One dissertation each...

Palin (right): *"Ph.D defence? - good, out of the door, line on the left, one dissertation each..."*

Idle (left): *"Er no, freedom, freedom for me, they said I hadn't done anything and I could go free and live on an island somewhere"*

Palin: *"Oh, oh. jolly good, well off you go"*

Idle: *"Nah, only pulling your leg, its a Ph.D defence for me!"*

Quality aspects of Internet telephony

Ian Marsh

Ph.D defence
June 5th 2009

The structure of the presentation

The structure of the presentation

- A very brief history of Internet telephony development

The structure of the presentation

- A very brief history of Internet telephony development
- Some words about deployment

The structure of the presentation

- A very brief history of Internet telephony development
- Some words about deployment
- Problems facing Internet telephony quality

The structure of the presentation

- A very brief history of Internet telephony development
- Some words about deployment
- Problems facing Internet telephony quality
- A very simplified walk-through of a Voice over IP system

The structure of the presentation

- A very brief history of Internet telephony development
- Some words about deployment
- Problems facing Internet telephony quality
- A very simplified walk-through of a Voice over IP system
- What quality *really* means

The structure of the presentation

- A very brief history of Internet telephony development
- Some words about deployment
- Problems facing Internet telephony quality
- A very simplified walk-through of a Voice over IP system
- What quality *really* means
- The dissertation's contributions

The structure of the presentation

- A very brief history of Internet telephony development
- Some words about deployment
- Problems facing Internet telephony quality
- A very simplified walk-through of a Voice over IP system
- What quality *really* means
- The dissertation's contributions
- The dissertations' conclusions

A historical time-line of packetised voice

A historical time-line of packetised voice

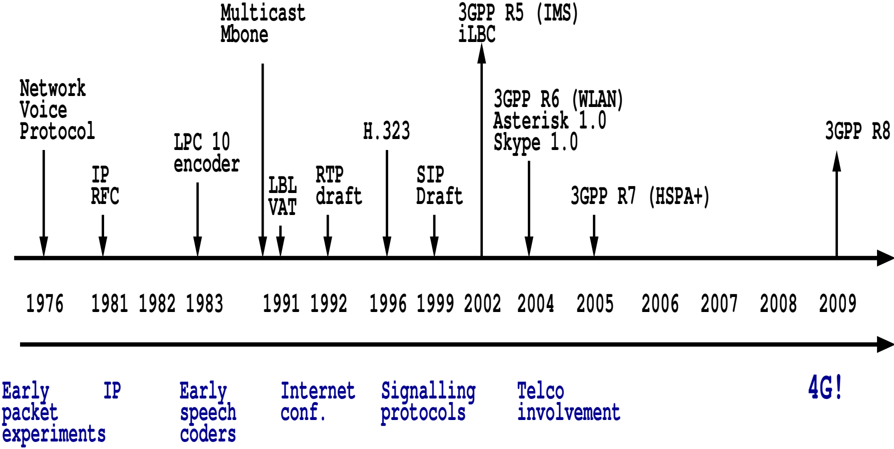

From a graphical perspective (1992 and 2005)

From a graphical perspective (1992 and 2005)

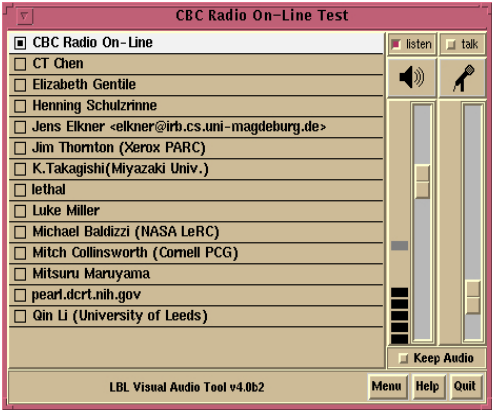

From a graphical perspective (1992 and 2005)

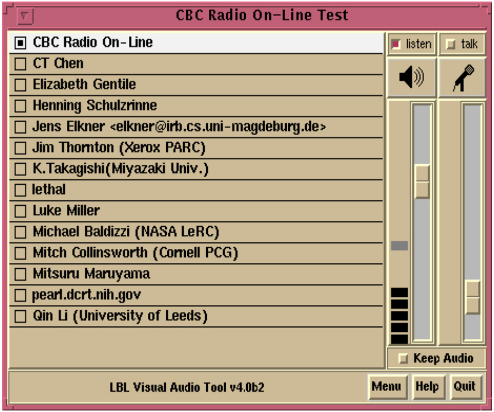

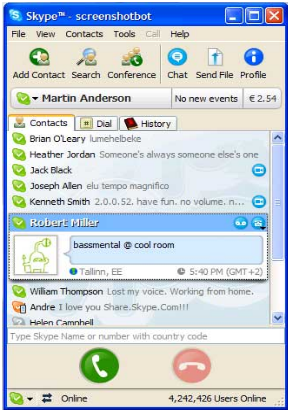

IP telephony deployment

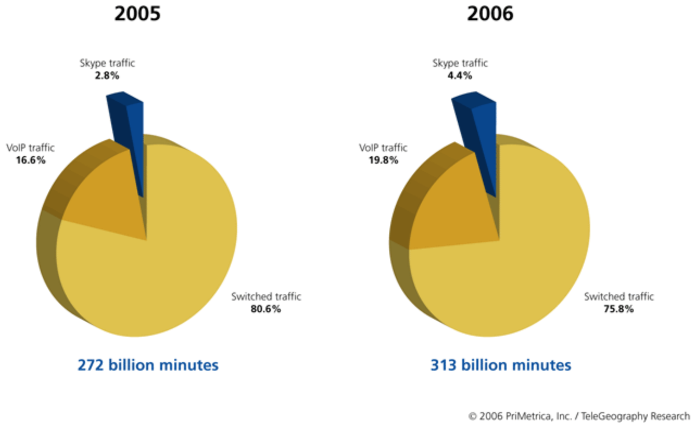

IP telephony deployment

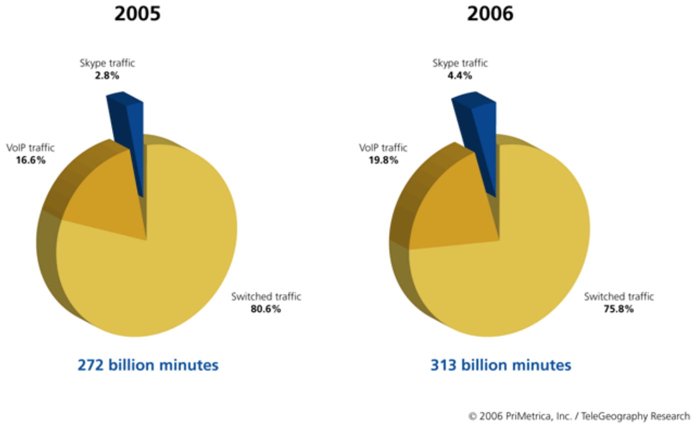

- Telephony grew 15% (in total) from 2005 to 2006

IP telephony deployment

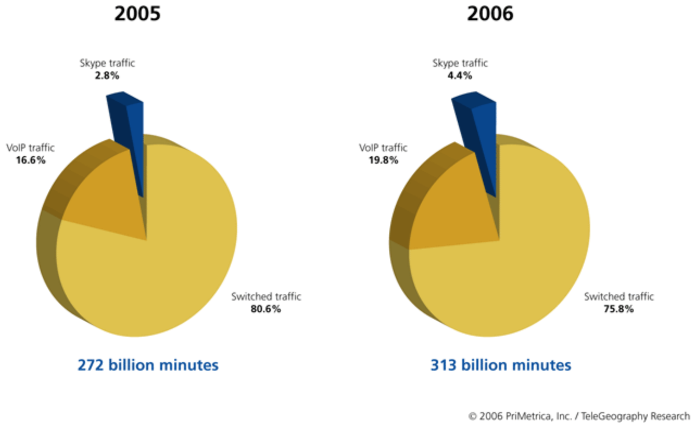

- Telephony grew 15% (in total) from 2005 to 2006
- VoIP grew from 19.4% in 2005 to 24.2% during the same period

IP telephony deployment

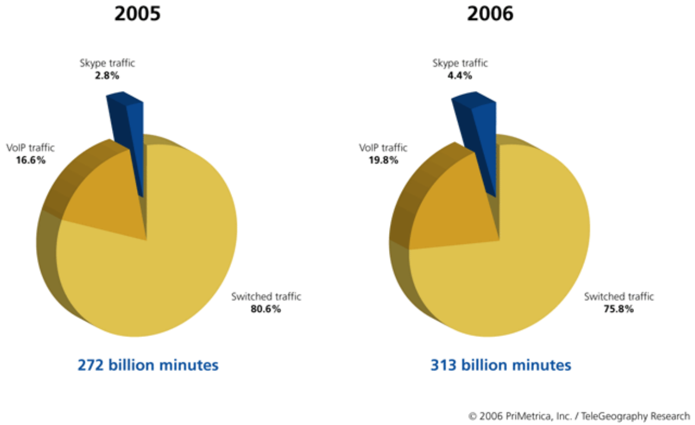

- Telephony grew 15% (in total) from 2005 to 2006
- VoIP grew from 19.4% in 2005 to 24.2% during the same period
- This still leaves 75% for migration

IP telephony deployment

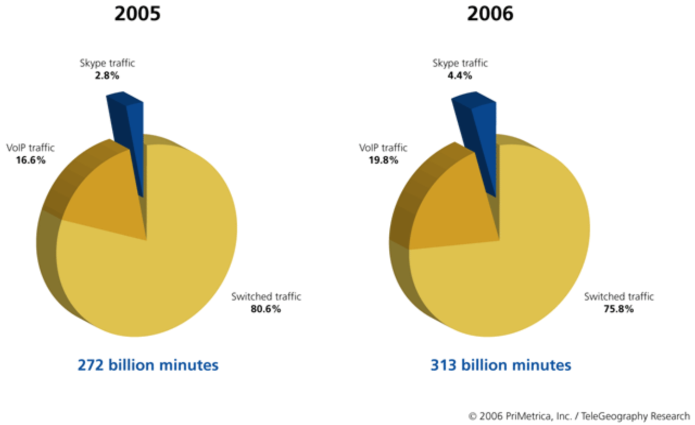

- Telephony grew 15% (in total) from 2005 to 2006
- VoIP grew from 19.4% in 2005 to 24.2% during the same period
- This still leaves 75% for migration
- And motivation for this work!

Problems facing VoIP quality

First a *non-problem*. There is nothing inherently problematic with sending real-time voice across a packet switched network, however, there is

Problems facing VoIP quality

First a *non-problem*. There is nothing inherently problematic with sending real-time voice across a packet switched network, however, there is

1. Disruption caused by competing traffic
   - At the single flow level, TCP induces loss by increasing its sending rate until a loss event
   - Increasing volumes of traffic, specifically, P2P, video streaming and IPTV, *without* appropriate capacity changes

Problems facing VoIP quality

First a *non-problem*. There is nothing inherently problematic with sending real-time voice across a packet switched network, however, there is

1. Disruption caused by competing traffic
   - At the single flow level, TCP induces loss by increasing its sending rate until a loss event
   - Increasing volumes of traffic, specifically, P2P, video streaming and IPTV, *without* appropriate capacity changes
2. 802.11-based communication
   - Interference, environment & mobility can induce problems using *unlicensed* spectrum
   - Poor quality radios in some wireless terminals

Problems facing VoIP quality

First a *non-problem*. There is nothing inherently problematic with sending real-time voice across a packet switched network, however, there is

1. Disruption caused by competing traffic
   - At the single flow level, TCP induces loss by increasing its sending rate until a loss event
   - Increasing volumes of traffic, specifically, P2P, video streaming and IPTV, *without* appropriate capacity changes
2. 802.11-based communication
   - Interference, environment & mobility can induce problems using *unlicensed* spectrum
   - Poor quality radios in some wireless terminals
3. Additional problems due to
   - Poor quality infrastructure in some countries
   - Poorly designed end-systems
   - Individual human tolerances

Simple VoIP walk-through

- On the next slide I am going to explain the path that real-time conversations take across the Internet from mouth to ear

Simple VoIP walk-through

- On the next slide I am going to explain the path that real-time conversations take across the Internet from mouth to ear
- I will describe:
    1. The important processing steps
    2. How loss, delay and jitter are introduced into a VoIP stream

Simple VoIP walk-through

- On the next slide I am going to explain the path that real-time conversations take across the Internet from mouth to ear
- I will describe:
  1. The important processing steps
  2. How loss, delay and jitter are introduced into a VoIP stream
- The illustration will be used to describe, in turn, each of the dissertations' contributions later

A simple VoIP walk-through

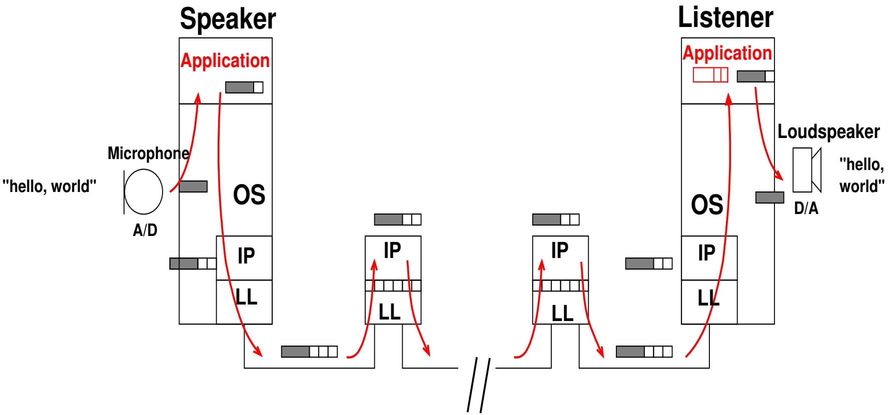

What quality *really* means...

What quality *really* means...

According to an ITU standard, **one-way delay** can be quantified as:

- "good" if the delay is less than 180 ms
- "acceptable" if the delay lies between 180 ms and 400 ms
- "unacceptable" if the delay exceeds 400 ms

What quality *really* means...

According to an ITU standard, **one-way delay** can be quantified as:

- "good" if the delay is less than 180 ms
- "acceptable" if the delay lies between 180 ms and 400 ms
- "unacceptable" if the delay exceeds 400 ms

Loss

- For telephony voice, acceptable **losses** can range from 1-10%
- Losses can be "hidden" using concealment techniques

What quality *really* means...

According to an ITU standard, **one-way delay** can be quantified as:

- "good" if the delay is less than 180 ms
- "acceptable" if the delay lies between 180 ms and 400 ms
- "unacceptable" if the delay exceeds 400 ms

Loss

- For telephony voice, acceptable **losses** can range from 1-10%
- Losses can be "hidden" using concealment techniques

What you are about to hear are 3 speech samples with:

What quality *really* means...

According to an ITU standard, **one-way delay** can be quantified as:

- "good" if the delay is less than 180 ms
- "acceptable" if the delay lies between 180 ms and 400 ms
- "unacceptable" if the delay exceeds 400 ms

Loss

- For telephony voice, acceptable **losses** can range from 1-10%
- Losses can be "hidden" using concealment techniques

What you are about to hear are 3 speech samples with:

1. The original sample

What quality *really* means...

According to an ITU standard, **one-way delay** can be quantified as:

- "good" if the delay is less than 180 ms
- "acceptable" if the delay lies between 180 ms and 400 ms
- "unacceptable" if the delay exceeds 400 ms

Loss

- For telephony voice, acceptable **losses** can range from 1-10%
- Losses can be "hidden" using concealment techniques

What you are about to hear are 3 speech samples with:

1. The original sample
2. The original corrupted by 2% loss

What quality *really* means...

According to an ITU standard, **one-way delay** can be quantified as:

- "good" if the delay is less than 180 ms
- "acceptable" if the delay lies between 180 ms and 400 ms
- "unacceptable" if the delay exceeds 400 ms

Loss

- For telephony voice, acceptable **losses** can range from 1-10%
- Losses can be "hidden" using concealment techniques

What you are about to hear are 3 speech samples with:

1. The original sample
2. The original corrupted by 2% loss
3. The original corrupted by 5% loss

What quality *really* means...

According to an ITU standard, **one-way delay** can be quantified as:

- "good" if the delay is less than 180 ms
- "acceptable" if the delay lies between 180 ms and 400 ms
- "unacceptable" if the delay exceeds 400 ms

Loss

- For telephony voice, acceptable **losses** can range from 1-10%
- Losses can be "hidden" using concealment techniques

What you are about to hear are 3 speech samples with:

1. The original sample
2. The original corrupted by 2% loss
3. The original corrupted by 5% loss

Jitter

- Buffers are required to handle jitter in VoIP systems
- Therefore they result in **additional** delay and/or loss

Now, presentation of the contributions

- Each contribution will now be shown in red using the previous figure

Now, presentation of the contributions

- Each contribution will now be shown in red using the previous figure
- Presentations will be brief, please consult the dissertation for full details

Protecting VoIP traffic (paper A)

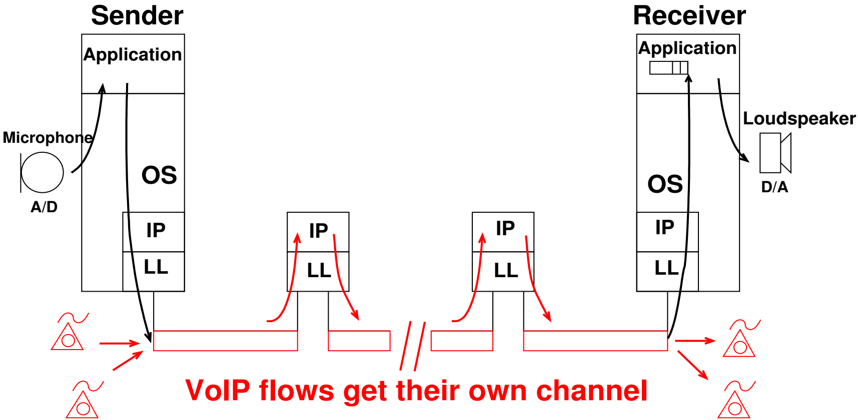

Protecting VoIP traffic

- Contribution: Three methods to dimensioning a link for statistically multiplexed VoIP flows

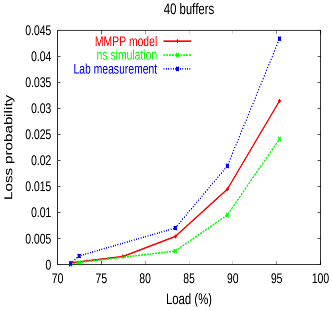

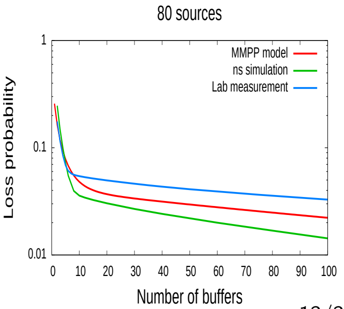

Protecting VoIP traffic

- Contribution: Three methods to dimensioning a link for statistically multiplexed VoIP flows
- The three colours show:
  1. A FreeBSD implementation (blue)
  2. A Markov-Modulated Poisson arrival model (red)
  3. A ns-2 simulation (green)

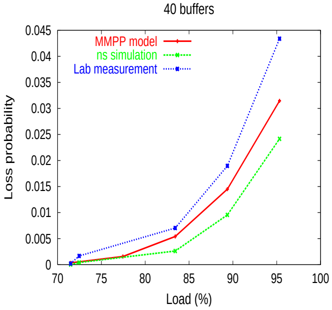

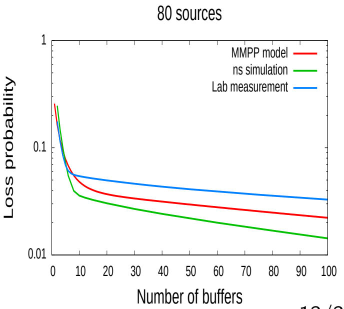

Protecting VoIP traffic

- Contribution: Three methods to dimensioning a link for statistically multiplexed VoIP flows
- The three colours show:
  1. A FreeBSD implementation (blue)
  2. A Markov-Modulated Poisson arrival model (red)
  3. A ns-2 simulation (green)
- Top figure is the loss probability against differing loads, calculated as follows:

$$Load = \frac{\#Sources \cdot ProbOn \cdot PeakRate}{Link \;\; capacity}$$

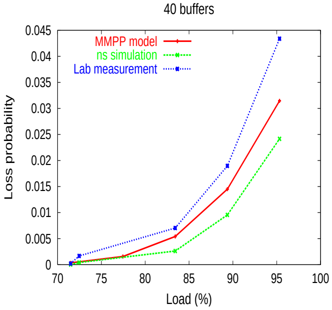

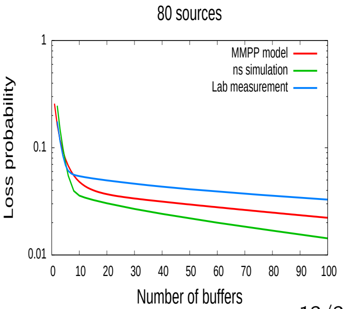

Protecting VoIP traffic

- Contribution: Three methods to dimensioning a link for statistically multiplexed VoIP flows
- The three colours show:
  1. A FreeBSD implementation (blue)
  2. A Markov-Modulated Poisson arrival model (red)
  3. A ns-2 simulation (green)
- Top figure is the loss probability against differing loads, calculated as follows:

$$Load = \frac{\#Sources \cdot ProbOn \cdot PeakRate}{Link \;\; capacity}$$

- Peak rate allocation would yield 25 sources for 85% load, whereas statistical multiplexing allows 72 sources

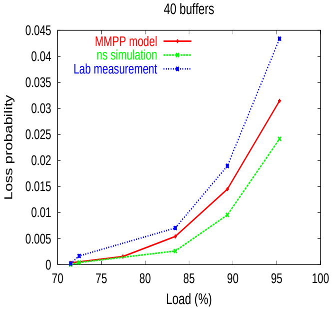

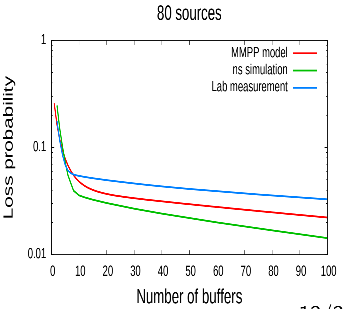

Protecting VoIP traffic

- Contribution: Three methods to dimensioning a link for statistically multiplexed VoIP flows
- The three colours show:
  1. A FreeBSD implementation (blue)
  2. A Markov-Modulated Poisson arrival model (red)
  3. A ns-2 simulation (green)
- Top figure is the loss probability against differing loads, calculated as follows:

$$Load = \frac{\#Sources \cdot ProbOn \cdot PeakRate}{Link\ \ capacity}$$

- Peak rate allocation would yield 25 sources for 85% load, whereas statistical multiplexing allows 72 sources
- Lower plot show loss against buffer sizes

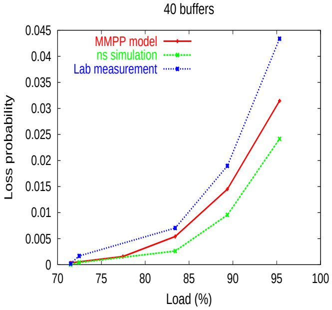

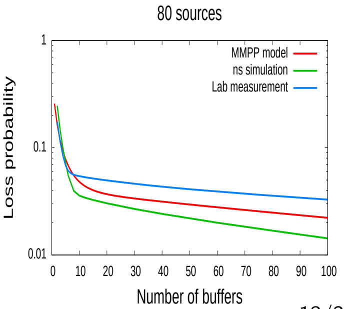

Mixing VoIP and data traffic (paper B)

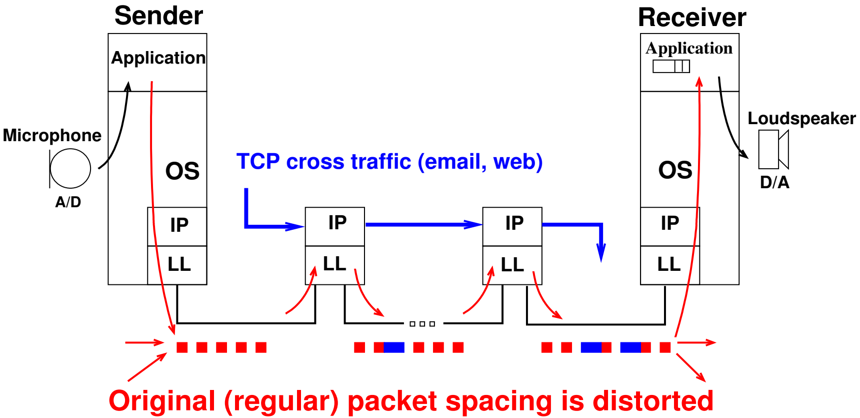

Mixing VoIP and data traffic (paper B)

- Contribution: Extensible model describing the arrival process of VoIP packets

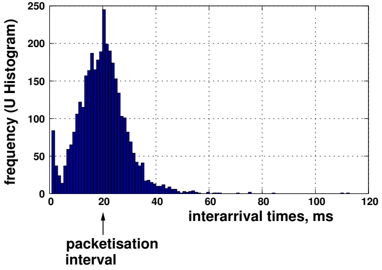

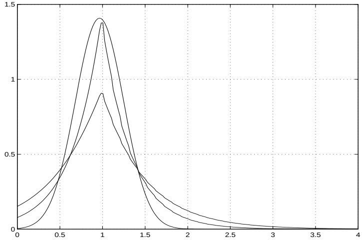

Mixing VoIP and data traffic (paper B)

- Contribution: Extensible model describing the arrival process of VoIP packets
- Most of the traffic on the Internet is TCP, therefore VoIP must coexist with bulk data

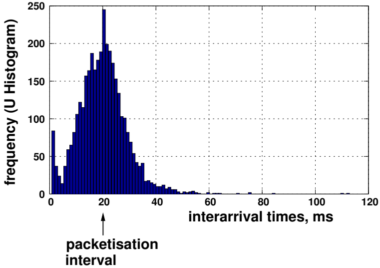

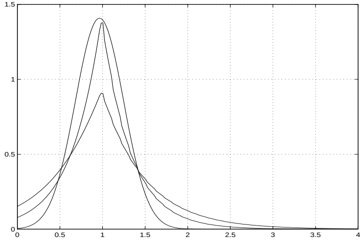

Mixing VoIP and data traffic (paper B)

- Contribution: Extensible model describing the arrival process of VoIP packets
- Most of the traffic on the Internet is TCP, therefore VoIP must coexist with bulk data
- Top plot shows the interarrival times for a single VoIP session

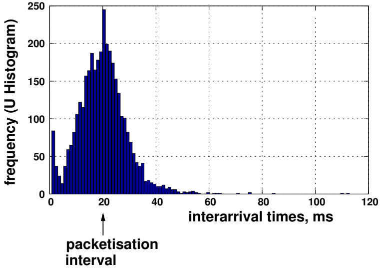

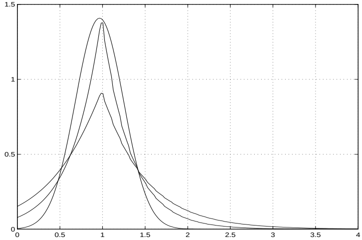

Mixing VoIP and data traffic (paper B)

- Contribution: Extensible model describing the arrival process of VoIP packets
- Most of the traffic on the Internet is TCP, therefore VoIP must coexist with bulk data
- Top plot shows the interarrival times for a single VoIP session
- The interarrival times of a packet delay *model* have been generated in the bottom plot

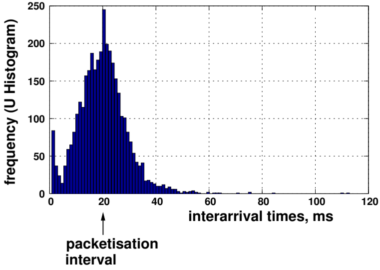

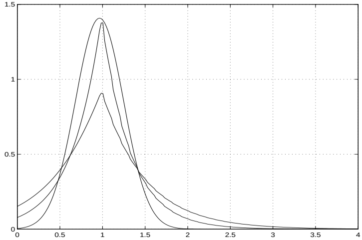

Mixing VoIP and data traffic (paper B)

- Contribution: Extensible model describing the arrival process of VoIP packets
- Most of the traffic on the Internet is TCP, therefore VoIP must coexist with bulk data
- Top plot shows the interarrival times for a single VoIP session
- The interarrival times of a packet delay *model* have been generated in the bottom plot
- Features:
  - $x = 0$ (origin) back to back packets
  - $x = 1$ (maximum) packet. interval
  - $x > 1$ (tail) of the distribution

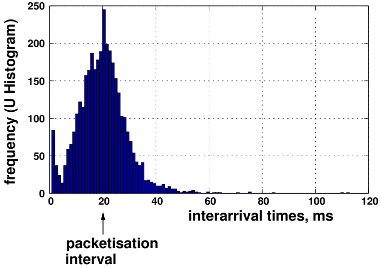

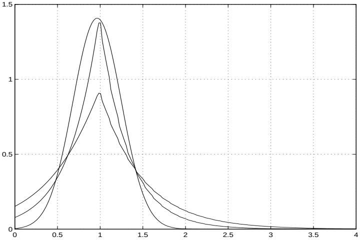

Mixing VoIP and data traffic (paper B)

- Contribution: Extensible model describing the arrival process of VoIP packets
- Most of the traffic on the Internet is TCP, therefore VoIP must coexist with bulk data
- Top plot shows the interarrival times for a single VoIP session
- The interarrival times of a packet delay *model* have been generated in the bottom plot
- Features:
  - $x = 0$ (origin) back to back packets
  - $x = 1$ (maximum) packet. interval
  - $x > 1$ (tail) of the distribution
- The model requires an *estimate* of the networks delay distribution (3 shown)

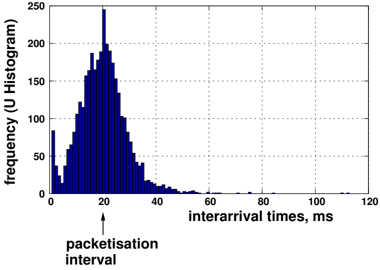

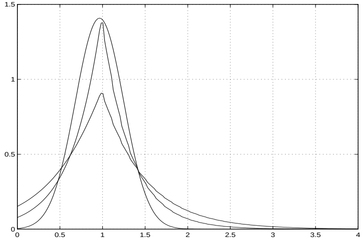

Measuring wide-area VoIP quality (papers D and E)

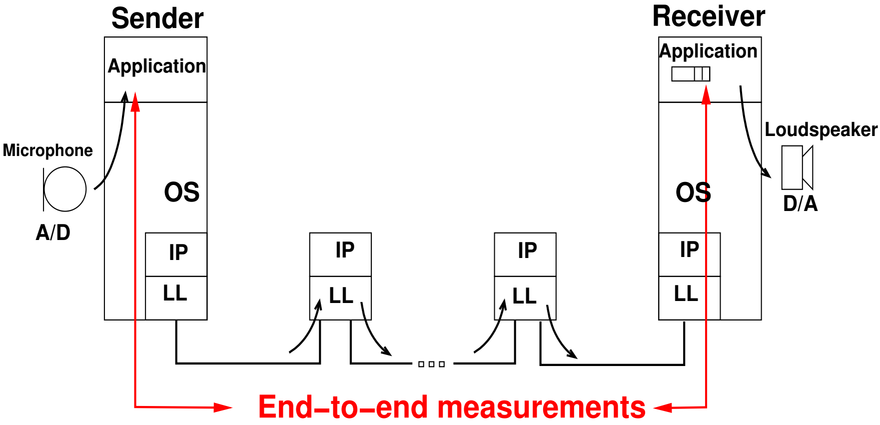

Measuring wide-area VoIP quality (papers D and E)

- Contribution: comprehensive measurement VoIP effort from 1998 and 2002

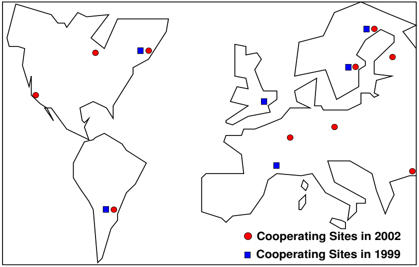

Transmitted session	
Call duration	70 seconds
Payload size	160 bytes
Packetisation time (ms)	20ms
Data rate	64kbits/sec
Without silence suppression	3653 packets
With silence suppression	2043 packets
Coding	8 bit PCM
Recorded call size	584480 bytes
Obtained data	
Number of hosts used	9
Number of traces obtained	22436
Number of data packets	32,771,021
Total data size (compressed)	411 Megabytes
Measurement duration	12 weeks

Measuring wide-area VoIP quality (papers D and E)

- Contribution: comprehensive measurement VoIP effort from 1998 and 2002
- Measured:
  - Loss
  - Delay
  - Jitter
  - Hops

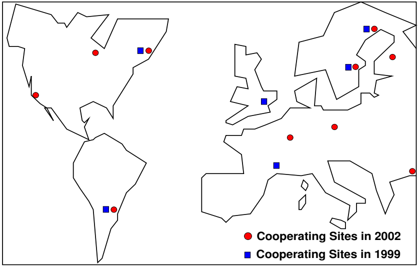

Transmitted session	
Call duration	70 seconds
Payload size	160 bytes
Packetisation time (ms)	20ms
Data rate	64kbits/sec
Without silence suppression	3653 packets
With silence suppression	2043 packets
Coding	8 bit PCM
Recorded call size	584480 bytes
Obtained data	
Number of hosts used	9
Number of traces obtained	22436
Number of data packets	32,771,021
Total data size (compressed)	411 Megabytes
Measurement duration	12 weeks

Measuring wide-area VoIP quality (papers D and E)

- Contribution: comprehensive measurement VoIP effort from 1998 and 2002
- Measured:
  - Loss
  - Delay
  - Jitter
  - Hops
- Hosts located at universities

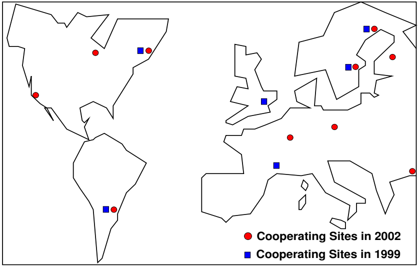

Transmitted session	
Call duration	70 seconds
Payload size	160 bytes
Packetisation time (ms)	20ms
Data rate	64kbits/sec
Without silence suppression	3653 packets
With silence suppression	2043 packets
Coding	8 bit PCM
Recorded call size	584480 bytes
Obtained data	
Number of hosts used	9
Number of traces obtained	22436
Number of data packets	32,771,021
Total data size (compressed)	411 Megabytes
Measurement duration	12 weeks

Measuring wide-area VoIP quality (papers D and E)

- Contribution: comprehensive measurement VoIP effort from 1998 and 2002
- Measured:
  - Loss
  - Delay
  - Jitter
  - Hops
- Hosts located at universities
- Most *sessions* showed good quality

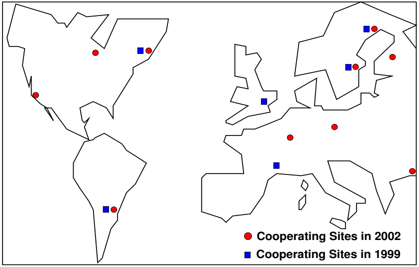

Transmitted session	
Call duration	70 seconds
Payload size	160 bytes
Packetisation time (ms)	20ms
Data rate	64kbits/sec
Without silence suppression	3653 packets
With silence suppression	2043 packets
Coding	8 bit PCM
Recorded call size	584480 bytes
Obtained data	
Number of hosts used	9
Number of traces obtained	22436
Number of data packets	32,771,021
Total data size (compressed)	411 Megabytes
Measurement duration	12 weeks

Measuring wide-area VoIP quality (papers D and E)

- Contribution: comprehensive measurement VoIP effort from 1998 and 2002
- Measured:
  - Loss
  - Delay
  - Jitter
  - Hops
- Hosts located at universities
- Most *sessions* showed good quality
- One *site* did not provide adequate quality

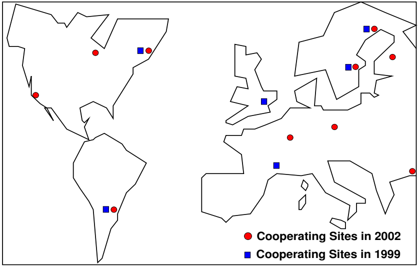

Transmitted session	
Call duration	70 seconds
Payload size	160 bytes
Packetisation time (ms)	20ms
Data rate	64kbits/sec
Without silence suppression	3653 packets
With silence suppression	2043 packets
Coding	8 bit PCM
Recorded call size	584480 bytes
Obtained data	
Number of hosts used	9
Number of traces obtained	22436
Number of data packets	32,771,021
Total data size (compressed)	411 Megabytes
Measurement duration	12 weeks

Measuring wide-area VoIP quality (papers D and E)

- Contribution: comprehensive measurement VoIP effort from 1998 and 2002
- Measured:
  - Loss
  - Delay
  - Jitter
  - Hops
- Hosts located at universities
- Most *sessions* showed good quality
- One *site* did not provide adequate quality
- Quality has slightly improved for the same sites since 1998

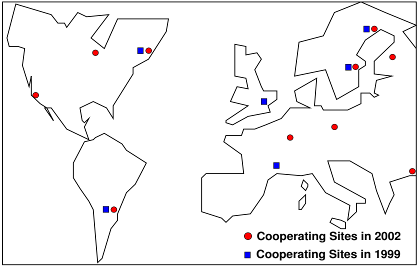

Transmitted session	
Call duration	70 seconds
Payload size	160 bytes
Packetisation time (ms)	20ms
Data rate	64kbits/sec
Without silence suppression	3653 packets
With silence suppression	2043 packets
Coding	8 bit PCM
Recorded call size	584480 bytes
Obtained data	
Number of hosts used	9
Number of traces obtained	22436
Number of data packets	32,771,021
Total data size (compressed)	411 Megabytes
Measurement duration	12 weeks

Measuring wide-area VoIP quality (papers D and E)

- Contribution: comprehensive measurement VoIP effort from 1998 and 2002
- Measured:
  - Loss
  - Delay
  - Jitter
  - Hops
- Hosts located at universities
- Most *sessions* showed good quality
- One *site* did not provide adequate quality
- Quality has slightly improved for the same sites since 1998
- Data used in papers B and F

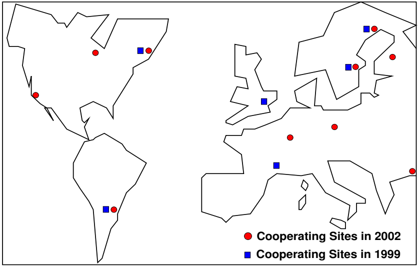

Transmitted session	
Call duration	70 seconds
Payload size	160 bytes
Packetisation time (ms)	20ms
Data rate	64kbits/sec
Without silence suppression	3653 packets
With silence suppression	2043 packets
Coding	8 bit PCM
Recorded call size	584480 bytes
Obtained data	
Number of hosts used	9
Number of traces obtained	22436
Number of data packets	32,771,021
Total data size (compressed)	411 Megabytes
Measurement duration	12 weeks

Self-admission control (paper F)

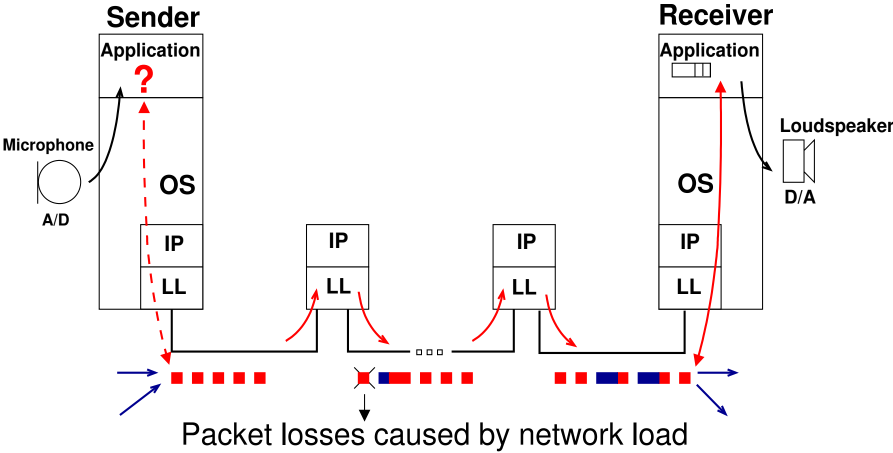

Self-admission control (paper F)

- Contribution: Free-standing endpoint admission control mechanism

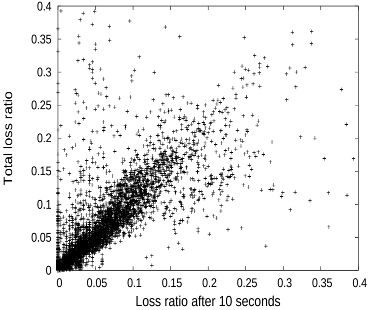

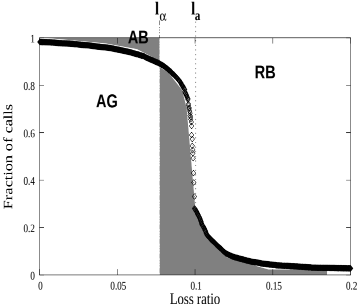

Self-admission control (paper F)

- Contribution: Free-standing endpoint admission control mechanism
- Goal is to estimate the quality of a call from its initial 'few' seconds

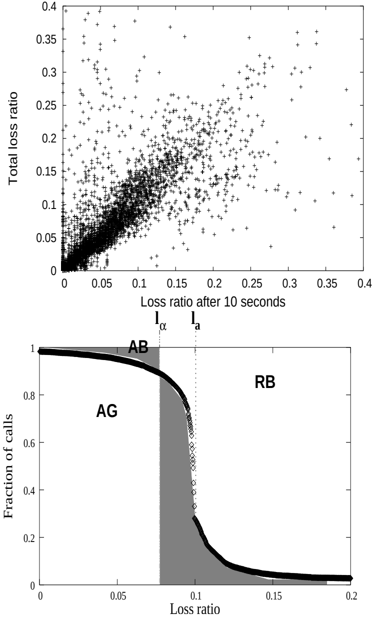

Self-admission control (paper F)

- Contribution: Free-standing endpoint admission control mechanism
- Goal is to estimate the quality of a call from its initial 'few' seconds
- Top plot shows the loss process for a number of calls

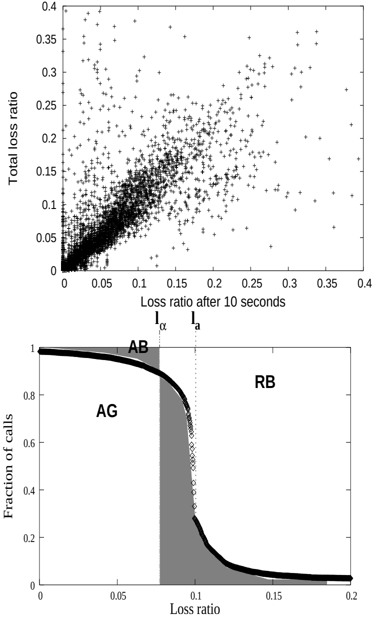

Self-admission control (paper F)

- Contribution: Free-standing endpoint admission control mechanism
- Goal is to estimate the quality of a call from its initial 'few' seconds
- Top plot shows the loss process for a number of calls
- Given the decision to admit or reject a flow, there are four possible outcomes:
  - Accepted and good quality (AG)
  - Accepted but bad quality (AB)
  - Rejected and bad quality (RB)
  - Rejected but good quality (RG)

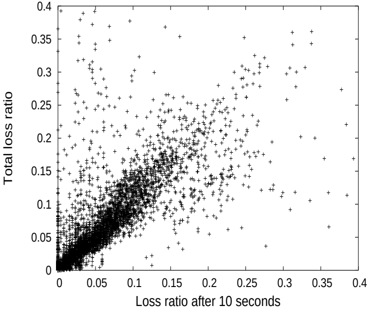

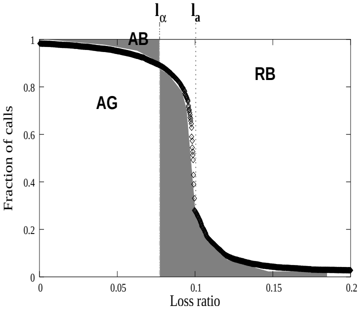

Self-admission control (paper F)

- Contribution: Free-standing endpoint admission control mechanism
- Goal is to estimate the quality of a call from its initial 'few' seconds
- Top plot shows the loss process for a number of calls
- Given the decision to admit or reject a flow, there are four possible outcomes:
  - Accepted and good quality (AG)
  - Accepted but bad quality (AB)
  - Rejected and bad quality (RB)
  - Rejected but good quality (RG)

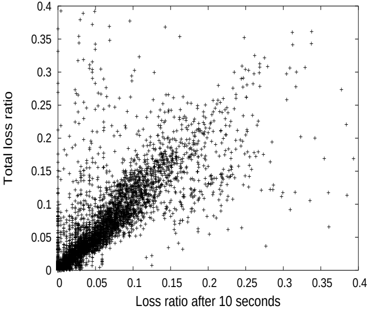

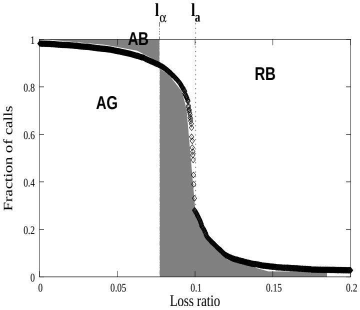

Self-admission control (paper F)

- Contribution: Free-standing endpoint admission control mechanism
- Goal is to estimate the quality of a call from its initial 'few' seconds
- Top plot shows the loss process for a number of calls
- Given the decision to admit or reject a flow, there are four possible outcomes:
  - Accepted and good quality (AG)
  - Accepted but bad quality (AB)
  - Rejected and bad quality (RB)
  - Rejected but good quality (RG)
- Our solution allows for strict and relaxed admission control strategies (strict shown)

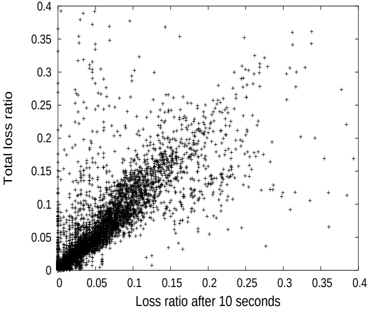

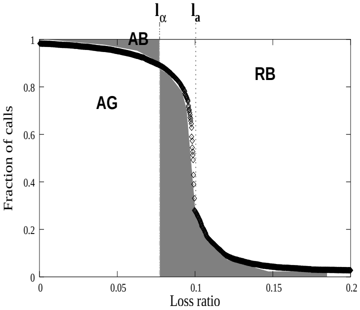

Self-admission control (paper F)

- Contribution: Free-standing endpoint admission control mechanism
- Goal is to estimate the quality of a call from its initial 'few' seconds
- Top plot shows the loss process for a number of calls
- Given the decision to admit or reject a flow, there are four possible outcomes:
    - Accepted and good quality (AG)
    - Accepted but bad quality (AB)
    - Rejected and bad quality (RB)
    - Rejected but good quality (RG)
- Our solution allows for strict and relaxed admission control strategies (strict shown)
- Strict policy minimises AB & maximises RG

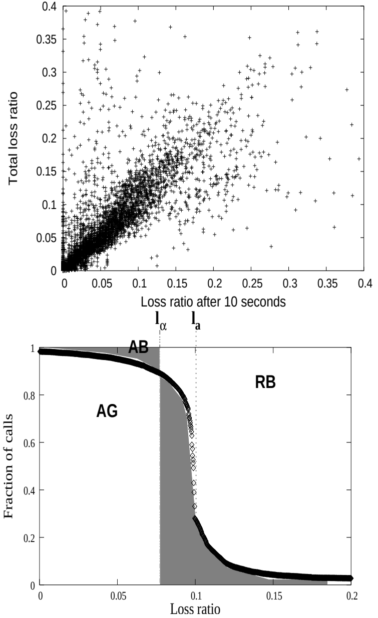

802.11 wireless access (paper G)

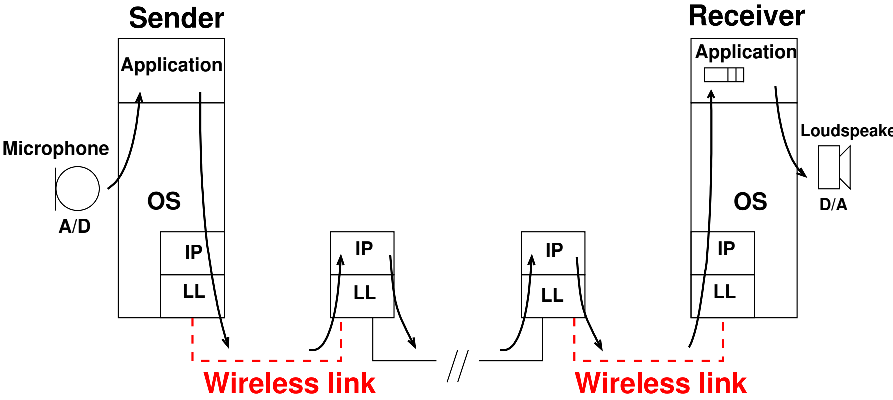

Investigate VoIP quality with local wireless access

802.11 access for VoIP

- Contribution: A comprehensive study of 802.11b networks for voice

Pos.	Loss % [min, mean, max]	RTT Delay (ms)
O_2	[0, 0.0, 0.1]	[1.9, 2.2, 2.4]
C	[0, 0, 0]	[1.9, 2.1, 4.0]
D	[0, 0, 0]	[2.1, 2.6, 3.1]
E_1	[0, 0.2, 2.6]	[2.8, 3.2, 5.4]
E_2	[9.4, 36.3, 89.1]	[5.3, 12.2, 24.3]
E_3	[0, 0.0, 0.2]	[2.8, 2.8, 4.0]
F_1	[20.1, 54.9, 88.7]	[5.4, 13.4, 24.6]
F_2	No signal	No signal
F_3	[1.8, 22.8, 84.9]	[4.6, 11.7, 13.7]
G	[0, 0.2, 1.9]	[1.9, 2.2, 3.8]
H_1	[0.4, 5.4, 30.2]	[3.5, 3.9, 8.1]
H_2	[3.4, 11.0, 28.3]	[5.9, 6.1, 11.7]
H_3	[0, 0.2, 2.7]	[3.4, 3.4, 5.4]]

802.11 access for VoIP

- Contribution: A comprehensive study of 802.11b networks for voice
- Studied different scenarios in a methodical manner:

Pos.	Loss % [min, mean, max]	RTT Delay (ms)
O_2	[0, 0.0, 0.1]	[1.9, 2.2, 2.4]
C	[0, 0, 0]	[1.9, 2.1, 4.0]
D	[0, 0, 0]	[2.1, 2.6, 3.1]
E_1	[0, 0.2, 2.6]	[2.8, 3.2, 5.4]
E_2	[9.4, 36.3, 89.1]	[5.3, 12.2, 24.3]
E_3	[0, 0.0, 0.2]	[2.8, 2.8, 4.0]
F_1	[20.1, 54.9, 88.7]	[5.4, 13.4, 24.6]
F_2	No signal	No signal
F_3	[1.8, 22.8, 84.9]	[4.6, 11.7, 13.7]
G	[0, 0.2, 1.9]	[1.9, 2.2, 3.8]
H_1	[0.4, 5.4, 30.2]	[3.5, 3.9, 8.1]
H_2	[3.4, 11.0, 28.3]	[5.9, 6.1, 11.7]
H_3	[0, 0.2, 2.7]	[3.4, 3.4, 5.4]]

802.11 access for VoIP

- Contribution: A comprehensive study of 802.11b networks for voice
- Studied different scenarios in a methodical manner:
  1. Line of sight (outside)

Pos.	Loss % [min, mean, max]	RTT Delay (ms)
O_2	[0, 0.0, 0.1]	[1.9, 2.2, 2.4]
C	[0, 0, 0]	[1.9, 2.1, 4.0]
D	[0, 0, 0]	[2.1, 2.6, 3.1]
E_1	[0, 0.2, 2.6]	[2.8, 3.2, 5.4]
E_2	[9.4, 36.3, 89.1]	[5.3, 12.2, 24.3]
E_3	[0, 0.0, 0.2]	[2.8, 2.8, 4.0]
F_1	[20.1, 54.9, 88.7]	[5.4, 13.4, 24.6]
F_2	No signal	No signal
F_3	[1.8, 22.8, 84.9]	[4.6, 11.7, 13.7]
G	[0, 0.2, 1.9]	[1.9, 2.2, 3.8]
H_1	[0.4, 5.4, 30.2]	[3.5, 3.9, 8.1]
H_2	[3.4, 11.0, 28.3]	[5.9, 6.1, 11.7]
H_3	[0, 0.2, 2.7]	[3.4, 3.4, 5.4]]

802.11 access for VoIP

- Contribution: A comprehensive study of 802.11b networks for voice
- Studied different scenarios in a methodical manner:
    1. Line of sight (outside)
    2. Non-line of sight (inside an office)

Pos.	Loss % [min, mean, max]	RTT Delay (ms)
O_2	[0, 0.0, 0.1]	[1.9, 2.2, 2.4]
C	[0, 0, 0]	[1.9, 2.1, 4.0]
D	[0, 0, 0]	[2.1, 2.6, 3.1]
E_1	[0, 0.2, 2.6]	[2.8, 3.2, 5.4]
E_2	[9.4, 36.3, 89.1]	[5.3, 12.2, 24.3]
E_3	[0, 0.0, 0.2]	[2.8, 2.8, 4.0]
F_1	[20.1, 54.9, 88.7]	[5.4, 13.4, 24.6]
F_2	No signal	No signal
F_3	[1.8, 22.8, 84.9]	[4.6, 11.7, 13.7]
G	[0, 0.2, 1.9]	[1.9, 2.2, 3.8]
H_1	[0.4, 5.4, 30.2]	[3.5, 3.9, 8.1]
H_2	[3.4, 11.0, 28.3]	[5.9, 6.1, 11.7]
H_3	[0, 0.2, 2.7]	[3.4, 3.4, 5.4]]

802.11 access for VoIP

- Contribution: A comprehensive study of 802.11b networks for voice
- Studied different scenarios in a methodical manner:
  1. Line of sight (outside)
  2. Non-line of sight (inside an office)
  3. With & without background load

Pos.	Loss % [min, mean, max]	RTT Delay (ms)
O_2	[0, 0.0, 0.1]	[1.9, 2.2, 2.4]
C	[0, 0, 0]	[1.9, 2.1, 4.0]
D	[0, 0, 0]	[2.1, 2.6, 3.1]
E_1	[0, 0.2, 2.6]	[2.8, 3.2, 5.4]
E_2	[9.4, 36.3, 89.1]	[5.3, 12.2, 24.3]
E_3	[0, 0.0, 0.2]	[2.8, 2.8, 4.0]
F_1	[20.1, 54.9, 88.7]	[5.4, 13.4, 24.6]
F_2	No signal	No signal
F_3	[1.8, 22.8, 84.9]	[4.6, 11.7, 13.7]
G	[0, 0.2, 1.9]	[1.9, 2.2, 3.8]
H_1	[0.4, 5.4, 30.2]	[3.5, 3.9, 8.1]
H_2	[3.4, 11.0, 28.3]	[5.9, 6.1, 11.7]
H_3	[0, 0.2, 2.7]	[3.4, 3.4, 5.4]]

802.11 access for VoIP

- Contribution: A comprehensive study of 802.11b networks for voice
- Studied different scenarios in a methodical manner:
  1. Line of sight (outside)
  2. Non-line of sight (inside an office)
  3. With & without background load
  4. Infrastructure mode (with an access point)

Pos.	Loss % [min, mean, max]	RTT Delay (ms)
O_2	[0, 0.0, 0.1]	[1.9, 2.2, 2.4]
C	[0, 0, 0]	[1.9, 2.1, 4.0]
D	[0, 0, 0]	[2.1, 2.6, 3.1]
E_1	[0, 0.2, 2.6]	[2.8, 3.2, 5.4]
E_2	[9.4, 36.3, 89.1]	[5.3, 12.2, 24.3]
E_3	[0, 0.0, 0.2]	[2.8, 2.8, 4.0]
F_1	[20.1, 54.9, 88.7]	[5.4, 13.4, 24.6]
F_2	No signal	No signal
F_3	[1.8, 22.8, 84.9]	[4.6, 11.7, 13.7]
G	[0, 0.2, 1.9]	[1.9, 2.2, 3.8]
H_1	[0.4, 5.4, 30.2]	[3.5, 3.9, 8.1]
H_2	[3.4, 11.0, 28.3]	[5.9, 6.1, 11.7]
H_3	[0, 0.2, 2.7]	[3.4, 3.4, 5.4]]

802.11 access for VoIP

- Contribution: A comprehensive study of 802.11b networks for voice
- Studied different scenarios in a methodical manner:
  1. Line of sight (outside)
  2. Non-line of sight (inside an office)
  3. With & without background load
  4. Infrastructure mode (with an access point)
  5. Using request to send/clear to send

Pos.	Loss % [min, mean, max]	RTT Delay (ms)
O_2	[0, 0.0, 0.1]	[1.9, 2.2, 2.4]
C	[0, 0, 0]	[1.9, 2.1, 4.0]
D	[0, 0, 0]	[2.1, 2.6, 3.1]
E_1	[0, 0.2, 2.6]	[2.8, 3.2, 5.4]
E_2	[9.4, 36.3, 89.1]	[5.3, 12.2, 24.3]
E_3	[0, 0.0, 0.2]	[2.8, 2.8, 4.0]
F_1	[20.1, 54.9, 88.7]	[5.4, 13.4, 24.6]
F_2	No signal	No signal
F_3	[1.8, 22.8, 84.9]	[4.6, 11.7, 13.7]
G	[0, 0.2, 1.9]	[1.9, 2.2, 3.8]
H_1	[0.4, 5.4, 30.2]	[3.5, 3.9, 8.1]
H_2	[3.4, 11.0, 28.3]	[5.9, 6.1, 11.7]
H_3	[0, 0.2, 2.7]	[3.4, 3.4, 5.4]]

802.11 access for VoIP

- Contribution: A comprehensive study of 802.11b networks for voice
- Studied different scenarios in a methodical manner:
  1. Line of sight (outside)
  2. Non-line of sight (inside an office)
  3. With & without background load
  4. Infrastructure mode (with an access point)
  5. Using request to send/clear to send
- Cross-layer techniques to infer VoIP quality at application proved useful

Pos.	Loss % [min, mean, max]	RTT Delay (ms)
O_2	[0, 0.0, 0.1]	[1.9, 2.2, 2.4]
C	[0, 0, 0]	[1.9, 2.1, 4.0]
D	[0, 0, 0]	[2.1, 2.6, 3.1]
E_1	[0, 0.2, 2.6]	[2.8, 3.2, 5.4]
E_2	[9.4, 36.3, 89.1]	[5.3, 12.2, 24.3]
E_3	[0, 0.0, 0.2]	[2.8, 2.8, 4.0]
F_1	[20.1, 54.9, 88.7]	[5.4, 13.4, 24.6]
F_2	No signal	No signal
F_3	[1.8, 22.8, 84.9]	[4.6, 11.7, 13.7]
G	[0, 0.2, 1.9]	[1.9, 2.2, 3.8]
H_1	[0.4, 5.4, 30.2]	[3.5, 3.9, 8.1]
H_2	[3.4, 11.0, 28.3]	[5.9, 6.1, 11.7]
H_3	[0, 0.2, 2.7]	[3.4, 3.4, 5.4]]

802.11 access for VoIP

- Contribution: A comprehensive study of 802.11b networks for voice
- Studied different scenarios in a methodical manner:
  1. Line of sight (outside)
  2. Non-line of sight (inside an office)
  3. With & without background load
  4. Infrastructure mode (with an access point)
  5. Using request to send/clear to send
- Cross-layer techniques to infer VoIP quality at application proved useful
- Need to do perform radio pre-planning and measurements before deployment

Pos.	Loss % [min, mean, max]	RTT Delay (ms)
O_2	[0, 0.0, 0.1]	[1.9, 2.2, 2.4]
C	[0, 0, 0]	[1.9, 2.1, 4.0]
D	[0, 0, 0]	[2.1, 2.6, 3.1]
E_1	[0, 0.2, 2.6]	[2.8, 3.2, 5.4]
E_2	[9.4, 36.3, 89.1]	[5.3, 12.2, 24.3]
E_3	[0, 0.0, 0.2]	[2.8, 2.8, 4.0]
F_1	[20.1, 54.9, 88.7]	[5.4, 13.4, 24.6]
F_2	No signal	No signal
F_3	[1.8, 22.8, 84.9]	[4.6, 11.7, 13.7]
G	[0, 0.2, 1.9]	[1.9, 2.2, 3.8]
H_1	[0.4, 5.4, 30.2]	[3.5, 3.9, 8.1]
H_2	[3.4, 11.0, 28.3]	[5.9, 6.1, 11.7]
H_3	[0, 0.2, 2.7]	[3.4, 3.4, 5.4]]

802.11 access for VoIP

- Contribution: A comprehensive study of 802.11b networks for voice
- Studied different scenarios in a methodical manner:
  1. Line of sight (outside)
  2. Non-line of sight (inside an office)
  3. With & without background load
  4. Infrastructure mode (with an access point)
  5. Using request to send/clear to send
- Cross-layer techniques to infer VoIP quality at application proved useful
- Need to do perform radio pre-planning and measurements before deployment
- Immediate rate changes beneficial for VoIP

Pos.	Loss % [min, mean, max]	RTT Delay (ms)
O_2	[0, 0.0, 0.1]	[1.9, 2.2, 2.4]
C	[0, 0, 0]	[1.9, 2.1, 4.0]
D	[0, 0, 0]	[2.1, 2.6, 3.1]
E_1	[0, 0.2, 2.6]	[2.8, 3.2, 5.4]
E_2	[9.4, 36.3, 89.1]	[5.3, 12.2, 24.3]
E_3	[0, 0.0, 0.2]	[2.8, 2.8, 4.0]
F_1	[20.1, 54.9, 88.7]	[5.4, 13.4, 24.6]
F_2	No signal	No signal
F_3	[1.8, 22.8, 84.9]	[4.6, 11.7, 13.7]
G	[0, 0.2, 1.9]	[1.9, 2.2, 3.8]
H_1	[0.4, 5.4, 30.2]	[3.5, 3.9, 8.1]
H_2	[3.4, 11.0, 28.3]	[5.9, 6.1, 11.7]
H_3	[0, 0.2, 2.7]	[3.4, 3.4, 5.4]]

802.11 access for VoIP

- Contribution: A comprehensive study of 802.11b networks for voice
- Studied different scenarios in a methodical manner:
  1. Line of sight (outside)
  2. Non-line of sight (inside an office)
  3. With & without background load
  4. Infrastructure mode (with an access point)
  5. Using request to send/clear to send
- Cross-layer techniques to infer VoIP quality at application proved useful
- Need to do perform radio pre-planning and measurements before deployment
- Immediate rate changes beneficial for VoIP
- Work proved an important precursor for the next contribution

Pos.	Loss % [min, mean, max]	RTT Delay (ms)
O_2	[0, 0.0, 0.1]	[1.9, 2.2, 2.4]
C	[0, 0, 0]	[1.9, 2.1, 4.0]
D	[0, 0, 0]	[2.1, 2.6, 3.1]
E_1	[0, 0.2, 2.6]	[2.8, 3.2, 5.4]
E_2	[9.4, 36.3, 89.1]	[5.3, 12.2, 24.3]
E_3	[0, 0.0, 0.2]	[2.8, 2.8, 4.0]
F_1	[20.1, 54.9, 88.7]	[5.4, 13.4, 24.6]
F_2	No signal	No signal
F_3	[1.8, 22.8, 84.9]	[4.6, 11.7, 13.7]
G	[0, 0.2, 1.9]	[1.9, 2.2, 3.8]
H_1	[0.4, 5.4, 30.2]	[3.5, 3.9, 8.1]
H_2	[3.4, 11.0, 28.3]	[5.9, 6.1, 11.7]
H_3	[0, 0.2, 2.7]	[3.4, 3.4, 5.4]]

802.11-based voice with alternative access (paper H)

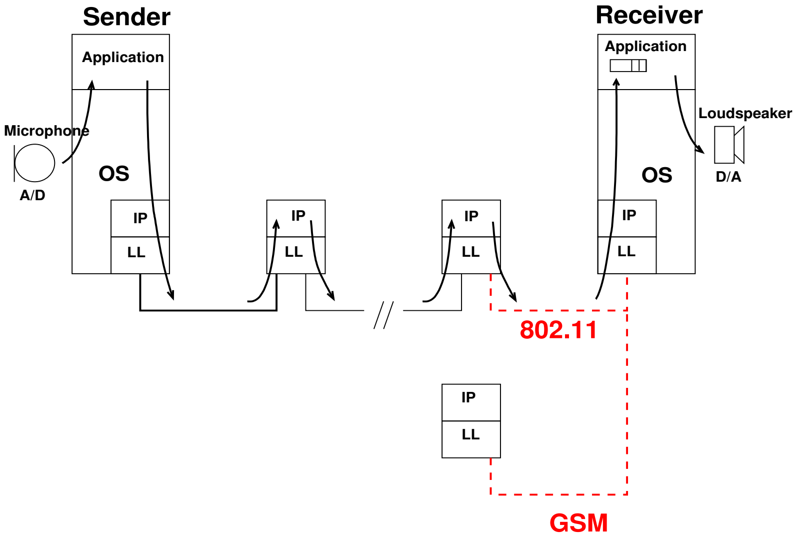

A handover solution

- Contribution: Fully automated handover mechanism from 802.11 to GSM

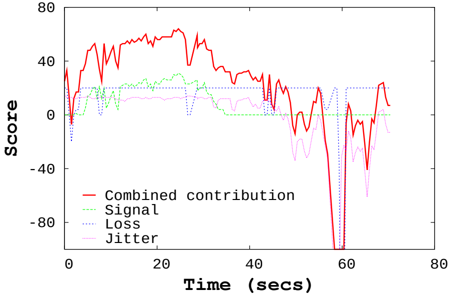

Quality started good and became bad	Timely HO **68**	Late HO 10
Quality started good and remained good	Unnecessary HO 7	No HO **15**

A handover solution

- Contribution: Fully automated handover mechanism from 802.11 to GSM
- Implemented a function to estimate the network parameters in a receiver

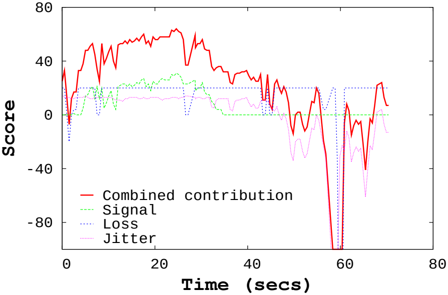

Quality started good and became bad	Timely HO **68**	Late HO 10
Quality started good and remained good	Unnecessary HO 7	No HO **15**

A handover solution

- Contribution: Fully automated handover mechanism from 802.11 to GSM
- Implemented a function to estimate the network parameters in a receiver
- When the quality drops below a given threshold, we switch from 802.11 to GSM

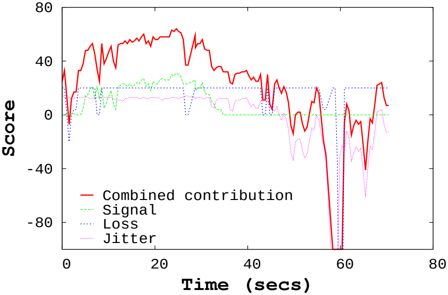

Quality started good and became bad	Timely HO **68**	Late HO 10
Quality started good and remained good	Unnecessary HO 7	No HO **15**

A handover solution

- Contribution: Fully automated handover mechanism from 802.11 to GSM
- Implemented a function to estimate the network parameters in a receiver
- When the quality drops below a given threshold, we switch from 802.11 to GSM
- The difficulty is in the prediction needed (i.e. the PSTN call setup)

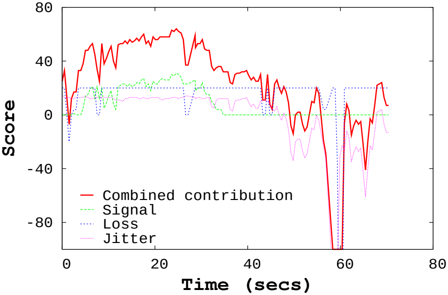

Quality started good and became bad	Timely HO **68**	Late HO 10
Quality started good and remained good	Unnecessary HO 7	No HO **15**

A handover solution

- Contribution: Fully automated handover mechanism from 802.11 to GSM
- Implemented a function to estimate the network parameters in a receiver
- When the quality drops below a given threshold, we switch from 802.11 to GSM
- The difficulty is in the prediction needed (i.e. the PSTN call setup)
- Tested the handover in a real implementation and conducted 100 user tests in an office environment (bottom table)

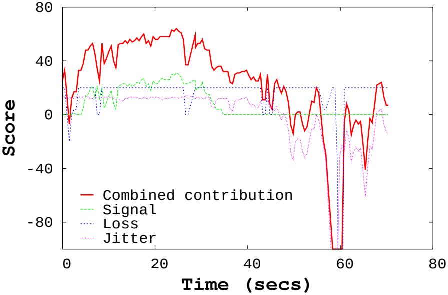

Quality started good and became bad	Timely HO **68**	Late HO 10
Quality started good and remained good	Unnecessary HO 7	No HO **15**

Single-sided measure based on PESQ (paper I)

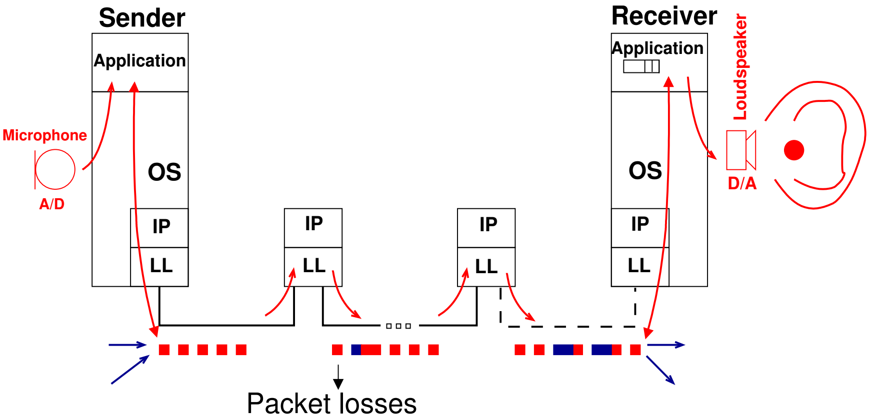

Single-sided measure based on PESQ

- Contribution: Single-sided real-time estimation for loss-based quality prediction

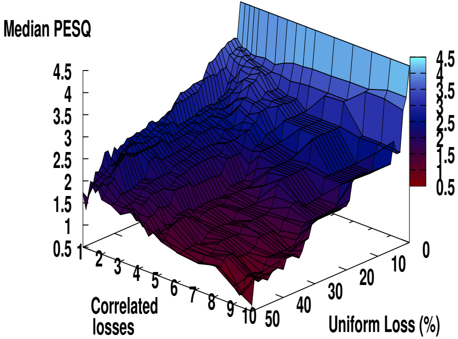

Single-sided measure based on PESQ

- Contribution: Single-sided real-time estimation for loss-based quality prediction
- Idea is to derive an off-line estimation of voice quality

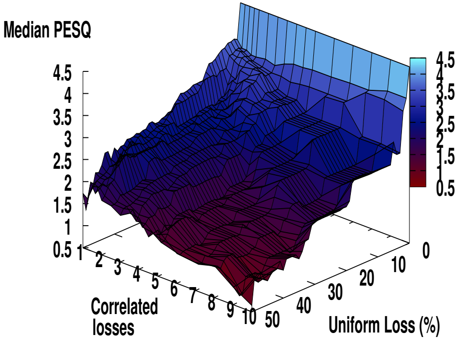

Single-sided measure based on PESQ

- Contribution: Single-sided real-time estimation for loss-based quality prediction
- Idea is to derive an off-line estimation of voice quality
- Compute off-line response to typical network losses using PESQ

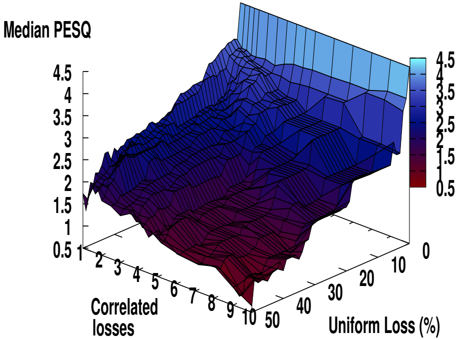

Single-sided measure based on PESQ

- Contribution: Single-sided real-time estimation for loss-based quality prediction
- Idea is to derive an off-line estimation of voice quality
- Compute off-line response to typical network losses using PESQ
    1. PESQ = Perceptual Evaluation of Speech Quality

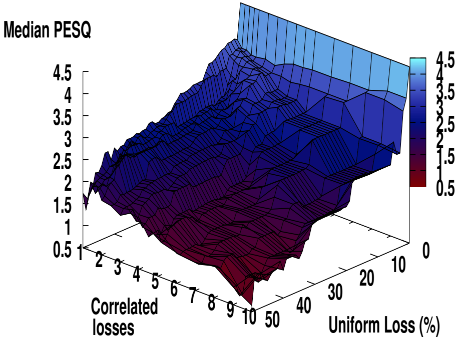

Single-sided measure based on PESQ

- Contribution: Single-sided real-time estimation for loss-based quality prediction
- Idea is to derive an off-line estimation of voice quality
- Compute off-line response to typical network losses using PESQ
    1. PESQ = Perceptual Evaluation of Speech Quality
    2. Compares a reference and degraded signal

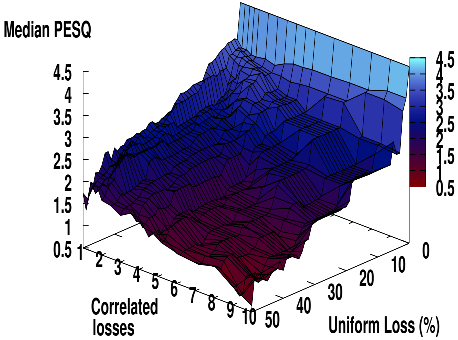

Single-sided measure based on PESQ

- Contribution: Single-sided real-time estimation for loss-based quality prediction
- Idea is to derive an off-line estimation of voice quality
- Compute off-line response to typical network losses using PESQ
    1. PESQ = Perceptual Evaluation of Speech Quality
    2. Compares a reference and degraded signal
    3. Outputs a value between 0.5 and 4.5

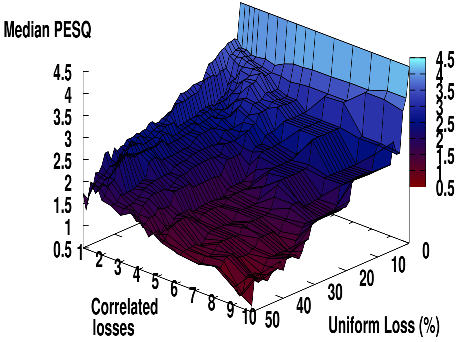

Single-sided measure based on PESQ

- Contribution: Single-sided real-time estimation for loss-based quality prediction
- Idea is to derive an off-line estimation of voice quality
- Compute off-line response to typical network losses using PESQ
  1. PESQ = Perceptual Evaluation of Speech Quality
  2. Compares a reference and degraded signal
  3. Outputs a value between 0.5 and 4.5
- 3D plot shows response to correlated and uncorrelated losses

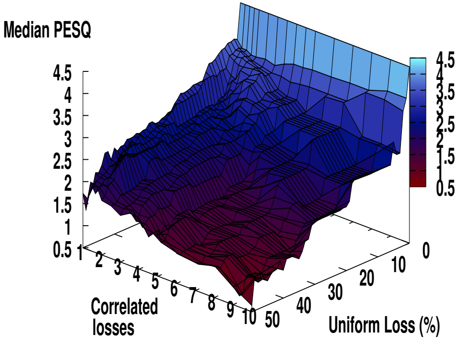

Single-sided measure based on PESQ

- Contribution: Single-sided real-time estimation for loss-based quality prediction
- Idea is to derive an off-line estimation of voice quality
- Compute off-line response to typical network losses using PESQ
    1. PESQ = Perceptual Evaluation of Speech Quality
    2. Compares a reference and degraded signal
    3. Outputs a value between 0.5 and 4.5
- 3D plot shows response to correlated and uncorrelated losses
- PESQ tends to underestimate the impact of bursty losses

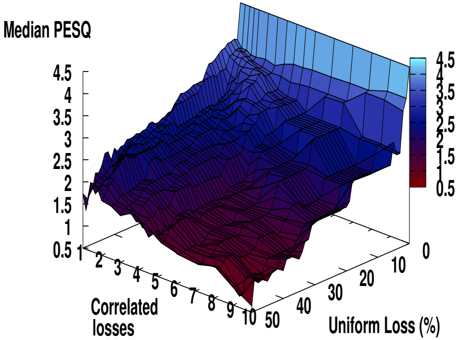

End-system adaption to network jitter (paper C)

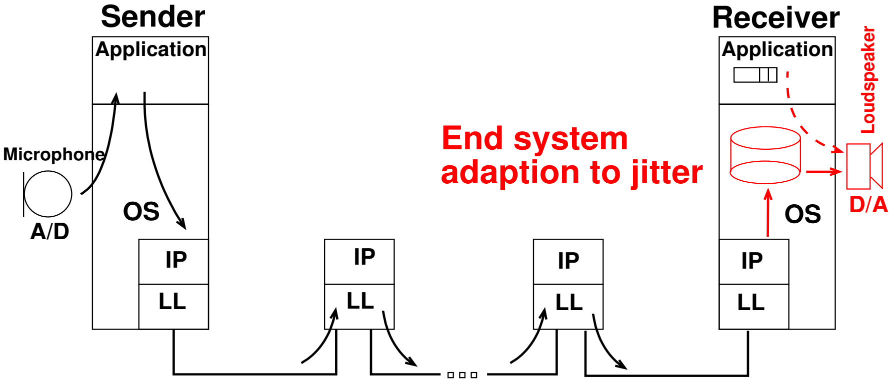

End-system adaption to network jitter

- Contribution: General guidelines in delay reduction techniques for end-systems

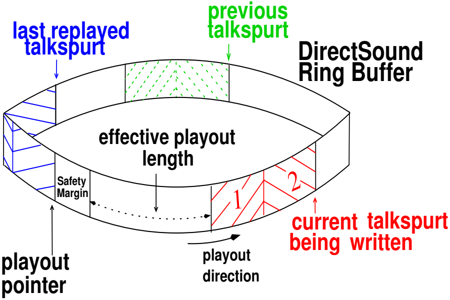

End-system adaption to network jitter

- Contribution: General guidelines in delay reduction techniques for end-systems
- End systems can add considerable delays to network stream

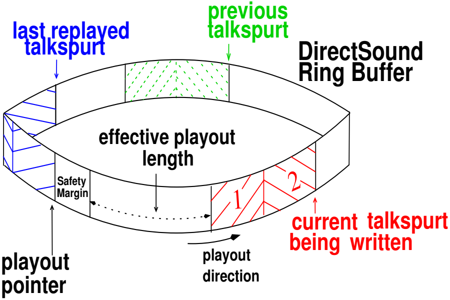

End-system adaption to network jitter

- Contribution: General guidelines in delay reduction techniques for end-systems
- End systems can add considerable delays to network stream
- Carefully engineered solutions make end-system optimisations possible

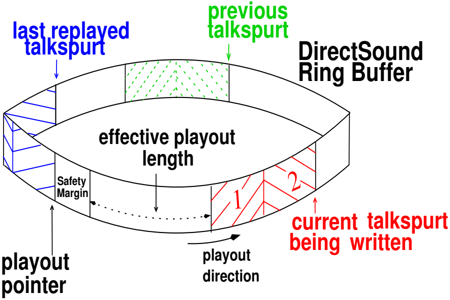

End-system adaption to network jitter

- Contribution: General guidelines in delay reduction techniques for end-systems
- End systems can add considerable delays to network stream
- Carefully engineered solutions make end-system optimisations possible
    1. Remove unnecessary data copying

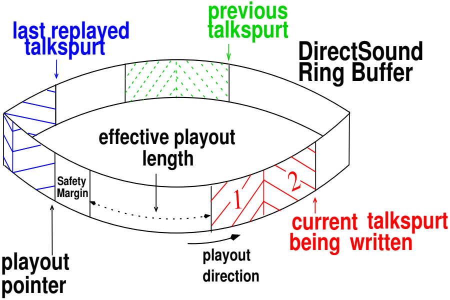

End-system adaption to network jitter

- Contribution: General guidelines in delay reduction techniques for end-systems
- End systems can add considerable delays to network stream
- Carefully engineered solutions make end-system optimisations possible
    1. Remove unnecessary data copying
    2. Setup memory transfers in advance (DMA)

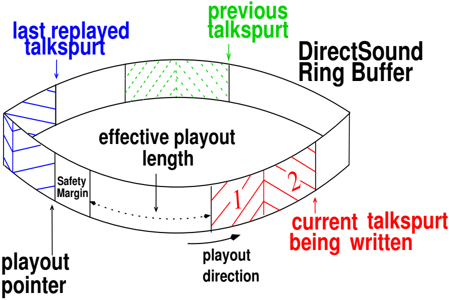

End-system adaption to network jitter

- Contribution: General guidelines in delay reduction techniques for end-systems
- End systems can add considerable delays to network stream
- Carefully engineered solutions make end-system optimisations possible
  1. Remove unnecessary data copying
  2. Setup memory transfers in advance (DMA)
  3. Optimise buffer playout for appropriate coding

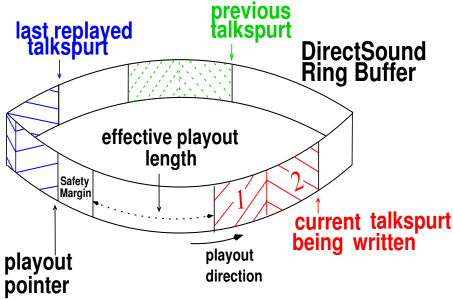

End-system adaption to network jitter

- Contribution: General guidelines in delay reduction techniques for end-systems
- End systems can add considerable delays to network stream
- Carefully engineered solutions make end-system optimisations possible
  1. Remove unnecessary data copying
  2. Setup memory transfers in advance (DMA)
  3. Optimise buffer playout for appropriate coding
  4. Require as few hardware interrupts as possible

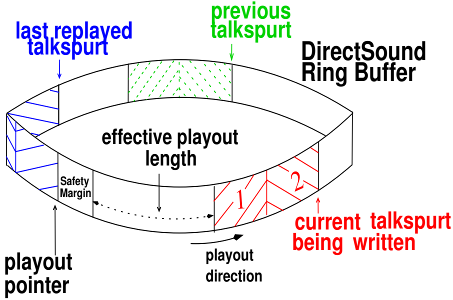

End-system adaption to network jitter

- Contribution: General guidelines in delay reduction techniques for end-systems
- End systems can add considerable delays to network stream
- Carefully engineered solutions make end-system optimisations possible
  1. Remove unnecessary data copying
  2. Setup memory transfers in advance (DMA)
  3. Optimise buffer playout for appropriate coding
  4. Require as few hardware interrupts as possible
- 100's of milliseconds can potentially be saved

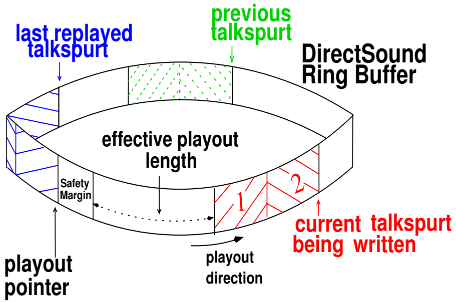

End-system adaption to network jitter

- Contribution: General guidelines in delay reduction techniques for end-systems
- End systems can add considerable delays to network stream
- Carefully engineered solutions make end-system optimisations possible
    1. Remove unnecessary data copying
    2. Setup memory transfers in advance (DMA)
    3. Optimise buffer playout for appropriate coding
    4. Require as few hardware interrupts as possible
- 100's of milliseconds can potentially be saved
- Paper describes a 3-tier solution to inhibit fluctuations

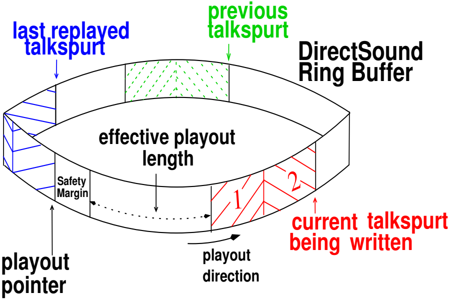

Dissertation conclusions

- A near complete system study for Voice over IP

Dissertation conclusions

- A near complete system study for Voice over IP
- Carefully engineered systems can provide good quality voice

Dissertation conclusions

- A near complete system study for Voice over IP
- Carefully engineered systems can provide good quality voice
    1. Controlled load for VoIP traffic, using dimensioning

Dissertation conclusions

- A near complete system study for Voice over IP
- Carefully engineered systems can provide good quality voice
  1. Controlled load for VoIP traffic, using dimensioning
  2. End system self-admission

Dissertation conclusions

- A near complete system study for Voice over IP
- Carefully engineered systems can provide good quality voice
  1. Controlled load for VoIP traffic, using dimensioning
  2. End system self-admission
  3. Monitor and measure the network performance

Dissertation conclusions

- A near complete system study for Voice over IP
- Carefully engineered systems can provide good quality voice
  1. Controlled load for VoIP traffic, using dimensioning
  2. End system self-admission
  3. Monitor and measure the network performance
     - For good engineering designs

Dissertation conclusions

- A near complete system study for Voice over IP
- Carefully engineered systems can provide good quality voice
  1. Controlled load for VoIP traffic, using dimensioning
  2. End system self-admission
  3. Monitor and measure the network performance
     - For good engineering designs
     - To gain insight into the problem

Dissertation conclusions

- A near complete system study for Voice over IP
- Carefully engineered systems can provide good quality voice
  1. Controlled load for VoIP traffic, using dimensioning
  2. End system self-admission
  3. Monitor and measure the network performance
     - For good engineering designs
     - To gain insight into the problem
     - Collate data for further investigation

Dissertation conclusions

- A near complete system study for Voice over IP
- Carefully engineered systems can provide good quality voice
    1. Controlled load for VoIP traffic, using dimensioning
    2. End system self-admission
    3. Monitor and measure the network performance
        - For good engineering designs
        - To gain insight into the problem
        - Collate data for further investigation
    4. Rate control within 802.11 networks

Dissertation conclusions

- A near complete system study for Voice over IP
- Carefully engineered systems can provide good quality voice
  1. Controlled load for VoIP traffic, using dimensioning
  2. End system self-admission
  3. Monitor and measure the network performance
     - For good engineering designs
     - To gain insight into the problem
     - Collate data for further investigation
  4. Rate control within 802.11 networks
  5. Where alternatives are available, use them, e.g. handover

Dissertation conclusions

- A near complete system study for Voice over IP
- Carefully engineered systems can provide good quality voice
  1. Controlled load for VoIP traffic, using dimensioning
  2. End system self-admission
  3. Monitor and measure the network performance
     - For good engineering designs
     - To gain insight into the problem
     - Collate data for further investigation
  4. Rate control within 802.11 networks
  5. Where alternatives are available, use them, e.g. handover
  6. *Estimate* the quality as near to the user as possible

Dissertation conclusions

- A near complete system study for Voice over IP
- Carefully engineered systems can provide good quality voice
  1. Controlled load for VoIP traffic, using dimensioning
  2. End system self-admission
  3. Monitor and measure the network performance
     - For good engineering designs
     - To gain insight into the problem
     - Collate data for further investigation
  4. Rate control within 802.11 networks
  5. Where alternatives are available, use them, e.g. handover
  6. *Estimate* the quality as near to the user as possible
  7. Delay aware end-systems

Dissertation conclusions

- A near complete system study for Voice over IP
- Carefully engineered systems can provide good quality voice
  1. Controlled load for VoIP traffic, using dimensioning
  2. End system self-admission
  3. Monitor and measure the network performance
     - For good engineering designs
     - To gain insight into the problem
     - Collate data for further investigation
  4. Rate control within 802.11 networks
  5. Where alternatives are available, use them, e.g. handover
  6. *Estimate* the quality as near to the user as possible
  7. Delay aware end-systems
- Used many different techniques:

Dissertation conclusions

- A near complete system study for Voice over IP
- Carefully engineered systems can provide good quality voice
  1. Controlled load for VoIP traffic, using dimensioning
  2. End system self-admission
  3. Monitor and measure the network performance
     - For good engineering designs
     - To gain insight into the problem
     - Collate data for further investigation
  4. Rate control within 802.11 networks
  5. Where alternatives are available, use them, e.g. handover
  6. *Estimate* the quality as near to the user as possible
  7. Delay aware end-systems
- Used many different techniques:
  1. Analysis and modelling

Dissertation conclusions

- A near complete system study for Voice over IP
- Carefully engineered systems can provide good quality voice
    1. Controlled load for VoIP traffic, using dimensioning
    2. End system self-admission
    3. Monitor and measure the network performance
        - For good engineering designs
        - To gain insight into the problem
        - Collate data for further investigation
    4. Rate control within 802.11 networks
    5. Where alternatives are available, use them, e.g. handover
    6. *Estimate* the quality as near to the user as possible
    7. Delay aware end-systems
- Used many different techniques:
    1. Analysis and modelling
    2. Simulation

Dissertation conclusions

- A near complete system study for Voice over IP
- Carefully engineered systems can provide good quality voice
  1. Controlled load for VoIP traffic, using dimensioning
  2. End system self-admission
  3. Monitor and measure the network performance
     - For good engineering designs
     - To gain insight into the problem
     - Collate data for further investigation
  4. Rate control within 802.11 networks
  5. Where alternatives are available, use them, e.g. handover
  6. *Estimate* the quality as near to the user as possible
  7. Delay aware end-systems
- Used many different techniques:
  1. Analysis and modelling
  2. Simulation
  3. Measurements

Dissertation conclusions

- A near complete system study for Voice over IP
- Carefully engineered systems can provide good quality voice
  1. Controlled load for VoIP traffic, using dimensioning
  2. End system self-admission
  3. Monitor and measure the network performance
     - For good engineering designs
     - To gain insight into the problem
     - Collate data for further investigation
  4. Rate control within 802.11 networks
  5. Where alternatives are available, use them, e.g. handover
  6. *Estimate* the quality as near to the user as possible
  7. Delay aware end-systems
- Used many different techniques:
  1. Analysis and modelling
  2. Simulation
  3. Measurements
  4. Implementation

Dissertation conclusions

- A near complete system study for Voice over IP
- Carefully engineered systems can provide good quality voice
  1. Controlled load for VoIP traffic, using dimensioning
  2. End system self-admission
  3. Monitor and measure the network performance
     - For good engineering designs
     - To gain insight into the problem
     - Collate data for further investigation
  4. Rate control within 802.11 networks
  5. Where alternatives are available, use them, e.g. handover
  6. *Estimate* the quality as near to the user as possible
  7. Delay aware end-systems
- Used many different techniques:
  1. Analysis and modelling
  2. Simulation
  3. Measurements
  4. Implementation
  5. Subjective tests

Dissertation conclusions

- A near complete system study for Voice over IP
- Carefully engineered systems can provide good quality voice
  1. Controlled load for VoIP traffic, using dimensioning
  2. End system self-admission
  3. Monitor and measure the network performance
     - For good engineering designs
     - To gain insight into the problem
     - Collate data for further investigation
  4. Rate control within 802.11 networks
  5. Where alternatives are available, use them, e.g. handover
  6. *Estimate* the quality as near to the user as possible
  7. Delay aware end-systems
- Used many different techniques:
  1. Analysis and modelling
  2. Simulation
  3. Measurements
  4. Implementation
  5. Subjective tests
- A byproduct of our research is a suite of software tools, some in commercial systems (handover module) and a large measurement data repository

Collaborators

Henrik Abrahamsson
Anders Gunnar
Bengt Ahlgren
Björn Grönvall
Olof Hagsand
Florian Hammer
Kjell Hanson
Ingemar Kaj

Gunnar Karlsson
Daniel Lorenzo
Fengyi Li
Gerald Q. Maguire Jr.
Ignacio Más
Victor Yuri Diogo Nunes
Juan Carlos Martín Severiano
Martín Varela
Thiemo Voigt

Q & A next...