NTNU

Innovation and Creativity

A Low-Latency Full-Duplex Audio over IP Streamer

Asbjørn Sæbø and U. Peter Svensson

Centre for Quantifiable Quality of Service in Communication Systems (Q2S)

LAC2006, 2006-04-27

What is LDAS?

— A Low Delay Audio Streamer in software
  - basically an audio to UDP/IP adapter
  - developed for Linux, using the ALSA sound system
  - currently at prototype stage

— A research tool
  - distributed multimedia interaction
  - perceived quality of service

— Aimed at demanding applications
  - low latency, high quality multichannel audio

Innovation and Creativity

Overview of the presentation

- Background and motivation
- Requirements and specification
- Main points of LDAS implementation
- Latency and latency measurements
- Conclusion

Q2S: Audio over IP Networks

- Quality beyond voice over IP:
  - Lower latency
  - Higher audio quality $\Rightarrow$ higher bit rates
  - Multiple channels $\Rightarrow$ higher bit rates
- Existing solutions
  - No access to source code, and/or
  - Not fulfilling requirements
  - Much to learn from “rolling our own”
- Goal: Fully open software, suitable as research tool

NTNU
Innovation and Creativity

Worldview

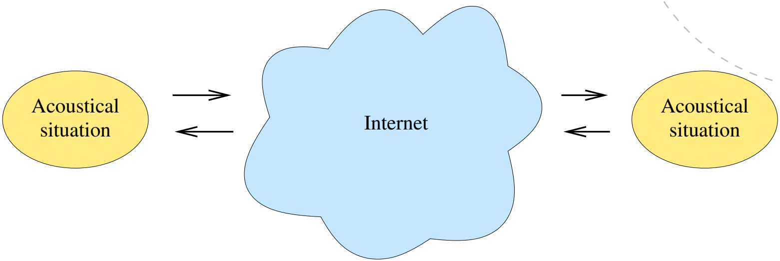

Quality of service, as perceived from the *endpoints*, by the *users*.
Connected by the net: Virtual presence and true interaction

NTNU
Innovation and Creativity

Networked ensemble playing

Other applications:

— Transmission of acoustic environments

— Advanced videoconferencing

Related work at Q2S/NTNU

Ola Strand:

— Transmission of audio and video

Håkon Liestøl Winge:

— Measured influence of latency

Otto Wittner / Sigurd Saue:

— Low-delay Windows Streamer

Trond Iver Røste Pedersen:

— Streaming in Java

Snorre Balliere Farner et.al.:

— Influence of latency and reverberation

NTNU
Innovation and Creativity

Requirements

- Low latency:
  Less than 20 - 30 ms for ensemble playing
- Multiple channels:
  For realistic transmission of
  acoustic environments and multi-channel 3D audio
- High quality:
  CD-quality and upwards

NTNU
Innovation and Creativity

Latency: How low can it be?

— Trondheim - New Zealand: 170 ms

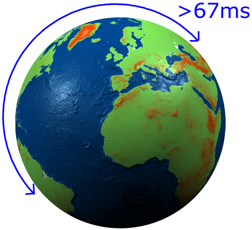

— The speed of light – too low?

— Only “near” parts of the world within reach

Specification (minimum)

Audio: 44.1kHz/48kHz, 16bit, two-channel PCM, possibility for coding.

Network: UDP/IP, with precautions against UDP shortcomings. No retransmission.

Latency: Less than 20 ms (analog to analog) over LAN.

Synchronisation:

- Receiver and sender must be kept in sync.
- Buffering, to minimise the effects of network transmission time jitter

NTNU
Innovation and Creativity

Implementation overview

Three threads: Recorder/sender, receiver and playback

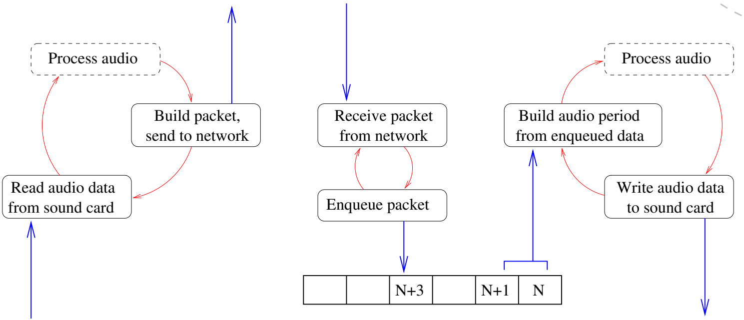

Packet format

The payload of the UDP packet:

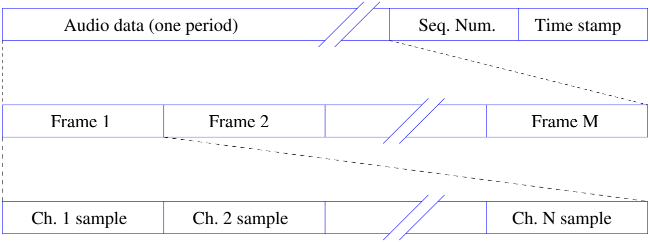

Packet stream positions:
Sequence number plus frame-level offset into packet

Packet stream control

Protocol implicitly defined by packet format (sequence numbering) and receiver enqueueing combination.
Handles UDP shortcomings and network transmission problems.

— Order packets
— Detect lost/missing packets
— Reject duplicates
— Reject late packets

Missing packets replaced by dummy data. (Possibility: Error concealment.)

NTNU
Innovation and Creativity

Synchronisation

— The receiver queue

- A sliding window onto the packet stream
- Ringbuffer of pointers to packets

— "Large scale" synchronisation

- "Early" packet (outside window) $\Rightarrow$ resynchronisation
- "Too many" late packets recently $\Rightarrow$ resynchronisation

— Drift adjustment by watermark algorithm

- Queue length too high $\Rightarrow$ skip single samples (frames)
- Queue length too low $\Rightarrow$ reuse single samples

Innovation and Creativity

Latency measurements

Audio delivered in *periods*

- Sender sound card buffer: One period of latency
- Receiver sound card buffer: Up to one period
- Receiver queue: Adjustable
- Network transmission time
- A/D and D/A: 2 – 3 ms
- Processing latency (software, OS)

Measurement setup

— Two Linux computers, 2.6.12 multimedia kernels
— M-Audio Delta44 audio interfaces
— Connected through a switch
— Full duplex stereo transmission with LDAS
— Impulse response measurements
— Transmission monitored by listening to audio

NTNU
Innovation and Creativity

Impulse response measurements

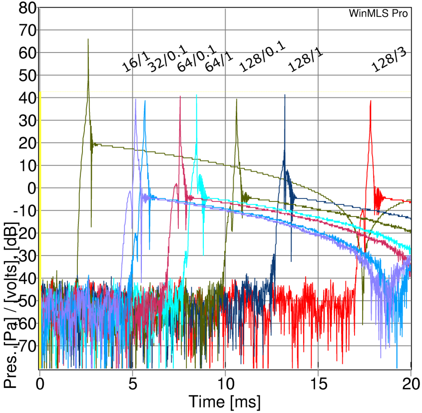

NTNU
Innovation and Creativity

Frequency response

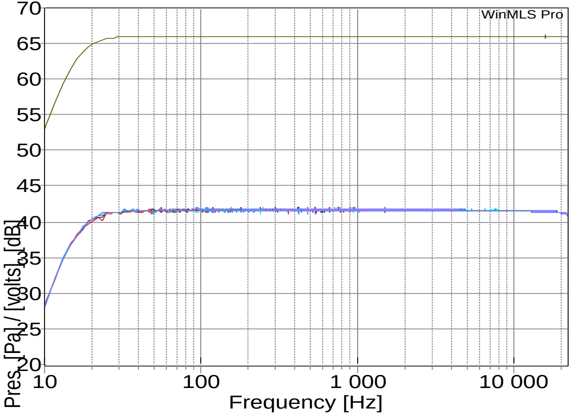

NTNU
Innovation and Creativity

Latency results

P. size	Queue len.	Mean latency	St. dev.
128	3	15.2	0.8
128	1	11.0	0.8
128	0.1	8.2	0.6
64	1	5.8	0.4
64	0.1	4.9	0.4
32	0.1	3.1	0.2
16	1	2.6	0.1

Latency in milliseconds for audio transmission with LDAS.
(Measured from analog input to analog output. Fs = 48kHz.)

NTNU
Innovation and Creativity

Similar solutions

— netjack (Torben Hohn)

— jack.udp (Rohan Drape)

— streamBD (SoundWIRE group, CCRMA)

— llcon (Volker Fischer)

NTNU
Innovation and Creativity

Thanks!

In particular:

— Lee Revell, Mindpipe Audio

Advice and discussions:

— Svein Sørsdal, SINTEF
— Georg Ottesen, SINTEF
— Jon Kåre Hellan, Uninett
— Paul Calamia, Rensselaer Polytechnic Institute

Availability

- Open Source, GNU General Public License
- Download: http://www.q2s.ntnu.no/˜asbjs/ldas/ldas.html
- Mailing list: https://pat.q2s.ntnu.no/mailman/listinfo/ldas-dev

NTNU
Innovation and Creativity