HD Multipoint Conference

iGrid 2005 CZ101

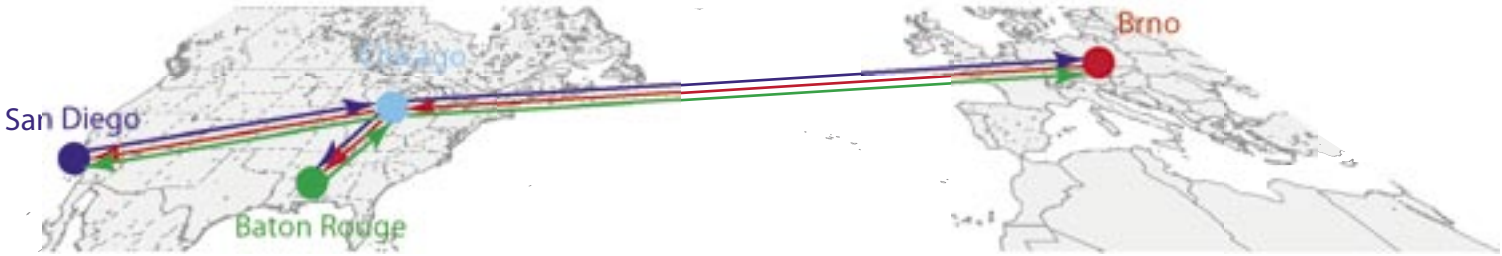

Contacts:
Luděk Matyska
CESNET/Masaryk University (CZ)
ludek@ics.muni.cz

Petr Holub
CESNET/Masaryk University (CZ)
hopet@ics.muni.cz

Collaborators:
CESNET/Masaryk University (CZ):
Luděk Matyska, Petr Holub, Miloš Liška, Lukáš Hejtmánek, Jiří Denemark, Tomáš Rebok

CESNET (CZ):
Stanislav Šíma, Jan Radil, Jan Gruntorád

Louisiana State University (USA):
Ed Seidel, Andrei Hutanu, Ravi Paruchuri

We present a virtual, interactive environment usable for high-quality lecturing and discussion, linking Brno (Czech Republic), LSU, and San Diego. It creates an illusion of a single meeting place for all. A lecture will be given at one of the sites, followed by a live discussion and interaction among the audience of all the three sites. To minimize delay and support ability to real time speach interruptions and to create a nearly immersive look-and-feel environment, uncompressed high definition video is required. One raw stream takes about 1.5 Gbps, so the three-site configuration needs around 5 Gbps and a larger multi-point configuration is likely to overload a single 10 Gbps link. There have been previous demonstrations of low-latency communication based on raw high-definition video between two sites, but this project demonstrates the strength of a multi-point communication based on virtual (simulated) multicast environment.

Networking Setup Overview

For the three-point HD transmission, we need a powerful networking infrastructure capable of distributing one egress 1.5 Gbps stream and two ingress 1.5 Gbps streams to each site. We use both dedicated and shared lambda lines for connecting each participating site. The networking part of the demo consists of the following items:

- 10GE networking interfaces for PCs based on Chelsio T110/T210 cards,
- Cisco 6505, 15454, and Force10 switch(the 6505 with LAN-PHY and WAN-PHY XENPAKs for framing conversion)
- interconnect at StarLight, Chicago (Figure 1).

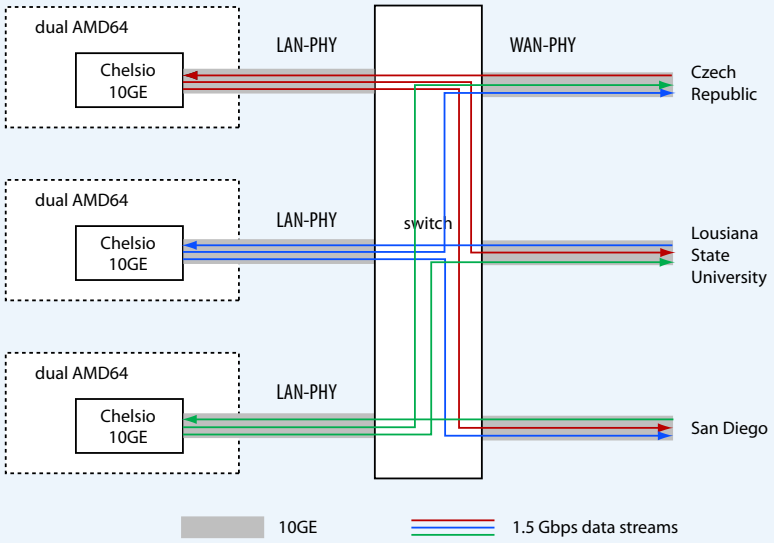

Figure 1: Reflector-based replication setup at StarLight, Chicago

Video Multiplication Overview

For multiparty videoconferencing, the video needs to be multiplicated to all the participants. Our long-term experiences with native multicast direct us to use distinct data replication mechanism. The following possibilities were considered and tested for the iGrid demo:

- software reflector (centralized or distributed),
- L1 optical splitter.

Reflector. Virtual multicast technology based on software reflector has been developed at the CESNET and Masaryk University. We have several implementations ranging from simple ultra-fast reflectors capable of replicating stream over 2 Gbps on a single machine, to complex active router based implementation with modular architecture. Also there is a distributed implementation that allows replication of data streams with bandwith beyond capability of any single machine. The ultra-fast implementation is used for the demo. The reflector is located at StarLight and runs on dual AMD Opteron computers with Chelsio T110/T210 NICs.

Optical Splitter. The optical splitter should serve in place of the reflector. It should provide far least latency, it should be speed independent and in some cases it could work over multiple lambdas (for broad-spectrum splitters). This scenario has been successfuly demonstrated in CESNET laboratory, but it was too complicated for deployment during this event.

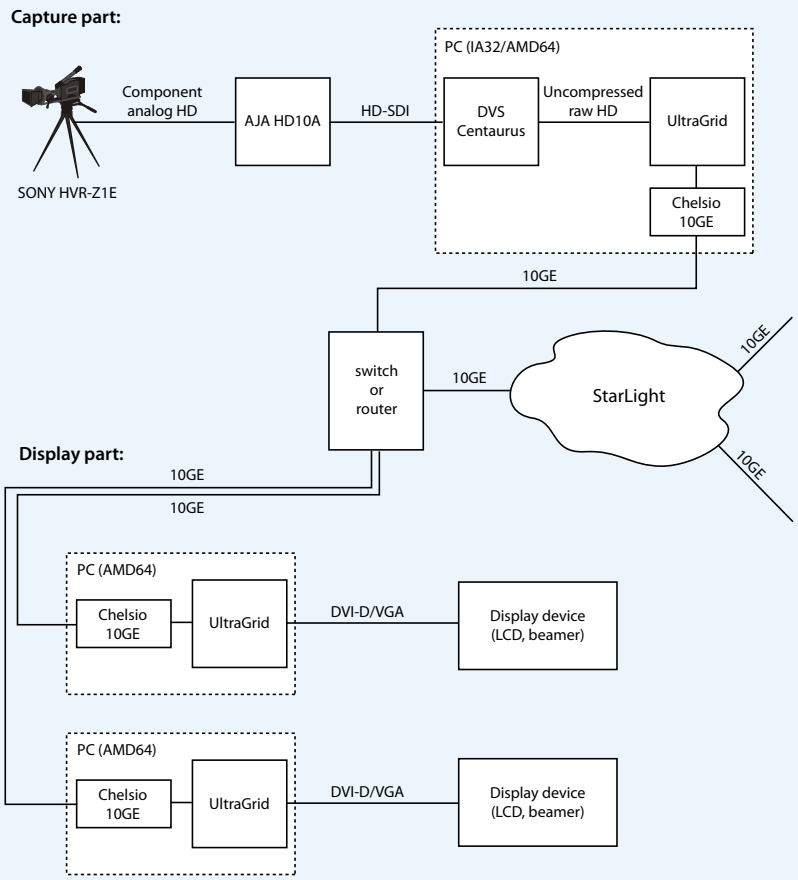

Figure 2: HD capture – send and receive – display workflows for each site

End-User Applications and Hardware Overview

Overview of the processing workflow is shown in the Figure 2.

Capture and sending part:
- HD camera SONY HVR-Z1E or SONY HDR-FX1 with AJA HD10A for analog component (YUV) to SDI conversion,
- hardware capture of HD-SDI (DVS Centaurus),
- software encapsulation into IP using custom UltraGrid-based software,
- sending using 10GE NIC.

Receiving and display part:
- reception using 10GE NIC,
- software decapsulation from IP and software display on AMD Opteron based machine using custom UltraGrid-based software (including software 10b to 8b color depth reduction and AMD64 optimized deinterlace algorithm implementation).