

## UltraGrid: Updates & Plans

Petr Holub  
CESNET z.s.p.o., Prague/Brno, Czech Republic

<Petr.Holub@cesnet.cz>



*Internet2 Collaboration SIG*  
2013-04-22



# UltraGrid Platform

- Technology

- an affordable platform for high-quality interactive image transmissions
- use of commodity hardware
  - ◆ PC (Linux, Windows **NEW!**) and Mac (MacOS X) platforms
  - ◆ commodity video capture cards
  - ◆ commodity GPU cards
  - ◆ 10GE is a plus but not necessary
- as low latency as possible on commodity hardware
- open-source software, BSD license
- a platform for validating research results (not just ours! :) )
  - ◆ compression & image processing, FEC, scheduling, congestion control...



# Applications of UltraGrid

- Generic scientific visualization
- Medicine
  - X-ray imagery, cardiology, pathology



# Applications of UltraGrid

- Education
  - remote education



# Applications of UltraGrid

- Cinematography

Detached BaseLight consoles at CinePost (Barrandov, CZ)

Mac Pro

Mac Pro

10GbE

Kona 3

Kona 3

dual-link HD-SDI

dual-link HD-SDI



BaseLight Four



SONY SXR4K



# Applications of UltraGrid

- Arts
  - distributed performances: music, theater



# UltraGrid Platform

- Supported formats
  - HD, 2K
  - 4K – tiled or native
  - 8K **NEW**
  - multichannel video (e.g., 3D HD, 4K)
- Uncompressed vs. compressed
  - low-latency compression
  - GLSL-accelerated DXT1, DXT5-YCoCg
  - CUDA-accelerated JPEG, DXT5-YCoCg
  - CPU-based DXT1, ffmpeg (e.g., H.264)
- Supported audio formats
  - uncompressed, multi-channel
  - Opus codec **NEW**



# UltraGrid Platform

- I/O
  - capture/playback cards: HD-SDI, SDI, HDMI, analog HD and SD
    - ◆ manufacturers' SDKs, Video4Linux2, QuickTime
  - screen capture input (up to 4K)
  - computer screen output (OpenGL, SDL)
  - SAGE output
  - specialized display filters
  - HDMI 1.4a: stere-HD, 4K
- Image composer **NEW**
- Full-duplex operation
- Simple GUI
  - QT-based, native MacOS
  - permanent storage of configuration
  - simple startup + advanced configuration dialog

Line-interlaced stereoscopic video





# UltraGrid Platform

## GUI on MacOS X

The screenshot displays the UltraGrid GUI on a Mac OS X desktop. The main window shows a live video feed of a server rack in a room with a clock and a sign that says "MNL". Overlaid windows include a terminal window with system logs, a "UltraGrid GUI" connection window with "Remote IP address: h64.R.manit.ca", and "UltraGrid Advanced Settings" with fields for MTU (1500), Display (qt), Capture, and Audio play (coreaudio 42). The dock at the bottom shows various Mac OS X icons.

# UltraGrid Platform

## GUI on Linux

The screenshot displays the UltraGrid GUI on a Linux desktop environment. The window title is "ultraGrid". The interface includes several panels:

- About:** Contains fields for "Remote IP address" and "URL/hostname", with "Terminal" and "Stop" buttons.
- Advanced settings:** A dialog box for "ultraGrid advanced settings" with fields for "MTU" (1500), "Display" (gl), "Capture" (testcard), "Audio capt" (embedded), and "Audio play" (alsa default). It also has "Save to config file", "OK", and "Cancel" buttons.
- Terminal:** A window titled "minerva" showing a terminal session with various commands and output, including "ultraGrid" and "ultraGrid - OpenGL Display".
- Output Log:** A window titled "ultraGrid output" displaying a list of performance metrics, such as "123 Frames in 5.8253 seconds = 24.6743 FPS".
- 3D Scene:** A window titled "ultraGrid - OpenGL Display" showing a 3D rendered scene of a desk with a computer monitor, keyboard, and chair. The monitor displays "MNL".



# UltraGrid Platform

- Audio

- balanced, unbalanced, HD-SDI, HDMI
- various system interfaces including JACK
- PortAudio, ALSA, CoreAudio, JACK
- embedded HD-SDI/HDMI
- simple mono software echo canceler based on Speex
- channel mixer/duplicator



# GPU-Accelerated Compression

- Available compression schemes
  - DXT1: CPU-based (FastDXT library from EVL)
  - DXT1, DXT5: OpenGL Shader Language (GLSL) based
  - JPEG: NVidia CUDA based
  - DXT5: NVidia CUDA based (for 8K)

SAGE display with various compressions





# GPU-Accelerated Compression

- Performance numbers (including transfer to/from GPU)
  - DXT1 GLSL: 798 Mpix/s (Nvidia 580GTX), 593 Mpix/s (ATI 6990)
  - DXT5 GLSL: 349 Mpix/s (Nvidia 580GTX), 305 Mpix/s (ATI 6990)
  - JPEG CUDA: up to 1.580 Mpix/s = 4.740 MB/s (Nvidia 580GTX, 4:4:4, Q=60)
  - DXT5 CUDA:  $\geq 1.580$  Mpix/s (Nvidia 580GTX)






# GPU-Accelerated Compression

- Performance of JPEG stages for 2160p video



(a) for JPEG encoder



(b) for JPEG decoder

# Forward Error Correction

- LDGM
  - CPU and GPU implementations
  - CPU (SSE optimized) is used because of CPU↔GPU transmissions overhead
  - packet loss up to 10% can be mitigated with reasonable overhead
  - can make JPEG survive up to 25% packet loss
  - performance issues above 2 Gbps
- Simple method: shifted multiplication

# Recent Updates

Since October 2012

- ffmpeg support – low latency H.264
  - **if linked with X264, UltraGrid becomes GPL** (GPL is viral)
  - starts at 150% CPU core for HD (settings-dependent) , well usable at >18 Mb/s
  - 4K being examined
  - due to licensing issues, we don't interface directly to X264 and leave it up to the user via ffmpeg/libavcodec
  - ultrafast vs. superfast quality settings
  - low-latency (“zero-latency”) mode
  - I-frames distributed in time to reduce bursts





# Recent Updates

Since October 2012

- Windows port
  - OpenGL, SDL displays
  - native BlackMagic SDK
  - DirectShow capture



# Recent Updates

Since October 2012


- Support for DELTACAST DVI-I/DVI-D grabbers
  - ideal for content capture, computer screen resolutions
  - supports multiple cards (e.g., 6x DVI-I in a single PC)
- File-based I/O
  - input/output of raw data
  - can be piped into mencoder (but not very convenient)
  - planned integration with further processing (e.g., GStreamer) for lecture/event/experiment recording, etc.
- Transcoding reflectors
  - change of formats “along the way”, as a part of multi-point data distribution
  - implemented using UltraGrid as backend
  - intended for automated setup with CoUniverse (later in 2013)



# Recent Updates

Since October 2012

- Integration of 2-camera GColl
  - group-to-group communication with partial gaze awareness



# Recent Updates

Since October 2012



- BlueFish444 capture card support
  - sub-frame I/O: a frame may be split up into 4 pieces
  - HD, 4K capture
- Audio compression based on Opus codec  
<http://www.opus-codec.org/>
  - uncompressed audio typically uses 1.5 Mbps ×3 for redundancy reasons
  - features both narrowband (voice) and fullband (music) compressions
  - includes SILK codec developed by Skype



# Recent Updates

Since October 2012

- Opus quality comparison



# Recent Updates

Since October 2012

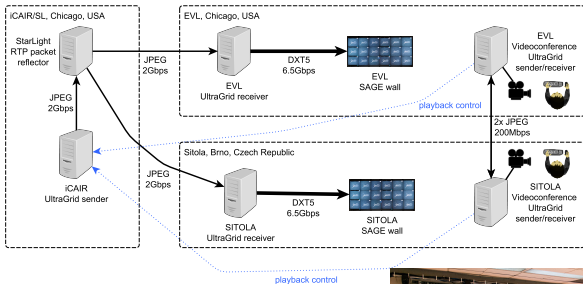
- Multichannel video processor (composer)
  - composition of images up to 4K
  - utilizes either GPU or CPU
  - allows logo overlay
  - allows black window overlay (for information removal, such as in medicine)



- composition is done typically on the sender

# World Firsts... 8K on Commodity PC

- 2012 – GPU-JPEG Transatlantic Multi-Point 8K



- from pre-rendered sources
- JPEG → DXT5-YCoCg on a single machine
- useful also as 16× HD (multi-camera setups)



# Award by ACM Multimedia SIG

- ACM Best Open-Source Software Competition Award





# Now... what bandwidth do I need?

(just rough estimates)

HD: 1080i50/59.94, 720p50/60

<i>Coding</i>	<i>Minimum [Mbps]</i>	<i>Optimum [Mbps]</i>
Uncompressed 4:2:2	1,500	1,500
DXT5	500	500
JPEG	60	200
H.264	5	30

4K: 2160p25/29.97

<i>Coding</i>	<i>Minimum [Mbps]</i>	<i>Optimum [Mbps]</i>
Uncompressed 4:2:2	6,000	6,000
DXT5	2,000	2,000
JPEG	150	500
H.264	15	80



# Latency

- Latency limits
  - <150 ms for interactivity: ITU-T rec G.114
- End-to-end latency
  - in a local network
  - measured using video (1/60 s quantization)
  - depends substantially on hardware cards used (2.0–5.0 frames)
  - Bluefish444 should get us a bit lower: sub-frame API
  - application-level traffic shaping to control bursts
- Uncompressed for DeckLink HD → DeltaCast 3G
  - 2.5 frames (83 ms)
- Impact of compressions
  - 2.5 frames (+<16.7 ms) for CUDA JPEG
  - 3.5 frames (+33.3 ms) for GLSL DXT1/5
  - 5 frames (+66.6 ms) for H.264 on a CPU (preliminary measurements)

# Future Plans

- Short-term:
  - Advanced multi-point
  - Self-organization with scheduling (release with CoUniverse)
  - Piping to GStreamer for further processing
- Long-term:
  - Acceleration of low-latency H.264/H.265
  - New compression formats for specific purposes (e.g., SAGE)
- Collaboration with izcat
  - shared small-scale subproject within the VisionAir 7FP EU project on H.264 performance and usability for 4K video
  - integration of izcat development results
  - MCU
  - SIP signaling support

# Availability

- UltraGrid
  - <http://ultragrid.sitola.cz/>
  - <http://sourceforge.net/p/ultragrid/>
  - <https://www.facebook.com/UltraGrid>
- GPUJPEG library
  - available both as part of the UltraGrid and as a standalone library
  - <http://gpujpeg.sourceforge.net/>



## Selected Papers

- [1] HOLUB, Petr - MATYSKA, Luděk - LIŠKA, Miloš - HEJTMÁNEK, Lukáš - DENEMARK, Jiří - REBOK, Tomáš - HUTANU, Andrei - PARUCHURI, Ravi - RADIL, Jan - HLADKÁ, Eva. High-definition multimedia for multiparty low-latency interactive communication. *Future Generation Computer Systems*, Amsterdam, The Netherlands, Elsevier Science, The Netherlands. ISSN 0167-739X, 2006, vol. 22, no. 8, pp. 856–861.
- [2] MATELA, Jiří - RUSŇÁK, Vít - HOLUB, Petr. Efficient JPEG2000 EBCOT Context Modeling for Massively Parallel Architectures. In *Data Compression Conference (DCC)*, 2011. Washington, DC, USA : IEEE Computer Society, 2011. ISBN 978-0-7695-4352-9, pp. 423–432. 2011, Snowbird, Utah, USA.
- [3] MATELA, Jiří - ŠROM, Martin - HOLUB, Petr. Low GPU Occupancy Approach to Fast Arithmetic Coding in JPEG2000. *Mathematical and Engineering Methods in Computer Science, Lecture Notes in Computer Science*, Heidelberg, Springer Berlin / Heidelberg, Germany. ISSN 0302-9743, 2011, vol. 2012, no. 7119, pp. 136–145.
- [4] HOLUB, Petr - ŠROM, Martin - PULEC, Martin - MATELA, Jiří - JIRMAN Martin. GPU-Accelerated DXT and JPEG Compression Schemes for Low-Latency Network Transmissions of HD, 2K, and 4K Video. *Future Generation Computer Systems*. Submitted 2012.

# Thank you for your attention!

<petr.holub@cesnet.cz>  
<ultragrid-dev@cesnet.cz>

This work is supported by LM2010005 project.