Evolution of Optical Networking Toward Rich Digital Media Services

Optical networking, burst switching to handle diverse and dynamic traffic, and media applications such as scientific visualization and consumer-driven media flows are discussed in this paper.

By Admela Jukan and Joe Mambretti

ABSTRACT | This paper discusses the evolution of optical networks from foundational data transport utilities and infrastructure to true enablers of next-generation services and facilities, specifically diverse rich digital media services, which are being produced, processed, and consumed by an increasingly sophisticated variety of end users. This evolution has been enabled by several major recent advances in optical networking, including the deployment of high-speed optical connections from core networks to enterprise and residential users; the dramatic increase in optical signal speed and reach; and the improvement in control and management planes to support increasingly diverse types of traffic and services. We describe two distinctive areas of broad impact of optical network technologies related to rich media services: high-resolution scientific visualization and high-quality, real-time consumer-driven media production and distribution. Finally, we discuss key trends that may develop next as optical networks further evolve shaped by both rich digital media service requirement and technology innovation.

KEYWORDS | Computer network management; digital multimedia broadcasting; optical fiber communication; optical interconnections; scientific computing; streaming media

A. Jukan is with the Institute of Computer and Network Engineering, Technische Universität Carolo-Wilhelmina zu Braunschweig, 38106 Braunschweig, Germany (e-mail: jukan@ida.ing.tu-bs.de).
J. Mambretti is with the International Center for Advanced Internet Research, Northwestern University, Evanston, IL 60208 USA (e-mail: j-mambretti@northwestern.edu).

I. INTRODUCTION

For many years, bandwidth-intensive applications, including rich digital media services, have evolved in concert with their supporting delivery platforms and underlying communication infrastructures. Driver applications and support technologies have challenged and enabled each other in an unprecedented progression. Along the way, there has been a distinct synergy between the requirements of rich digital media and next-generation optical networks, architecture, services, and technology. Today, multimedia content has not only become the main component of the overall traffic, which is expected to continue doubling every 12–18 months [1], [2], but is also requiring the network to be highly adaptive to meet the widely fluctuating bandwidth demands of such applications.

Four major characteristics have driven the advancement of optical networks, which have been increasingly deployed by communication service providers worldwide: 1) *large and flexible capacity*: a range of choices between 40 and 100 Gb/s channels and a very large number of smaller sized channels, enabled by the advances in optical transmission technologies, including variable channel spacing and new modulation schemes (optical phase in addition to its amplitude and code); 2) *reach*: opportunity to transmit channels all-optically for several thousands kilometers, with fewer amplifiers and regenerators; 3) *framing options*: support of multiple framing options ranging from packet streaming to specialized bursting to end-to-end circuits; and 4) *programmability and reconfigurability*: infrastructure that is highly software and hardware programmable, with a broad portfolio of network control and management options, enabling dynamic channel setup and infrastructure virtualization.

While optical networks continued to progress, information sent over the Internet has shifted from traditional data services (e.g., e-mail, Web pages) to multimedia content, based on complex media content creation, transport, and consumption paradigms. For instance, the demand for high-definition (HD) multimedia content and 3-D media delivery technologies, requiring gigabit per second processing and transmission speeds, is evolving from those applications used by a few specialized entertainment industries toward a consumer-level "HD service" accessible to individuals in multiple public and private sectors. With major increases in end-device compute power, which are enabling the production and consumption of media on virtually any smart device, everybody has become empowered to be a producer and/or consumer of rich media. The trend towards *democratization* of media production and consumption has emerged especially through the combination of powerful end devices and cloud computing services.

To analyze the achievements to date related to media and optical networks and suggest what may appear next, in this paper, we will look at the evolution of optical networking toward rich digital media services from both application and technology perspectives. To this end, we will first present general requirement considerations for digital media services and implications for network architecture and technology. Here, we will set the stage for a more detailed consideration of these and other requirements for digital media services, and implications for optical-centric network architecture and technology, which we will present in the sections that follow. Then, we will shift our focus to two specific areas of interest: high-resolution scientific visualization and high-quality, real-time consumer-driven media production and distribution. These two focus areas are chosen not only because of the level of broad impact they have had in this evolution, but also because they have special requirements that are not well matched by today's networks based on the traditional operation system interface (OSI)

model [3]. We will conclude the paper with a brief review of open issues, challenges, and our vision for future progress.

II. GENERAL REQUIREMENT CONSIDERATIONS FOR DIGITAL MEDIA SERVICES AND IMPLICATIONS FOR NETWORK ARCHITECTURE AND TECHNOLOGY

In digital media applications, *high quality* refers to multiple measurable parameters, including frame rate, pixel resolution, a wide range of colors, and other quantifiable characteristics that provide for a user experience far superior to what is possible with common broadcast technologies today [4]. Table 1 shows some of these parameters along with resulting data rates required. In addition to three high-quality digital media services, i.e., 3-D HDTV, digital cinema and ultrahigh-definition video (UHD), for comparison, Table 1 also shows the standard values for HDTV available in today's broadcast systems. As indicated, the frame rate (in frames per second) and color (in bits per pixel) increase with the quality of the digital media service, with the overall better resulting portrayal of motion, as well as representation of red, green, and blue colors with a larger number of bits [5], [6]. The digital cinema and HD remote visualization services are denoted as 2 k, 4 k, and 3-D 4 k; for instance, 4 k corresponds to approximately 4000 pixels horizontal and 2000 vertical, whereas 3-D 4 k has twice the number of pixels, etc.; an even higher resolution format called 8 k is being developed today. All data shown have been calculated for a single-view and multiview video (MVV) transmission, where the latter incorporates multiple videochannels into a single data stream. Finally, we show the data rates for compressed and uncompressed videos, where the compression ratio was assumed 20 : 1 with a JPEG2000 codec.

As shown in Table 1, the resulting data rates are orders of magnitude higher than those used in HD video

Table 1 Multiple Measurable Parameters and Bandwidth Requirements of HDTV and High-Quality, Rich Digital Media. [1][84], [2][85], [3][86], [4][87], [5][88]

Digital Media Service		Frame Rate (Frames/s)	Resolution (HxV in pixels)	Color (bits/pixel)	Data Rate in Gb/s Single View Video		Data Rate in Gb/s Multi-View Video (16 channels)	
					Uncompressed data rates (Gb/s)	Compressed data rates in Gb/s (20:1)	Uncompressed data rates (Gb/s)	Compressed data rates in Gb/s (20:1)
HDTV[1]		25	1920x1080	30	1.56	0.07778	24.89	1.24
		50			3.11	0.15552	49.77	2.49
3D HDTV[2] (Stereo HDTV)		50	2x1920x1080	30	6.22	0.31104	99.52	4.98
Digital Cinema/ HD remote visualization[3]	2K	24	2048x1080	36	1.91	0.09555	30.56	1.53
		48			3.82	0.1911	61.12	3.06
	4K	48	4096x2160	36	15.2	0.76441	244.61	12.23
	3D 4K	48	2x4096x2160	36	30.58	1.53	489.22	24.46
Ultra High Definition Video [4,5] (8K and beyond)		60	7680x4320	36	71.66	3.58	1.15Tb/s	57.33
		120			143.33	7.17	2.29Tb/s	114.66

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

Jukan and Mambretti: Evolution of Optical Networking Toward Rich Digital Media Services

transport, especially for MVV applications. To this end, various compression techniques have been proposed. Of note is the H.264/MPEG-4 AVC standards family, which has been shown to provide an excellent performance in compressing 2-D/3-D images and MVV video steams [7]; the latter is also known as multiview video coding (MVC), which can significantly reduce the bandwidth requirement for HD MVV streaming. Despite the effectiveness of the compression techniques, video compression may lead to several issues related to either the audience satisfaction or network transmission. First, compression techniques remove pixels from an HD image in order to reduce the image size, which reduces the resolution, and may decrease the user perception of quality of experience (QoE) as a consequence. Second, from the network perspective, as the compression techniques may eliminate the redundancy of the video streams, the inherent resilience to losses also decreases from a transmission viewpoint [8]. On the other hand, significant compression gain can be obtained in MVV by eliminating the overlapping between different views based on inter-view prediction. However, this requires an inter-view dependency, where even a small error in one view stream may cause larger video distortions after decompression.

Currently, single-view HDTV services are being delivered across multiple network technologies and architectures, including the traditional Internet Protocol (IP) packet networks, connection-oriented networks (SONET), and content distribution networks (CDNs). However, none of them will be able to efficiently meet the bandwidth and other service level requirements of high-quality digital media services.

Traditional IP networks were designed to support the transit of discrete small packets of information in an end-to-end, host-based configuration. While relying on several basic transport protocols (e.g., UDP, TCP), IP network are best effort and "nonguaranteed" [9]. Many potential solutions have been suggested to address this issue, ranging from specialized alternative transport protocols (e.g., Fast TCP, Grid TCP) [10], [11], to traffic differentiation methods (e.g., DiffServ) to, most promising, connection-oriented approaches (e.g., MPLS). However, in connection-oriented networks, there is a challenge in designing services for digital media distribution, since classic methods of load balancing do not work well for rich media services. This is because one large flow in the network, such as 6–7 Gb/s for uncompressed UHD signals [12], cannot be load balanced based on the traditional principle of shared link capacities. In addition, traditional network circuit-switched layers cannot provide the level of dynamic capacity required by digital media services either. Primarily designed for analog telephone applications, traditional SONET and WDM circuits are inadequate in supporting bursty data traffic without overprovisioning, thus making the use of bandwidth for digital media services inefficient with regard to resource utilization [13].

Finally, in order to efficiently deliver static multimedia content, network and content providers currently use widespread content-delivery networks (CDNs), which have the inherent ability to cache multimedia content near the network edge, and then deliver this content on demand. CDNs represent a potentially important technique for delivering SHD and UHD multimedia [14]. Presently, these services rely on a set of geographically distributed proxies/gateways connected over the Internet. Streaming media are cached in the dedicated proxy servers, which are statically deployed beforehand. These proxies are currently connected with an overlay IP network. The closest proxy server is then used for streaming video, instead of the original source. For multicast services, these servers presently use Internet multicasting to distribute media to their immediate group of users. This type of multicasting is highly scalable because an arbitrarily large number of users can efficiently share a single channel. However, as it is implemented by IP, it cannot efficiently or easily support the needs of high-capacity users with stringent QoS requirements [15]. Therefore, there is an emerging need for a more scalable solution that does not only rely on Internet multicast.

Optical networks, while not as flexible as IP especially for low-bandwidth content, are more cost efficient at delivering high-bandwidth content with QoS guarantees. Due to their capability to support a variety of available transmission formats, switching technologies and spectrum management options (e.g., OTDM, WDM, OPS), optical networks are becoming common practice for high-definition video conferencing, high-performance computing as well as high-definition entertainment content distribution, scientific visualization, live broadcast, etc. (Fig. 1) [16]–[18]. Optical networks can satisfy several other challenging design and operational requirements such as the dynamic allocation of network resources, multicasting, effective integration of multiple layers (physical, Ethernet, Internet), and new mechanisms for network management and control. The widely fluctuating demands of rich digital media applications can be enabled by the deployment of dynamic Ethernet services over optical networks (typically from 1 Gb/s and up to a current maximum of 100 Gb/s), or large lambda services (from a minimum of 10 Gb/s and beyond 1 Tb/s). Much progress has been made over the last few years in transitioning Ethernet from a local area network (LAN) technology to an enabler of high-quality services in metro and core networks, and further progress in this area is expected with the completion of the IEEE 100 GE standard.

III. RECENT ADVANCES IN OPTICAL NETWORKING

Three major recent advances in optical networks have marked a new frontier for rich digital media services. First, the demand for high-speed optical network connections is
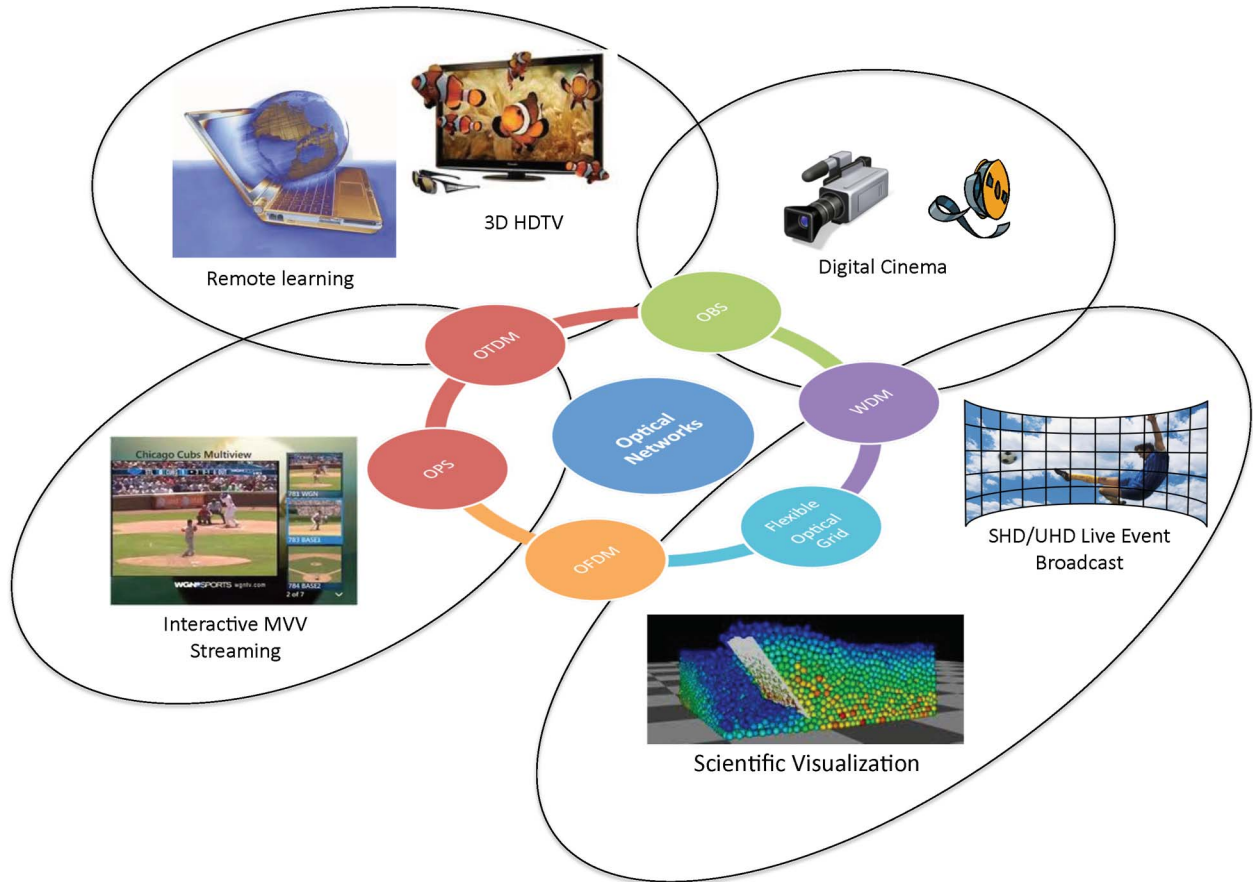
This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

Jukan and Mambretti: Evolution of Optical Networking Toward Rich Digital Media Services

Fig. 1. *Example digital media applications supported by a variety of optical networks.*

extending from core networks to enterprise and residential users, transforming the traditional access network segment and creating two new distinctive market and technology trends, i.e., cellular backhaul and private access. Second, optical technologies evolved both in speed and reach, with 100 s of gigabit per second transmission speed per wavelength, and thousands of kilometers of optical signal transmission without regeneration, with flexible spectrum allocation and fiber sharing in the time and frequency domains. Finally, traditional control and management planes for optical networks are also evolving to meet the requirements of increasingly diverse types of traffic, services, and communities while using shared—not separate—infrastructure, traditional control, and management planes for optical networks are also evolving to meet the needs of new requirements. We will address these three trends in more detail in the following sections.

A. Optical System Proliferation in Various Network Segments

Fig. 2 illustrates the influx of optical networking in the traditional segments of telecommunication networks, i.e., access, metro, and core.

The most prominent fiber-to-the-X access network technology (X = home, building, etc.) is the passive optical network (PON), which encompasses a number of its variants, including long-reach PONs and a combination of various PON architectures referred to as hybrid PONs [19], [20]. The efficient sharing of fiber capacity by a large number of residential users is achieved by the deployment of unpowered optical splitters (hence the name "passive"). The unpowered operation is critical to the evolution of access network segments where installing fiber optic infrastructure is a dominant cost factor. While 10-G PONs have been standardized [21], [22], transmission bit rates as large as 100 Gb/s per wavelength can be achieved, over the distance of 100 km between the central office and homes [19].

Cellular backhaul, which facilitates the transmission between cellular network base stations and their switching center, is today also migrating to fiber-based solutions [23]. The traditional copper-based, low-bandwidth solutions known as T1/E1 offer capacity of only 2 Mb/s and have a cost structure that increases linearly with capacity, which clearly does not economically scale with increasing number of users, base stations, and mobile network video services [24], [25].
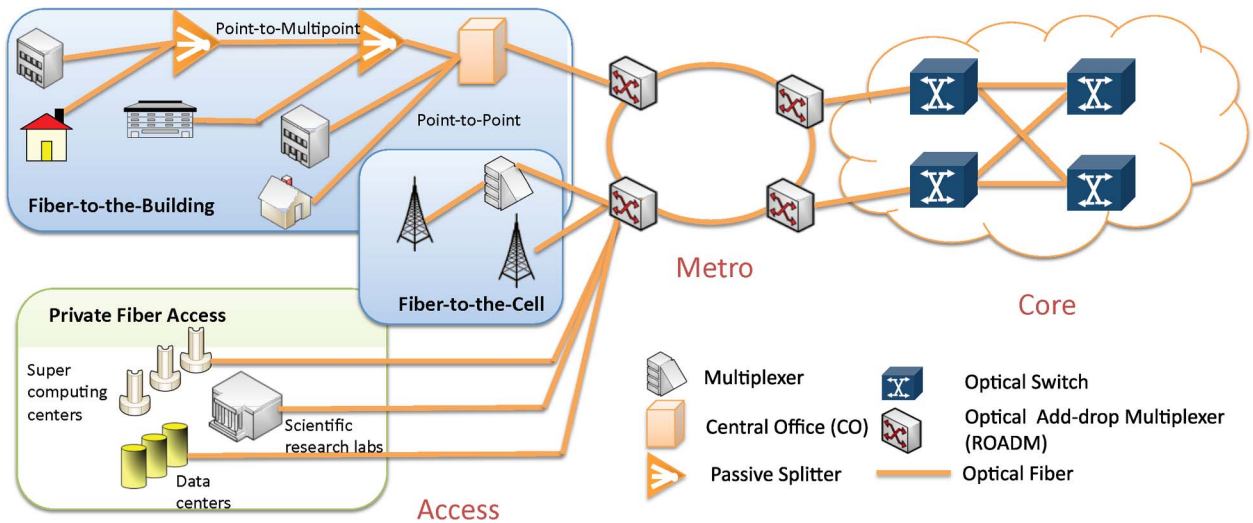
Fig. 2. *Fiber optic service proliferation in traditional telecommunication network segments, and beyond.*

Finally, private access is a new paradigm in the network access segment where the fiber connection, typically at gigabit per second of speed, is exclusively used by and privately dedicated to a specific user (not shared), e.g., a user who requires a large-scale media stream, such as science labs and large data centers. This concept describes a major growing trend, a transition from the one-size-fits-all network architecture, to one that must increasingly take into consideration diverse types of traffic and service requirements. While the private access connection paradigm has evolved from and is fostered by eScience applications and experimental network trials [26]–[29], it is broadly accepted that cloud computing, data center interconnection, distributed computing, and visualization in finance, health, and government sectors will soon deploy these types of access networks architectures [30], [32], [33].

Also shown in Fig. 2 are the metro and core network segments. The metro segment has been traditionally designed to meet requirements for flexible traffic aggregation on the one side (access), and high bandwidth on the other (core). However, metro and core networks continue to converge and blend in terms of both architecture and optical system design, the topic discussed next.

B. Advances in Optical Network System Design

Fig. 3 illustrates a typical optical transmission and switching system, and emphasizes a few recent and noteworthy advances. In an optical transmission system, multiple wavelengths are multiplexed into a single fiber, which is referred to as wavelength division multiplexing (WDM) transmission, where all-optical amplifiers are used to compensate for power losses in a long-haul transmission. Fig. 3 also illustrates other important components of the system, including optical switches, reconfigurable optical

add–drop multiplexer (ROADM), as well as optical receiver, transmitter, multiplexers/demultiplexers (MUXs/DMUXs), and wavelength-selective switches (WSS). Each of the subcomponents has technologically advanced to support multiple electronic layers, large capacity per channel, flexible wavelength (frequency) allocation, and fast optical switching and hardware processing. The Internet packets, Ethernet frames, SONET/SDH, OTN containers, etc., can all be modulated onto a wavelength today, and transmitted over the same fiber.

Especially remarkable has been the recent offering of Ethernet technologies as a service on optical-centric networks. Even as traditional SDH/SONET is being improved through advanced framing schemes, it is being increasingly replaced in many areas by Ethernet solutions, which are designed to make direct use of WDM optical networks. These major trends are expected to continue for the foreseeable future, when the implementation of 100-Gb/s Ethernet is expected to bring about high levels of parallelism for digital media application processing. The parallelism allows a single user to access the network via multiple parallel wavelengths, or by inverse multiplexing of single-user digital media frames. Enabling parallelism is a paradigm shifting capability for digital media services, because there are physical limitations of serial transmission in bit rates per lambda, and the use of parallel ports, transmission channels and resources with multilambda/fiber network interfaces can be applied to improve the total performance [34]. Parallel optical interfaces are not only embodied in the Ethernet standards, but are also driving the evolution of optical networking toward parallelism at terabit scale [30]. Another attraction of 100-Gb/s transmission for digital media services is the fact that it is coherent, which gives the power to compensate for large amounts of dispersion as well as
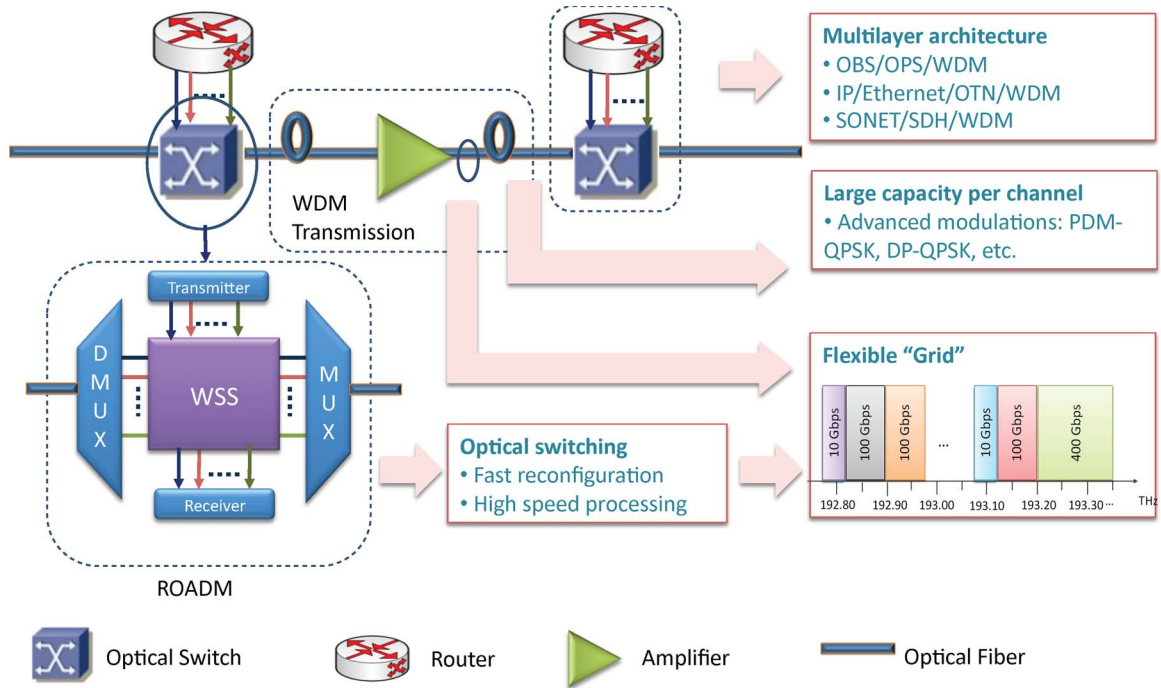
Fig. 3. *Advances in optical transmission and switching system design.*

having a much better optical signal-to-noise ratio (SNR) performance.

In the ongoing quest for large amounts of transmission and switching capacity, *flexible grids* play the key role. In a flexible grid, as opposed to the fixed one (as defined by ITU-T), high data rate channels can coexist with low data rate channels in optical spectrum [31]. For instance, while 10-Gb/s channels require 50-GHz spacing, 100- and 400-Gb/s channels can be allocated with 75- and 150-GHz spacing, respectively, etc. Also of note are advances in hardware programmability, fast processing, advanced coding, and modulation techniques. Field-programmable gate array (FPGA) technologies for optical transmitters and receivers have advanced, being able to produce line rates at 100-Gb/s speeds, thus enabling a true real-time digital signal processing [35], [36]. Fast reconfigurable ROADMs have been reported with integrated optical label readers and channel selectors [37], [38]. The concept of flexible add–drop bandwidth has been also used in ROADMs and wavelength selective switches with liquid-crystal-on-silicon (LCOS) technologies [39], [40]. New ROADMs making use of LCOS arrays can give very fine spectral resolution (1 GHz) and allow a wide range of add–drop spectral shapes, which is especially beneficial for UHD media services. Numerous results from field trials and experiments have been reported, such as in [41], with a ten-channel WDM transmission system with PDM 16-quadrature amplitude modulation (16-QAM) and 224 Gb/s per channel, over 2000 km of fiber.

C. Enhanced Control Planes, Dynamic Provisioning, Automation, and Programmability

Today, multiple architectural design trends are shaping the structure and function of next-generation control and management planes. Several of these major trends are listed here.

Unified Versus Multiple Control Planes: Traditionally, each network layer has had its own separate control and management plane. More recently, several standards organizations, including the Internet Engineering Task Force (IETF), have been examining the potential to provide architecture for a single unified control and management plane that could address the needs of any layer. The capability to manage traffic flows that may traverse multiple different layers (Internet, optical) at various times in an end-to-end path is especially important for high-resolution digital media. Providing options for keeping digital media flows within optical domains for as long as possible enhances efficiency and helps ensure high quality.

Enhanced and IP-Based Control and Management Planes: There are many activities directed at enhancing control and management plane functionality to enable more degrees of dynamic control over optical paths [42]. Two major control plane standards have already reached commercial maturity: Automatically Switched Optical Networks (ASON) standardized in ITU-T, and Generalized Multi-Protocol Label Switching (GMPLS) developed by

the IETF. GMPLS has gained popularity in the past years. With the Internet-specific extensions (GMPLS-TE [43]), the use of GMPLS-based is not only confined to the optical layers, but it extends to the packet switching networks. Also, of note is a novel management plane paradigm under the umbrella of Multi-Technology Operation System Interface (MTOSI), as standardized by the Telemanagement Forum (TMF) [44]. MTOSI is expected to facilitate the interaction between vendor-specific network management systems, which is a critical feature in today's systems.

Dynamic Versus Static Provisioning: Traditionally, the majority of optical paths have been designed and implemented as static components within a static infrastructure. Today, however, major providers of optical switching equipment provide an ability to dynamically establish optical paths between switches at time scales that are on the order of hundreds of milliseconds. These capabilities are being further enhanced to allow for increasingly granulated optical component control and management. Such innovative capabilities are already being implemented in advanced communications exchange facilities [45].

Decentralization Through Edge Subsystems and Signaling: As networks expand in the number of services provided, service attributes, location served, and edge devices accessed, it is necessary to remove bottlenecks forced by the reliance on centralized processes to determine dynamic changes related to standard adjustments, enhancements, and responses to fault conditions. By providing options for decentralized control and management, new architectures are enabling edge systems to directly adjust their network resources through specialized signaling. These decentralized capabilities can be accessed through policy-based processes by independent organizations, individuals, and even applications. Various third party control and management subsystems have been proposed, most notably the path computation element (PCE) framework, and latest extensions further enhance the path computation and setup functions in core networks [46]. Another example is a monitoring framework called perfSONAR, which can be used for network performance monitoring in multiple layers [47].

Network Programmability: One of the major areas of network innovation over the last decade has been the design and development of capabilities for programming networks, including the dynamic provisioning of lightpaths and virtualization [48], [49]. It has been demonstrated that such programmability can extend to enabling dynamic provisioning for optical multicasting, including to support high-quality, large-volume digital media [50]. Some of the early work in this area was derived from Grid networking activities, which enabled these distributed computing environments to utilize networks as "first class" resources directly programmable as opposed to external resources

that could not be managed [51]. As networks become programmable environments managed by multiple parties, it has also become important to develop common, standardized methods for describing individual programmable network resources, especially to enable the design of processes based on web services, which in turn can enable the virtualization of network services and infrastructure [52].

Automation with Analytics: To enhance dynamic provisioning, configuration, fault detection, and response in complex multiservice networks, a combination of manual programmability and automation of many common functions is critical. Both the manual programmability and automation can be combined with sophisticated analytics, so that many elements in the network, both resources and traffic behaviors, can be continually monitored and analyzed. Any types of special conditions can be detected and responses can be provided depending on requirements. Some of these types of techniques are emerging, such as the OpenFlow framework, which is based, in part, on the concept of extending flow tables to servers external to switches, allowing individual flow monitoring, and providing mechanisms for responding to specific types of flow behaviors [53].

IV. HIGH-RESOLUTION VISUALIZATION IN SCIENCE

Data-intensive science applications tend to be significant drivers of technology innovation because they experience technical barriers many years before such limitations are experienced by other communities. Consequently, the technology solutions used by such applications provide a window through which it is possible to predict future service and technology requirements. At this time, many computational science communities have already designed "science data networks" that are optimized for optical transport. Among the most important discovery tools used by data-intensive science are innovative techniques for viewing large volumes of data. Scientific visualization requires the rendering and transport of large-volume, high-quality digital media among multiple sites, including 3-D digital media, UHD digital media, augmented reality environments, and photo realistic virtual reality. In this section, we examine high-resolution data-intensive scientific visualization, which is a type of rich digital media application.

In general, media are prepared and transmitted primarily to provide a passive viewing experience. However, increasingly digital media environments are being created to allow for immersive, photorealistic, interactive experiences, such as 3-D environments. These environments allow direct interactivity with digital media objects including direct manipulation, sometimes incorporating physical feedback using haptic techniques. These types of environments require continuous high-speed signaling to

and from the digital media environment as well as rapid continuous changes in the environments. Such applications enable virtual worlds to be displayed and directly experienced within specialized spaces. For these types of media applications bidirectional high performance, network designers must give consideration to many more issues that are common with traditional source destination stream management requirements. For example, considerations may include those related to general workflow management, accessing multiple large remote data repositories, selecting data, merging data, transporting the resulting composite information to a rendering site, and interactively visualizing the results.

For large-scale, high-resolution images, the end site may require a "graphics card" that comprises a large-scale computer cluster. Extremely high-resolution images today cannot be displayed on any single display device because none provide a capability for streaming the large number of bits (pixels) required. Consequently, tiled displays are used, composed of individual high-resolution screens, each capable of displaying extremely large numbers of bits. For ultrahigh-resolution scientific images, the network requirements for such displays can far exceed the capacity of metro core networks. Furthermore, these environments require more than mere arrays of displays; they must enable viewers to interact with digital media objects, as if they existed in physical reality. For instance, groups of scientists may be examining, and interacting with, the same 3-D protein molecular structure from multiple sites on several continents. In addition, they must also be integrated with extremely large data sets, which must be transported as part of scientific workflows. Early experiments with large-scale data transport over optical networks demonstrated the utility of large-scale dynamic lightpath providing for multiple data-intensive applications [48], [54], [55].

A. Designing the Network for Scientific Visualization

Designing networks for scientific visualization requires consideration of multiple components, processes, and tim-

ing among processes. For these types of high-capacity networks, the data plane is separated from the control plane. The control plane comprises multiple processes, which at the highest level can be based on web services. A high-level edge process is a method of signaling into the network to request resources. Such signaling, which can be initiated directly by applications and instruments, is managed by policy-based access and authorization procedures. Within the network, various types of processes are used to provide for the interactivity between and among request processes and resource discovery and utilization processes, which determine what resources are available at what times, and how they can be allocated. These processes include those related to optical channel routing, topology and device discovery, path selection, maintaining databases of physical links, creating new paths, optimizing path selections, providing for traffic engineering, implementing constraint-based routing when required, supporting interworking, and implementing protection and restoration tools. In other words, these types of processes require much more state information than is common in traditional static networks, as well as more interprocess communication and messaging.

Given these types of processes for dynamic provisioning, one of the key issues is the relative timings among individual processes. Illustration in Fig. 4 provides a relative description of such timings, which are derived from testing on a large-scale metro area network research testbed with 24 individually addressable 10-Gb/s lightpaths using an experimental control plane, integrated with GMPLS. None of the processes are optimized. The time can be significantly reduced. The issue here is not the actual process times but the times relative to all the processes. The timings are in seconds and describe a 20-GB file transfer using FTP. The path allocation request is almost instantaneous (0.0023% of the process). The control plane processing is 0.0167% of the process. The path ID is returned quickly (0.0023%). The longest relative timing period related to provisioning is the network reconfiguration (0.116%). The FTP setup time (0.00065%), path

End-To-End Transfer Time

Function	Time (s)	Percent
Path Allocation Request	0.5	0.002325
Server Processing	3.6	0.016741
Path ID Returned	0.5	0.002325
Network Reconfiguration	25	0.116257
FTP Setup Time	0.14	0.000651
Data Transfer 20 GB	174	0.809152
Path Deallocation Request	0.3	0.001395
Server Processing	11	0.051153
Totals	215.04	1

Fig. 4. *Path setup times.*

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

Jukan and Mambretti: Evolution of Optical Networking Toward Rich Digital Media Services
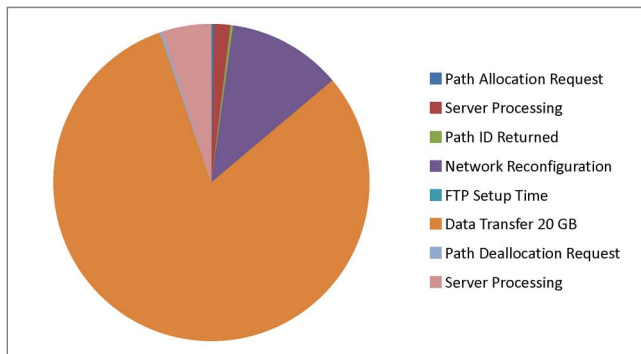
deallocation (0.001395%), and final control plane processes are minimal (0.051%).

In sum, when designing and operating a dynamic optical network for scientific visualization applications, a set of system management processes is required to provide for locating devices, implementing configurations, communicating management instructions, initiating and stopping various modules, providing for resource balancing, adjusting interfaces, and other functions. In addition, there are other processes that directly monitor and adjust physical components. More detailed descriptions on the emerging requirements and tools for networks interconnecting supercomputing centers, and scientific labs can be found in series of formal reports developed by federal agencies [26]–[28].

B. Experimentation With Rich Digital Media in HPDMnet

Over the last few years, a number of large-scale (global) distributed environments have been designed and implemented as platforms to support next-generation digital media, especially large-scale scientific visualization, 3-D displays, and UHD digital media. One of the first such large-scale, distributed environments based on dynamically provisioned lightpaths, which was designed for data-intensive science and visualization, was the OptIPuter, a national fabric with international extensions based on the international Global Lamda Integrated Facility (GLIF) [29], [57], [58]. A more recent example is the High-Performance Digital Media Network (HPDMnet), a global distributed experimental dynamically provisioned Layer1/Layer2 (L1/L2) end-to-end network spanning four continents that can support extremely high-resolution scientific visualization, medical imaging, virtual spaces, stereo 3-D movies and 4 k and 8 k streams [59]. The HPDMnet initiative comprises three areas of activity, one related to content, one to services and intelligent middleware, and one to an underlying network research testbed.

A typical HPDMnet infrastructure topology, based on the GLIF, is shown in Fig. 5. Note that the GLIF has not been designed as a network, but as a large-scale distributed facility within which it is possible to create multiple specialized networks with individual specialized characteristics. To this end, the GLIF is based on international open communications exchanges, such as the StarLight international/national communications exchange in Chicago, IL, that enables peering at all network layers, not merely traditional L3 peering, which has become common at almost all network exchanges today. StarLight provides support for multiple advanced distributed facilities, including the National Science Foundation (NSF)-funded TransLight [56], [60]. In the HPDMnet cloud, multiple techniques exist or a under investigation for digital media transport, using L1/L2 paths, including direct live streaming end-to-end, network caching, edge caching, file transfer end to end, media fusion, and optical multicasting. We will here present a couple of such HPDMnet experiments, one related to optical multicasting and the other to end-system design.

In the HPDMnet cloud, optical multicasting has been demonstrated, and has proven to be a particularly efficient technique for large-scale digital media streams. Fig. 6 shows one of the first HPDMNet optical multicasting demonstrations that was staged a few years ago for an international supercomputer conference in Tampa, FL. At each of three sites, i.e., at the conference in Tampa, at Louisiana State University in Baton Rouge, and at Masaryk University in Brno, Czech Republic, a 750-Mp/s stream of high-quality, high-resolution digital media was sent using the StarLight international/national communications exchange on an L1/l2 path to the NetherLight exchange in Amsterdam, The Netherlands. For optical multicast, two techniques can be used. One uses optical frame duplication, within optical networking core equipment, enabling wavelength path (lighpath)-based streams to be sent to remote locations to be duplicated and redistributed. Another technique is based on L2 frame duplication using server-based computer clusters. Each site was sending one 1.5-Gb/s stream and receiving two from the other sites allowing for many to many streaming. This technique has subsequently been used to support a number of digital media showcases for conferences, workshops, collaborative activities, and other specialized events.

Although this network uses IP addressing, all paths avoid routers to ensure that large data streams can be transported with high performance and high quality. To achieve this performance and quality, various configuration parameters along the paths must be addressed. For example, standard maximum transmission unit (MTU) sizes are inadequate to support these large flows sufficiently. An IEEE standard Ethernet frame size is 1518 B, which includes the payload, L2 header, and frame check sequence (FCS). Because this frame size does not adequately support large-scale, long-duration flows, HPDMnet uses jumbo frames, which are approximately 9000 B. Jumbo frames are not actually part of an IEEE standard but are a convention, initially suggested by designers and implementers of high-performance networks supporting data-intensive science research and then incorporated by manufacturers into high-performance switches and network interface cards (NICs). Because jumbo frames are not a standard there is some variance among implementations. However, in general, the 9000-B size is widely used, and it is important for these types of networks to have jumbo frames end-to-end on all paths. It is notable that as 40- and 100-Gb/s paths are being implemented, discussions have begun about creating a much larger frame size for these channels.

A typical end-system setup in HPDMnet is shown in Fig. 7. This specific setup uses the open source scalable adaptive graphics environment (SAGE) system that supports visualization [61]. SAGE is a technology for visualization of multiple high-definition video streams in a
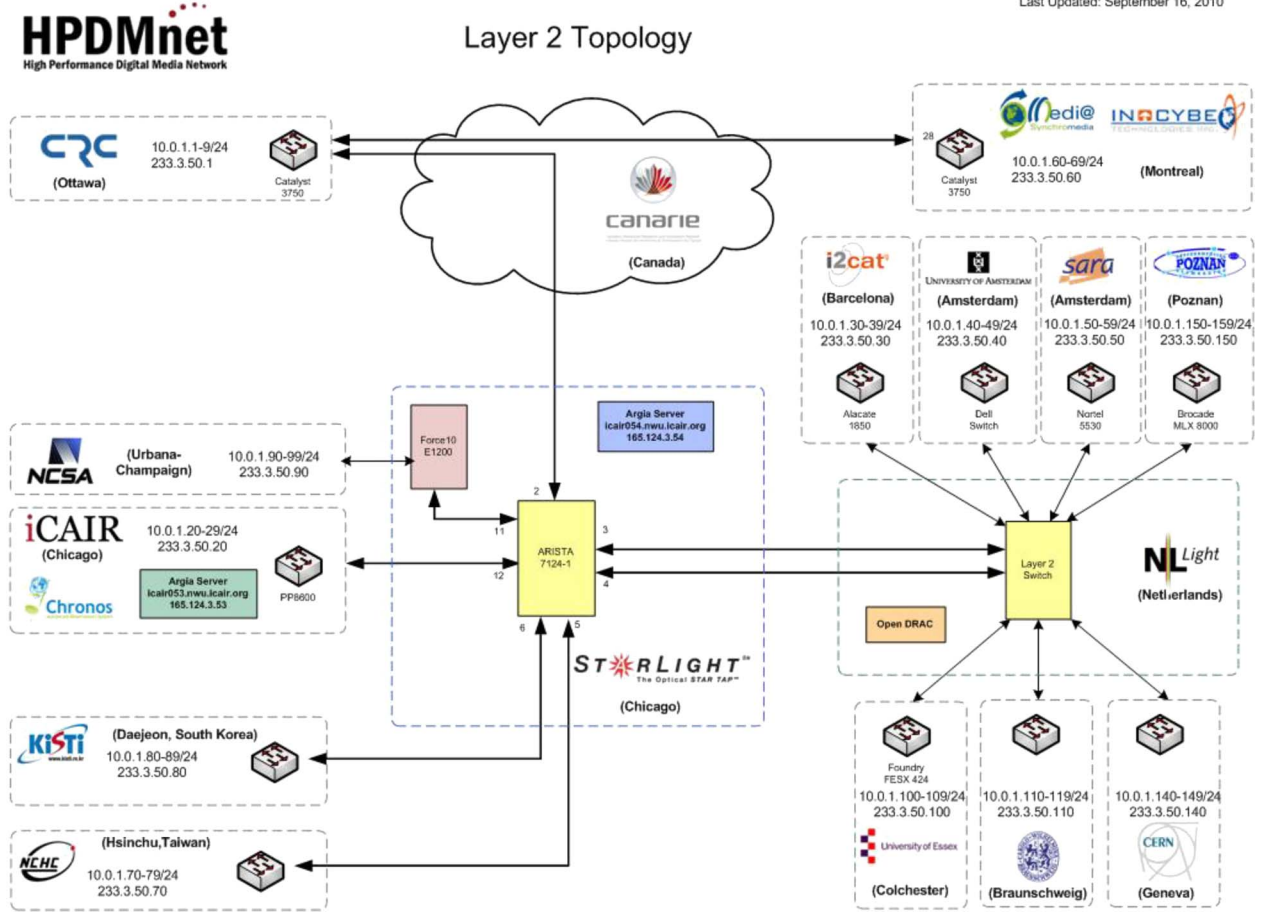
This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

Jukan and Mambretti: Evolution of Optical Networking Toward Rich Digital Media Services

Fig. 5. *HPDMnet global networks topology.*

tiled screen. Each local lab in HPDMnet furthermore connects to the HPDMnet cloud. The tiled display wall (here: 3 × 3) is used for visualization, where each display is controlled by a single SAGE Receiver, a part of the SAGE Receiver cluster. Every SAGE server is running a FreeSpace Manager server responsible for the control of the whole visualization system, as well as an application server with the SAGE application interface library (SAIL) supporting streaming via different applications to the receivers. Finally, the SAGE user interface (UI) is used, which is a user-friendly graphical tool for configuration of the visualization parameters for the FreeSpace Manager. As previously mentioned, this kind of end system uses a "graphics card" that comprises a computer cluster (here: nine machines). Multiple and high-resolution images could not be displayed on any single-display device, and therefore, tiled displays are used. In the example shown in Fig. 7, four individual streams are received, each capable of being displayed in foreground or background, and in various sizes.

When the number of streams increases, so do the bandwidth and processing power, which raises the issue of network utilization and end-system scalability. To address

the bandwidth- and end-system performance challenge, the potential to implement a media composing middleware was investigated, which, based on the receiver requirements, can compose multiple media streams into a single stream for visualization [59], [62]. For instance, assume that three streams are received, and the user would like to see two of them in the foreground in high resolution, and the third one in the background, in low resolution. In this scenario, one could media-compose the two high-resolution streams and transmit them together over HPDMnet, and add the third stream through the Internet. In [62], it was measured and showed that while transmission of the two streams independently would require two times the bandwidth of each stream, if these streams are composed, based on the percentage of overlap of the videos, the total bandwidth requirement would decrease significantly. For only 25% of video overlap, the bandwidth saving of 12% was already measured, etc. Note that bandwidth savings in HPDMnet were not in fact important for network resource optimizations, since the network was mostly underutilized, but it had a large impact on end-system delay

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

Jukan and Mambretti: Evolution of Optical Networking Toward Rich Digital Media Services

and processing. With networked media composition, the end-system performance can better scale with the number of streams visualized. In the future, a scenario is envisioned, in which many media composers are distributed throughout the network, and only streams chosen by users are composed and transmitted in the desired quality.

V. HIGH-QUALITY REAL-TIME CONSUMER VIDEO APPLICATIONS OVER OPTICAL NETWORKS

With the increasing demand of HD multimedia content and the emergence of 3-D media technology, delivering media with quality assurances to end users is becoming progressively bandwidth, processing, and memory intensive [4], [63]. This evolution is much like the one that has been experienced in science networks, as discussed in Section IV. For instance, in a high action sports game, audiences would like to have a comprehensive view of the game to obtain a precise judgment on scores. In such scenarios, recording an event with a multiview video greatly enhances the visualization experience, which is one of the critical factors in consumer video applications. Since multiview video requires considerably larger bandwidth,

using optical networks for rich media transmission is emerging as a viable solution to serve high-quality consumer video applications.

In this section, we are mainly concerned with the evolution of high-quality consumer video applications and their potential deployment over optical networks. To this end, we analyze the candidate optical network services from the point of view of required capacity, scalability, and real-time content delivery. Since consumer digital media services are in their early stages of wider deployment over optical networks, we discuss how media content can be moved closer to the user, with the applications of optical bypass between IP routers.

A. Optical Network Services for High-Quality Consumer Media Applications

Optical networks have been already deployed for digital cinema distribution with demonstrated cost efficiency and superb performance [64], [65]. In Japan, several movie theaters have conducted a trial with 4-K digital cinema distribution over optical networks, and demonstrated that network-based digital cinema distribution is both cost efficient and secure [66]–[68].
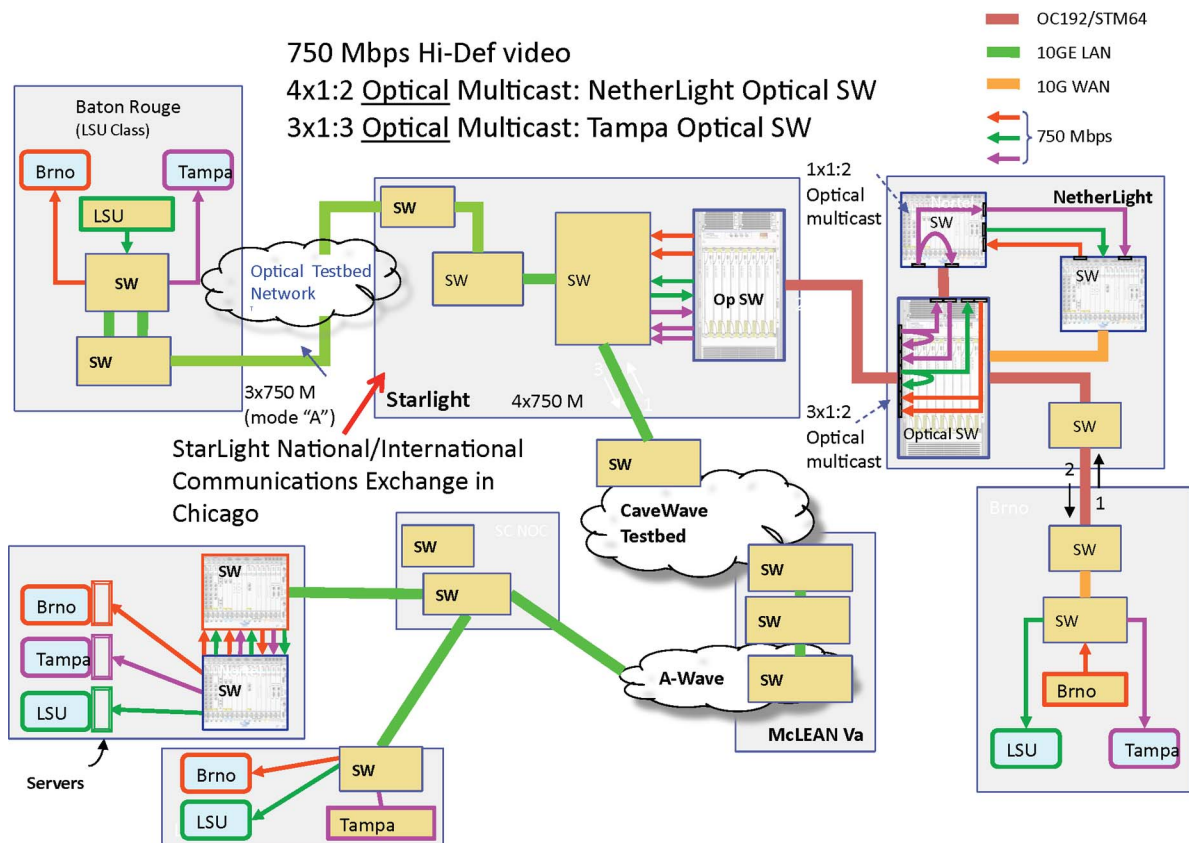

Fig. 6. An HPDMnet topology for L1 dynamic optical multicast demonstration that was staged for an international supercomputer conference (Op SW is an Optical Switch; SW is an L2 switch).
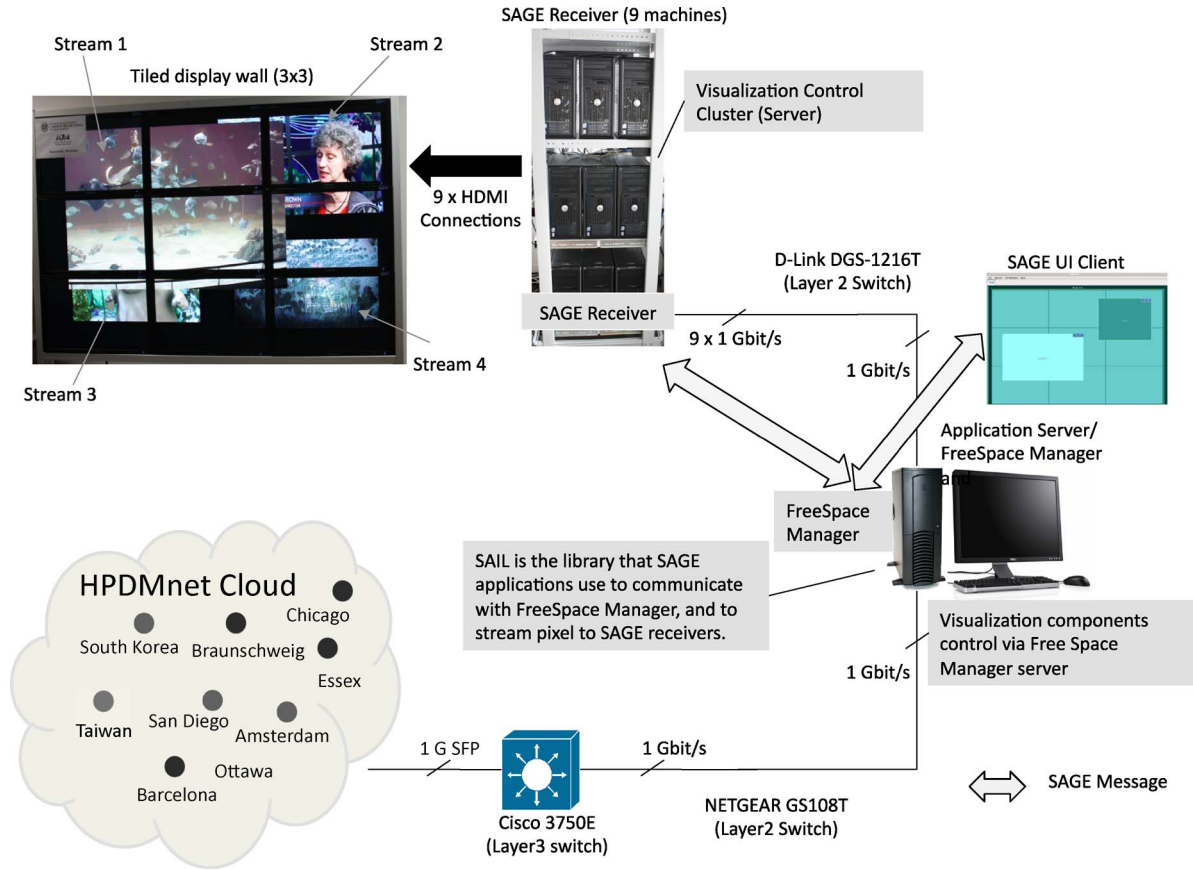
This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

Jukan and Mambretti: Evolution of Optical Networking Toward Rich Digital Media Services

Fig. 7. *A typical HPDMnet end-system configuration with tiled displays controlled by the SAGE software. This setup was used for an HPDMnet any-cast demo at ACM Supercomputing in 2009, and the screenshots were taken from real traffic received at the TU Braunschweig from Taiwan (Stream 1), Chicago (Stream 2), Amsterdam (Stream 3), and one local sender in Germany (Stream 4).*

Optical burst switching (OBS) network is a good candidate to serve digital cinema applications, and bulk data transfer in general, which typically requires a high capacity connection for a specific time period, between a limited number of known users and sources of content. Compared to optical circuit-switched WDM network, OBS can achieve better capacity utilization and a high degree of statistical multiplexing while also accommodating the bursty nature of media applications in an efficient fashion [69], [70].

Optical circuit-switched WDM networks, on the other hand, are best suited for more static connections. A typical use case is the connection between video sources and a production house, where statically preconfigured optical networks can be used to support the high bandwidth needs for video transmission from different video sources. For extreme data rate requirements, coherent and parallel transmission techniques at 100 Gb/s can be applied. As previously mentioned, the main attractions of 100 G for real-time SHD/UHD applications that require low jitter and synchronization are not only the high bandwidth, but also coherent transmission and parallelism.

Another important family of optical transmission and switching techniques includes optical time division multiplexing (OTDM), and optical packet switching (OPS) and optical slot switching (OSS).

OTDM is intended for the creation of optical multiplexed streams at speeds significantly in excess of the maximum speed of electronics [70]. Similar to SONET/SDH and OTN networks, which deploy the concepts of virtual concatenation and inverse multiplexing, OTDM can be used to transmit each view stream in an MVV separately. Each MVV stream can either be encapsulated as an individual OTDM stream, or decomposed and transmitted in multiple OTDM streams. The latter relies on the multiresolution encoding technique to encode high-definition images in 3-D scenes or MVV scenes into multiple objects with different resolutions and transfer them in multiple streams [71]. An SHD/UHD video stream can therefore be encoded into multiple streams and transferred with multiple paths. All video streams must be synchronized with the timing information embedded in the videos captured from different cameras used together as an MVV stream. When using OTDM for multiview

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

Jukan and Mambretti: Evolution of Optical Networking Toward Rich Digital Media Services

video, since all the views are transmitted on the same lightpath (wavelength), full synchronization between individual streams can be maintained, since all the streams experience the same transmission delay and attenuation.

OPS is the optical realization of its electronic counterpart, and has unique capabilities to provide high bandwidth, low latency, and fine granularity, which are essential for rich digital media. OPS systems are still difficult to implement because practical optical buffers are not yet available. However, it has been shown that OPS network services carry significant potential to serve the real-time streaming in consumer video applications. The recent paper by Kataoka et al. [72] presented a successful field trial of 4-K uncompressed streaming over multicast-capable 80-Gb/s OPS network using SOA-based broadcast-and-select switch and stacked optical-code label processing. OSS, on the other hand, has also been demonstrated [73] as a new optical packet switching architecture for HD contents transfer. In this architecture, the control plane is separated from the data plane and provides the signaling for the content delivery network, while the OSS layer consists of a nanosecond-range fast optical switch and a Ethernet/optical slot protocol converter.

To scale to a large number of end users, next-generation media infrastructures will have to build on technologies able to handle not only circuits but also time-shared granularities in the optical domain [74], [75]. To this end, not only hybrid node architectures have been proposed [76]–[78], consisting of a parallel arrangement of optical or electronic packet and circuit switches, but also more flexible gridless arrangements are on the horizon [79]. These new methods enable flexible allocation of optical spectrum and are expected to be the next important step in the evolution of optical networking towards rich digital media services.

B. Moving Content Closer to the Users

Today, the delivery of real-time multimedia content comprises three primary functions: media production/composition, adaptation, and delivery. The high-quality media production and composition function produces multimedia content using one or more incoming video feeds from different locations (e.g., different camera views in a sporting event or dual camera view in a news report). Specifically, the composition process requires operations such as interlacing, overlaying, and merging, and uses raw, uncompressed UHD media, which requires very high processing power and storage facilities. The latter presents a significant entry barrier into live, real-time UHD media production and composition.

When the media content needs to be delivered to multiple end devices, it typically requires adaptation to different formats [80]. This process may include adapting media to 3-D or 2-D formats, change in resolution and frame size, video quality or compression and encoding format to match the user's end devices and bandwidth. Note that media adaptation must also support real-time processes to support delivery of live content to the end users, which again requires significant processing when considering the number of concurrent requests and the different end device and network constraints. Currently, media adaptation is limited to support a fixed number of end devices and encoding formats, typically by preprocessing and storing digital content. Most of the media adaptation happens at either the transmitter (e.g., video source) or receiver (end-user hosts). For UHD adaptation, on the other hand, the required high processing power is provided by large cluster facilities owned by production houses of large content producers.

After adaptation, content distribution networks (CDNs) are typically deployed for media delivery, as intermediaries between media content production companies and end users. As mentioned in Section II, CDNs are distributed networks of storage and servers close to the edge of the metropolitan or regional networks, and end users, thus minimizing overall bandwidth cost [81]. The actual multimedia content is delivered to the end users using multicast over the traditional IP network infrastructure, or alternatively, a static optical connection can be configured for a high-end user, such as a large-screen viewing provider of a major event. Note that over an IP layer, a unique multicast tree is generated for each media adaptation function (i.e., each supported end device).

It should be noted that among the media functions of production, composition, adaptation, and delivery discussed above, only "delivery to the end user" has successfully deployed advanced networking solutions, e.g., CDNs. Indeed, in UHD applications, such as digital cinema, traditional postal services and not networks are often used to "deliver" large storage discs between media composition and adaptation sites, and to the movie theaters. This is where optical networks can play an important role in the future.

Fig. 8 depicts the possible influx of optical networking into the existing media delivery platforms. For easier understanding, we simplify the representation of media flows from content creators (video sources) to end users (media consumers). We show that video from multiple sources to be used for production and composition can be delivered to the "production house in the cloud" by using configurable optical networks to support the high bandwidth needs. The production functions can be deployed dynamically, by using cloud computing services [83]. Similarly, the composition/production, adaption, and delivery functions can use optical networks for their interconnections.

Optical networks can also provide significant advantages to media delivery to end users. As shown in Fig. 8, the delivery functions can be "moved" closer to the users, by dynamically setting up connections in the optical layer, which bypass the IP routers otherwise traversed [82]; see dashed lines between the edge routers. In this way, we can significantly reduce the bandwidth consumed in the IP
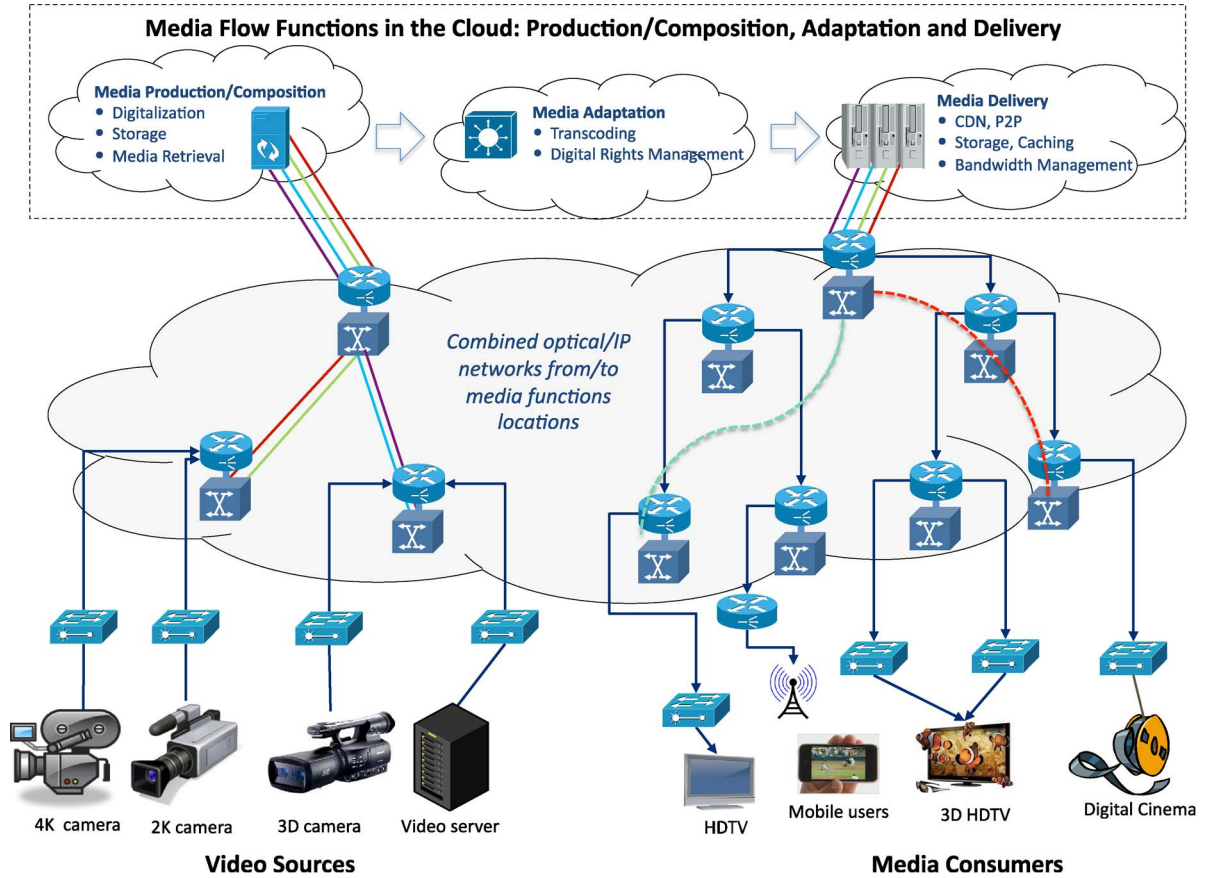
This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

Jukan and Mambretti: Evolution of Optical Networking Toward Rich Digital Media Services

Fig. 8. *Media flow functions in the cloud: composition, adaptation, delivery, and increasingly many new capabilities with combined optical/IP network connectivity and optical bypass (dashed line).*

network infrastructure, and as a consequence serve a larger customer base, with improved user experience. In addition, localization of CDN operations can be dynamically changed, thus significantly decreasing the strain on the network infrastructure.

VI. SUMMARY AND OUTLOOK

In this paper, we discussed the evolution of optical networks from foundational data transport utilities and infrastructure to true enablers of next-generation applications and facilities, especially rich digital media services. The paper described the recent achievements of optical networking in various segments, from the network core to aggregation and access sites, to edge devices. We focused on a few distinctive areas of broad impact of optical network technologies related specifically to rich media services: high-resolution scientific visualization and high-quality real-time consumer video applications.

It is not difficult to predict some key aspects of future optical networks, as they transition from supporting millions of small data flows to the foundation of complex digital media centric environments. For example, the need

for increasingly larger capacity at all levels has been, remains today, and will be a very long-term trend. The innovations in optical switching and transmission will also continue. Of all of the trends, perhaps the most significant are those that will further enable optical network services and infrastructure to be highly programmable, allowing new services to be added, enhanced, and reconfigured instantaneously and dynamically. A wide range of new programmability and virtualization capabilities are enabling close integration of networks with other types of resources, including computing and storage systems, instruments, and sensors. It is envisioned that in the near future, the evolution of optical networking toward rich digital media services will also be closely integrated with the rapidly evolving cloud-based media services, which are expected to be transformative to scalable multimedia processing and delivery.

Thus far, the pioneering rich digital media applications for optical technologies have primarily been those designed for scientific research. High-quality, real-time media production and distribution are now emerging and becoming democratized, as they migrate from being based on expensive capabilities owned and managed by a few large organizations to being based on common, inexpensive

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

Jukan and Mambretti: Evolution of Optical Networking Toward Rich Digital Media Services

ubiquitous resources that anyone can use. However, this is just the beginning. Engineering design, architectural modeling, medical procedures, and crisis response are just a few examples of many emerging rich digital media services and applications that will challenge and take advantage of further great optical networking advances. ∎

REFERENCES

[1] M. Novell, V. Vusirikala, and R. Hays, "Overview of requirements and applications for 40 gigabit and 100 gigabit ethernet," Ethernet Alliance White Paper, Jun. 2010.

[2] T. B. Zahariadis, G. Pau, and G. Camarillo, "Feature topic on future media internet," *IEEE Commun. Mag.*, vol. 49, no. 3, Mar. 2011.

[3] H. Zimmerman, "OSI reference model, the ISO model of architecture for open systems interconnection," *IEEE Trans. Commun.*, vol. COM-28, no. 4, pp. 425–432, Apr. 1980.

[4] J. Leigh, L. Renambot, A. Johnson, B. Jeong, R. Jagodic, N. Schwarz, D. Svistula, R. Singh, J. Aguilera, X. Wang, V. Vishwanath, B. Lopez, D. Sandin, T. Peterka, J. Girado, R. Kooima, J. Ge, L. Long, A. Verlo, T. DeFanti, M. Brown, D. Cox, R. Patterson, P. Dorn, P. Wefel, S. Levy, J. Talandis, J. Reitzer, T. Prudhomme, T. Coffin, B. Davis, P. Wielinga, B. Stolk, G. Bum Koo, J. Kim, S. Han, J. Kim, B. Corrie, T. Zimmerman, P. Boulanger, and M. Garcia, "The global lambda visualization facility: An international ultra-high-definition wide-area visualization collaborator," *Future Generat. Comput. Syst.*, vol. 22, no. 8, pp. 964–971, Oct. 2006.

[5] M. Armstrong *et al.,* "High frame-rate television," presented at the BBC Research, Sep. 2008, WHP 169.

[6] E. Nakasu *et al.,* "Technical development towards implementation of extremely high resolution imagery system with more than 4000 scanning lines," in *Proc. IBC Conf.,* 2006.

[7] International Communications Union. [Online]. Available: http://www.itu.int/rec/T-REC-H.264

[8] L. Onural, T. Sikora, J. Ostermann, A. Smolic, M. Civanlar, and J. Watson, "An assessment of 3DTV technologies," in *Proc. NAB Broadcast Eng. Conf.*, 2006, pp. 456–467.

[9] V. Jacobson B. Braden *et al.,* "TCP extensions for high performance," *RFC 1323*, May 1992.

[10] C. Jin, D. Wei, S.H. Low, J. Bunn, H.D. Choe, J.C. Doylle, H. Newman, S. Ravot, S. Singh, F. Paganini, G. Buhrmaster, L. Cottrell, O. Martin, and W.-c. Feng, "FAST TCP: From theory to experiments," *IEEE Network*, vol. 19, no. 1, pp. 4–11, Jan. 2005.

[11] E. He, P. Vicat-Blanc Primet, and M. Welzl, "A survey of transport protocols other than "Standard" TCP open grid forum," Document GFD.55.

[12] A. Masuda *et al.,* "Application-network collaborative bandwidth on-demand for uncompressed HDTV transmission in IP-optical networks," in *Proc. IEEE/IFIP Network Oper. Manage. Symp. Workshops*, 2010, pp. 177–180.

[13] C. R. Lima *et al.,* "High definition video broadcast over core IP networks," in *Proc. Opt. Fiber Commun. Conf.*, 2006, DOI: 10.1109/OFC.2006.215779.

[14] W. P. K. Yiu, X. Jin, and S. H. G. Chan, "Challenges and approaches in large-scale P2P media streaming," *IEEE Multimedia*, vol. 14, no. 2, pp. 50–59, Apr.–Jun. 2007.

[15] H. H. Yen, S. S. W. Lee, and B. Mukherjee, "Traffic grooming and delay constrained multicast routing in IP over WDM networks," in *Proc. Int. Conf. Commun.*, Beijing, China, May 2008, pp. 5246–5251.

[16] R. L. Grossman, Y. Gu, M. Sabala, C. Bennet, J. Seidman, and J. Mambretti, "The open cloud testbed: A wide area testbed for cloud computing utilizing high performance network services," abs/0907.4810, 2009.

[17] R. Grossman, Y. Gu, M. Sabala, J. Mambretti, A. Szalay, and K. White, "The open science data cloud," in *Proc. Sci. Cloud*, Jun. 2010, DOI: 10.1145/1851476.1851533.

[18] R. Pereira, M. Azambuja, K. Breitman, and M. Endler, "An architecture for distributed high performance video processing in the cloud," in *Proc. IEEE 3rd Int. Conf. Cloud Comput.*, 2010, pp. 482–489.

[19] D. Qian, N. Cvijetic, Y.-K. Huang, J. Yu, and T. Wang, "100 km long reach upstream 36 Gb/s-OFDMA-PON over a single wavelength with source-free ONUs," in *Proc. Eur. Conf. Opt. Commun.*, 2009, pp. 1–2.

[20] Y. Tanaka, S. Yoshima, N. Kataoka, J. Nakagawa, N. Wada, and K. Kitayama, "100-Km uplink transmission of 10 G- and 1 G-ONU co-existing TDM-OCDMA-PON system using dual-rate burst-mode receiver," in *Proc. Conf. Opt. Fiber Commun.*, 2011, pp. 1–3.

[21] ITU-T PON standards of G. 987 (XG-PON).

[22] IEEE PON standards IEEE 802.3av (10G-EPON).

[23] O. Tipmongolsilp, S. Zaghloul, and A. Jukan, "The evolution of cellular backhaul technologies: Current issues and future trends," *IEEE Commun. Surv. Tut.*, vol. 13, no. 1, pp. 97–113, First Quarter 2010.

[24] G.-K. Chang, Z. Jia, J. Yu, and A. Chowdhury, "Super broadband optical wireless access technologies," *Proc. Conf. Opt. Fiber Commun./Nat. Fiber Opt. Eng. Conf.*, 2008, DOI: 10.1109/OFC.2008.4528441.

[25] Network Strategy Partners, *The Business Case for RAN Backhaul Using Mobile Transport Over Packet*, May 2009.

[26] K. Bergman, J. Mambretti, J. St Sauver, and B. Wing, "Networking Research Challenges Workshop Report," Large Scale Networking Coordinating Group, Networking and Information Technology Research and Development Program, (NITRD) (Department of Energy: Office of Science, National Science Foundation Directorate for Computer and Information Science and Engineering), 2008.

[27] T. Ndousse, N. Ghani, J. Mambretti, D. Petravick, B. Wing, N. Rao *et al.,* "Workshop report on advanced networking for distributed petascale science: Research challenges and opportunities," 2008.

[28] J. Mambretti and T. Aoyama, "Report of the Interagency Optical Networking Testbed Workshop 3," Large Scale Networking Coordination Group, Jul. 2007.

[29] J. Mambretti, "OptIPuter: Enabling advanced applications with novel optical control planes and backplanes," *Future Generat. Comput. Syst.*, vol. 25, no. 2, pp. 137–141, Feb. 2009.

[30] O. Ishida and A. Araki, "Challenging tera bit class LAN over wide area networks," *J. Lightw. Technol.*, vol. 27, no. 12, Jun. 15, 2009.

[31] S. Gringeri, B. Basch, V. Shukla, R. Egorov, and T. J. Xia, "Flexible architectures for optical transport nodes and networks," *IEEE Commun. Mag.*, vol. 48, no. 7, pp. 40–50, Jul. 2010.

[32] A. Bobyshev, S. Bradley, M. Crawford, P. DeMar, D. Katramatos, K. Shroff, M. Swany, and D. Yu, "A collaborative network middleware project by Lambda Station, TeraPaths, and Phoebus," *J. Phys., Conf. Ser.*, 062034, 2010.

[33] B. Gibbard, D. Katramatos, and D. Yu, "TeraPaths: End-to-end network path QoS configuration using cross-domain reservation negotiation," in *Proc. 3rd Int. Conf. Broadband Commun. Netw. Syst.*, 2006, pp. 1–9.

[34] S. Kobayashi, Y. Yamada, K. Hisadome, and O. Ishida, "Packet-based lane bundling for terabit-LAN," in *Proc. Asia-Pacific Conf. Commun.*, 2008, pp. 1–5.

[35] R. Schmogrow, M. Winter, B. Nebendahl, D. Hillerkuss, J. Meyer, M. Dreschmann, M. Huebner, J. Becker, C. Koos, W. Freude, and J. Leuthold, "101.5 Gbit/s real-time OFDM transmitter with 16 QAM modulated subcarriers," in *Proc. Opt. Fiber Commun. Conf.*, 2011, pp. 1–3.

[36] K. Fukuchi, D. Ogasahara, J. Hu, T. Takamichi, T. Koga, M. Sato, E. L. T. de Gabory, Y. Hashimoto, T. Yoshihara, W. Maeda, J. Abe, T. Kwok, Y. Huang, K. Hosokawa, Y. Yano, M. Shigihara, Y. Ueki, Y. Saito, Y. Nomiyama, K. Kikuchi, A. Noda, S. Shioiri, M. Arikawa, T. Wang, and T. Tajima, "112 Gb/s optical transponder with PM-QPSK and coherent detection employing parallel FPGA-based real-time digital signal processing, FEC and 100 GbE ethernet interface," in *Proc. Eur. Conf. Exhibit. Opt. Commun.*, 2010, DOI: 10.1109/ECOC.2010.5621114.

[37] H. Hu, H. Ji, M. Galili, M. Pu, H. C. H. Mulvad, L. K. Oxenlowe, K. Yvind,

J. M. Hvam, and P. Jeppesen, "320 Gb/s phase-transparent wavelength conversion in a silicon nanowire," in *Proc. Opt. Fiber Commun. Conf. Exhibit./Nat. Fiber Opt. Eng. Conf.*, 2011, pp. 1–3.

[38] A. Rohit, A. Albores-Mejia, N. Calabretta, X. J. M. Leijtens, D. J. Robbins, M. K. Smit, and K. A. Williams, "Fast remotely reconfigurable wavelength selective switch," in *Proc. Opt. Fiber Commun. Conf. Exhibit./Nat. Fiber Opt. Eng. Conf.*, 2011, pp. 1–3.

[39] R. Ryf, Y. Su, L. Moller, S. Chandrasekhar, L. Xiang, D. T. Neilson, and C. R. Giles, "Wavelength blocking filter with flexible data rates and channel spacing," *J. Lightw. Technol.*, vol. 23, no. 1, pp. 54–61, Jan. 2005.

[40] S. Frisken, G. Baxter, D. Abakoumov, H. Zhou, I. Clarke, and S. Poole, "Flexible and grid-less wavelength selective switch using LCOS technology," in *Proc. Opt. Fiber Commun. Conf. Exhibit./Nat. Fiber Opt. Eng. Conf.*, 2011, pp. 1–3.

[41] A. H. Gnauck, P. J. Winzer, S. Chandrasekhar, X. Liu, B. Zhu, and D. W. Peckham, "Spectrally efficient long haul WDM transmission using 224-Gb/s polarization-multiplexed 16-QAM," *J. Lightw. Technol.*, vol. 29, no. 4, pp. 373–377, Feb. 15, 2011.

[42] *General Principles and General Reference Model for Next Generation Networks*, ITU-T Y. 2011, Oct. 2004.

[43] T. D. Nadeau and A. Farrel, *IETF Draft: Generalized Multiprotocol Label Switching (GMPLS) Traffic Engineering Management Information Base*. [Online]. Available: http://tools.ietf.org/html/draft-ietf-ccamp-gmpls-te-mib-16

[44] *Telemangement Forum*. [Online]. Available: http://www.tmforum.org/browse.aspx

[45] J. Mambretti, T. DeFanti, and M. Brown, "StarLight: Next generation communication services, exchanges, and global facilities," in *Advances in Computers*, vol. 80. Amsterdam, The Netherlands: Elsevier, 2010, pp. 191–207.

[46] A. Farrel, J.-P. Vasseur, and J. Ash, *rfc4655: A Path Computation Element (PCE)-Based Architecture*. [Online]. Available: https://datatracker.ietf.org/doc/rfc4655/

[47] perfSONAR. [Online]. Available: http://www.perfsonar.net/

[48] J. Mambretti, D. Lillethun, and J. Weinberger, "Optical Dynamic Intelligent Network Services (ODIN): An experimental control plane architecture for high performance distributed environments based on dynamic lightpath provisioning," *IEEE Commun.*, vol. 44, no. 3, pp. 92–99, Mar. 2006.

[49] J. Mambretti, "Designing 21st century communications: Architecture, services, technology, and facilities *Networks for Grid Applications*, vol. 25. Berlin, Germany: Springer-Verlag, Sep. 2009, pp. 1–6, ser. Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering.

[50] E. Grassa, S. Figueroela, A. Fornsb, G. Junyenta, and J. Mambretti, "Extending the Argia software with a dynamic optical multicast service to support high performance digital media," *Opt. Switching Netw.*, vol. 6, no. 2, pp. 120–128, Apr. 2009.

[51] F. Travestino, J. Mambretti, and G. Karmous-Edwards, *GridNetworks: Enabling Grids With Advanced Communication Services*. Chichester, U.K.: Wiley, 2006.

[52] Web Services Resource Framework (WSRF) Technical Committee, *Organization for the Advancement of Structured Information Standards*. [Online]. Available: www.oasis-open.org

[53] OpenFlow. [Online]. Available: http://www.openflow.org/

[54] D. B. Hoang, T. Lavian, S. Figueira, J. Mambretti, I. Monga, S. Naiksatam *et al.*, "DWDM-RAM: An architecture for data intensive services enabled by next generation dynamic optical networks," in *Proc. IEEE Global Telecommun. Conf. Workshops*, Dec. 2004, pp. 400–409.

[55] Y. Gu, R. Grossman, and J. Mambretti, "A peer-to-peer infrastructure for distributing large scientific data sets over wide area high-performance networks," in *Proc. 1st Int. Conf. Netw. Grid Appl.*, 2007.

[56] T. DeFanti, C. De Laat, J. Mambretti, and B. St Arnaud, "TransLight: A global scale lambda grid for e-science," *Commun. ACM*, vol. 46, *Special Issue on Blueprint for the Future of High Performance Networking*, no. 11, pp. 34–41, Nov. 2003.

[57] Global Lambda Interconnecting Facility. [Online]. Available: www.glif.is

[58] T. DeFanti, M. Brown, J. Leigh, O. Yu, E. He, J. Mambretti, D. Lillethun, and J. Weinberger, "Optical switching middleware for the OptIPuter," *IEICE Trans. Commun.*, vol. E86-B, *Special Issue on Photonic IP Network Technologies for Next-Generation Broadband Access*, no. 8, pp. 2263–2272, Aug. 2003.

[59] J. Mambretti, M. Lemay, S. Campbell, H. Guy, T. Tam, E. Bernier, B. Ho, M. Savoie, C. de Laat, R. van der Pol, J. Chen, F. Yeh, S. Figuerola, P. Minoves, D. Simeonidou, E. Escalona, A. Gonzalez, A. Jukan, W. Bziuk, D. Kim, K. Cho, H. Lee, and T. Liu, "High Performance Digital Media Network (HPDMnet): An advanced international research initiative and global experimental testbed," in *Future Generation Computer Systems*. Amsterdam, The Netherlands: Elsevier, 2011.

[60] J. Mambretti, T. DeFanti, and M. Brown, "StarLight: Next generation communication services, exchanges, and global facilities," in *Advances in Computers*, vol. 80. Amsterdam, The Netherlands: Elsevier, 2010, pp. 191–207.

[61] SAGE. [Online]. Available: http://www.evl.uic.edu/cavern/sage/index.php

[62] O. Tipmongolsilp, M. Chamania, and A. Jukan, "Network-centric media composition under extreme bandwidth requirements for high quality video streaming," *ACM Supercomput.*, New Orleans, LA, Nov. 2010, (Poster).

[63] M. Sugawara, K. Masaoka, M. Emoto, and Y. Matsuo, "Research on human factors in ultrahigh-definition television (UHDTV) to determine its specifications," *SMPTE Motion Imaging J.*, pp. 23–29, Apr. 2008.

[64] CineGrid. [Online]. Available: http://www.cinegrid.org

[65] T. Fujii, "Recent progress of digital cinema over optical networks," in *Proc. Opt. Fiber Commun. Conf.*, 2005, DOI: 10.1109/OFC.2005.193169.

[66] T. Shimizu *et al.*, "International real-time streaming of 4K digital cinema," *Future Generat. Comput. Syst.*, vol. 22, no. 8, pp. 929–939, Oct. 2006.

[67] T. Fujii *et al.*, "Digital cinema over optical network-status of super HD development," in *Proc. Opt. Fiber Commun. Conf. Expo./Nat. Fiber Opt. Eng. Conf.*, 2011, pp. 1–3.

[68] H. Sakamoto, K. Minami, K. Shirakawa, T. Fujii, Y. Saito, and H. Yamane, "The '4K pure cinema' joint digital cinema trial," *NTT Tech. Rev.*, vol. 4, no. 7, pp. 18–22, Jul. 2006.

[69] Y. Yin, W. Zhang, H. Guo, T. Tsuritani, X. Hong, J. Wu, H. Jiang, X. Wang, X. Cao, L. Liu, T. Otani, and J. Lin, "Demonstration of QoS-classified survivability schemes in transparent OBS/GMPLS networks using streaming media transmission," in *Proc. Eur. Conf. Opt. Commun.*, 2009, pp. 1–2.

[70] D. Simeonidou, D. K. Hunter, M. Ghandour, and R. Nejabati, "Optical network services for ultra high definition digital media distribution," in *Proc. 5th Int. Conf. Broadband Commun., Netw. Syst.*, U.K., Sep. 2008, pp. 165–168.

[71] D. Tian and G. AlRegib, "Multistreaming of 3-D scenes with optimized transmission and rendering scalability," *IEEE Trans. Multimedia*, vol. 9, no. 4, pp. 736–745, Jun. 2007.

[72] N. Kataoka, K. Sone, K. Kaneko, T. Aoyama, K. Shimizu, T. Kawano, M. Maruyama, N. Wada, Y. Aoki, S. Kinoshita, H. Otsuki, H. Harai, H. Onaka, T. Miyazaki, and K.-i. Kitayama, "4K uncompressed streaming over multicast-capable80 $(8\lambda \times 10)$ Gbps colored optical packet switching network using SOA switch and stacked OC-label processing," in *Proc. Conf. Opt. Fiber Commun.*, 2009, pp. 1–3.

[73] T. Kasahara, K. Tokuhashi, K. Kikuta, D. Ishii, Y. Arakawa, S. Okamoto, and N. Yamanaka, "P2P HD contents transport network based on PLZT high-speed optical slot switching (OSS)," in *Proc. Int. Conf. Photon. Switching*, 2008, DOI: 10.1109/PS.2008.4804165.

[74] Y. Qin, G. Zervas, V. Martini, M. Ghandour, M. Savi, F. Baroncelli, B. Martini, P. Castoldi, C. Raffaelli, M. Reed, D. Hunter, R. Nejabati, and D. Simeonidou, "Service-oriented multi-granular optical network testbed," in *Proc. Conf. Opt. Fiber Commun.*, San Diego, CA, Mar. 2009, pp. 1–3.

[75] G. Zervas, L. Sadeghioon, D. Klonidis, Y. Qin, R. Nejabati, and D. Simeonidou, "Demonstration of novel multi-granular switch architecture on an application-aware end-to-end multi-bit rate OBS network testbed," in *Proc. Eur. Conf. Opt. Commun.* Germany, Sep. 2007, pp. 1–2.

[76] L. Hui, M. Takagi, H. Imaizumi, and H. Morikawa, "Preliminary demonstration of hybrid optical switching node with dynamic wavelength resource allocation using SOA switch," in *Proc. Opt. Electron. Commun. Conf.*, Jul. 13–17, 2009, DOI: 10.1109/OECC.2009.5222735.

[77] T. Miyazawa, H. Furukawa, K. Fujikawa, N. Wada, and H. Harai, "Partial implementation and experimental demonstration of an integrated optical path and packet node for new-generation networks," in *Proc. Opt. Fiber Commun. Conf.*, 2010, pp. 1–3.

[78] M. Takagi, H. Li, K. Watabe, H. Imaizumi, T. Tanemura, Y. Nakano, and H. Morikawa, "400 Gb/s hybrid optical switching demonstration combining multi-wavelength OPS and OCS with dynamic resource allocation," in *Proc. Opt. Fiber Commun. Conf.*, 2009, pp. 1–3.

[79] G. Zervas *et al.*, "Multi-granular optical cross-connect: Design, analysis and demonstration," *J. Opt. Commun. Netw.*, vol. 1, no. 1, pp. 69–84, Jun. 2009.

[80] J.-C. Gregoire and A. Jukan, "On the support of media functions within the IMS," in *IMS Handbook: Concept, Technology and Services*. Boca Raton, FL: CRC Press, 2009.

[81] G. Pallis and A. Vakali, "Insight and perspectives for content delivery networks," *Commun. ACM*, vol. 49, no. 1, pp. 101–106, Jan. 2006.

[82] M. Chamania, M. Caria, and A. Jukan, "Effective usage of dynamic circuits for IP routing," in *Proc. IEEE Int. Conf. Commun.*, South Africa, May 2010, DOI: 10.1109/ICC.2010.5501805.

[83] R. Pereira, M. Azambuja, K. Breitman, and M. Endler, "An architecture for distributed high performance video processing in the cloud," in *Proc. IEEE 3rd Int. Conf. Cloud Comput.*, 2010, pp. 482–489.

[84] EBU-TECH 3299, *High Definition (HD) Image Formats for Television Production*, Geneva, Jan. 2010.

[85] DVB BlueBook A151, *DVB Commercial Requirements for 3D-TV*, Jul. 2010.

[86] Digital Cinema Initiatives, *Digital Cinema System Specification, v1.2*, Mar. 7, 2008.

[87] Open House, *Experience the Future of Digital Broadcasting in Japan*, 2011.

[88] M. Kanazawa, K. Mitani, K. Hamasaki, M. Sugawara, F. Okano, K. Doi, and M. Seino, "Ultrahigh-definition video system with 4000 scanning lines," *SMPTE Mot. Imag. J.*, vol. 112, no. 10-11, pp. 339–346, Oct. 1, 2003.

ABOUT THE AUTHORS

Admela Jukan received the B.Sc. (Dipl.-Ing.) degree from the Faculty of Electrical Engineering and Computing (FER), University of Zagreb, Zagreb, Croatia, the M.Sc. degree in information technologies from the Politecnico di Milano, Milan, Italy, and the Ph.D. degree *(cum laude)* in electrical and computer engineering from the Technische Universitaet Wien, Vienna, Austria.

She is Chair Professor of Communication Networks at the Electrical and Computer Engineering Department, Technische Universitaet Carolo-Wilhelmina zu Braunschweig (Brunswick), Braunschweig, Germany. Prior to coming to Brunswick, she was a member of the research faculty at the Institut National de la Recherche Scientifique (INRS), University of Illinois at Urbana-Champaign (UIUC), and Georgia Tech (GaTech). From 2002 to 2004, she served as Program Director in Computer and Networks System Research at the National Science Foundation (NSF), Arlington, VA.

Dr. Jukan has chaired and cochaired several international conferences, including IFIP ONDM, IEEE ANTS, IEEE ICC, and IEEE GLOBECOM. She serves as an Associate Editor for the IEEE Communications Magazine and IEEE Network. She is co-Editor-in-Chief of the Elsevier *Optical Switching and Networking Journal*. She is elected Vice Chair of the IEEE Optical Network Technical Committee (Chair in 2014). She is a member of the HPDMnet initiative, and she also leads the European Union project ONE, focusing on coordinated network management of optical networks and the Internet. She was awarded an "IBM Innovation Award" for applications of parallel computing and optical networking for rich digital media services.

Joe Mambretti received the Ph.D. and MBA degrees from the University of Chicago, Chicago, IL, in 1977 and 1989, respectively.

He is Director of the International Center for Advanced Internet Research at Northwestern University, Evanston, IL, which is focused on developing digital communications for the 21st Century. The Center, which was created in partnership with a number of major high tech corporations (www.icair.org), designs and implements large-scale infrastructure and applications (metro, regional, national, and global). He is also Director of the Metropolitan Research and Education Network (MREN, http://www.mren.org), an advanced high-performance network interlinking organizations in seven upper midwest states. With its research partners, iCAIR has established multiple major network research testbeds to develop new architecture and technology for dynamically provisioned communication services and networks, including those based on lightpath switching. iCAIR has partnered with the University of Illinois at Chicago (UIC) to create StarLight (www.startap.net/starlight), an advanced global communications exchange based on leading-edge optical technologies. He is a Principal Investigator (PI) of the National Science Foundation (NSF)-funded International Global Environment for Network Innovations (iGENI) initiative, a PI for the national TeraFlow Network, which supports an Open Cloud computational science testbed. He is one of the PIs of HPDMnet, an international optical testbed, a Co-Director of the Open Cloud Consortium, a founding member of the Global Integrated Lambda Facility, a world wide distributed optical communications infrastructure.