


When Latency Matters and How to Optimize It

Andreas Petlund

Simula Research Laboratory AS

Copenhagen, June 18th 2012



The speed illusion

Confusing bandwidth with speed – a business model

Fast, Reliable, and Co 

Choose a broadband package



Kontakt Om Telenor Job English Tilbudssavis Selvcenter Webmail

Privat Erhverv

Søg

Brug Verden Mobiltelefoner Mobilabonnementer Mobilt Bredbånd **Bredbånd & Telefoni** Mit Telenor Kundeservice

Bredbånd & Telefoni

Fordele som kunde

Skift til Telenor

Hjemmetelefon

Services og bredbåndsstyr

Hastigheder og dækning

Udlandspriser for telefoni

Bredbåndsfakta

Vilkår

Ring mig op

Hurtigt Bredbånd



Trådløs router til 0 kr.

0,-
trådløs router

0 kr. i oprettelse

0,-
i oprettelse



Services Shop Support

Community

Search

How fast is your Internet?

Check your connection now and discover how fast you're really going.

TEST SPEED



FiOS
A NETWORK AHEAD

Is your home ready for America's Fastest Internet?

Check Availability

The speed illusion

Confusing bandwidth with speed – a business model



Get commercial video

The speed illusion

Confusing bandwidth with speed – a business model



My private Get connection induced a 2 second delay when the link was saturated

The modern stock market



Hosting centre



Stock exchange

Information Week (2007): “A 1-millisecond advantage in trading applications can be worth \$100 million a year to a major brokerage firm”.

The modern stock market



Hosting centre

Stock exchange



Information Week (2007): “A 1-millisecond advantage in trading applications can be worth \$100 million a year to a major brokerage firm”.


The money at stake

Hibernia Atlantic



\$300 million investment
Shaves off 6ms from the 65 ms transmission time

Russian cable:
The cost of cutting
London-Tokyo
latency by 60ms:
\$1.5 billion



©NewScientist

Latency and online gaming

Claypool latency video

The cost of delay in web search

- In 2009 Microsoft Bing found that introducing 500ms extra delay translated to 1.2% less advertising revenue.
- Google experimented with injecting 400ms delay before they returned their search page. Initial searches declined linearly over time, dropping by 0.76% after 6 weeks and continuing linearly after that. Interestingly, once the artificially introduced delay was removed, it took a similar period to linearly regain the original demand.
- Performing this experiment on their whole user base would have cost Google **\$75M** per year.

400ms delay = \$75.000.000

What's the problem?

Why interactive applications is challenging

Dept. of wild claims: Interactive applications produce data patterns that are fundamentally different from what the networks were designed for.

The nature of interactivity:


- Transmit when it happens.
- The traffic is “hiding”
 - Very small share of traffic,
 - but large share of streams.

The previous focus on throughput:

- Unbalanced technology

The shared medium of the Internet :

- No guarantees.




Interactive traffic patterns

application	payload size (bytes)			packet interarrival time (ms)						avg bandwidth used	
	avg	min	max	avg	med	min	max	percentiles		(pps)	(bps)
								1%	99%		
Casa (sensor network)	175	93	572	7287	307	305	29898	305	29898	0.137	269
Windows remote desktop	111	8	1417	318	159	1	12254	2	3892	3.145	4497
VNC (from client)	8	1	106	34	8	< 1	5451	< 1	517	29.412	17K
Skype (2 users) (TCP)	236	14	1267	34	40	< 1	1671	4	80	29.412	69K
SSH text session	48	16	752	323	159	< 1	76610	32	3616	3.096	2825
Anarchy Online	98	8	1333	632	449	7	17032	83	4195	1.582	2168
World of Warcraft	26	6	1228	314	133	< 1	14855	< 1	3785	3.185	2046
Age of Conan	80	5	1460	86	57	< 1	1375	24	386	11.628	12K
YouTube stream	1446	112	1448	9	< 1	< 1	1335	< 1	127	111.111	1278K
HTTP download	1447	64	1448	< 1	< 1	< 1	186	< 1	8	> 1000	14M
FTP download	1447	40	1448	< 1	< 1	< 1	339	< 1	< 1	> 1000	82M

Event-driven, not congestion-control restricted.

Extra vulnerable to excess delay and retransmissions.

The extra delay



How to get there?

Alternative 1: The perfect solution, but oh so unlikely

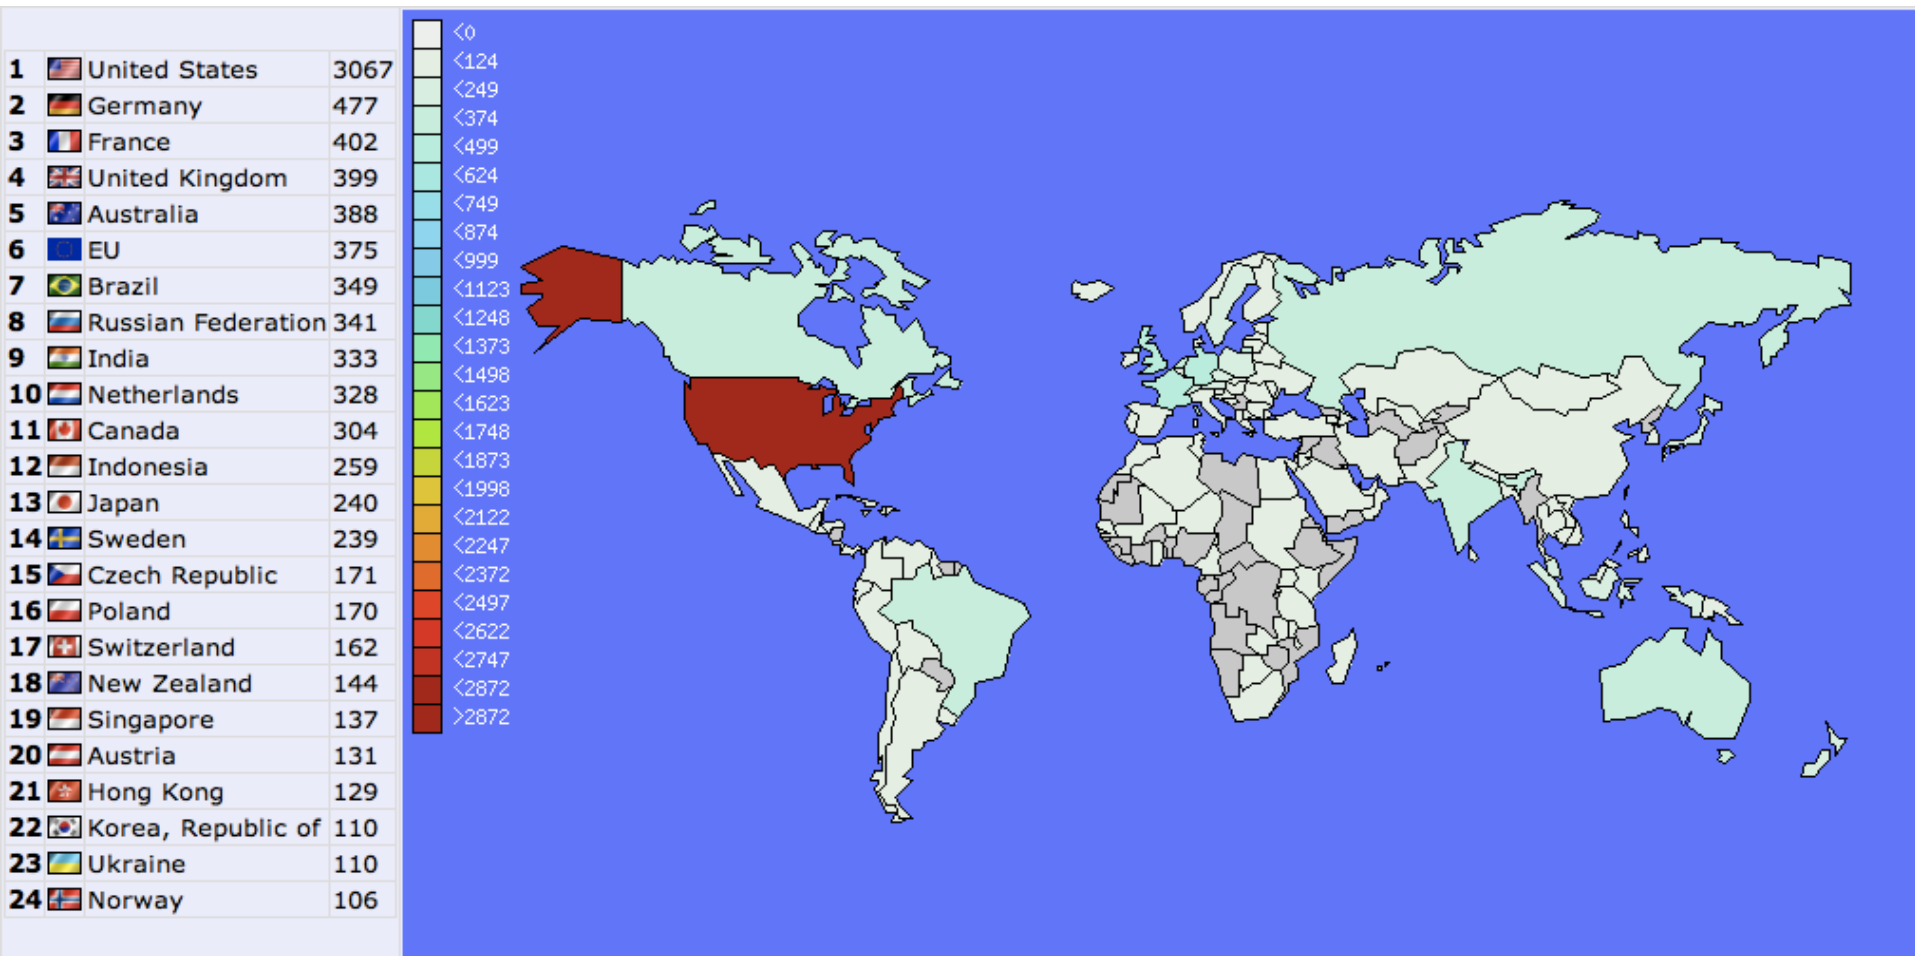
Get everybody to reserve dedicated links for low-latency traffic:

- More of a political than technical problem.
- Will need worldwide support to produce proper results.
- May well come at the expense of robustness.
- Terribly expensive.
- Prognosis: very unlikely.



Ease of deployment for all-in solutions

A critical use case: IPv6



IPv6 prefixes per country June 2012.

The incremental approach

- Awareness-raising:
 - Bufferbloat-project: low-latency evangelism –well placed
 - Much can be done by configuring systems right.
- Incremental approaches:
 - Create mechanisms that will have a positive effect also if only some nodes support them.
- AQMs, self-configuring AQMs and network hints.
 - ECN
 - DCTCP
 - Van Jacobson / Kathleen Nichols – “Codel”
 - +++
- Sender-side only modifications.



Some ongoing standardisation efforts

TCP efforts:

- Sending data with SYN packets: draft-ietf-tcp-fastopen
- Avoid excess retransmission delays:
 - Early retransmit: RFC5827
 - draft-hurtig-tcpm-rtorestart-01
 - Thin-stream retransmissions: Upcoming
- AB/IRTF Workshop on Congestion Control for Interactive Real-Time Comm.
- RTCWeb
- Google – Increased initial window (short-lived streams initiative): draft-hkchu-tcpm-initcwnd-01
- SPDY: draft-mbelshe-httpbis-spdy-00
- Multipath efforts: MPTCP, SCTP
- Network and AQM initiatives.

What can be done at the providers?

Business model approaches:

- Nuance the bandwidth=speed wording in campaign.
 - Big help in awareness-raising!
- Create a market for latency-driven services?
 - Provide a low-latency option to customers as a chargeable extra service

Technical approaches:

- Reduce buffer sizes.
- Use self-configuring AQM schemes for low latency.
- Learn how to configure static AQM schemes for low latency.
- Support network hint systems and network/end-host collaborative schemes.

Help yourself to lower latency

Tune your buffers:

- Minimise host buffers (where applicable)
- Minimise router buffers:
 - Lower link utilisation for a single stream.
 - Much better latency.
 - Many routers have QoS options.

Thin-stream tuning:

- Optimize gateways: Byte-based queue, not multiple of MSS*
- Use thin-stream optimizations (Linux: `sysctl -a | grep tcp_thin`)
- Early retransmit.
- Turn off delayed ACKs
- Turn off Nagle's algorithm.

* Multimedia-unfriendly TCP Congestion Control and Home Gateway Queue Management – MMSYS 2011

A retransmission use-case: Funcom

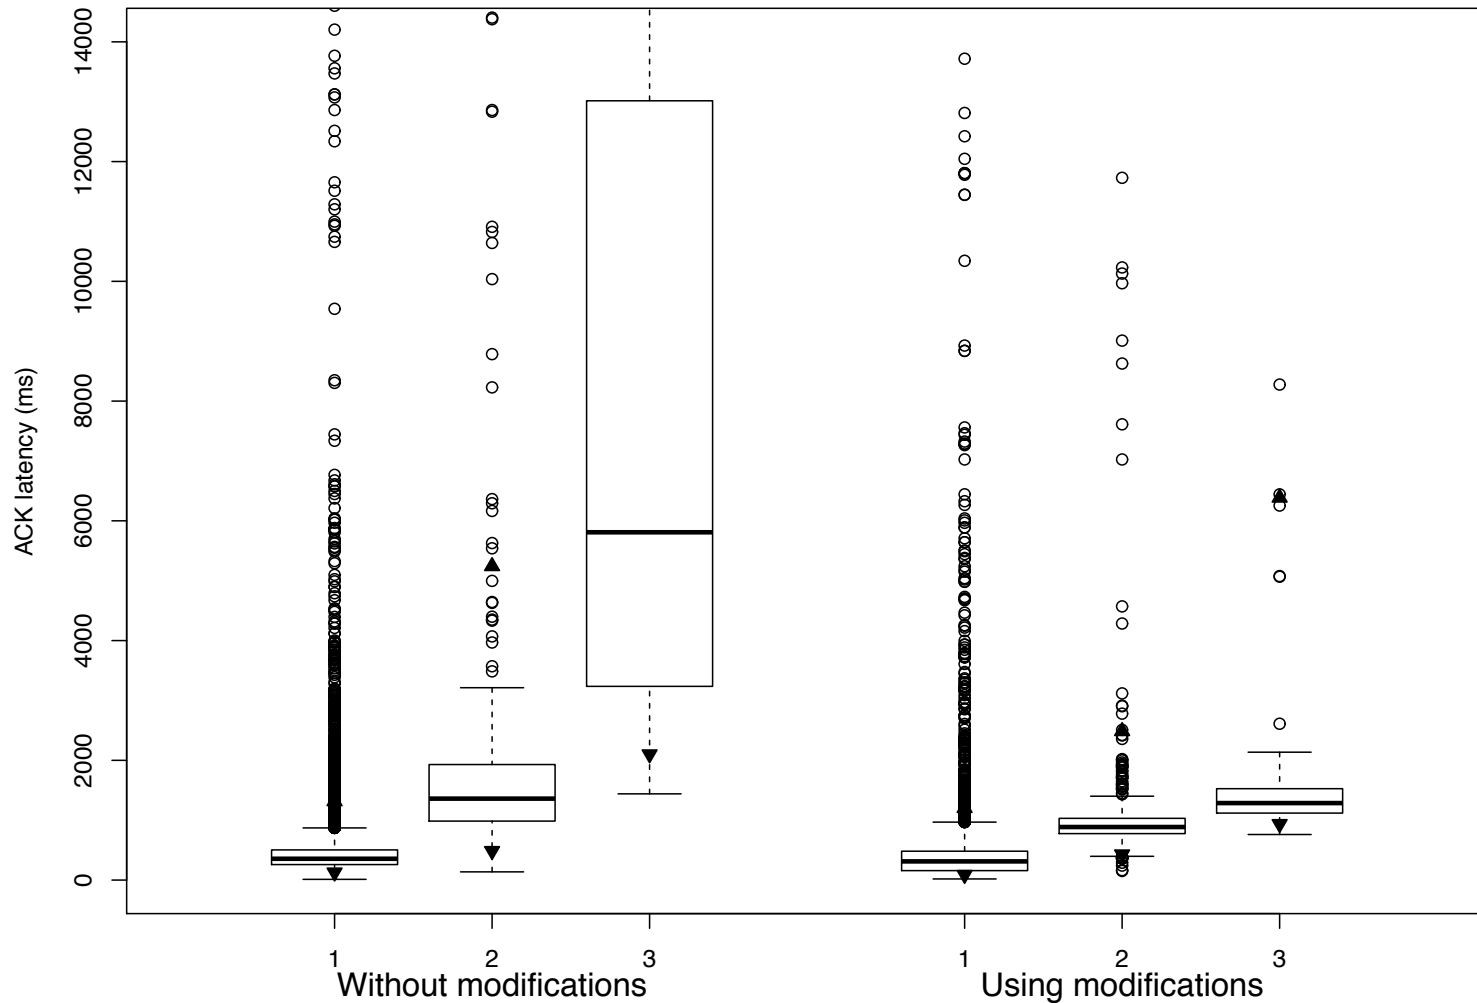
Age of Conan deployment of thin-stream mechanisms



A retransmission use-case: Funcom

Age of Conan deployment of thin-stream mechanisms

Application latency – AoC 1hr trace



Where's the killer apps?

Dpt. of wild speculations – Mobile interactivity



Today:

- Turn-based games
- Latency-resistant games
- Networked interaction (Google docs ++)
- 2-participant video chat

Tomorrow:

- Proper MMOs?
- 1st person shooters?
- Multi-participant virtual interaction?
- +++

Where's the killer apps?

Dpt. of wild speculations – reliable low-delay Internet



Q&A

