High-Quality Media Production

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

<michal.krsek@cesnet.cz>, <petr.holub@cesnet.cz>

14th Annual Global LambdaGrid Workshop 2014
Queenstown, New Zealand, 2014–09–30

Why should we consider media production using high speed networks?

Movie Studios

- Solving 21st century problems with early 20th century paradigms:
  - splitting post production work among external suppliers to reduce costs
  - production is digital, media is still transferred on hard drives…
  - … because it is part of approved standards
  - "Internet is *the dangerous place* where pirates will steal your data" mentality

Broadcasting Companies

- Coax cables inside building plus satellite feeds
  - centralized production
  - most of content comes from “cans” – playout from hard drives
  - local networks – SDI-based
  - live production – satellite feeds
- There is ongoing but slow shift to online media processing
  - typically remote editing of recorded content

“Our” Goals in the Past

- Use of uncompressed media as much as possible
  - to avoid delay,
  - to preserve maximum quality.
- Focus on lightpaths
  - to create private networks,
  - to guarantee the bandwidth.
- Adaptation of existing applications to run over long-distance high-bandwidth networks

What CESNET Has Done

- Examined and challenged most of the steps of the media production and post-production:
  - dailies processing,
  - debayer,
  - color grading,
  - stereography,
  - remote VFX ,
  - media (film) restoration.

  workflow

- Over the long-distance high-bandwidth networks, often photonic.

Feedback from the User Community

- Holy Grail for the community is working with uncompressed media in its original format …
  - … but budgetary and engineering constraints prevent use of uncompressed media ☹
- Environment needs to be adjusted to budget
- We need support for the complex workflows

Time has come to shift our focus

Focus Shift

- Adaptation of media to available networks
  - target user community is not willing and often unable to get access to the high-bandwidth long-distance networks,
  - we need to maximize utilization of available infrastructure,
  - and enforce upgrades of only those bits that are severely limiting production quality.
- Transparent and content agnostic
- Self-organizing

Demo: Broadcasting from Sport Events

- Motivation scenario: live broadcasting from multiple sports events
- Goals:
  - utilize *available* network capacity to achieve optimum broadcasting quality,
  - network and processing have to be integrated,
  - no administrative rights over the network are needed,
  - no frames should be lost during switching,
  - support for blending images,
  - support for live & pre-recorded content.
- … and that all without deteriorating current performance of their workflows (in any aspect).

CESNET

Demo: Broadcasting from Sport Events

- Typical structure of the network
  - SDI infrastructure (possibly converted to optics or Cat5/6/7),
  - high-bandwidth on-site (local) network,
  - medium-capacity long-distance links.
- Today, a control center is needed at each event, plus there is another coordinating control center at the broadcaster.

Demo: Broadcasting from Sport Events

- Practical workflow for the future:
  - previews from all cameras should be delivered to the control center,
  - the control center directs what signal is delivered to the distribution center,
  - the distribution center delivers the signals to the cable/satellite distribution.
- Optimizations compared to state of the art:
  - no control center at each event location, only one control center remains,
  - dynamic adaptation of the image quality w.r.t. the current network capacity.

CESNET

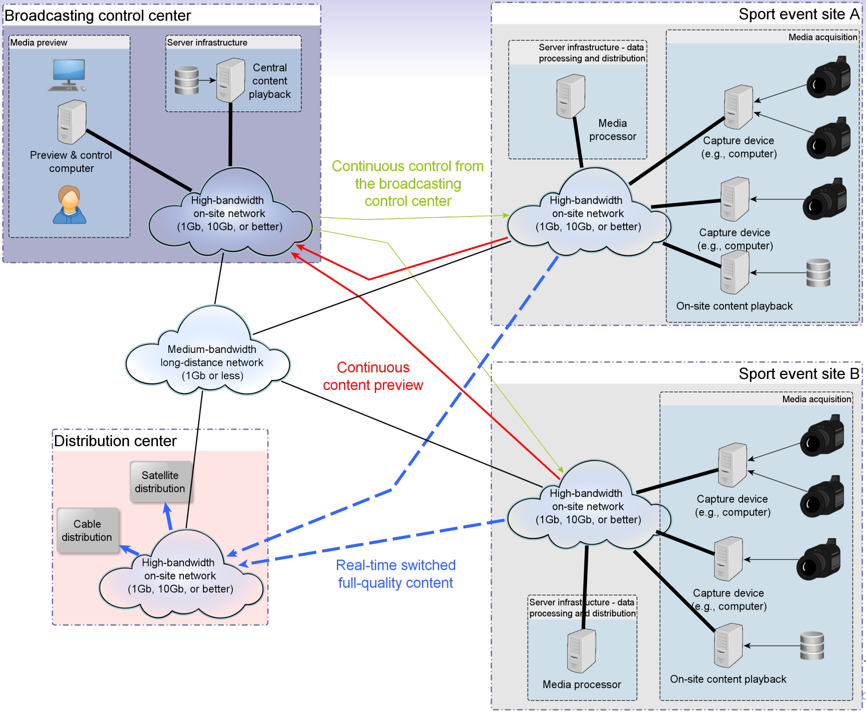
Broadcasting control center
Media preview
Preview & control computer
Server infrastructure
Central content playback
High-bandwidth on-site network (1Gb, 10Gb, or better)
Continuous control from the broadcasting control center
Sport event site A
Server infrastructure - data processing and distribution
Media processor
Media acquisition
Capture device (e.g., computer)
Capture device (e.g., computer)
On-site content playback
High-bandwidth on-site network (1Gb, 10Gb, or better)
Medium-bandwidth long-distance network (1Gb or less)
Continuous content preview
Sport event site B
Media acquisition
Capture device (e.g., computer)
Capture device (e.g., computer)
On-site content playback
High-bandwidth on-site network (1Gb, 10Gb, or better)
Server infrastructure - data processing and distribution
Media processor
Distribution center
Satellite distribution
Cable distribution
High-bandwidth on-site network (1Gb, 10Gb, or better)
Real-time switched full-quality content

Interesting (Research) Problems

- User has very limited control of the network
  - never resolved problem with consistent multi-domain multicast deployment
    - multicast is also “second class citizen” in common networks
  - programmability of networks in multi-domain environment is still an open problem
    - being researched also in GLIF community
  - application-level data distribution
    - creates an overlay network,
    - enables content processing as a part of the distribution – very powerful.

Interesting (Research) Problems

- Scheduling of data flows comparable to the capacity of network links in an overlay network
  - optimization of visual quality and latency
  - even overlay multicast routing with latency optimization is an $\mathcal{NP}$-hard problem[1]
  - maximum visual quality $\neq$ maximum bandwidth
    - visual quality improvements beyond certain level is not perceptible
    - uncompressed video rarely brings *visual* improvement

[1] Mokhtarian, Kianoosh, and H-A. Jacobsen. "Minimum-Delay Multicast Algorithms for Mesh Overlays." ACM/IEEE Transactions on Networking, April 2014. As well as other papers on overlay multicast optimizations.

Interesting (Research) Problems

- Problem of "no frame must be lost" and scene blending
  - should not rely on synchronous characteristics of the network,
  - transitional period where both streams are received at the same time
    $\Longrightarrow$ has to be supported by the scheduling model,
  - possible transient image quality degradation.

Demo: Media Switching from Events

Proof-of-concept

- Streaming from two sports events: Brno & Prague, CZ
  - content provided by DAITE, industrial partner of CESNET
  - multi-camera synchronous recordings from hockey matches
- Control center in Queenstown, NZ
- Distribution center in Queenstown, NZ

Demo: Infrastructure Scheme

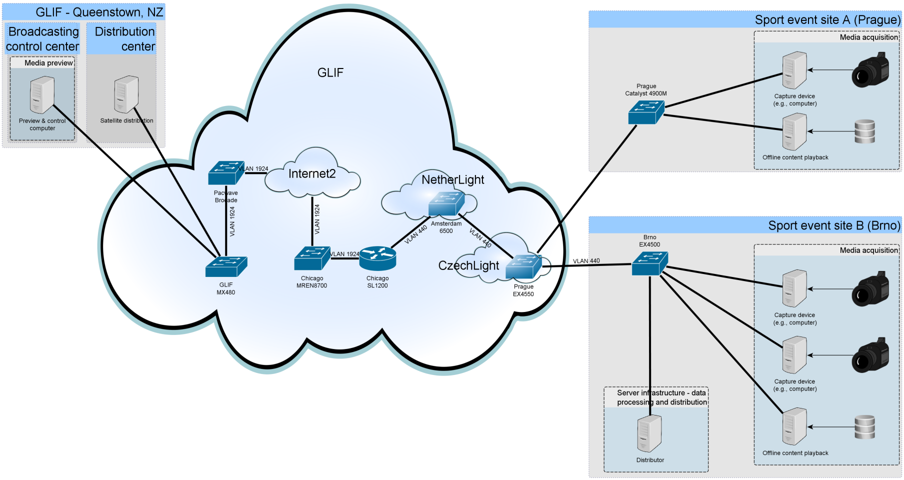

Demo: Media Switching from Events

Proof-of-concept

- CoUniverse
  - open-source framework:
    `https://www.sitola.cz/CoUniverse`
  - development version is used for the demo
  - framework for self-organization and orchestration of high-performance media applications
    - JXTA overlay network for control-plane communication
    - data plane runs directly on the native network
    - concept of data*producers, consumers, distributors*
    - data flow scheduling tries to satisfy wishes of consumers and optimizes given objective function
    - support for bandwidth-on-demand services (CoUniverse+NSIv2 demoed during SC'13 with AIST and SURFsara)

Demo: Media Switching from Events

Proof-of-concept

- CoUniverse
  - testbed for network stream schedulers
    - for streams with bitrates comparable to capacities of networks links
    - multi-criteria optimization: latency minimization, quality maximization, possibly other criteria (e.g., latency equalization)
    - from constraint-based and mixed integer programming to ant-colony optimization algorithms
    - can deal with transcoding and partially-known topologies[1]
    - can deal with bandwidth estimate uncertainties[2]
    - direct communication with applications – namely UltraGrid

[1] TROUBIL, Pavel - RUDOVÁ, Hana - HOLUB, Petr. Media Streams Planning with Transcoding. In 12th IEEE International Symposium on Network Computing and Applications (NCA 2013). USA : IEEE, 2013. ISBN 978-0-7685-5043-6, pp. 41-48. 22.8.2013, Cambridge, Massachusetts, USA.

[2] TROUBIL, Pavel - RUDOVÁ, Hana - HOLUB, Petr. Media Streams Planning with Uncertain Link Capacities. In IEEE 13th International Symposium on Network Computing and Applications NCA 2014. USA : IEEE, 2014. ISBN 978-1-4799-5393-6, pp. 197-204. 21.8.2014, Cambridge, MA, USA.

Demo: Media Switching from Events

Proof-of-concept

- UltraGrid[1]
  - open-source high-performance media streaming: `http://www.ultragrid.cz/`
  - optimized for quality maximization and latency minimization on commodity hardware
    - uncompressed, M-JPEG compressed, H.264, …– spanning 40 Mbps–10 Gbps range for 4K video
    - end-to-end latency as low as 75 ms for video, even lower for audio – depends on HW setup
  - GPU-accelerated image (de)compression and processing
    - GPU JPEG and DXT compressions[2]

[1] HOLUB, Petr - MATELA, Jiří - PULEC, Martin - ŠROM, Martin. UltraGrid: Low-Latency High-Quality Video Transmissions on Commodity Hardware. In Proceedings of the *20th ACM international conference on Multimedia*. New York, NY, USA : ACM, 2012. ISBN 978-1-4503-1089-5, p. 1457-1460. 2012, Nara, Japan.

[2] 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*, Amsterdam, The Netherlands, Elsevier Science, The Nederlands. ISSN 0167-739X, 2013, vol. 29, no. 8, 1991–2006-16 pp.

Demo: Media Switching from Events

Proof-of-concept

- UltraGrid
  - CPU- and GPU-based LDGM FEC
    - GPU is needed >800 Mbps
  - new developments for the demo:
    - on-the-fly configuration interfaces (for CoU): configuration of sending destinations and compression parameters, reporting available formats and their processing requirements
    - on-the-fly configuration changes without restart: destination, compression
    - sharp switching between sources and blend/dissolve effect
  - full transcoding reflector with (almost) identical feature set

UltraGrid

CESNET

Demo: Media Switching from Events

Proof-of-concept

- How does it work?
  - the operator sees videos from all the sources coming to the control center computer
  - the operator decides, which video should be sent to the distribution compiuter in full quality
    - this propagates to the scheduler, which finds a new plan with the transition period, if there was a previous plan
    - scheduler notifies the nodes which are required to implement the change
  - quality preference is biased toward the distribution

Conclusions

This demo is just first proof-of-concept prototype to demonstrate viability of distributed media routing and processing for live media production.

- Open questions (selection of ☺)
  - will the users react positively to this concept? (we have some positive feedback already)
  - how can be programmability of the network effectively utilized, given the multi-domain environments?
  - are there any other hidden show-stoppers for the adoption of the concept?

CESNET

Thank you for your attention!

Michal Krsek: <michal.krsek@cesnet.cz>
Petr Holub: <petr.holub@cesnet.cz>

Big props to other members of development & support team:
Martin Pulec, Jiří “Maara” Marek, Pavel Troubil, Lukáš Ručka

This work is supported by LM2010005 project.

MINISTERSTVO ŠKOLSTVÍ, MLÁDEŽE A TĚLOVÝCHOVY