



TRANSMITTING LIVE UNCOMPRESSED 8K VIDEO AND STILL NOT BUILDING CORRESPONDING SDN NETWORKS

Miloš Liška

Sitola Seminar

Oct. 24th 2018

Brno

■ Technology


- an affordable platform for very high-quality interactive video (up to 8K) and audio transmissions
- use of commodity (gaming) hardware
 - Linux and Windows PC and Mac OS platforms
 - commodity video capture cards
 - commodity GPU cards
 - commodity sound cards
 - any reasonable network
- as low latency as possible on commodity hardware
- open-source software, BSD (GPL) license

■ Community, user support

- New QT GUI
- BitFlow capture cards support
- Complete Spout and Syphon support
- NVidia and Intel based video decoding acceleration
- AJA cards support on Windows and MacOS
- Multiple improvements of Blackmagic cards support (multi-link, 3D, passthrough)
- Performance and latency improvements of OpenGL display
- Support for 64-bit Windows builds
- Support for Comprimato J2K codec



- We have already successfully demonstrated compressed 8K video transmission in 2012 and 2016



- Uncompressed transmissions are much more challenging

- 8K 60 fps video = $7680 \times 4320 \times 60 \times 16$ (4:2:2 subsampling with 8b) =
31.85 Gbps
- Implementation based completely on commodity HW
- 8K video pipeline, zero-copy operations as much as possible
- RTP/UDP data transmission, jumbo frames!
- As low latency as possible

- Problem of fitting 8K video capture card(s), NIC, NVME Raid and GPU(s) into a PC physically and connecting all subsystems to the CPU with sufficient PCI-E lanes/bandwidth
 - Intel architecture, PCI-E 3.0 usually limited up to 40 lanes
 - 4 lanes always reserved for south bridge
 - 8K video capture card = 8 lanes
 - 40GE NIC = 8 lanes, 100GE NIC = 16 lanes
 - NVME RAID = at least 8 lanes (has to bypass south bridge!)
 - GPU = 16 lanes (can work with 8)
- Limited physical space on the motherboard together with semi-static assignment of PCI-E lanes to physical slots

- Traditional video interfaces still lacking enough bandwidth
- 16x 3G-SDI, now 4x 12G-SDI
 - AJA CorVid88, Blackmagic 8K Pro
- 4x HDMI 2.0 could work, but not supported by any existing video capture card, using 4 cards is not practical
- HDMI 2.1 is only emerging these days

- Video capture → Packetization → Payload reconstruction → Video display
- Video capture and sending
 - BM: BM → **BM double buffer** → **Framebuffer** → **Packets payload**
 - AJA: AJA low-level Double/Ring buffer → **Framebuffer** → **Packets payload**
- Receiving and video display
 - NIC → kernel net buffer → **framebuffer copy** (packets still in L3 cache) → **OpenGL texture copy** (directly to Texture Object bypassing the Pixel Buffer Object)

■ Bursts!

- RTP/UDP transmission with no flow control/congestion control
- One sends packets at the link speed using UDP (i.e., 40Gbps)
- With 9000B jumbo frames 1.8us/packet at 40Gbps (0.7us at 100Gbps) – interrupts handling!
- Frame based operations, single frame is 5ms at 100Gbps which potentially leads to congestion

■ Other reason is the limited PCI-E lanes availability, 100GE NIC requires 16 lanes

- Tiled 4K displays (also Sharp LV-70X500E)
 - 4x HDMI 2.0 or DP 1.1
 - Nvidia Quadros or Radeon Pro SSG
 - Gaming cards feature 3xDP + 1xHDMI2.0 with different displaying latencies
- Samsung QLED 8K
 - HDMI2.1 → NVidia RTXs
- Dell UP3218K UltraSharp
 - 2x DP 1.2 → Nvidia GTX (Pascal and newer)

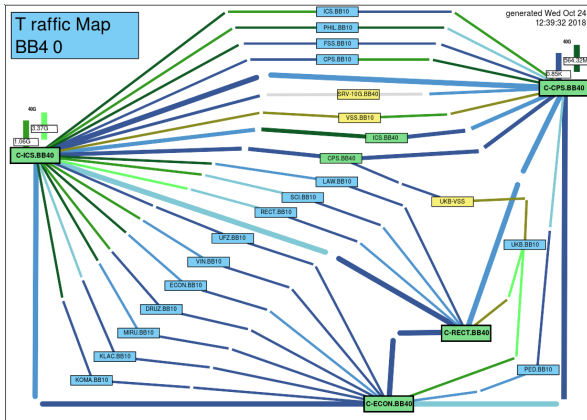
■ Sender update

- Asus Z170-A Motherboard + i7-6700K CPU + 32GB RAM
- **Sharp 8C-B60A** 8K camera with commodity quad 12G-SDI interface
- Single **Blackmagicdesign DeckLink 8K Pro** capture card
- **Mellanox ConnectX 40/56 GbE NIC**
- Allows to add additional GPU such as NVidia Quadro P4000 or better to support live 8K H.265/HEVC compressed transmissions
- Allows to add **ASUS HYPER M.2 x16 PCIe VROC RAID card** + 2× Samsung 960 M2 for uncompressed 8K video disk IO

■ Receiver update


- MSI X299 Gaming Motherboard + i9-7900X CPU + 64GB RAM
- Memory modules overclocked to 2900 MHz to achieve necessary video data throughput in UltraGrid pipeline
- Mellanox ConnectX 40/56 GbE NIC
- NVidia Quadro M6000/P4000 or better with **Sharp LV-70X500E** or tiled 4K displays setup
- NVidia GTX 1080 or better with native 8K **Dell UP3218K UltraSharp** display (or NVidia RTX to support upcoming Samsung 8K displays)

- End-to-end network availability (esp. on campus)
- One does not simply go to Mordor and connect to 40GE (or even worse 100GE)

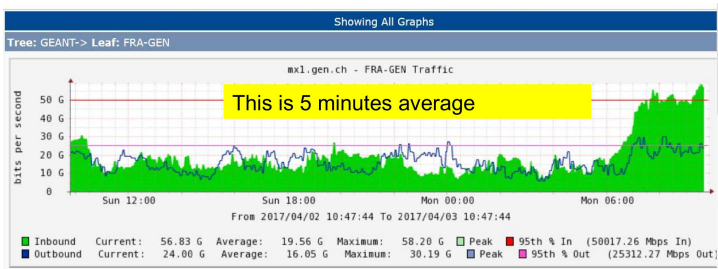


-
- The diagram illustrates the network topology of the Czech Republic, showing various cities and their interconnections. The legend indicates the following link types and router types:
- 100 Gbps DWDM (Orange line)
 - 100 Gbps (Blue line)
 - 40 Gbps (Red line)
 - nx 10 Gbps (Green line)
 - 10 Gbps (Light green line)
 - 1 Gbps (Dark green line)
 - DWDM circuits (Dashed line)
 - P router (Red icon)
 - PE router (Blue icon)
 - access L2/L3 switch (Blue icon with 'X')
 - PE router Alcatel-Lucent 7750 SR-12e (Purple icon)
- Key cities and routers shown include: Liberec (R139), Hradec Králové (R129, R130), Opava (Cat67), Ostrava (R144), Karviná (Cat54), Pardubice (R145), Olomouc (R132, R133), Zlín (R143), Brno I (R142), Brno II (R121), Jihlava (R146), Jindřichův Hradec (Cat48), České Budějovice (R141), Plzeň (R140), Písek (R149), Cheb (R117), Ústí n. L. (R147), Děčín (Cat53), Most (R117), TIS 1 AS 6762, TIS 2 AS 6762, GEANT3 backup, SANET, ACOnet, VIX, SID, AMS-IX, NIX4 CE COLO, Google CE COLO, Akamai cache, LHCONE, GEANT3, AMS-IX, NIX1 CRA, Google, and various CRS-X routers.


■ Network is heavily shared



Barcelona – Prague transmission through most utilised GEANT network link: Gen-Fra resulting in framedrops/audio artifacts.



Link capacity with 50% utilisation
during 1 s




25 fps transmission, i.e., 40 ms
to transmitt a single frame
(approx 3 Mbyte/frame)


Frames delayed in the packet mix
resulting in > 25 % framedrops.

- Based partially on recent discussions at I2 Technology Exchange
- Dedicated Circuits
- QoS
- SDN-enabled Networks with (guaranteed) Bandwidth on Demand Services

■ Going back to roots – d'oh



- Over-provisioning paradigm
- No one likes the Elephant streams on their network!
- Quality of Service
 - Very old concept
 - Implemented using mostly “ancient” protocols and not really supported in computer networks today
 - Except for ToS in IP (L3 only) and 3 Exp (as in Experimental) bits in MPLS (which are mostly defaulting to 0 anyways)
 - Ruled out by the over-provisioning paradigm in the past
 - No legitimate purpose is served by offering QoS services, it even breaks net neutrality!
 - But over-provisioning is not an answer for transmissions requiring bandwidth comparable to network speeds (on a shared infrastructure)



- And it does not look better on a larger scale
 - Géant AutoBahn
 - Bandwidth on demand service
 - Never went fully into production, still more or less a testbed
 - Allows to request a dynamic circuits with bandwidth up to 10Gbps – d'oh
 - Physically uses the same shared network infrastructure – d'oh
- So the future of high-speed and low-latency video transmissions is apparently only to be seen again...

Thank you for your attention!

`<milos.liska@cesnet.cz>`
`<ultragrid-dev@cesnet.cz>`

`http://www.ultragrid.cz/`

This work is supported by the CESNET E-Infrastructure (LM2015042) project.