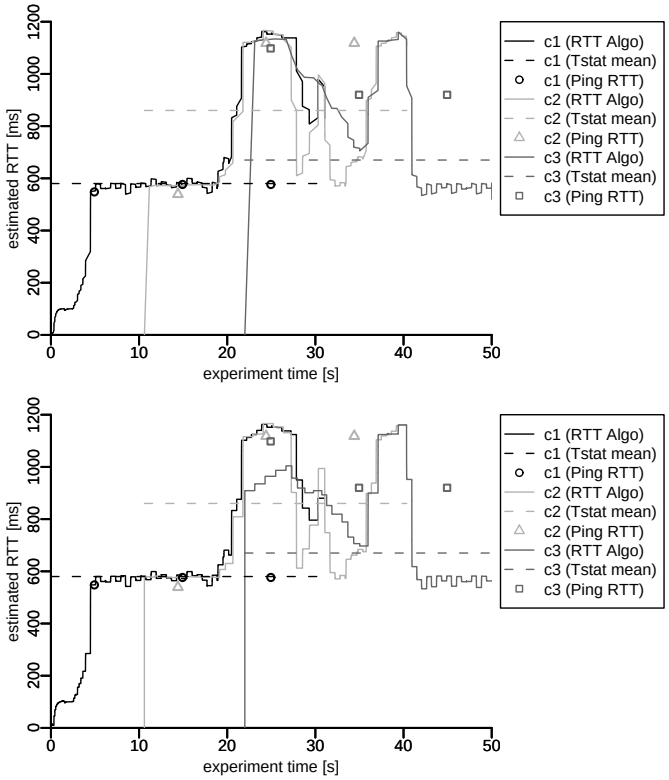

Conclusion

- Talk abstract
- Talk outline

Review

Paper 2: ACE2005

Paper 1: RTT

Conclusion

Questions

Thanks

- Very good and interesting time
- Improved my knowledge of computer networking
- Improved my C++ skills (I come from the Java world ...)
- Got to know quite some knew tools
- Have gained insight in an academic work environment
- I'm top on Google for the search key "Urs Keller" (but only in Australia) :-)

- Talk abstract
- Talk outline

Review

Paper 2: ACE2005

Paper 1: RTT

Conclusion

Questions

Thanks

Questions?

SWINBURNE CENTRE FOR ADVANCED INTERNET ARCHITECTURES

- Talk abstract
- Talk outline

Review

Paper 2: ACE2005

Paper 1: RTT

Conclusion

Questions

Thanks

Thanks