УДК 621.395

High Speed WAN Data Transmission for Globalized Content Production Environments

E. Siemens, C. Herpel, R. Köhler, A. Aust

Audiovisual content creation often involves actors at many distributed locations. Tighter integration of these actors, among others by means of an efficient electronic content exchange, is urgently needed to improve the economics and to speed up especially high end, digital film production workflows. In this paper, we present the vision for a networked and managed audiovisual content production environment that may be distributed around the world. Necessary steps and major challenges towards the realization of this vision are presented here and solution approaches are sketched. A focus of this paper is on the specific challenge of fast wide-area data transport at multi-gigabit data rates. Currently a discrepancy exists between the available network bandwidth and the ability to "fill this pipe" in order to transport uncompressed imagery at 4k resolution. A new, but already deployed, solution for this issue is elaborated. The paper concludes with an outlook on the next steps towards our vision to make globalized content production more efficient.

Keywords: globalized content production, post production, fast wide area data transfer.

1. Globalizing Content Production

For a long time, the production and repurposing of audiovisual content for all its customers along the content exploitation chain has already been a distributed business with many actors. Film post production involves editors, visual effects houses and the actual post facility assembling the final deliverable, possibly all of them at different locations. Film derivatives along the exploitation chain are also typically produced at various locations. With the advent of "digital intermediate", all stages of the post production workflow are now executed in the digital domain. Only shooting and distribution still mostly take place on celluloid with a strong trend towards digital media as well.

Today globalization of cinematographic content production workflows takes place since specific equipment, artistic talent or low labor rates are only available at specific locations. Media *exchange* between workflow actors remains a challenge due to delays introduced by shipping physical media ranging from celluloid, video tape, data tape to portable hard disks. Fortunately, with increasing bandwidth of IT networks, it becomes more and more feasible to migrate to electronic content exchange.

Notably, for a long time electronic content exchange has been an option in the broadcast world, where dedicated "contribution links" have been used to deliver content in real-time from a supplying broadcaster to the receiving broadcaster(s) either for live transmission to the end user or for later repurposing. These links have always been very limited in data rate, starting with 34 Mbps for SDTV 1 many years ago and developing into a multitude of formats and bitrates up to 100 Mbps currently. For a live production of a show or soccer game the multitude of different feeds together with special signals, like those from slow-motion cameras, easily add up to an overall data rate budget beyond 1 Gbit/sec.

Today digital cinematographic production-quality content has resolutions between HDTV and 4k (4096x3116x24 fps). In this realm, compression is typically not accepted or has at least to be visually lossless. This leads to data rates between 1,5 Gbit/s and more than 10 Gbit/s for a *single* real-time contribution stream between post production facilities.

Both broadcasters and digital post production houses are currently using file-based transfer of content between their facilities; however, this is often done just via ftp at average data rates of several 10 Mbit/s and almost certainly without specific guarantees for timely delivery. Management of such inter-facility workflows is mostly manual and time consuming.

The availability of well-integrated and managed data transfer solutions supporting very high data rates enables streamlining of collaboration between distributed facilities. Figure 1 sketches networked facilities grouped according to their functions in the production workflow, while many other subdivisions are conceivable.

In this paper we expose the key challenges for efficient and reliable WAN data transfer (section 2) and propose a solution (section 3). Furthermore we sketch the elements of a complete solution for making truly globalized content production facilities possible (section 4), completed by a look at the next steps on this way.

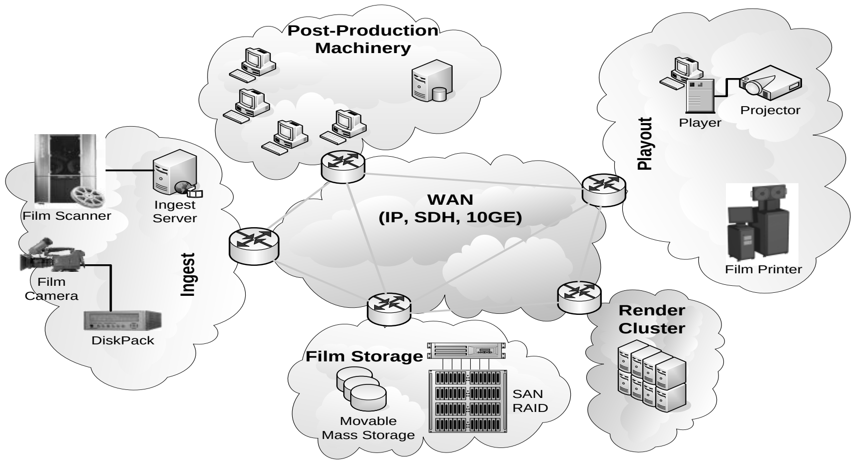

Figure 1. Overall concept of a globalized 4K production facility

2. Fast WAN Data Transmission Challenges

As described in the previous section, the data transmission between particular domains of data acquisition, processing and play-out is a key component in the globalizing world of content production. Each facility must be able to transmit uncompressed image data together with meta-data very fast – preferably in real-time or even faster than real-time. Depending on the facility design, such a data transmission system could be implemented as a stand-alone box - e.g. as one of many bricks in a service oriented multimedia enterprise - or as an integrated part of other image processing software or hardware. But unnecessary store-and-forward copy operations to dedicated WAN transfer stations, as they occur today, need to be avoided.

Regardless of the chosen integration approach, a number of major challenges have been identified on the way towards the real-time image data transmission at multi-gigabit rates.

1. The major precondition is a suitable protocol for reliable data transport. Such a transport protocol has to sustain multi-gigabit rates in spite of network delay, jitter and packet losses. From the user point of view, the protocol has to provide an error-free transmission of data between different sites.

Commodity reliable transport protocols on top of IP, such as TCP and its derivates, fail on this task due to their Bandwidth-Delay-Product (BDP) limitations 2. Considering transcontinental uncompressed 4k data transmission, we have to support BDPs in the order of 100 Mbytes. Dealing with packet losses, care has to be taken of the fact that packet losses on internet links usually occur in bursts. So, we need packet loss recovery mechanisms, which do not fail even in case of hundreds or thousands of consecutive packet losses.

2. Another challenge lies in the selection of hardware platforms that incorporate the functionality to pass image data at multi-gigabit rates from an acquisition device or a storage system through the network protocol stack to the interface hardware. Since a single CPU of a high-end PC can easily be fully loaded just with data transfers, multi-core computer systems are needed, where a performance gain can be achieved by distributing individual data processing tasks to different cores. However, in this case, usually each access to commonly used data must be synchronized, leading to significant performance penalties.

3. As outlined in Figure 1, a high-efficient data transmission system has to be integrated into a diversity of environments. Sometimes the transmission service can be provided by dedicated hardware, in other cases it must be an integral part of other tools within the post-production chain. Consequently, we have to design the overall transmission system in a modular fashion in order to be able to integrate it into a variety of hardware- and software platforms.

3. Reliable High-Speed WAN Transmission Solution

Addressing the key issues pointed out in the previous section, a fast data transmission system with the goal to provide a high-performance and flexible transmission software stack was successfully developed. The software stack had to be executable on a commodity hardware and operating system and also to be easily integrated into third party post-production solutions. Due to the wide availability of Linux as a server platform and even as a platform for a diversity of high-performance embedded devices and appliances, Linux was identified as the most suitable operating system. Another point in favor of Linux is the good driver support from a diversity of NIC manufacturers. Nevertheless during design phase it was foreseen to be able to port the developed solution to other operating systems like Windows with limited effort.

The particular steps of performance improvement related to the challenges mentioned above are depicted in more detail in the next subsections.

3.1. RWTP – Reliable WAN Transfer Protocol

The lack of existing transport protocols, which meet the requirements of reliable data transmission over long distances 3 led to a new transport protocol being developed - the Reliable WAN Transfer Protocol (RWTP). It is optimized for streaming uncompressed 4k image and audio data in a reliable fashion. At this point it is worth to mention two transport protocol proposals with a similar scope:

- While Reliable Blast Transfer Protocol (RBUDP) 4 is designed especially for fast transmission of huge files, rather than a sequence of smaller files, it provides data rates in the range of our interest. The main idea of RBDUP is to transmit all data from sender with full link speed. At the end of the transmission the data receiver reports to the sender the list of received datagrams. On reception of this receiver report, all lost datagrams have to be repeated. The main constraint of RBUDP is that all retransmissions are started only when all the data packets of the actual iteration ("blast") are received or timeouted. This implies the existence of the notion of "all data". However, streaming data don't have this notion, since usually the amount of streaming data is not known during the streaming. Moreover, packet retransmissions at ends of particular "blasts" would introduce unacceptable delays into the media streaming application.

- The other approach, called UDP-based Data Transfer Protocol (UDT) 5 provides a reliable stream service with data rates up to 1 Gbit/s. Although this rata rate is much higher than the data rates of various TCP flavors, this is insufficient for an uncompressed real-time transmission of 4k data.

On this background, the retransmission mechanisms were improved by introducing proactive positive and negative data acknowledgments. Furthermore, a new congestion control algorithm was tailored for multi-gigabit streaming of audiovisual data. While this congestion control algorithm does behave more aggressive than TCP and even UDT, with these extensions, the transmission system is able to recover from packet losses without end-to-end data stream disruptions – even when packets are lost in burst of hundreds of datagrams.

Obviously, the lack of fairness of RWTP leads to a need for a very careful network planning and network design, when deploying RWTP. One effective QoS measure would be flow isolation in accordance to the IntServ or the DiffServ frameworks. So, we strongly believe that a co-existence of aggressive protocols like RWTP and fair-behaving data flows like TCP or UDT can be provided by means of network QoS like IntServ of DiffServ or by means of L2-VPNs.

In order to run protocol throughput tests, a test client-server application was implemented which transfers several gigabytes of data from sender to receiver without file-system or other IO interactions and without any image processing. Network impairment on an Anue network emulator 6 device was configured to insert delays up to 100 ms and equally distributed packet loss with a mean loss rate of 0.3 %. This scenario roughly emulates a transcontinental link with sporadic congestion occurrence.

With the described setup, we revealed a significant performance boost compared to other rate-based transport protocols 7. Depending on used 10GE NICs, sustained data rates above 6 Gbit/s has been reached.

3.2. Efficient Protocol Buffer

As mentioned in section 0, fast data exchange between application and network stack and interface board is the second big challenge when implementing fast data transmission. For this purpose, sender and receiver buffers have to be implemented that are used by the application and the transport protocol stack. Concurrent access to the sender and receiver buffers from different processing threads has to be performed very frequently – at least once for each packet sent or received. For example for a rate of 6 Gbit/s and using jumbo frames of approximately 9 kBytes this would lead to a packet rate of roughly 83.000 packets per second. Hence, approximately one thread synchronization would be required every 12 microseconds. Even using recent hardware, such a synchronization rate would overload the system!

To improve this situation, an optimized buffer system comprising a number data buffers has been implemented. A simplified structure diagram of this data buffer system is shown in Figure 2. Besides sender buffer and receiver buffer, both organized as ring buffers, one more buffer with control information about sent and retransmitted data packets has been introduced. Each of the buffers is accessed via a specific set of pointers, acting as delimiters of different buffer areas. With a very specific access pattern to each of the three buffers it is possible to assure access to data of send buffer and receive buffer by multiple threads without having to use a lot of CPU intensive locking/synchronization mechanisms.

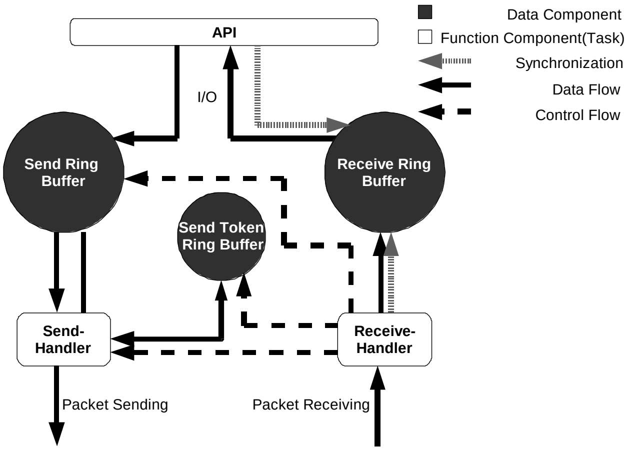

Figure 2. Simplified structure of a buffer system for fast data transmission

Only when sending data packets smaller than the path MTU, synchronization is still required. This occurs essentially only if the sender buffer contains less data than MTU size of the corresponding interface – a very rare case in the steady state of image data transmission.

Comparative performance tests of the outlined buffer system against other buffer systems, synchronized with mutexes revealed a stable data rate boost of approx. 3 Gbit/s.

3.3. The Modular Software Concept

The third challenge, the need of integration of a fast data path into a variety of applications and systems is addressed by modularization.

Before the image data at the sender side can be passed to the RWTP protocol buffer, this data must pass different other processing nodes between the data origin and the RWTP stack. The origin of data can be a SAN or NAS based file system or an image acquisition device, such as a digital camera or a film scanner. At the receiver side, the data travels in the opposite direction as depicted in Figure 3.

Moreover, film data may be acquired, stored or played out in a different format than the one used for data transmission, due to constraints of the equipment used. As an extreme example, 4k image data can be stored being decomposed into four 2k-quadrants, transported as 4k DPX images and played out using 4 HD-SDI links. In this case, an on-the-fly assembling/reassembling of image data is necessary. Many other image format conversions, before or after the transmission, are conceivable. From these considerations, the constraint for the overall transmission application design can be deduced that the RWTP stack must operate with best performance in parallel to other CPU consuming tasks on a flexible number of processor cores.

These considerations lead to an overall application design as described in Figure 3. The main concept of data flow is based on the pipes and filters design pattern 8. Additionally data copying is avoided wherever possible.

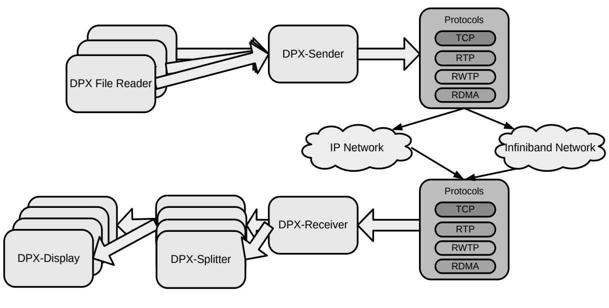

Figure 3. Software concept of the 4k data transmission system

In the application design, image data is fetched from the data source by a DPX File Reader. A number of different file reading strategies have been implemented, which can gather image data from SAN or NAS storage systems as well as from HD-SDI real-time inputs. The reader blocks can be executed in thread pools and perform data gathering synchronously or asynchronously. The DPX-Sender performs necessary format transformations and passes all the image data to the transport protocol stack. Depending on the available network topologies, the transport protocol engine is able to instantiate one of the legacy protocols (UDP, TCP) or, in case of a WAN transmission, the RWTP.

Since all building blocks provide a unique interface to their neighbors within the processing chain, the concrete instantiation with the appropriate synchronous or asynchronous file reader, DPX processor or transport protocol can be done at initialization time of the application. The extensive use of the strategy design pattern makes the system highly configurable.

Putting the finishing touch on the application design, the distribution of individual tasks to CPU cores has been optimized, so queue underruns between processing blocks could be minimized and consequently, the data throughput maximized.

3.4. Performance Tests and Their Results

Performance tests have been performed within the Thomson Corporate Research 10G-testbed located in Hanover, Germany. Besides a diversity of 10Gbit routers and switches, the testbed comprises a productive 10 Gigabit-Ethernet WAN-loop of the German carrier LambdaNet. Additional network impairments like packet delays, packet delay jitter and packet losses have been inserted using a network emulator Anue XGEM 10.

With the described example application, end-to-end data rates of approximately 4.5 Gbit/s in file-transfer and streaming modes could be reached in our test bed, chaining the real WAN connection with an emulated further 8000 km connection and insertion of delay jitter and equally distributed packet loss with a mean of 0.3 %. During the transmission, an additional on-the-fly image splitting from 4k to four 2k outputs has been performed. Additional tests have been performed by insertion of packet loss bursts with burst length of several hundred packets. Even with bursty losses, the target data rate has been sustained.

4. Assembling the Global Facility

The efficient usage of network resources on a WAN is an important factor for speeding up distributed workflows by eliminating shipping delays. Remote talent can now be integrated seamlessly even in live or near-live interactive production scenarios. This enables further de-localization towards cost optimization, be it by reducing travel time of the highly wanted creative people or by shifting work to lower labor cost areas. However, several other aspects have to be considered before this proposition can fully prove its economic viability.

4.1. Global Facility Management

The management of a globally distributed content production facility and its resources is a major challenge. However, if this challenge can be overcome, expensive equipment such as film scanners, film printers, storage systems or compute clusters can be shared efficiently over multiple facilities thus enabling smooth workflows across teams at various locations.

In order to describe and enact business-oriented workflows, physical devices are abstracted to the services they are offering, requiring service discovery & monitoring. This can be done at a global level using grid middleware frameworks, such as Globus Toolkit 9 with the Web Services Resource Framework (WSRF) 10 as its basis.

In order to execute business workflows it is necessary to have workflow description, resource allocation, job scheduling & monitoring as well as billing functionalities in place. While again no comprehensive solution exists as of today, there are both interesting approaches such as the generic GRIA 11 and business-specific offerings such as Thomson's ContentShare2 that address part of these functionalities.. A time-slot-oriented workflow management is needed additionally, in order to have deterministic workflow behavior when resources are only made available at precisely defined times for cost or efficiency reasons.

4.2. Heterogeneous Networks

Existing network islands are currently being connected and legacy sneaker-nets (carrying hard disks / tape / film) start to disappear. The result, however, is not yet a homogeneous networking landscape, but a heterogeneous one with, most likely, three different network interconnects, all of them being standard IT-networks of the office / datacenter area rather than proprietary ones such as SDI 12.

Ethernet is the oldest technology but with a strong and backwards compatible development path together with the largest number of installed ports and broadest support of the industry. It is the most often used carrier for the TCP/IP protocol suite. Supported physics are copper and fiber up to the highest data rates 13. Strong QoS mechanisms are missing at hardware level today. However, some of such mechanisms exist as part of the TCP/IP protocol suite. Although not initially designed for high-bandwidth and hardware QoS, Ethernet survives because of its flexibility and currently ongoing standardization efforts that will add hardware QoS. With Ethernet networks of any size from LAN to WAN are possible. Today 10 gigabit speed is in the market and 40/100 Gbps are on the roadmap for the next 1-2 years 14. Considering that 10 GbE today is compatible with SO-NET/SDH telecom infrastructure and might eventually replace it, this “Datacenter Ethernet” is the strongest contender for the future single network link 15.

Fibre Channel – designed as a reliable high speed network technology with low overhead and QoS features dominates storage networks 16. Link speed is 4 Gbps today with 8 Gbps just emerging. FC networks are bound to local installations, with inter-network connectivity solutions existing via Ethernet. FC is not being discussed as a network for non-storage traffic at all. On the other hand, experimental or even product-level storage networks via Infiniband and 10GbE do exist but are not yet fully competitive.

Infiniband, the new comer of the 90’s captures the high-performance/cluster computing market because of its lowest latency and low costs per port. Similar to Fibre Channel it also contains QoS and reliability at hardware level and can only be used in LAN, again with wide area bridges via Ethernet existing. Although copper and fiber physics are specified, Infiniband is only available on copper networks; fiber connections are only available via port adapters and multi-ribbon fiber cables without broad support by the industry. With its 20 Gbps it is the fastest currently available link 17.

From the above analysis it may be concluded that total cost of ownership considerations will lead to Ethernet, with both copper and fiber physics, being the only network in the content production facility (and elsewhere). Activities of the IT industry like Datacenter Ethernet push this convergence.

However, until this day the heterogeneity has to be managed, among others with solution permitting a seamless QoS across such a network.

QoS will become an important mechanism because in contrast to standard IT-networks the bandwidth per connection is in the range of the maximum available bandwidth. This means, over-provisioning does not work. In addition, content transmission is time critical, at least for real-time streams and may not degraded or stopped in any case. So topology aware QoS Management is needed with full control over the network elements to ensure that specific connections are operated with 100% service guarantee.

4.3. Storage

With all celluloid production gradually being turned into digital production and most of the celluloid ending up being scanned onto storage arrays, the storage quantities for digital content production are huge and increasing steadily. The raw material for a feature film produced in 2K resolution can easily reach 10-100 TB, with the final deliverable amounting to 1 TB. Even taking into account a 1:10 visually lossless compression, these are huge capacities, knowing that facilities will often have multiple productions under processing simultaneously.

Hierarchical storage systems are needed to cope with this situation, with only the most performing “online” tier being able to support the online workflows including the fast data transfer solution exposed above. The online SAN solutions are usually rather expensive and use over-provisioning to ensure data rates, without providing true determinism and, hence, predictable QoS. Software QoS solutions like the distributed I/O scheduler demonstrated at NAB 2007 18 can be used to lower the storage price point or to enable second tier (“near-line”) storage to support our fast data transfer solution.

4.4. Security

Security is also a key requirement for globalized facilities. Because content and services are spread over multiple sites which are interconnected either with private or public IP networks, sniffing attacks or “man in the middle” attacks could compromise the secrets of the film and broadcast industry.

So a security model of a globalized workflow must be present at all levels inside and outside of each local facility. Here, the state-of-the-art IT-technology provides proven security mechanisms to authenticate network devices and users and to manage access to storage systems. More challenging are the encryption of time critical access to distributed services via Web Services and, most importantly, the encryption of high-bandwidth data for network traversal and for writing it on encrypted storage systems. The tremendous data rates rule out any standard software solution today.

Another hot topic is the global management of authentication and authorization. Hereby, the authentication and authorization information has to be transparently spread over the several involved locations in a secured way.

4.5. Migration to IP for Live Production

The heterogeneous network of a facility for non-live production is today complemented by a completely proprietary network infrastructure for broadcast facilities. All live feeds use SDI or even older standards. Even in this demanding area a replacement by standard Ethernet infrastructure is in reach, as detailed in a companion paper to this conference 19.This migration will eliminate the last remaining media breaks. Cameras with 10GbE interface will make it possible to seamlessly use them both in live or non-live productions, either directly broadcasting the signal, if needed via global links using our WAN transfer solution, or captured directly onto a disk system for further processing.

4.6. Metadata

The multitude of formats and networks today effectively very much limits the transparency of production-oriented metadata. Since multiple metadata standards with different scope exist, it will be a tremendous effort to merge them or transcode between them. However, the migration to a single transmission network is still likely to foster a wider availability of metadata.

5. The Next Steps

Cost and time pressure will continue to be a driving factor towards more efficient global collaboration in digital content production. Key solution bricks do exist, such as the WAN transfer solution exposed here. In future this solution will integrate security and will become available as a service. On the side of network providers a cost-efficient time-shared access to high speed networks is required and already being studied 20.

Work, relevant for the Global Facility Management, is taking place, among others, in the EU-funded IRMOS project 21, addressing QoS and real-time aspects of a distributed service-oriented infrastructure. Obviously, it is conceivable to incorporate the workflow management tools into applications like ContentShare2 22.

Once the assumed migration of broadcast studios to IP infrastructure will have taken place, a lot of interesting work remains to maximize the synergies between live and non-live audiovisual content production across the, currently very firm, boundaries between broadcast and (digital) film production.

References

1. ITU-T. Transmission of component-coded digital television signals for contribution-quality applications at the third hierarchical level of ITU-T Recommendation G.702, ITU-T recommendation J.81, 1993
2. Dina Katabi, Mark Handley, Charlie Rohrs. Congestion control for high bandwidth-delay product networks. In Proc. Of: ACM SIGCOMM Computer Communication Review, Oct. 2002
3. C. Anglano, M. Canonico. Performance analysis of high-performance file-transfer systems for Grid applications. Concurrency and Computation: Practice & Experience, Vol 18 pp 807-816, Jul. 2006
4. E. He, J. Leigh, O. Yu and T. DeFanti, Reliable Blast UDP: Predictable High Performance Bulk Data Transfer. IEEE Cluster Computing, 2002: p. 317
5. Y. Gu „UDT: A High Performance Data Transport Protocol“ PhD Thesis, University of Illinois at Chicago, 2005
6. Anue Systems. Network emulator Anue XGEM, Product information page. URL:http://www.anuesystems.com/GEM_XGEM_network_simulator.htm. 2008
7. Xinran (Ryan) Wu and Andrew A. Chien. Evaluation of rate-based transport protocols for lambdagrids. In HPDC '04: 13th IEEE International Symposium on High Performance Distributed Computing. 2004
8. F. Buschmann et. al. Pattern-Oriented Software Architecture, Vol 1. John Wiley & Sons,2000
9. Globus Toolkit Home Page. URL: http:// www.globus.org. 2008
10. The WS-Resource Framework, Home Page. URL:http://www.globus.org/wsrf/
11. GRIA Home Page. URL: http://www.gria.org. 2008
12. Society of Motion Picture and Television Engineers. “ Serial Digital Interface (SDI)”, SMPTE-259M, 1997
13. IEEE. IEEE 802 part 3: Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access, method and physical layer specifications, IEEE 802.3-2005, 2005
14. IEEE P802.3ba 40Gb/s and 100Gb/s Ethernet Task Force, Home Page URL: http://www.ieee802.org/3/ba/. 2008
15. M. Wadekar. “Enhanced Ethernet for Data Center: Reliable, Channelized and Robust. 15th IEEE Workshop on Local & Metropolitan Area Networks 2007, p.65
16. Fibre Channel Industry Organisation. Home Page URL:http://www.fibrechannel.org/. 2008
17. Infiniband Trade Association. Home Page URL: http://www.infinibandta.org/. 2008
18. C. Herpel. “Di/oS - Improve Cost and Time Efficiencyof SAN-Based Workflows” Thomson, Corporate Newsletter Nr. 1, 2007 URL:http://www.thomson.net/GlobalEnglish/research/research/Documents/Newsletter/Corporate_Research_summer2007.pdf

19. F. Le Bolzer et. al. Prodim@ges - A new Video Production Environment based on IP wireless and optical links. Accepted to 2008 NEM Summit.
20. MUPBED Project Home Page. URL: http://www.ist-mupbed.org/. 2008
21. IRMOS Project Home Page. URL: http://www.irmosproject.eu. 2008
22. ContenShare2 Product Home Page. URL: http://www.thomsongrassvalley.com/products/security/cs2/. 2008

Received December 7, 2010

Edward Siemens, Ph.D.
Thomson Corporate Research, Karl-Wiechert-Allee 74, 30625, Hanover, Germany;
phone: +49 511 418 2168, e-mail: eduard.siemens@thomson.net

Carsten Herpel
Thomson Corporate Research; e-mail: carsten.herpel@thomson.net

Ralf Köhler
Thomson Corporate Research; e-mail: ralf.köhler@thomson.net

Andreas Aust
Thomson Corporate Research; e-mail: andreas.aust@thomson.net

Высокоскоростная передача данных в глобальных сетях для систем производства цифровой аудиовизуальной информации

Э. Сименс, К. Херпель, Р. Кёлер, А. Ауст

В создание аудиовизуальных материалов часто вовлечены лица, находящиеся в пространственно распределенных местах. Для улучшения экономических показателей и ускорения производственных процессов создания фильмов на базе передовых цифровых технологий крайне необходима более тесная интеграция всех участников посредством эффективного электронного обмена информацией. В статье представлен взгляд авторов на сетевую и управляемую среду для производства цифровой аудиовизуальной информации, которая может быть распределена по всему миру. Представлены необходимые шаги и основные проблемы в реализации этой среды, а также обозначены подходы к их решению. Статья особенно сфокусирована на проблеме сверхбыстрой передачи данных в глобальных сетях, что необходимо для передачи несжатых изображений формата 4k HDTV. Исследовано новое, но уже реализованное решение для этой задачи. Приводится обзор следующих шагов, направленных на дальнейшее повышение эффективности распределённой системы производства цифровой аудиовизуальной информации.

Ключевые слова: распределённое производство мультимедийной информации, пост-обработка мультимедийной информации, сверхбыстрая передача данных в глобальных сетях.