EMERGING SONET-BASED SOLUTIONS FOR CABLE NETWORK EVOLUTION

Thomas Darcie, Eric Manning, Dale Shpak
University of Victoria
Nick Tzonev
Syscor R&D Inc

Abstract

Approaches for the design of cable transport networks to support emerging services are compared. While SONET is established firmly for wide-area network (WAN) transport, a diverse set of requirements drives a correspondingly diverse set of choices for metro (MAN) and secondary/hub networks. We evaluate the performance of next-generation (NG) SONET and contemporary alternatives in meeting these requirements, identifying strengths and weaknesses with both. Requirements or opportunities for new capabilities are identified and used as guidance for a cable-friendly third-generation SONET technology.

INTRODUCTION

Having made massive infrastructure investments in DOCSIS, 2-way plant, and digital video technology, the cable industry now holds an enviable position for the delivery of a broad range of emerging services. Yet with a growing pool of cable-modem and DTV subscribers, and the vast majority of homes passed or penetrated with broadband capability, impressively lucrative new-service revenues remain just out of reach. How can cable operators leverage the tremendous capability created by prior investments to improve return?

Popular directions for new revenue potential include various incarnations of voice, data, and digital video services. Each of the many possibilities present specific, often stringent, requirements for QoS (bandwidth, latency, delay, jitter, packet loss rate), and Wide-area or Metro-area Network (WAN or MAN) connectivity. In response, vendors have created a variety of transport network approaches, each of which has strengths and weaknesses, with no one approach serving all requirements well. Each approach presents a continually evolving progression of features.

A central point of debate within the networking community surrounds the efficacy of SONET and next-generation (NG) SONET technology, in comparison to evolving native Ethernet alternatives. Numerous business and technical factors have propelled a recent wave of enthusiasm for NG-SONET technologies. Once an expensive and static technology supporting only TDM services, improvements in cost, proven product stability, a pervasive embedded base, and data-friendly enhancements (RPR, GFP, VC, LCAS (described later)) have made Next-Generation SONET (NG-SONET) the vehicle of choice for major telecom service providers in both WAN and MAN. Perhaps the strongest argument in its favour is the ability to capitalize on the enormous existing embedded base of SONET technology.

However, the simplicity, flexibility, and utility of native Ethernet connectivity have driven its migration from access aggregation or local area network (LAN) to MAN. Special challenges like multicast video-on-demand have driven cable networks to transport-overlay alternatives (RPR, POS) or native Ethernet (GbE) in some portions of the network, but the resulting network may be less suitable for other high QoS services.

In this paper we discuss future directions for SONET, and how new, emerging SONET capabilities can address transport-network challenges within the cable industry. When we compare proposed networking approaches to delivery of key service opportunities, it becomes evident that NG-SONET plays a key role in providing transport for high-QoS and wide-area connectivity, while overlay techniques using other technologies are sometimes preferred for digital video delivery and local data aggregation. This unfortunate situation exacerbates network complexity, increasing both cost and technical risk.

Finally, we suggest enhancements to NG-SONET, defined collectively as third-generation (3G) SONET, that add additional networking capabilities within the transport network. An example is *SONET Frame Switching* (SFS) [Shpak], which provides dynamic provisioning, multicast distribution, statistical multiplexing, and various tiers of QoS support, while retaining compatibility with the existing SONET infrastructure and transparency to existing data formats. This capability can be added incrementally to legacy SONET rings by adding low-cost switch hardware and simple (transport) network control.

REFERENCE ARCHITECTURE

Recognizing that there is a diverse assortment of network topographies in use in modern cable systems, we base our discussion on the reference architecture shown in Figure 1. This captures the general notion of a hierarchy of interconnected networks leading from access aggregation (HFC) to a WAN network (regional, national, global) spanning independent systems.

SERVICE OPPORTUNITIES

To provide context for evaluating the suitability of various network approaches, we attempt to categorize voice, digital video, and data services that are most likely to have major market impact. We do this by service type, bandwidth, QoS requirements, MAN/WAN requirements, and other positive or negative attributes

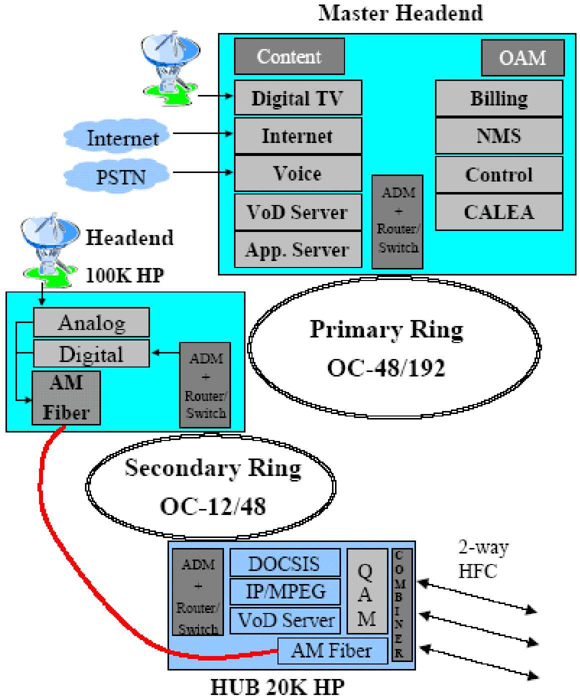

***Figure 1**: Reference architecture of hierarchical cable network with distribution and aggregation through HFC (fiber nodes, DOCSIS, QAM, etc.), MAN connectivity through hub ring, and MAN/WAN connectivity through head-end ring. Gateways to other territories and networks are through the master head-end.*

Voice

Voice services fall into various classes, depending on distribution (Voice-over-cable or DOCSIS) and network (PSTN, Internet, or managed IP network) methodologies. Voice-over-Cable (TDM-based voice modems) provides complete alignment with legacy PSTN approaches, but overlooks the opportunity to make good use of the features of DOCSIS and the PacketCable initiative. We consider here the continued requirement to transport the resulting aggregated DSO/T1-level streams through MAN and MAN. This is referred to as Voice-Circuit-PSTN (or **Voice-CP**). We do not consider the use of this circuit-based voice distribution approach within Internet or managed IP networks.

Voice over DOCSIS can be supported by all three network approaches, with or without invoking DOCSIS 1.1 QoS features. We consider only the specific important case of a high-QoS voice service that is desired through a managed IP network. This approach, referred to as Voice-DOCSIS-Managed IP (**Voice-DM**), is a next step towards higher quality that goes beyond the models of existing VoIP-over-cable providers (since DOCSIS 1.1 QoS capabilities are used.)

Video

Digital video services can be categorized as *broadcast* or *demand- based*. For delivery of broadcast digital video, a high channel count (e.g., 200) is distributed from a master head end to each head end, as illustrated in Figure 1. This is then converted to QAM channels and broadcast over analog fiber to hubs and distribution. Since this is a) business as usual and b) does not affect digital traffic on the hub ring, it is not considered further.

Demand-based video can be defined by the location of the server or cache (hub, head end, or master head end) and the transport or distribution method (MPEG/QAM or DOCSIS/QAM). The analysis is not sensitive to the distribution method. However, since the aggregate bandwidth in the hub network can become very large, especially as servers are pushed to the head end, cost per unit bandwidth becomes a key parameter.

We assume that the tradeoff between server and transport cost pushes the servers towards the hub for popular content (e.g. top 10 hits or channels) and towards the master head end for customized or archival content. [A large amount of research on *content networking* [Cameron] examines ways to optimize the placement of content based on access patterns and frequency.] To cover this tradeoff, we consider three alternatives for the delivery of demand-based video. The first (Video Centralized, or **Video-C**) involves a centralized (master head end) VOD server delivering streams of video to individual customers over the IP-MPEG/QAM path from the hub. The second uses a VOD server located in the hub (Video Distributed, or **Video D**) to provide transport-efficient service to all subscribers sharing the hub. The third (Video Head-end, or **Video H**) takes the middle ground, placing servers in each head end to serve all subscribers sharing the hub network

Data

Data services include a large set of applications, bandwidths, service quality attributes, and MAN/WAN requirements. DOCSIS-based consumer-oriented Internet access (Data - Broadband Access, **Data BA**), represents the existing service as supported by most cable providers today. The industry seeks to exploit this base and expand into other opportunities. These include, in order of increasing sensitivity to delay and jitter: VLAN (including Ethernet Virtual Private LAN - EVPLan), Network Attached Storage (NAS), Ethernet Virtual Private Line (EVPL), Storage Area Network (SAN), Ethernet Private Line (EPL), and traditional T1. EVPLan and EVPL represent the primary (in terms of likely demand) services recently defined by the Metro Ethernet Forum (MEF) [Clavenna1, Clavenna2]. It is important to note that these new Ethernet services need not be delivered by an Ethernet network: customer ports need only have Ethernet interfaces.

While numerous other definitions could be explored, we feel that these seven types span a representative set of likely requirements. These will be abbreviated to **Data-XXX** in what follows.

Many of the services discussed are supported on Layer 3 (L3) router-based networks. This adds another layer through which performance must be maintained. For the purposes of this transport-layer study, we recognize that these services require QoS from lower layers end-to-end, and do not consider the additional complexities of router evolution.

Data-VLAN represents the class of Virtual LAN services offered typically across LAN or MAN, based on the IEEE 802.1q specification. With DOCSIS as the access vehicle, such services are clearly an attractive complement to existing broadband services. Included here are Ethernet LAN Services (E-LAN) as defined recently by the MEF.

Data-T1 represents the vast majority of business services provided by the ILECs to provide Internet and PBX connectivity, and contributing roughly $23 B in high-margin revenue (In-Stat/MDR Research). A DOCSIS–based end-to-end means to attack this market would be a valuable capability.

Data-NAS describes remote data storage on centralized or shared network-based media. Peak transfer data rates can be high, while latency is not a prime consideration. Excessive latency is often an irritant rather than a failure to deliver the contracted service.

Data-SAN describes high-performance networked storage, typically using a dedicated network (e.g., Fiber Channel) for this purpose. Although these networks commonly exist in data centers, it is often desired to extend the SAN for remote backups or for servicing remote locations. To retain high performance, the network must have low latency and negligible packet loss.

In an attempt to define a data-centric evolution for T1 service, **Data-EVPL** specifies various service qualities, such as Committed Information Rate (CIR) and Excess Information Rate (EIR). This would support a combination of real time (voice) and data requirements across MAN and WAN. Also supported by EVPL is a wide class of interactive real-time applications (real-time interactive gaming or music, video conferencing/telephony, etc.), in which only a portion of the bandwidth must be allocated with high QoS.

Finally, **Data-EPL** specifies a committed rate over a dedicated channel (essentially Ethernet over a SONET channel), effectively providing TDM-like guarantees, but over a data network. Data-EPL service is much like Data-T1 but has Ethernet interfaces.

In summary, we have defined two classes of voice, three video, and seven data services, for a total of twelve spanning a wide variety of network requirements. All appear to be within the opportunity space available to cable service providers. We now explore the networking technologies that can be considered in support of these services.

NETWORK TECHNOLOGY CHOICES

A daunting assortment of evolving network choices is available. We summarize the main approaches in Figure 2. Included in "other" are those approaches that have specific but narrowly-defined applicability, (e.g., Ethernet over SONET, Fiber Channel).

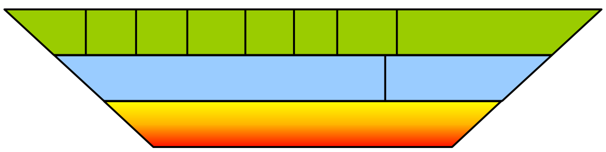

Figure 2: Layering of transport and overlay layer-2 networks on wavelength-division multiplexed (WDM) transmission.

Transport Layer

SONET is the backbone of telecom networks. SONET provides high reliability, QoS, scalability, and mature OAM capabilities. Accordingly, many carriers prefer to extend the lifetime of their SONET infrastructure.

Next-Generation SONET (NGS) is currently defined as the combination of Generic Framing Procedure (GFP), Virtual Concatenation (VCAT), and Link Capacity Adjustment Scheme (LCAS). These technologies enhance SONET/SDH by adding flexibility to payload types and by supporting finer-grained provisioning of SONET channels. GFP is a simple framing method that can be used to encapsulate many different types of data. VCAT improves channel utilization by enabling finer-grained provisioning of point-to-point channels. LCAS provides a mechanism for on-the-fly modification the capacity of these provisioned channels.

Although Next-Gen SONET improves SONET utilization, it is still a point-to-point technology, and therefore requires numerous point-to-point channels when interconnecting multiple nodes. Each of these channels must be provisioned to support the peak data rate. Since the channels are not shared, it does not provide any statistical multiplexing gain to improve utilization when transporting bursty traffic.

In recent years, ***Optical Ethernet (OE)*** has been proposed as a transport technology. The main argument supporting Optical Ethernet is that the large deployment of Ethernet in the enterprise helps to drive down the price of components. Ethernet has several attractive features, such as the ability to rapidly provision bandwidth with fine granularity and simple internetworking with enterprise networks.

However, there are several reasons why Ethernet has not made major inroads into carrier networks. Firstly, carriers already have substantial investments in their SONET/SDH and ATM networks. They would often prefer to improve the performance of these networks rather than cap their existing networks and grow a completely new transport network. Much of the Ethernet traffic that is carried over transport networks is not transported natively: it is encapsulated in SONET/SDH. There are other concerns about using Ethernet in transport networks, such as standardized support for end-to-end QoS (perhaps using MPLS), rapid protection switching (non-standard proprietary approaches with interworking and/or performance issues), mature OAM capabilities, and scalability [MEF].

Overlay Networks

The most common methods for transporting IP data over SONET/SDH are Packet over SONET (PoS), Asynchronous Transfer Mode (ATM), Resilient Packet Ring (RPR), Frame Relay (FR), and Ethernet. All these technologies have strengths and weaknesses that make them suitable for some applications and less suitable for others.

Packet over SONET (PoS) is a long-serving method for transporting packet data over SONET. A major weakness of PoS is that it uses provisioned point-to-point channels between each of the network elements. Because of the point-to-point nature of PoS, the interconnection of multiple nodes requires a mesh consisting of a provisioned channel between every node. This reduces network flexibility as each of these SONET channels must be provisioned to meet peak bandwidth requirements and there is limited opportunity to exploit statistical multiplexing to improve bandwidth utilization over the individual channels.

Frame Relay (FR) is a very important commercial packet service that is somewhat like X.25 with cut-through forwarding. The nodes in a FR network switch packets over provisioned paths in the network (Permanent Virtual Circuits). A FR network may also support Switched Virtual Circuits that can be established on demand. With FR, each unidirectional link is established having a guaranteed or Committed Information Rate (CIR) (which may be zero) and a Peak Information Rate (PIR). Since FR uses shared channels, it has statistical multiplexing gain.

An advantage of FR over VPNs is that a FR connection is truly private: no TCP/IP ports are exposed to hackers.

Asynchronous Transfer Mode (ATM) and SONET are the major transport technologies used by telcos. ATM can transport many different types of packet- and circuit-oriented traffic types. Notably, nearly all DSL traffic is aggregated using ATM-based DSLAMs. ATM uses a small, fixed cell size that results in two important hardware advantages: the resulting hardware has a simple and regular queue structure and ATM can have a low forwarding latency. However, ATM has a connection-oriented protocol that is not efficient for small data transfers that are typical for some Internet communications. It does not scale well and it is not well suited for carrying transient or best-effort traffic. It is more appropriate for establishing network services requiring higher levels of QoS (mainly voice). Other problems with ATM are high cost and excessive signaling overhead.

Resilient Packet Ring (RPR) was designed to address many of the shortcomings of PoS. RPR introduces oversubscription into SONET. With RPR, the nodes are addressable, can share the capacity on a SONET ring, and any node can insert traffic into the shared channel. This helps to overcome the poor channel capacity utilization in PoS that results from having to use dedicated point-to-point channels. In addition, RPR supports differentiated QoS which can be used to offer tiered Service Level Agreements (SLA). RPR was designed to be media-independent and can currently be carried over SONET or Ethernet. Therefore, it requires its own protection switching mechanism. RPR's most serious weakness is that it cannot natively span multiple rings. Typically, a router is used to interconnect RPR rings.

Ethernet: In a typical overlay deployment, Ethernet switches are connected over point-to-point SONET links by mapping Ethernet frames into the SONET payload.

Each Ethernet switch in the overlay network must be capable of determining the required output port for Ethernet frames destined to *any* Ethernet address *anywhere* on the Ethernet network. Unlike RPR nodes, which simply forward packets that are not destined for that node, each Ethernet switch must resolve the destination port for every Ethernet frame that arrives on any of its interfaces. The advantage of this additional processing is that an Ethernet switch has the capability to forward traffic any output port, thereby giving it the capability to interconnect rings or meshes (unlike RPR).

As mentioned above, there are issues with the scalability of Ethernet, including the speed of network recovery when using spanning tree protocols.

TECHNOLOGY SUITABILITY

In evaluating the alternative networking approaches, we compare key attributes of the transport-layer and overlay networks, for each of the 12 classes of service we have defined, and for applicability in WAN and MAN networks.

First, each service class has different degrees of tolerance for network impairments. Table 1 illustrates the tolerance to network outages, delay and/or jitter, and packet loss. There is a wide variation in tolerance for each impairment, ranging from low for Data-SAN, Data-EPL, Voice CP, and Data T1, to high for best effort data service (Data- BA).

Table 1: Tolerance to Network Impairments

	Outages	Delay/Jitter[1]	Loss
Voice-CP	Low	Low	Low
Voice-DM	Low	Mod	Low
Video-C	Low	Mod	Mod
Video-H	Low	Mod	Mod
Video-D	Mod	Mod	Mod
Data-BA	High	High	High[2]
Data-VLAN	Mod	High	High[2]
Data-T1	Low	Low	Low
Data-NAS	Mod	High	Mod[2]
Data-SAN	Low	Low	Low[2]
Data-EVPL	Mod	Mod	Mod[2]
Data-EPL	Low	Low	Low
1. Jitter is often accommodated by adding delay (jitter buffer) 2. Assumes higher-layer error recovery (TCP)			

These network impairments are mapped onto the three transport layer alternatives (SONET, NG SONET, and Ethernet) on Table 2. Also included are transport efficiency and scalability, characteristics that do not affect service directly, but are of direct concern to the service provider. Scalability is the ability to migrate to large networks and high speed, particularly across WAN. SONET has the clear advantage in providing protected, low-delay, low-jitter, and low-loss service that is scalable across WAN, but with low transport efficiency. NG SONET takes a substantial step in increasing this efficiency. Depending on the particular approach, Ethernet restoration time can be large or small, but difficulties can be encountered in scalability.

Table 2: Transport Network Characteristics

	SONET	NG SONET	Ethernet
Restoration Time	Low	Low	Low-High[1]
Delay	Low	Low	Mod
Jitter	Very Low	Very Low	High
Loss	Very Low	Very Low	Mod
Transport Efficiency	Low	Mod	Mod
Scalability	High	High	Low-Mod
1. Spanning Tree, Fast Spanning Tree, Proprietary Faster Methods			

Overlay Networks

Table 3 shows the strengths and weakness of the various overlay network approaches within the MAN. The point-to-point and static use of transport bandwidth limits the applicability of POS. ATM excels for most services, but has cost and scalability challenges. FR is not generally applicable for high bandwidth (video or broadband data) broadcast or multicast services. RPR satisfies a broad set of MAN requirements, but its inability to span rings limits applicability in WAN. Ethernet has broad applicability in MAN, but struggles or fails for delay or jitter sensitive services and services requiring high availability.

Table 3: Suitability of Overlay Technologies for MAN

	PoS	ATM	FR	RPR	Ethernet
Voice-CP	No[1]	Yes[3]	OK[1]	OK[1,3]	OK[1,3]
Voice-DM	No[2,4]	Yes	OK[4]	Yes	Yes
Video-C	No[4,5]	No[7,8]	No[5]	No[9]	Yes
Video-H	No[4,5]	No[7,8]	No[5]	No[9]	Yes
Video-D	No[4,5]	No[8]	No[5]	Yes	Yes
Data-BA	No[2,4]	Yes	No[4]	Yes	Yes
Data-VLAN	No[2,4]	Yes	Yes	Yes	Yes
Data-T1	No[1]	Yes[3]	Yes[3]	OK[1,3,9]	OK[1.3]
Data-NAS	No[4]	OK[8]	OK[4]	OK[9]	Yes
Data-SAN	No[1]	OK[8]	OK[1]	OK[1,9]	No[1]
Data-EVPL	No[A]	Yes	Yes	Yes	Yes
Data-EPL	No[2]	Yes	Yes	OK[1,]	OK[1]
1. Excessive latency or jitter 2. Poor network utilization 3. Requires emulation 4. Inherently point-to-point	5. No multicast 6. Channel is not shared 7. Scalability is an issue 8. Cost is an issue	9. Cannot span rings A. No QoS mechanism			

A similar evaluation of overlay alternatives for WAN reflects a combination of the performance of the overlay approach in aggregating traffic and the capabilities of the underlying transport network. The inherent point-to-point nature of PoS is again a limitation, as is the inability of RPR to span rings. Ethernet's inability to provide hard QoS and limited scalability are exacerbated in WAN.

FUTURE DIRECTIONS

Several conclusions are clear from the preceding discussion.

- No one solution fits all requirements.
- SONET and NG SONET, in conjunction with the appropriate overlay technology, are the means by which WAN requirements will be met now and in the future.
- Ethernet (overlay and transport) has a strong place in MAN, but must invoke L3 for WAN.

Unfortunately, QoS mechanisms for L3 WAN data services remain a key challenge. Complexity of various alternatives (e.g. MPLS, flow-based routing) continues to plague economics, operations and general acceptance of these solutions.

Therefore, we seek future directions for SONET and NG SONET that accomplish the following:

- **Increase the transport efficiency** for data, particularly in MAN (as per Table 3), while concurrently supporting low-latency TDM traffic. This requires increased flexibility in mapping different services into large (OC48-192) and small (OC3-12) streams, the ability to overprovision best-effort traffic on top of guaranteed, and broadcast/multicast support.
- Provide an efficient means (unlike PoS) to create **multipoint connections across MAN and WAN**. This requires increased flexibility in defining the connectivity between interconnected SONET nodes, preferably on a frame-by-frame basis.
- Provide a transport capability that dramatically **simplifies the challenge of QoS-controlled routing**. This requires sufficient flexibility, in terms of both bandwidth allocation and addressing, to be built into the SONET layer.
- Enable **monitoring, policing, and billing of rate- or usage-based service level agreements.**

An example of an approach to 3rd-Generation (3G) SONET that attempts to provide these features is SONET Frame Switching (SFS) [Shpak]. SFS introduces two main components:

1) a novel frame-based protocol, and
2) a network switching architecture that uses segmented, hierarchical MAC addresses.

A frame-oriented protocol is utilized in order to take advantage of SONET framing, rather than using an octet-stream-based protocol that requires additional synchronization to locate packet boundaries. Rather than layering legacy protocols on top of SONET, SFS uses a connectionless MAC-layer protocol. Hierarchical addressing provides minimum byte overhead while enabling straightforward switching and topology discovery. SFS is presented in the context of a ring-based architecture but other topologies such as linear or mesh can be readily realized.

To L3 protocols, all nodes connected through SFS appear to be on a contiguous L2 network (the SFS "cloud"). This allows for routing decisions to be made only at the ingress nodes. The traffic can then be switched between multiple SFS nodes without making further routing decisions. To a router connected to the SFS network, any other router connected to the SFS network appears to be separated by a single routing hop. SFS can be readily connected to legacy networks by extracting the IP packets and forwarding them to an existing router or, if the SFS node has a router, by making the routing decision within the SFS node and forwarding packets through one or more legacy-compatible ports on the SFS node.

SUMMARY

In summary, with its pervasive embedded base, unrivaled timing precision, and restorability, SONET provides the ideal foundation for transport of emerging cable services. NG-SONET and emerging 3rd-generation SONET capabilities like SONET Frame Switching (SFS) can provide the efficiency and flexibility previously derived from Ethernet and Layer-3 networks, but with the timing performance and reliability expected from SONET. It is anticipated that these capabilities will result in substantial simplification of cable networks and reduce dependence on complex layer-3 alternatives.

REFERENCES

[Cameron] C. Cameron. S. H. Lo, and D. Wei, "A Server Placement Algorithm for Content Distribution", Pro. IEEE Information Theory Workshop, 2002, pp. 26-28

[Clavenna1] S. Clavenna, "Standardizing Ethernet Services", Light Reading, Jan 9, 2004.

[Clavenna2] S. Clavenna, "Future Ethernet Standards Unveiled", Light Reading, Dec 8, 2003.

[MEF] "Metro Ethernet Networks A Technical Overview," Metro Ethernet Forum, 2003.

[Shpak] D.J. Shpak, P.-C. Tang, and E.G. Manning, "Supporting everything-on-demand in cable services," *19th National Fiber Optics Engineers Conference*, Orlando, Sept. 2003.