HD Multi-point Videoconferencing

Petr Holub
hopet@ics.muni.cz
CESNET and Masaryk University, Czech Republic

TERENA Networking Conference
Catania, Italy, 2006–05–18

Talk Overview

Virtual Collaborative Environments

- An environment for the least obtrusive remote collaboration
  - transmission of natural communication channels (image, sound, …)
  - elimination of disturbing elements
    - low-quality media (both digital and analog artifacts)
    - latency
    - unnatural physical elements

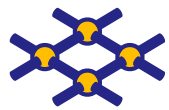

Problem Statement

- high image quality multi-party collaborative environment based on open-source products
  - (at least) full HD image with close to maximum quality
  - low latency
  - multi-party data distribution without need for network-native multicast
  - open-source implementation on open-source platform
  - for free

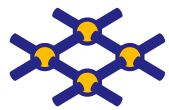

Proposed Solution

- 1920×1080 image resolution
- uncompressed video stream
- multi-party data distribution based on our UDP packet reflectors
- Linux-based implementation

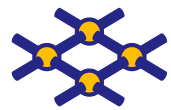

What does “high-quality” video mean?

- resolution
- frame rate
- progressive/interlaced video
- color space depth
- color space sampling
- end-to-end latency

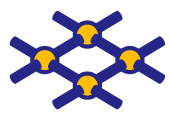

Compressed vs. uncompressed video

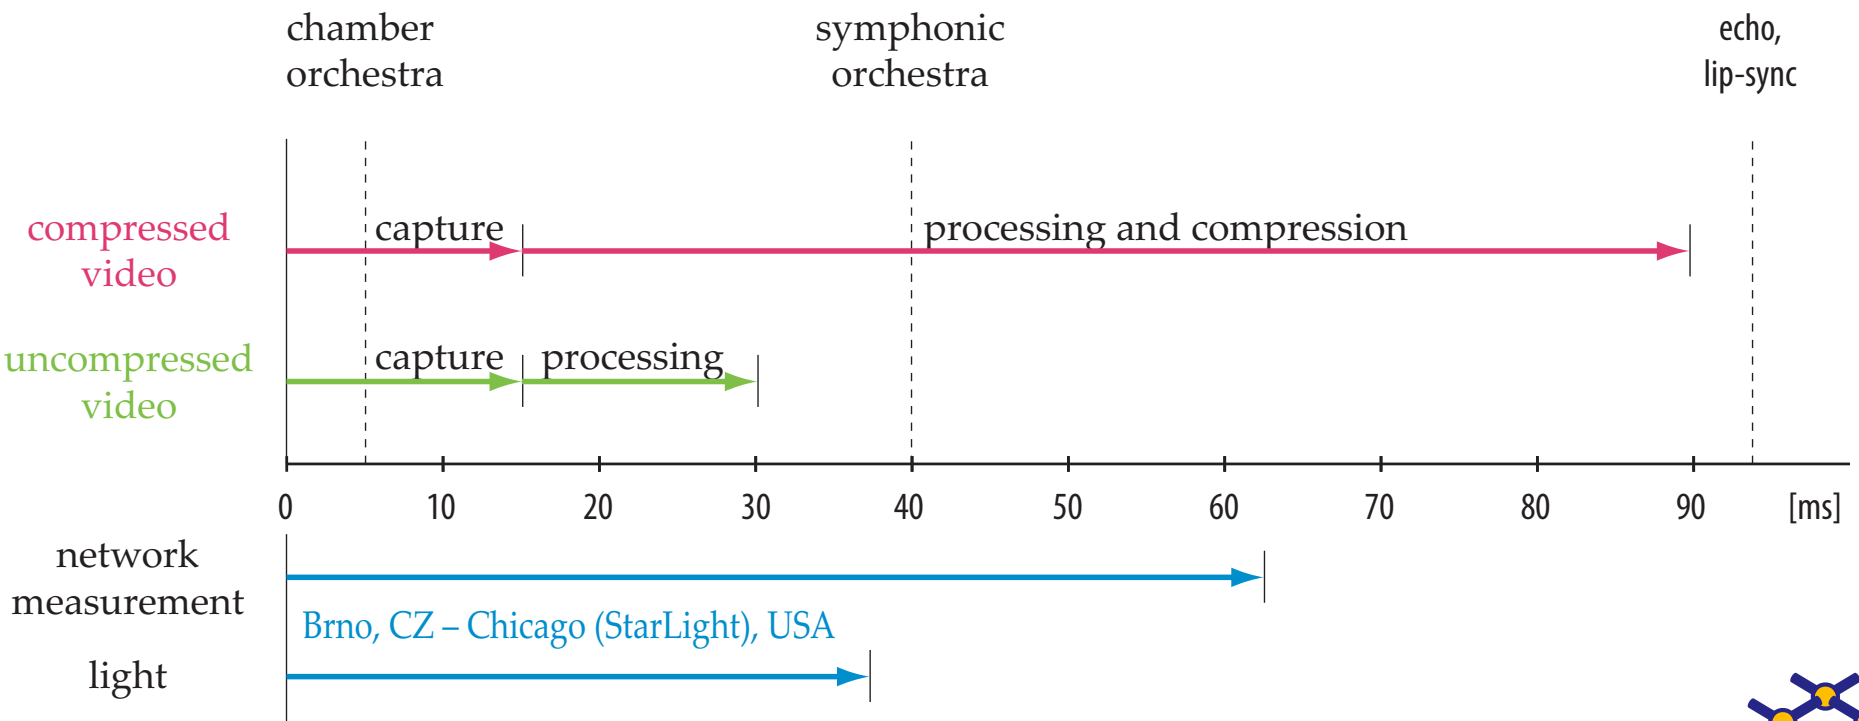

Compressed vs. uncompressed video

+ latency minimization
+ maximum quality preservation including pixel/field/frame independence
  - important esp. for further processing and (re)compression
+ fewer problems with artifacts (esp. in case of network dropouts)
  - DCT-based compression schemes have very characteristic response to data losses
– bandwidth wasting
– higher demands on capacity of video processing systems

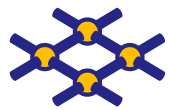

HDTV

- multiple supported formats: 720p, 1080i, 1080p
  - common 1920 × 1080 @ 60i
- support for high-quality audio (e.g. 16 bit @ 48kHz with 5.1 channels)

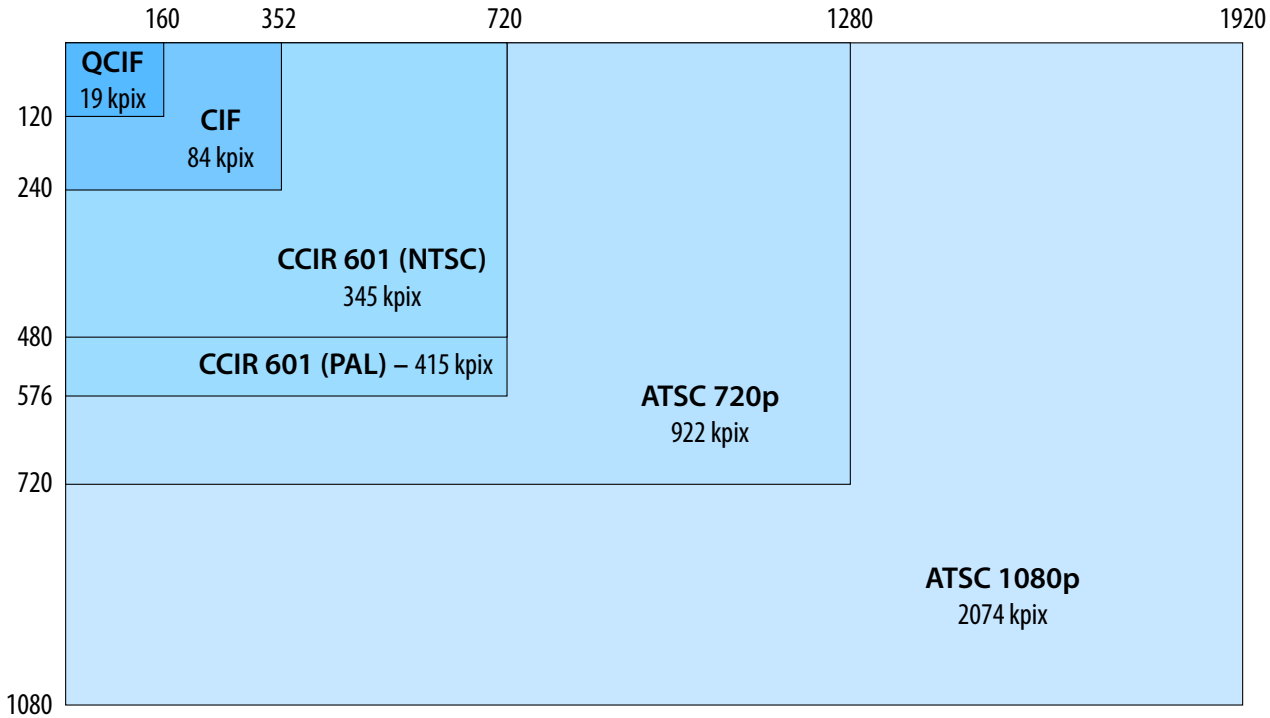

What Is Uncompressed HD About?

What is usually understood under uncompressed HD?
(1920 × 1080, 1.485 Gbps, transmitted over SDI, SMPTE 292M)

- Bandwidth calculation:

$$\underbrace{2200 * 1125}_{\text{total resolution}} * \underbrace{30}_{\text{bit/point}} * \underbrace{30}_{\text{fps}} * \underbrace{2/3}_{\text{4:2:2 sampling}} = 1.485.000.000 \text{ bps}$$

- Resolution: includes 1920 × 1080 of effective resolution, but also adds up blanking lines, totaling 2200 × 1125.
- Color depth: 10 bits/point/color plane $\Longrightarrow$ 30 bits/point
  - Computers are usually unable to render more than 8 bits/color plane.
- Frame rate: 24p, 25p, 29.97p, 30p, 50i, 59.94i, 60i
- Sampling: usually 4:2:2
  - 4:4:4 needed? You will also need double-link SDI...

What Is Uncompressed HD About?

What is usually understood under uncompressed HD?
(1920 × 1080, 1.485 Gbps, transmitted over SDI, SMPTE 292M)

- Bandwidth calculation:

$$\underbrace{2200 * 1125}_{\text{total resolution}} * \underbrace{30}_{\text{bit/point}} * \underbrace{30}_{\text{fps}} * \underbrace{2/3}_{\text{4:2:2 sampling}} = 1.485.000.000 \text{ bps}$$

- Resolution: includes 1920 × 1080 of effective resolution, but also adds up blanking lines, totaling 2200 × 1125.
- Color depth: 10 bits/point/color plane $\Longrightarrow$ 30 bits/point
  - Computers are usually unable to render more than 8 bits/color plane.
- Frame rate: 24p, 25p, 29.97p, 30p, 50i, 59.94i, 60i
- Sampling: usually 4:2:2
  - 4:4:4 needed? You will also need double-link SDI...

What Is Uncompressed HD About?

What is usually understood under uncompressed HD?
(1920 × 1080, 1.485 Gbps, transmitted over SDI, SMPTE 292M)

- Bandwidth calculation:

$$\underbrace{2200 * 1125}_{\text{total resolution}} * \underbrace{30}_{\text{bit/point}} * \underbrace{30}_{\text{fps}} * \underbrace{2/3}_{\text{4:2:2 sampling}} = 1.485.000.000 \text{ bps}$$

- Resolution: includes 1920 × 1080 of effective resolution, but also adds up blanking lines, totaling 2200 × 1125.
- Color depth: 10 bits/point/color plane $\Longrightarrow$ 30 bits/point
  - Computers are usually unable to render more than 8 bits/color plane.
- Frame rate: 24p, 25p, 29.97p, 30p, 50i, 59.94i, 60i
- Sampling: usually 4:2:2
  - 4:4:4 needed? You will also need double-link SDI...

What Is Uncompressed HD About?

What is usually understood under uncompressed HD?
(1920 × 1080, 1.485 Gbps, transmitted over SDI, SMPTE 292M)

- Bandwidth calculation:

$$\underbrace{2200 * 1125}_{\text{total resolution}} * \underbrace{30}_{\text{bit/point}} * \underbrace{30}_{\text{fps}} * \underbrace{2/3}_{\text{4:2:2 sampling}} = 1.485.000.000 \text{ bps}$$

- Resolution: includes 1920 × 1080 of effective resolution, but also adds up blanking lines, totaling 2200 × 1125.
- Color depth: 10 bits/point/color plane $\Longrightarrow$ 30 bits/point
  - Computers are usually unable to render more than 8 bits/color plane.
- Frame rate: 24p, 25p, 29.97p, 30p, 50i, 59.94i, 60i
- Sampling: usually 4:2:2
  - 4:4:4 needed? You will also need double-link SDI...

What Is Uncompressed HD About?

What is usually understood under uncompressed HD?
(1920 × 1080, 1.485 Gbps, transmitted over SDI, SMPTE 292M)

- Bandwidth calculation:

$$\underbrace{2200 * 1125}_{\text{total resolution}} * \underbrace{30}_{\text{bit/point}} * \underbrace{30}_{\text{fps}} * \underbrace{2/3}_{\text{4:2:2 sampling}} = 1.485.000.000 \text{ bps}$$

- Resolution: includes 1920 × 1080 of effective resolution, but also adds up blanking lines, totaling 2200 × 1125.
- Color depth: 10 bits/point/color plane $\Longrightarrow$ 30 bits/point
  - Computers are usually unable to render more than 8 bits/color plane.
- Frame rate: 24p, 25p, 29.97p, 30p, 50i, 59.94i, 60i
- Sampling: usually 4:2:2
  - 4:4:4 needed? You will also need double-link SDI…

HD-SDI over IP

- Packetization specified in RFC 4175 (“RTP Payload Format for Uncompressed Video”)
  - encapsulation: payload/RTP/UDP/IP
  - augmenting common RTP headers with payload headers (e.g., additional packet numbers because of fast RTP counter wrap-arounds—0.5 s for HD-SDI)
  - per video-line processing
- Reasonable to use jumbo frames (best $>$8,500 B)
  - decreases packetization size overhead
  - decreases host load due to decreasing number of pps

Implementation

- Linux based
- comes from UltraGrid by Perkins & Gharai
  - extended to support full-HD 1080i
  - support for SW display including color space down-sampling and field de-interlacing (assembly optimized for AMD64)
  - number of other enhancements
- used with DVS Centaurus HD capture cards
  - problems with latency, since the card doesn't support DMA and requires buffering at least 4 fields for reliable operation
  - quite expensive
  - there are other cards but not supported in Linux :(

Application Workflow

Capture part:

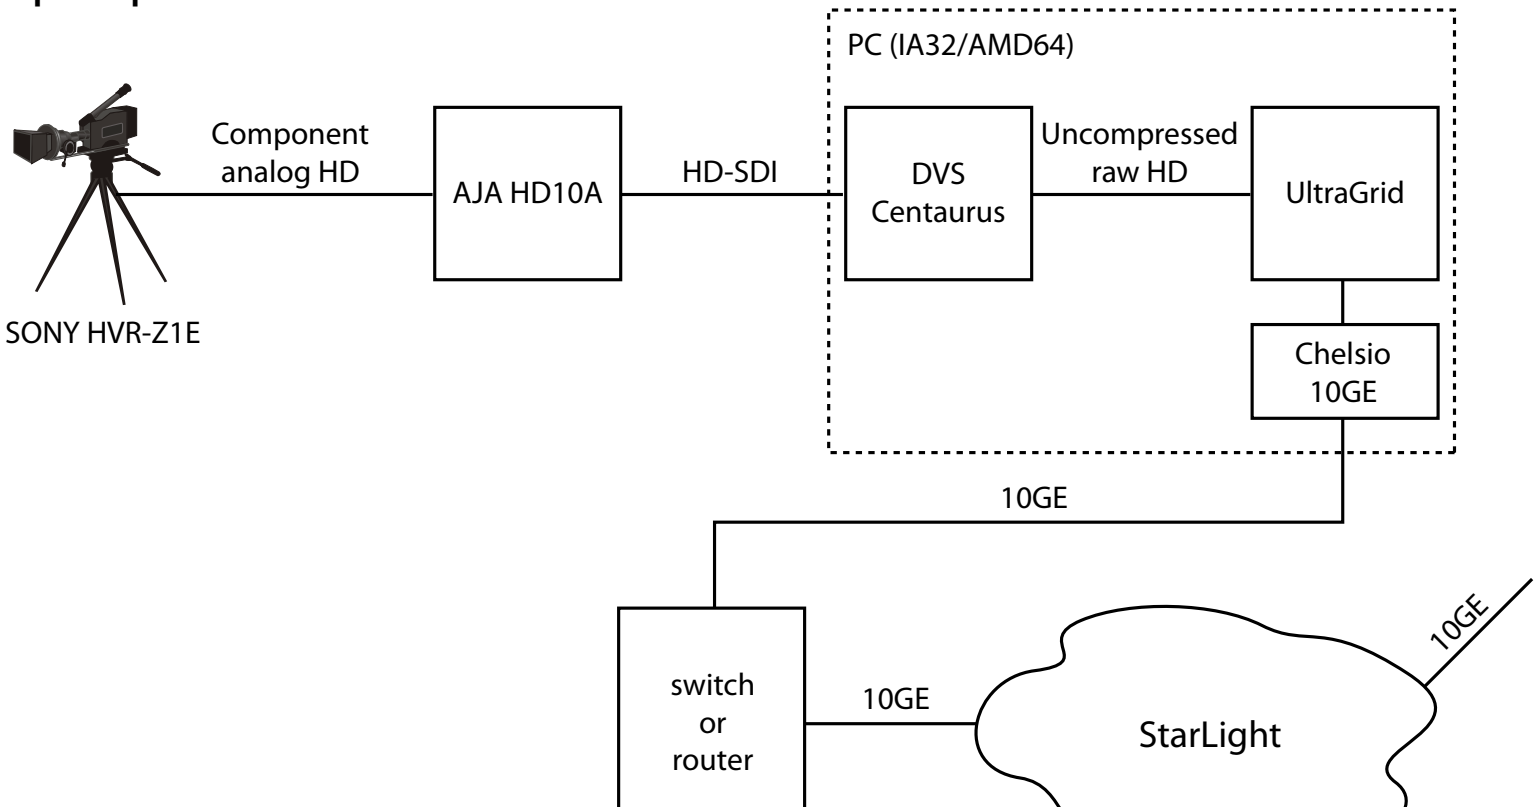

Application Workflow

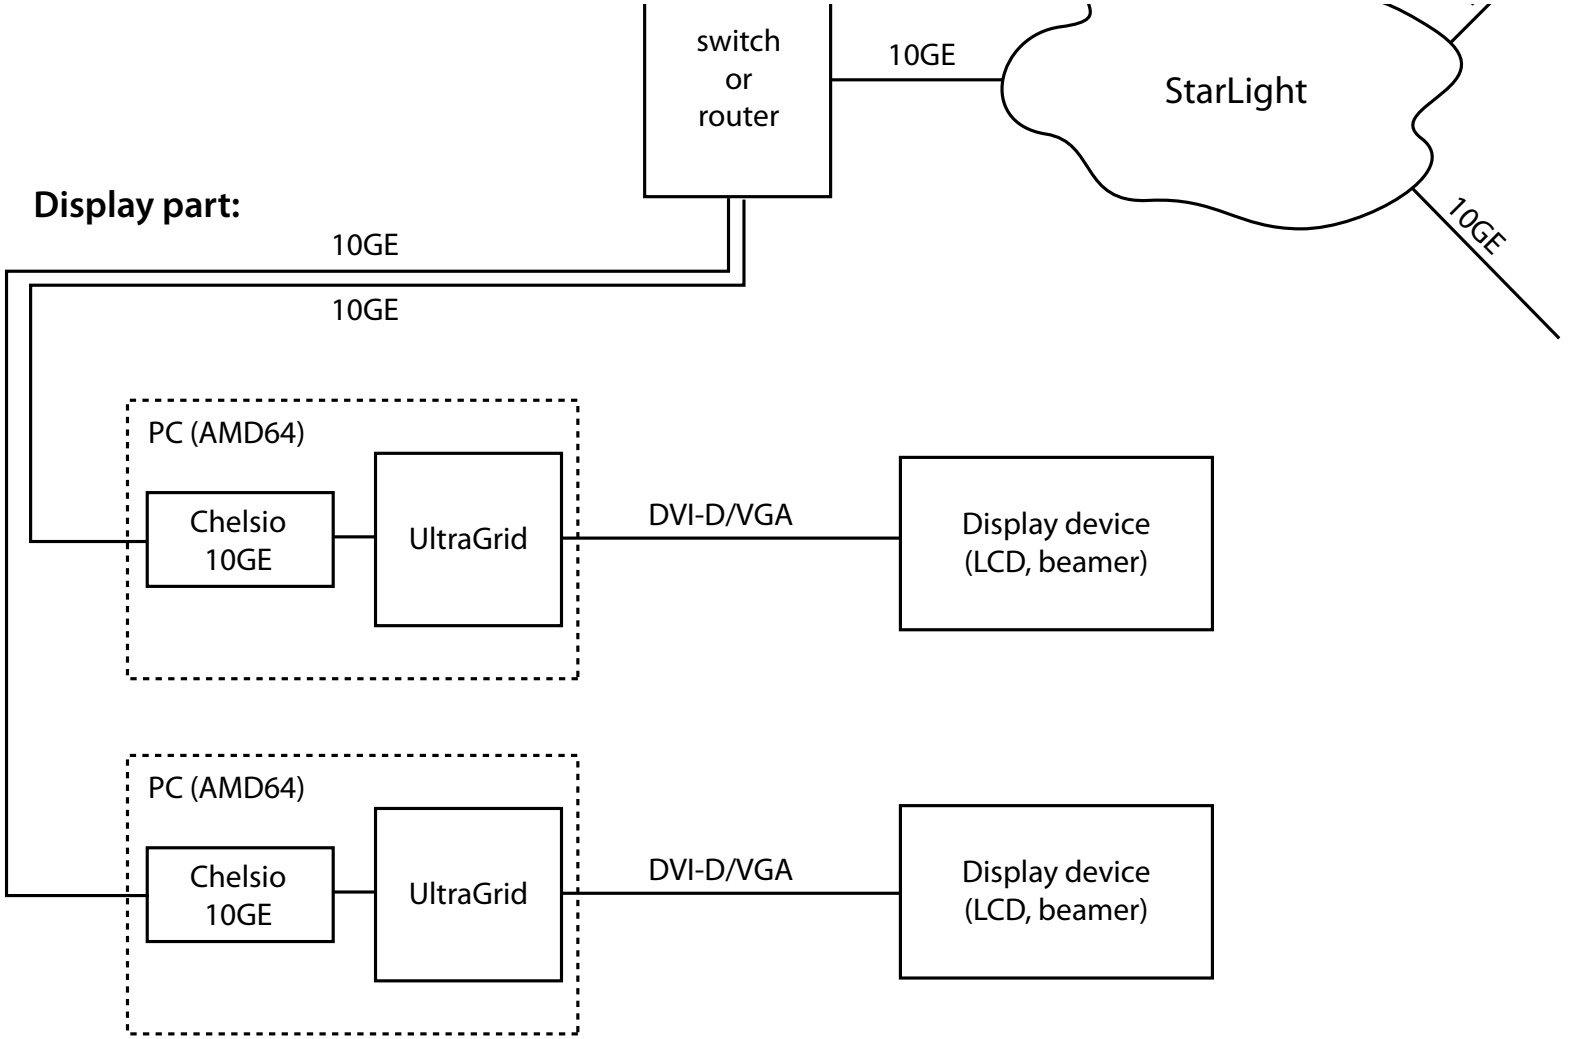

Measured E2E Latency

- unidirectional latency measurement
  - SONY Z1E → AJA HD10A → PC → PC → LCD

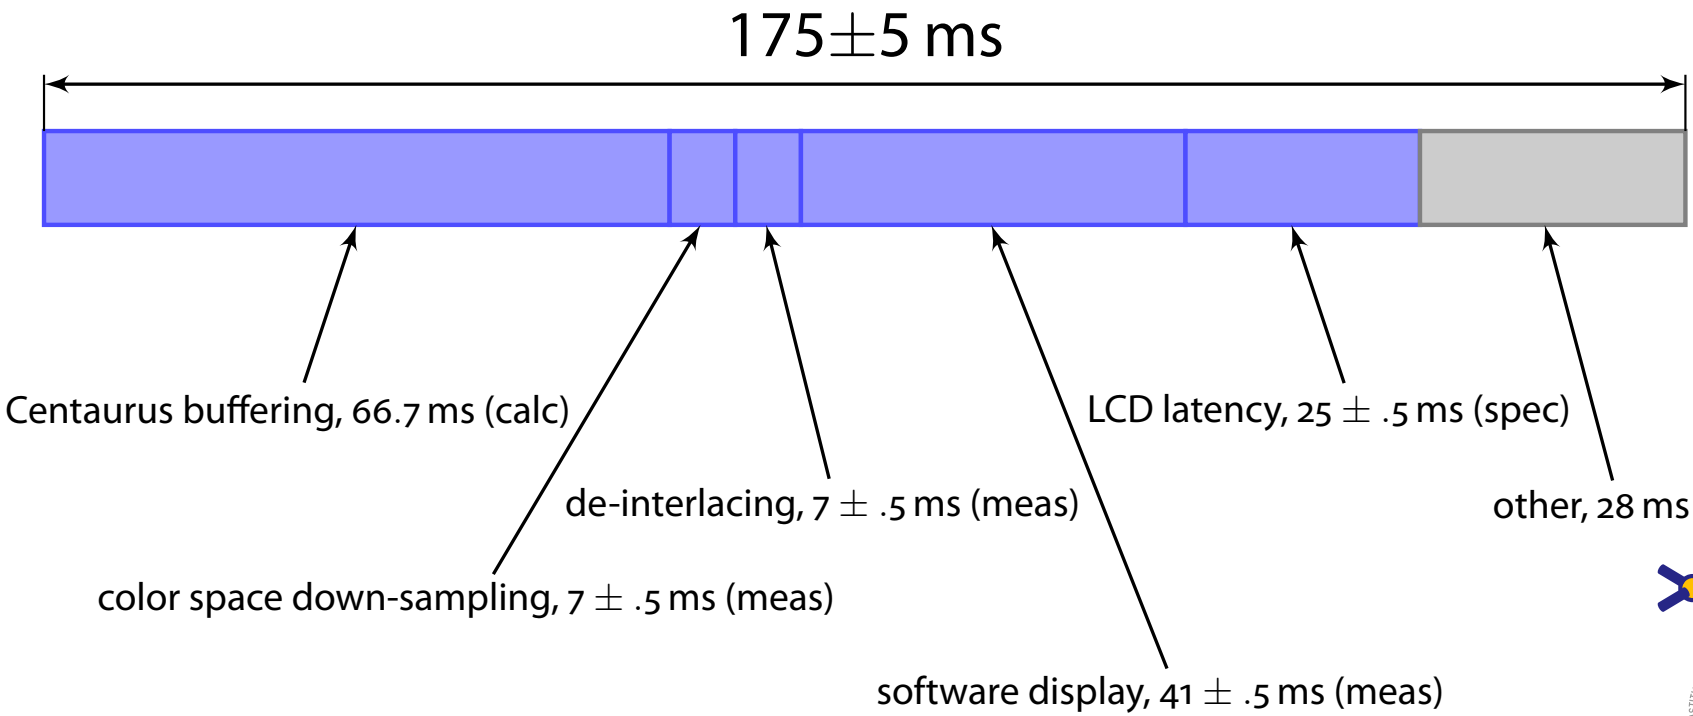

Network data distribution

- Point-to-point transmissions
  - natural for todays networks
- (Multi)Point-to-multipoint transmissions
  - requires distribution service either on network or application layer
  - IP multicast
    - native
    - UDP packet reflectors—Active Elements
    - optical splitters

iGrid 2005 Workshop

- 26.–29. 9. 2005, San Diego—La Jolla, CA, USA
- `http://www.igrid2005.org`
  - "iGrid 2005, the 4th community-driven biennial International Grid event, is a coordinated effort to accelerate the use of multi-10Gb international and national networks, to advance scientific research, and to educate decision makers, academicians and industry researchers on the benefits of these hybrid networks."
- the most bleeding edge applications
- organized by Global Lambda Integrated Facility (GLIF)

Our Demos

- CZ101—HD Multipoint Conference
  - demonstration of uncompressed HD-based collaborative environment
- US127—Interactive Remote Visualization across the Louisiana Optical Network and the National LambdaRail
  - demonstration of collaborative visualization based on the uncompressed HD video transmission
  - joint demo with Louisiana State University
- equivalent of US127 demo has been shown also at SuperComputing 2005

- Implementation:
  - three places communication using uncompressed HD streams
- problems:
  - for 3-way conference it means 1.5 Gb/s egress and 3 Gb/s ingress
  - Active Elements need to handle at least 4.5 Gb/s

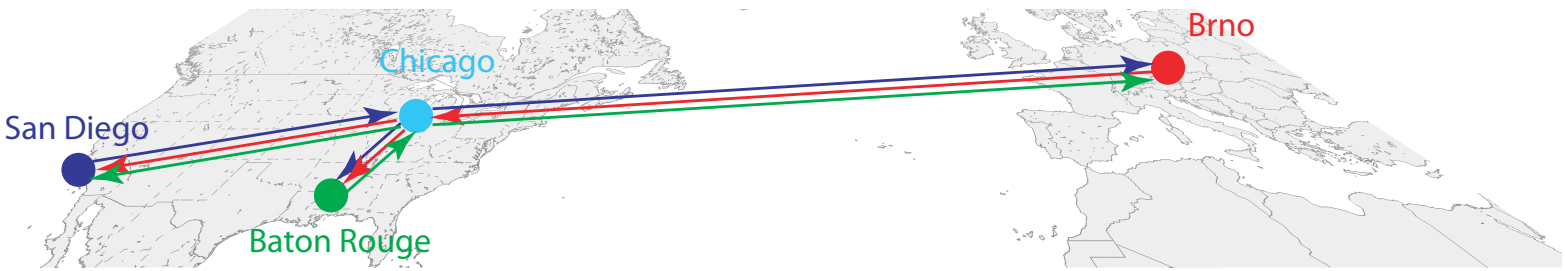

Active Element Setup

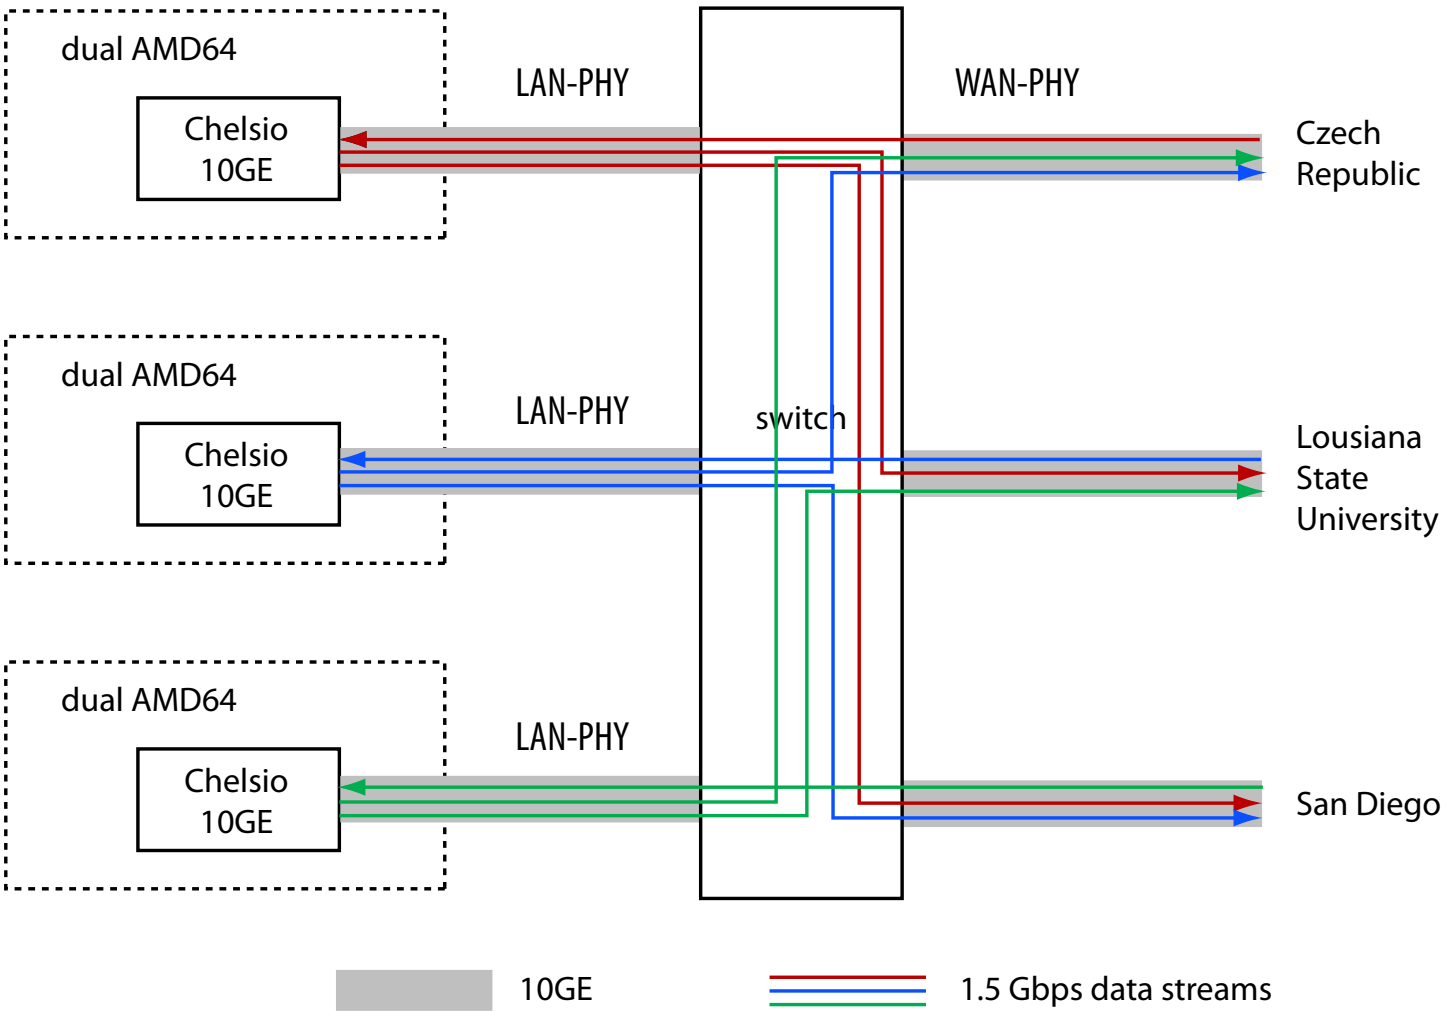

Experienced Latencies

- Reflector latency: 13±2 ms
- Circuit latencies:
  - San Diego ↔ StarLight: 78.2±.2 ms (routed)
  - Louisiana ↔ StarLight: 31.09±.04 ms (switched)
  - Brno (CZ) ↔ StarLight: 126.7±.3 ms (routed)

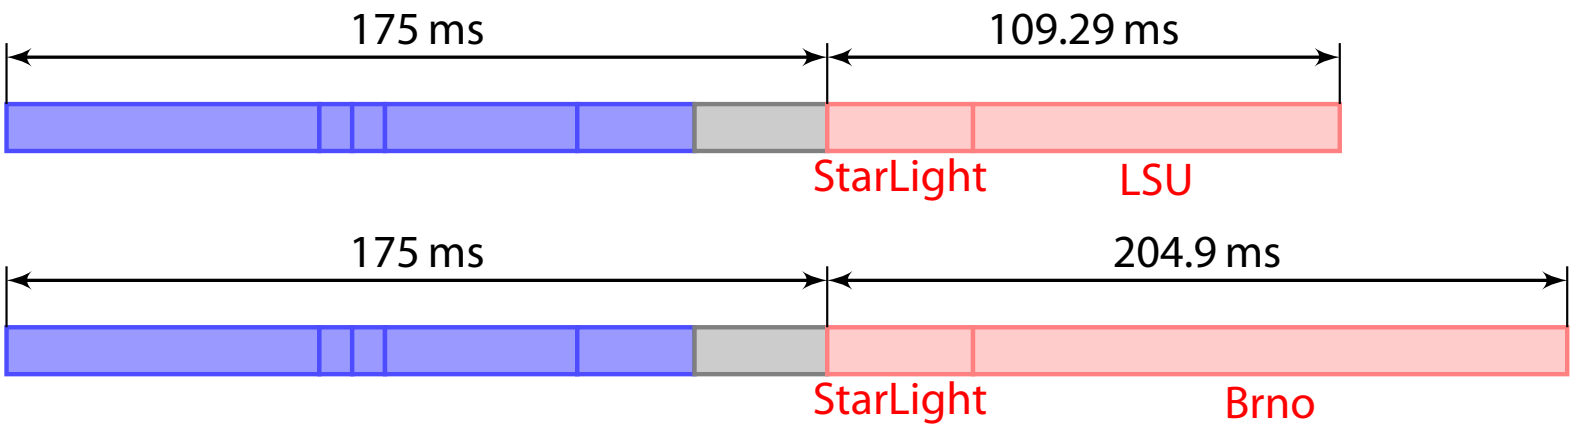


Brno

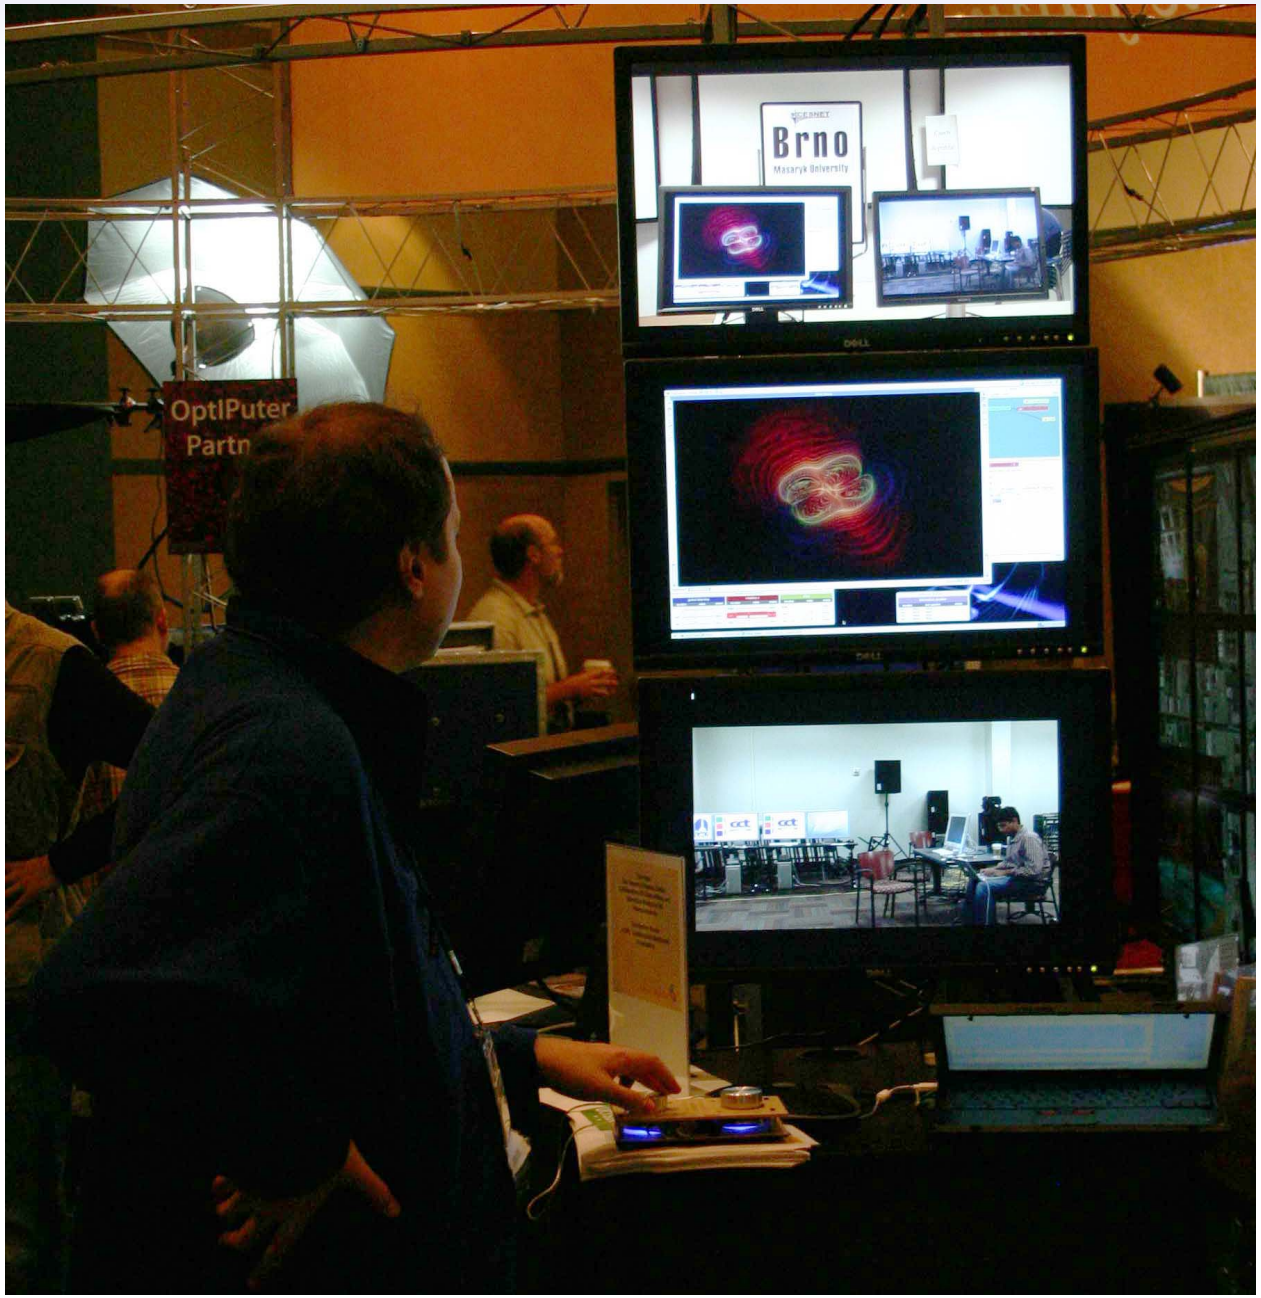
Brno
Masaryk University
OptIPuter

Future Work

- Going beyond the HD resolution
  - 2K, 4K, 6K resolutions and ultra-HD resolution

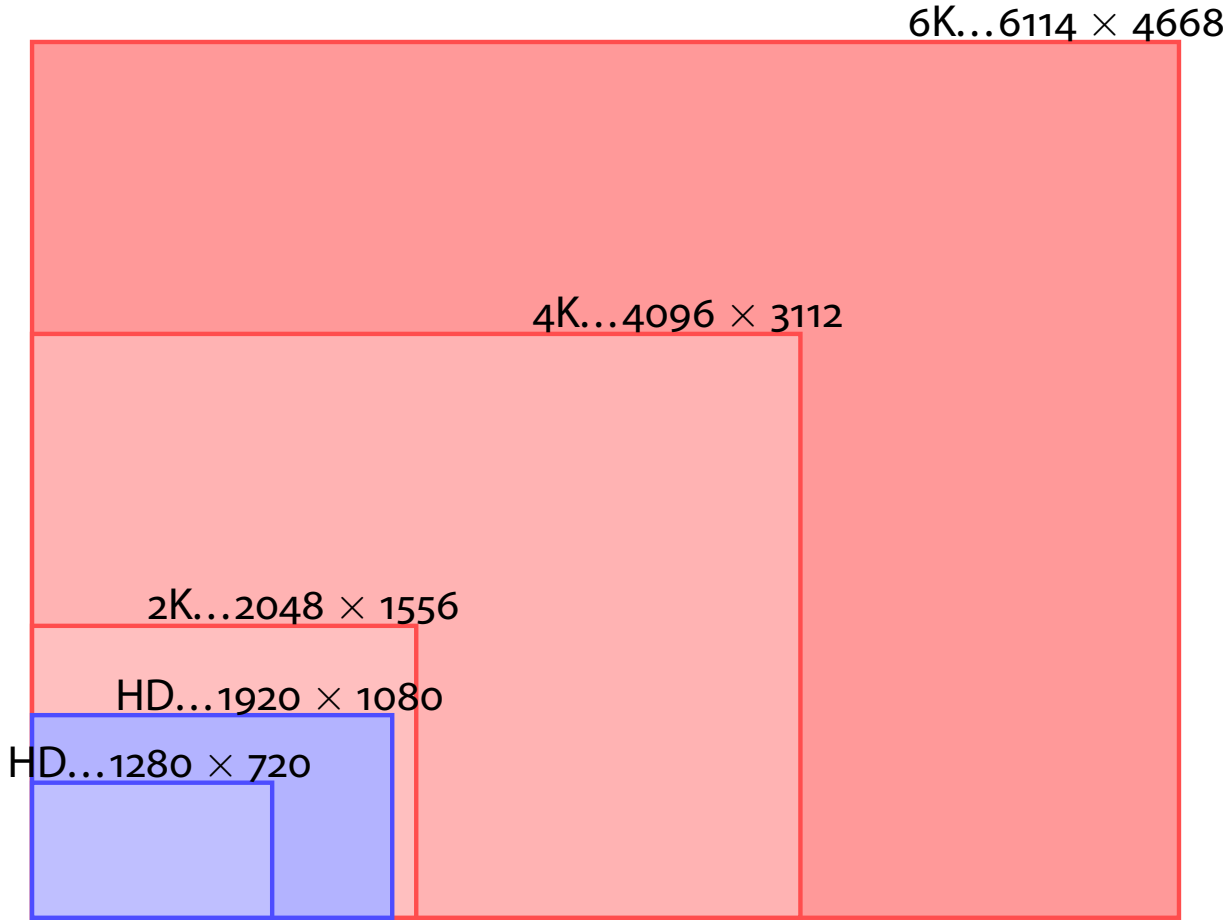

Future Work

- Stereoscopic/holographic projections
- Truly high-definition audio
  - most of the current problems are with physical/electrical/acoustic installation of the non-computer equipment
  - higher quantization and sampling resolutions (24 b @ 192 kHz)
- Incorporating other natural media for natural immersion

Future Work

- More efficient reflectors
  - optical splitters (not deployed for iGrid and SC demos as it is too experiemntal and requires L1 network access)
- We need a reasonable network control plane to interact with!
  - the control plane that is able to work across multiple administrative domains
  - manual optical currently requires tremendous effort by host of people
  - integration with reflectors (be it implemented on any level)

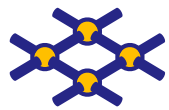

Related Projects

- UltraGrid (Colin Perkins, Ladan Gharai)
  - `http://csperkins.org/research/ultragrid/`
  - available under (pseudo-)BSD license
  - Ladan in incorporating our code (version 0.4.3)
- ResearchChannel (Michael Wellings et al.)
  - an implementation under MS Windows
  - availability of cheaper AJA XENA LH and Black-Magic DeckLink HD
  - currently switching to open-source
  - the data is not sent according to the RFC 4175


Conclusions

- Transmission of uncompressed HD 1080i @ 1,5 Gb/s
  - \+ low latency
  - \+ high image quality
  - − high demands on the network capacity
  - − high demands on scalability of the distribution network
  - − so far relatively expensive
- For more information:
  `https://sitola.fi.muni.cz/igrid/`

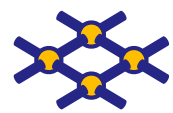

Who is involved…

- Leading: Petr Holub and Luděk Matyska
- Miloš Liška, Lukáš Hejtmánek, Jiří Denemark, Tomáš Rebok
- others for the demos…

More details are in the special issue of Future Generation Computer Science by Elsevier Science dedicated to iGrid 2005.

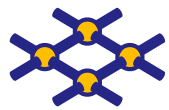


Thank you for your attention!

Q?/A!

hopet@ics.muni.cz
https://sitola.fi.muni.cz/igrid/