

Value-Added Services in Next-Generation SONET/SDH Networks

Nasir Ghani and Qing Liu, University of New Mexico

Ashwin Gumaste, IIT Bombay

John Lankford, Ciena

Abdallah Shami, University of Western Ontario,

Chadi M. Assi, Concordia University

Ahmad Khalil, City University of New York

Driss Benhaddou, University of Houston

ABSTRACT

Advances in next-generation SONET/SDH have introduced novel features for generic protocol framing/encapsulation, virtual concatenation, inverse multiplexing, dynamic circuit adjustment, and so on. In turn, these provisions have enabled much improved multi-tiered service provisioning and are viewed very favorably by carriers, particularly incumbents. This article looks at this evolved framework with a particular focus on value-added services creation. Results from a sample performance evaluation study also are presented to quantify some of the achievable gains.

INTRODUCTION

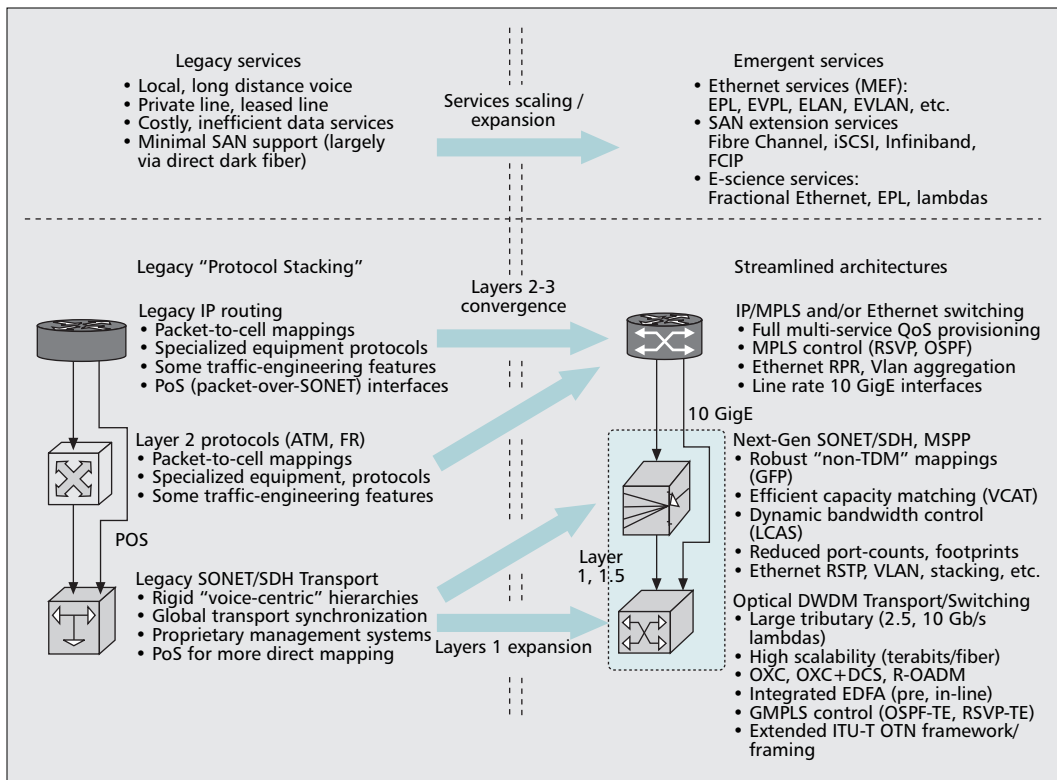
Sustained traffic growth has led to many advances in circuit-switched technologies. Foremost among these, optical dense wavelength division multiplexing (DWDM) has yielded unparalleled bandwidth-distance scalability in the core, with current systems supporting over 100 channels/fiber. Related control standards — such as Internet Engineering Task Force (IETF) generalized multi-protocol label switching (GMPLS) and International Telecommunication Union-Telecommunication (ITU-T) automatically-switched transport network (ASTN) — also have matured [1]. However, DWDM is generally not suitable for the edge owing to the prevalence of “sub-wavelength” demands, for example, Ethernet-based services, storage area network (SAN) extension, and legacy private leased line (PLL) [2]. Hence, advances in next-generation synchronous optical network/synchronous digital hierarchy (SONET/SDH NGS) time-division multiplexing (TDM) have provided much-needed multi-protocol grooming features here [1–4]. Indeed, NGS has become the premiere optoelectronic solution for fiber backbones and has redefined traditional transport hierarchies (Fig. 1).

NGS provides key enhancements for dynamic circuit provisioning and improved traffic mapping. These advances offer significant potential for supporting new value-added services cost-effectively over legacy and emergent infrastructures. Overall, NGS has gained strong traction with incumbents — owing to its inherent compatibility with the installed base — and most fiber backbones support overlying NGS grooming layers [1]. Hence, researchers are seeking to leverage this technology to develop improved multi-tiered services, particularly Ethernet-based [3, 5–9] services.

This article focuses on the application of NGS technology for value-added service creation and is organized as follows. First, we present a brief overview of the NGS framework. Subsequently, evolutions in various market sectors are detailed — corporate, residential, and e-science — and the applicability of NGS for improved multi-tiered service provisioning is surveyed. A sample study is then presented along with final thoughts and conclusions.

NEXT-GENERATION SONET/SDH: A BRIEF OVERVIEW

Legacy TDM was primarily built for voice and PLL services and leveraged intermediate layer 2 asynchronous transfer mode (ATM) and frame-relay (FR) systems for overlaying data services (Fig. 1). However, as data demands outpaced legacy services, the limitations of these multilayering set ups became apparent [2]. Even though various solutions were developed to improve Internet Protocol (IP)-TDM interfacing, for example, packet over SONET (PoS) [3], these techniques suffered from high bandwidth inefficiencies via rigid mappings to “next-highest” TDM carrier. Moreover, all-or-nothing SONET/SDH protection proved very problematic for diversified multitiered services. To address



Overall, the GFP-VCAT combination delivers significant improvements in TDM timeslot efficiency and port reduction, well beyond legacy centralized back-hauling. Additionally, inverse multiplexing and LCAS are very timely for improving resiliency.

Figure 1. Network technology and services evolutions.

these problems, new NGS standards evolved, most notably, generic framing procedure (GFP G.7041), virtual concatenation (VCAT G.707), and link capacity adjustment scheme (LCAS G.7042) [2, 3].

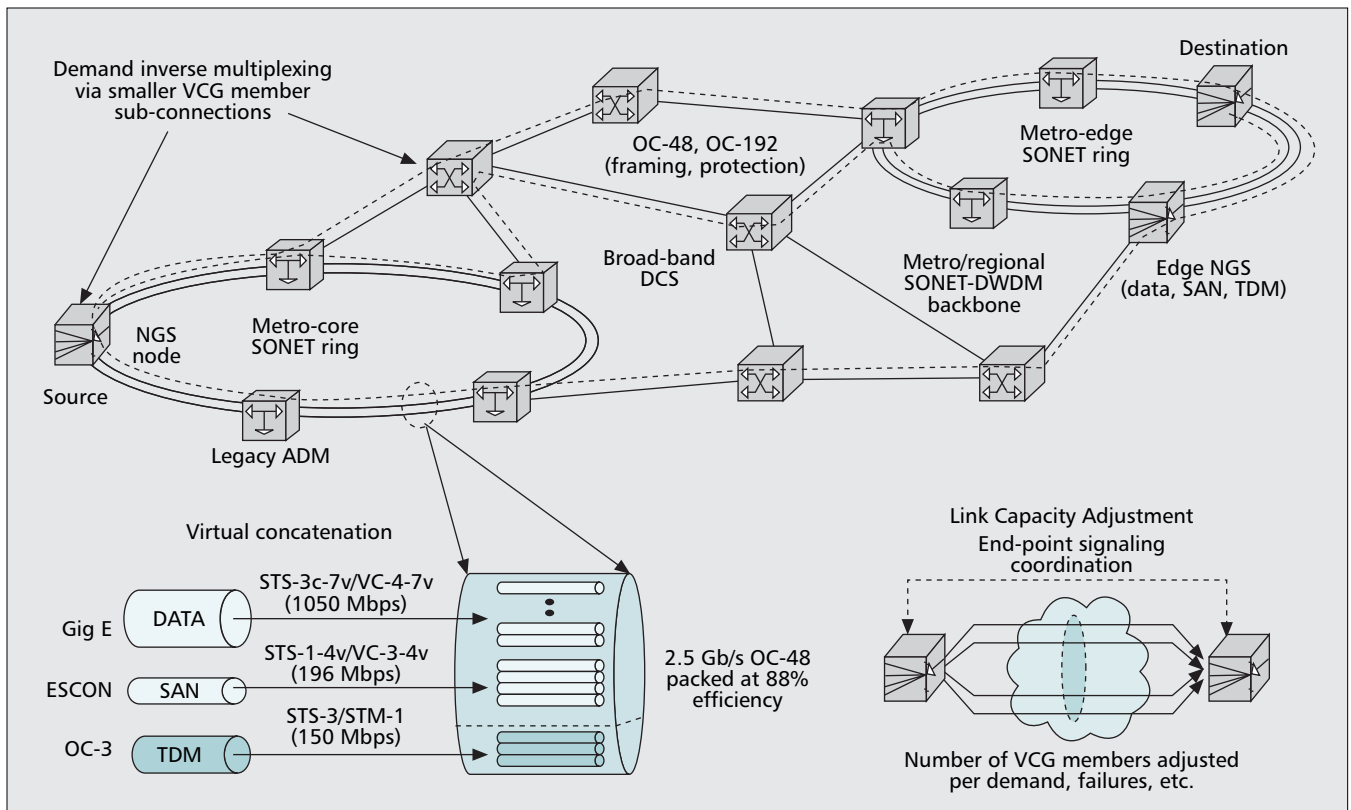
GFP provides efficient mappings of diverse protocols directly onto byte-synchronous TDM optical carrier-n (OC-n)/ synchronous transport mode (STM-n) channels, greatly improving data plane efficiencies [3]. This scheme uses robust error-controlled frame delineation (like asynchronous transport mode [ATM]) and supports two payload mappings, frame (GFP-F) and transparent (GFP-T). The former yields deterministic overheads for Ethernet medium access control (MAC) or IP packets, whereas the latter transparently maps 8b/10b encoded payloads (Fiber Channel, Enterprise Systems CONnection [ESCON], fiber connectivity [FICON]) with minimal packetization/buffering delays. Meanwhile, VCAT addresses the inherent tributary mismatch of legacy SONET/SDH by combining multiple slower-speed time-slots to "right-size" end-user tributaries, namely, 1.54 Mb/s VT1.5, 48.38 Mb/s STS-1, or 155 Mb/s STS-3c increments. SONET/SDH protection and operations administration and management (OAM) features are also extendible to these tributaries, yielding carrier-class management of data/SAN services. Although newer Ethernet line/path OAM standards are being developed [2], wide-scale adoption will take time, and related costs are unknown.

Another key VCAT provision is inverse multiplexing [3, 8], which allows OC-n and/or concatenated tributaries to be split into multiple sub-connections (Fig. 2). These entities consti-

tute a virtual concatenation group (VCG), where each VCG member can be separately provisioned through multipath routing over legacy SONET/SDH or NGS networks using add-drop multiplexers (ADMs) or digital cross-connects (DCS). These streams are recombined at the receiver using buffering to resolve multipath delays. As per [4, 9], current standards can handle up to 128 ms delay for up to 64 VCG members, which is adequate for global distances. In all, inverse multiplexing can achieve very good bandwidth efficiency in mesh networks.

Finally, the LCAS protocol [3] complements VCAT with "hitless" VCG trail readjustment — ideal for dynamic/time-varying or asymmetric demands. Namely, VCG end-points use two-way signaling to synchronize the addition/removal of channels from a VCG. LCAS is well-suited for designing multitiered services as it can provision pre/post-fault switchovers on a per-VCG member basis. For example, the VCG sink can monitor incoming members and notify the source of any trail failures within 64–128 ms, providing near SONET/SDH-like timescales [8]. Upon failure notification, the source can take various actions, for example, initiate protection for failed VCG members (high-end services), initiate slower signaling restoration for the failed VCG members (mid-tier services), or simply let the connection run at a lower rate (graceful degradation).

Overall, the GFP-VCAT combination delivers significant improvements in TDM timeslot efficiency and port reduction, well beyond legacy centralized back-hauling [1]. Additionally, inverse multiplexing and LCAS are very timely for improving resiliency and will help minimize



■ **Figure 2.** VCAT inverse multiplexing operation in mixed SONET-DWDM networks.

higher-layer IP routing table and/or Ethernet spanning tree disruptions. Hence, NGS has gained significant traction with incumbents, allowing them to recoup legacy investments and still resolve deficiencies in data/SAN support. In turn, these saliciencies are yielding many operational expenditure (OPEX) and capital expenditure (CAPEX) reductions. Particularly, with regard to the latter, only selected premise upgrades are required without costly changes to legacy core/management systems. Finally, GFP mappings also are compatible with newer ITU-T optical transport unit (OTU) framing formats, enabling smooth integration with future optical standards. As these market offerings continue to mature, the focus has shifted to building new service models.

SERVICES OVERVIEW

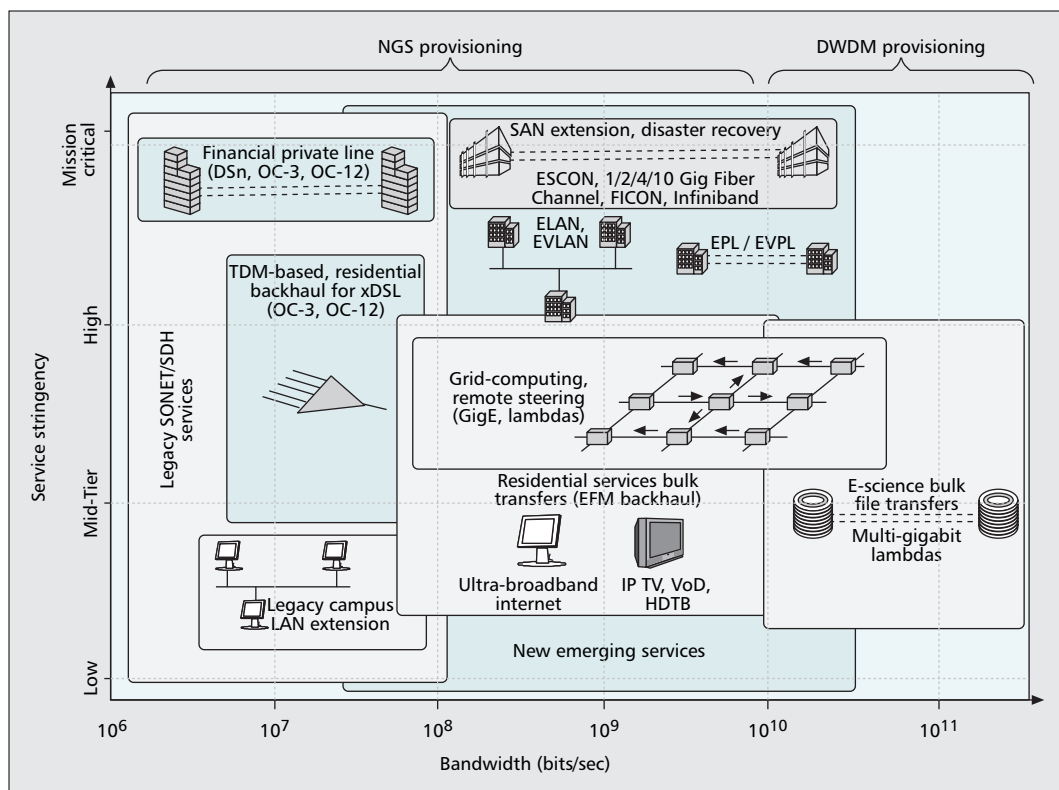
Network operators today are facing wide-ranging services growth across key sectors, including corporate (small, medium, large businesses), residential, and e-science (Fig. 3). Although the applications are different in each of these markets, the underlying requirements have many commonalities. Foremost, there is a continual push for bandwidth scalability and reduced price-per-bit through CAPEX and OPEX reduction [1, 2]. Equally important is the desire to deliver flexible multitiered services to offset the continued price-per-bit declines of legacy services. Specifically, carriers would like to support advanced service-level agreements (SLAs) [2] for diversified applications with quantifiable quality of service (QoS) parameters such as bandwidth,

delay/jitter, survivability, and so on. For example, delay and recovery requirements for mission-critical financial services are generally much tighter than those for campus LAN extension.

CORPORATE SERVICES

Legacy services are comprised of trunked voice and PLL — DS-3, T1, T3, and OC-*n* — with the latter primarily being used to overlay ATM and frame-relay devices for IP data services. Today, PLL still represents a sizeable, although declining, portion of commercial traffic. Nevertheless, with the massive proliferation of IP/Ethernet-based applications, there are now concerted efforts to extend Ethernet over *wide-area* domains by carrier-class features [2]. The main challenges here are to resolve the low reliability, slow protection, and lack of latency/packet loss guarantees of enterprise Ethernet. Hence, new data interfaces were developed supporting long-reach optics up to 100 km (10 Gigabit Ethernet). In addition, improved Ethernet address scalability, management, and wide-area bridging (IEEE 802.1a/d/h) standards also emerged.

For clients, a key requirement is the definition of new Ethernet service models. Here, the Metro Ethernet Forum (MEF) has tabled three broad categories — namely Ethernet private line (EPL) and Ethernet (private) LAN (E-LAN) and Ethernet tree (E-Tree) — along with associated user-network interface (UNI) standards [2]. EPL is based upon point-to-point Ethernet virtual connection (EVC) support and is a generalization of legacy PLL for point-to-point data interconnectivity. This service formalizes packet-level SLA parameters, such as committed infor-



■ **Figure 3.** Network services and applications.

The residential market has undergone numerous changes as carriers continue broadband rollouts and plan for future ultra-broadband capabilities. A defining trend here is bundled triple-play services, combining high-speed Internet, voice, and video services at reduced price points.

mation rate (CIR), committed burst size (CBS), frame delay/jitter, priority, and so on. Furthermore, Ethernet virtual private line (EVPL) is a variation of EPL that performs edge multiplexing at the UNI. A major evolution here is the ability to specify *fractional* 100 BaseT, Gigabit Ethernet, and 10 Gigabit Ethernet rates. This is very appealing for cost-sensitive small- and medium-sized businesses wanting to purchase incremental capacity without costly hardware upgrades. Moreover, carriers can resell EPL connectivity between corporate sites or as underlying bearer service for higher revenue offerings, such as packet voice (voice-over-IP), packet video (video conferencing), managed IP-based virtual private networks (IP-VPN), and so on. Note that the ITU-T Study Group 15 also has approved similar EPL (G.8011.1) and Ethernet virtual private line (EVPL, G.8011.2) standards, defining services from a carrier's perspective.

Meanwhile, newer E-LAN and E-Tree services support multipoint connectivity between dispersed corporate sites. Associated end-user costs are drastically lower because enterprises do not require specialized ATM or PoS interfaces. Akin to EVPL services, Ethernet virtual private LAN (EVP-LAN) and Ethernet virtual private tree (EVP-tree) services also allow multiple client LAN entities (virtual LAN identifiers) to share a multipoint E-LAN/E-Tree connection. At the SONET/NGS level, these services can be implemented using point-to-point EPL connections between switching end points, for example, either mesh or tree. In all, Ethernet services likely will challenge PLL dominance by the end of the decade.

In addition to Ethernet services, more spe-

cialized commercial services also will emerge. For example, mission-critical storage technologies have been widely adopted in the financial community. These solutions use specialized protocols (ESCON, Fibre Channel) and are driving growth in *SAN extension* for disaster recovery, data center consolidation, data mirroring, and so on [15]. Newly legislated business continuity requirements (e.g., Sarbanes-Oxley in the United States) also are pushing SAN extension trends. Currently, most SAN devices/fabrics use full-rate 1 or 2 Gb/s line-rate Fibre Channel interfaces, with newer standards supporting 4 and 10 Gb/s speeds. These interfaces have extremely stringent QoS and reliability requirements, with stated bit-error rates (BER) as low as 10^{-12} . In general, most of these interfaces can be directly mapped onto TDM tributaries using NGS GFP. Note that there is also much focus on IP-based storage, such as block-level Internet small computer system interface (iSCSI) and fiber channel over IP (FCIP), and these standards can use EPL bearer services.

RESIDENTIAL SERVICES

The residential market has undergone numerous changes as carriers continue broadband rollouts and plan for future ultra-broadband capabilities. A defining trend here is bundled triple-play services, combining high-speed Internet, voice, and video services at reduced price points. To deliver these packages in a cost-effective manner, providers are actively migrating toward all-IP access. Hence, Ethernet is again being adapted to provide scalable bandwidth delivery over a wide range of media: cable, copper, fiber, and even air. For example, Ethernet in the first mile

Another market experiencing rapid growth is e-science bandwidth services for large governmental and research organizations. For example, collaborative research in areas such as high-energy physics and astronomy is increasing the need for remote information distribution between dispersed sites.

(EFM IEEE 802.3ah) initiatives are actively developing copper and fiber solutions. EFM-copper standards use digital subscriber line (xDSL) schemes, for example, asymmetric (A)DSL, ADSL2, very high bit-rate (V)DSL, and so on, to send Ethernet frames over (bonded) twisted-pair lines. Meanwhile, EFM-fiber schemes are looking at new schemes for point-to-point and point-to-multipoint fiber plants. In particular, the IEEE 802.3ah multipoint control protocol (MPCP) supports gigabit-level connectivity in Ethernet passive optical networks (EPON). Note that the ITU-T G.984 Gigabit PON (GPON) standard also defines PON capabilities for faster 2.5 Gb/s SONET/SDH OC-48/STM-16 links. In all, EFM technologies are of key interest to incumbents competing with cable operators in the content delivery market. Meanwhile, cable operators have deployed data over cable service interface specification (DOCSIS) to support residential services. Currently, DOCSIS 2.0 has gained strong traction, supporting up to 20 Mb/s per user (although operational speeds are limited to smaller values), and the newer DOCSIS 3.0 standard is promising even higher speeds through channel bonding.

As these converged access infrastructures mature, client applications are being packetized rapidly. A foremost example is voice-over-IP (VoIP), which is eroding legacy fixed-voice services at a sustained, rapid pace. Furthermore, packet video represents a major growth segment in the residential sector with new applications such as TV/cable distribution and playback, home-office teleconferencing, remote learning, gaming, and so on. For carriers, these applications will propel new data services between edge content distribution sites. Hence, EPL and E-LAN provide a strong fit as underlying bearer services for residential backhaul over metro/wide-area domains (Fig. 3). These types can provide full packet voice/video QoS support along with flexible tributary resizing to match time-varying aggregates. More importantly, SLA prioritization can ensure high-availability/uptime for (emergency) voice services. Note that many carriers use TDM-based backhaul (OC-3, OC-12 PLL) for residential xDSL aggregation, which can readily be supported via NGS.

The previous discussions have focused largely on *wireline* services. Although detailed considerations for wireless markets are beyond the scope of this article, a myriad of new and emerging services therein (particularly video offerings) also are expected to increase metro core EPL backhaul demands.

E-SCIENCE SERVICES

Another market experiencing rapid growth is e-science bandwidth services for large governmental and research organizations. For example, collaborative research in areas such as high-energy physics, fusion, astronomy, and climate modeling is increasing the need for remote information distribution between dispersed sites. Here many new e-science applications are emerging, including large-scale file transfers, grid computing, remote steering and visualization, data mining, and so on. [14]. More importantly, associated data scales are far outpacing

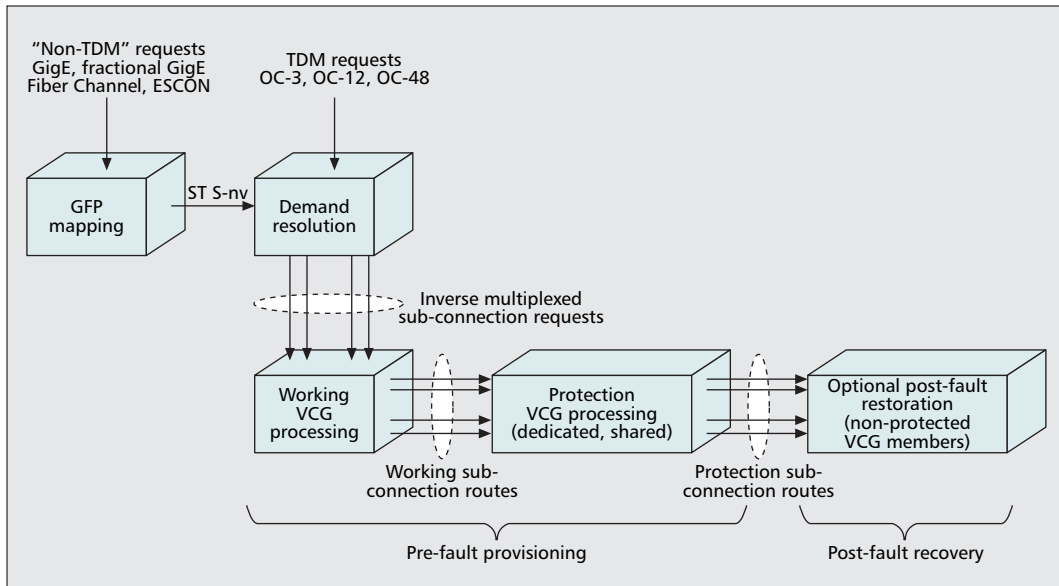
commercial demands, with current datasets ranging from terabytes to petabytes. Hence, research organizations are actively building their own dedicated infrastructures, many of them based upon underlying DWDM transport (Fig. 3).

Given that most scientific data originates in IP/Ethernet format, EPL services are very germane. Here NGS provides a good means to rapidly overlay a wide range of full/fractional rate data services — from 1 Mb/s all the way to 10 Gb/s — particularly over TDM or optical transport network (OTN)-framed optical trunks or lambdas. For increased scalability, larger TDM fabrics also can be used to provision coarser timeslots, for example, OC-3 (155 Mb/s), thereby simplifying VCG management. Along these lines, researchers are already experimenting with large-scale dataset mappings over TDM tributaries using specialized transport-layer protocols such as Hurricane, Tsunami, Real-time Transport Protocol (RTP), and so on (Ultra ScienceNet study [14]).

TIERED SERVICES SUPPORT

NGS offers many benefits for carriers to support a wide range of client services — data, SAN, legacy TDM — and compares favorably with alternative IP/multiprotocol label switching (MPLS) or Ethernet-based circuit emulation/pseudo-wire technologies [3]. Although some of these technologies can match SONET/SDH timescales, for example, MPLS fast reroute (FRR) and IEEE 802.17 resilient packet ring (RPR), these layer 2/3 offerings mandate costly core upgrades and careful overlay design/placements. Hence NGS-based multi-tiered services are starting to receive much focus. For example, the Optical Internetworking Forum (OIF) recently demonstrated its UNI 2.0 and network node interface (NNI) 1.0 standards in conjunction with multi-vendor operation support system (OSS) systems to support fractional EPL circuits using GFP and LCAS [13]. From a more algorithmic perspective, various studies were also conducted. For example, multipath inverse multiplexing routing strategies for SONET-DWDM networks are presented in [6]. The results here show that inverse multiplexing gives good blocking reduction, 10–15 percent lower than regular (non-inverse multiplexed) routing, and can compensate for reduced grooming inside the core. Meanwhile, [7] studies several iterative heuristics for multipath VCAT routing, namely, shortest-path first (SPF), widest-path first (WPF), and max-flow (MF). The findings here show very competitive performances between the schemes, with SPF performing best at higher loads. Also, [8] introduces the cumulative differential delay routing (CDDR) problem to minimize delays between VCG members and shows it to be NP-complete. A path pre-computation solution is developed to lower sink-side buffering — a key benefit for higher-speed links.

However, the more interesting aspect of NGS is its ability to improve survivability because inverse multiplexing offers much more flexibility versus legacy SONET protection, such as 1+1/1:1 span protection, bi-directional line-switched ring (BLSR), and so on. Herein, researchers also developed some advanced sub-



■ **Figure 4.** Tiered service survivability scheme for NGS networks.

However, the more interesting aspect of NGS is its ability to improve survivability because inverse multiplexing offers much more flexibility versus legacy SONET protection, such as 1+1/1:1 span protection, bi-directional line-switched ring (BLSR), and so on.

connection survivability schemes. Earlier, [9] proposed several low-overhead protection schemes for Ethernet over SONET (PESO) to increase VCG resiliency, namely, PESO α , β , γ schemes. These strategies try to achieve sufficient multipath diversity between VCG members so as to ensure sufficient immunity to single-link failures. Nevertheless, detailed performance results are not presented. More recently, [10] also tabled a similar strategy for achieving degraded service-aware bandwidth provisioning across multiple VCG paths. Specifically, sub-connection routing is performed to ensure that no one path carries more than a given fraction of the total flow. Multipath load distribution also is performed to minimize the maximum incremental link utilization using integer linear programming (ILP). The results here show good improvements in blocking and load balancing. However, the above schemes do not explicitly provision back-up sub-connections and hence are susceptible to reduced topological connectivity. For example, if only three link-disjoint routes are available between a source-destination pair, the PESO α and β schemes cannot guarantee less than 33 percent VCG traffic disruption for a single link failure.

Meanwhile, others have studied more direct sub-connection protection and restoration strategies. For example, [11] proposes two inverse multiplexing shared-protection schemes. The first approach, protecting individual VCG member (PIVM), allows back-up capacity sharing between link-disjoint VCG members, whereas the second approach, provisioning fast restorable VCG (PREV), only allows sharing between link-disjoint VCG members with the same source-destination. As expected, PIVM gives much higher efficiency, although recovery is generally slower and requires complex per-link VCG member conflict state. Hence, this scheme is only amenable to centralized implementations because GMPLS routing does not provide connection-level state. Meanwhile, PREV provides faster recovery because switchover routes are

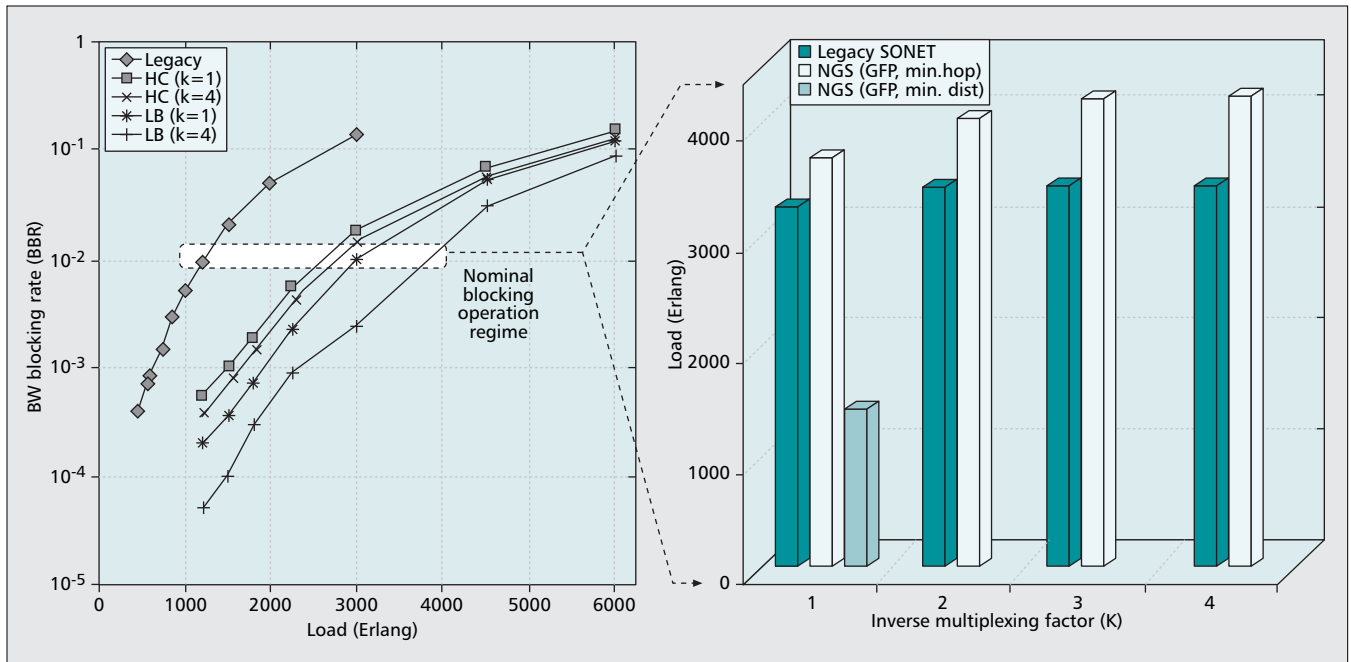
known in advance but requires complex min-cost flow pre-calculation. However, this scheme is more susceptible to networks with reduced connectivity/path diversity. Finally [12] develops a new *effective multipath bandwidth* metric that takes into account both link bandwidth and availability constraints. Furthermore, two multipath routing heuristics are developed to achieve desired availability levels, with both showing significant improvements versus single path provisioning strategies.

Although the above contributions represent a good start, additional avenues exist for improving multi-tiered service models. In particular, it is important to fully leverage inverse multiplexing for more selective *per-VCG member* recovery, and this is quite germane for a competitively-priced, low-mid range full/fractional EPL/EVPL service that can sustain degraded throughputs during outages. In addition, *post-fault* VCG restoration also can be used to boost low mid-range service reliability and efficiency owing to the lower delay stringency of data traffic. These issues are now considered.

SAMPLE PERFORMANCE STUDY

The survivability scheme in [15] leverages NGS features to build new services with differentiated, that is, tiered, levels of recovery. This algorithm uses shortest-path/ k -shortest path heuristics and readily can be integrated into the constraint-based routing (CBR) engines [4] of centralized OSS or distributed GMPLS control planes. From a high-level perspective, the scheme defines a protection factor, ρ ($0 \leq \rho \leq 1$), and routes at least ρn STS-1 units of back-up capacity for a n STS-1 demand. Namely, incoming requests are inverse multiplexed, and a minimum subset of working VCG sub-connections are protected (shared, dedicated) until this threshold is reached. Varying ρ here allows carriers to achieve tiered protection levels, and the overall scheme consists of several steps (Fig. 4).

Foremost, incoming STS-1 mapped requests



■ **Figure 5.** Mixed PLL and Gigabit Ethernet services with full protection (1 percent BBR).

are “evenly” split over K VCG sub-connections to generate multiple VCG member requests, where K is the *inverse multiplexing factor*. Although this value can be optimized per a given criterion (e.g., maximize load, minimize blocking), this is generally an NP-complete problem [7]. Hence, in practical settings, most carriers likely will pre-specify K in order to limit VCG set-up complexity. Next, working VCG sub-connection routes are sequentially computed by running Dijkstra’s shortest-path algorithm over the available capacity graph using either a minimum hop count or minimum distance link cost metric (the latter metric weights a link proportional to the inverse of its free capacity for load-balancing [15]). If a working route cannot be found for all VCG sub-connections, the request is dropped, otherwise the algorithm proceeds to the protection phase (Fig. 4). Note that additional metrics can be added for delay constraints on VCG members, as per [8].

The protection phase cycles through all the computed working VCG routes and attempts to route a *diverse dedicated protection* route for each. A running count of the aggregate routed protection capacity of working VCG members is checked after each successful protection sub-connection computation. If this count exceeds pn STS-1 units, the request is successful; otherwise it is blocked. The overall compute complexity here is $O(KNL)$ for a network with N nodes and L links. Overall, this scheme guarantees tiered recovery (pn STS-1 units) for single-fault events, as well as graceful degradation for multiple-fault events. Furthermore, *post-fault* restoration also can be used to reroute failed, non-protected VCG members, possibly yielding full (100 percent) bandwidth recovery (Fig. 4). Essentially such restoration allows carriers to offer lower-cost services with improved recovery without having to price in full protection guarantees. However, post-fault per-VCG restoration likely

entails notable path computation and signaling complexities and/or delays. Note that future studies also can consider shared protection between VCG members for improved resource efficiencies [9]. Nevertheless, associated compute complexities and bookkeeping overhead are much higher as one must consider sharing among VCG members in the same connection and also in different connections. This also complicates distributed GMPLS implementations.

The multitiered NGS scheme is tested using coding custom models in the *OPNET Modeler™* environment. Tests are performed using the National Science Foundation Network (NSFNET) topology comprising 16 nodes and 25 bi-directional links (node degree 3.125), with all nodes having STS-1 (50 Mb/s) switching granularity and OC-192 (10 Gb/s) link rates. Furthermore, connections have exponentially distributed holding and inter-arrival times, with the mean holding time set to 600 hours (scaled) and the mean inter-arrival time adjusted according to load. First, full protection is tested for mixed legacy PLL and EPL demands. The former comprise SONET OC-3 (155 Mb/s), OC-12 (622 Mb/s), and OC-48 (2.5 Gb/s) requests. Meanwhile, the latter comprise Gigabit Ethernet and fractional Gigabit Ethernet demands (uniform from 50 Mb/s to 1.0 Gb/s in 50 Mb/s STS-1 increments). All connection types are chosen with equal probability (20 percent) yielding an average request size of about 19 STS-1 units (Gigabit Ethernet resolved to 21 STS-1 increments via GFP). For the case of legacy SONET operation, however, all full/fractional Gigabit Ethernet demands are mapped to OC-48 tributaries.

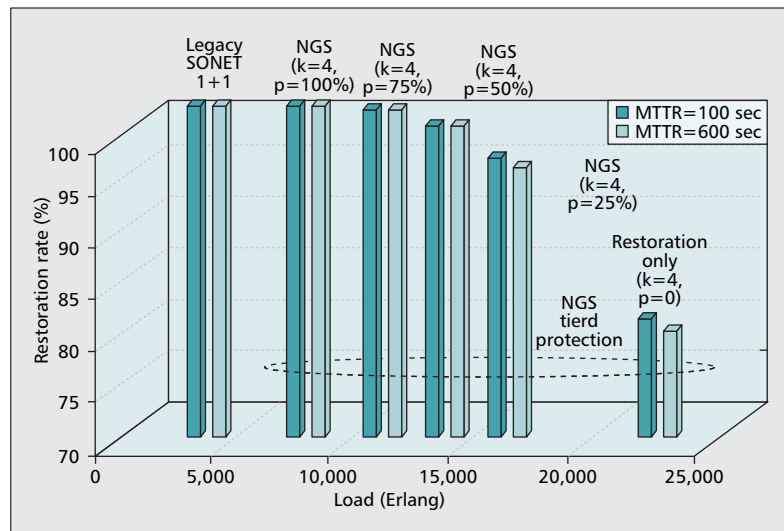
Figure 5 summarizes two key results for this full-protection scenario — bandwidth blocking rate (BBR) versus load and carried load for a normalized BBR of 1 percent. (Note that the BBR metric is a modified version of connection blocking that takes into account connection size

[7]). The latter plot is very germane for carriers as it represents the revenue-generation capability at a nominal “low-blocking” operating point. The findings clearly indicate that NGS inverse multiplexing provides much better performance over a wide load range and carries almost 250 percent higher load versus legacy SONET (1 percent BBR). More importantly, demand splitting yields very significant improvement versus non-inverse multiplexed GFP-only mappings and is particularly effective when coupled with the load-balancing metric. For example, minimum distance routing shows almost an order of magnitude blocking reduction (versus hop count without inverse multiplexing) and a 20 percent increase in carried load at 1 percent BBR (i.e., 3600 Erlang for $K = 1$ versus 4270 Erlang for $K = 4$ [Fig. 5]).

In general, load balancing routing is very beneficial for inverse multiplexing as it tends to distribute sub-connection routes over the network and limits excessive congestion at bottleneck links. Additional tests for lower node degrees (below 2.5) also yield decent blocking reductions. Finally, the results in Fig. 5 show that inverse multiplexing factors over $K = 4$ have minimal additional benefits for *full* protection, allowing carriers to limit overall management complexity. Note that the above gains are seen even with predominantly TDM line-rate demands (only 40 percent Gigabit Ethernet or fractional Gigabit Ethernet) and the inclusion of other “non-TDM” demands (Fiber Channel, ESCON) likely will yield higher efficiencies.

Next, tiered survivability is tested for fractional EPL services, as these services are expected to gain significant traction in coming years. Namely, different protection thresholds are tested ($\rho \leq 1$) in conjunction with *post-fault* restoration (Fig. 4) to gauge carried load and connection recovery rates for single-link failures. In the event of a link failure, all non-protected VCG members traversing a failed link are identified and sub-connection rerouting is performed for each. The goal here is to achieve *full* recovery for partially protected services. All link failures are randomly generated using exponentially distributed inter-arrival times with mean 600 hours. Here, two different mean-time-to-repair (MTTR) values are tested: 100 and 600 hours (exponential). These times are typically measured in hours, that is, truck roll repairs, and are much larger than SONET/SDH switchover times.

To stress restoration recovery, performance is gauged using the connection recovery rate for carried loads up to 20 percent BBR, namely, high loads. This metric is defined as the percentage of failed connections that fully recover 100 percent of their throughput after a link failure. Note that for inverse multiplexed connections, this implies recovery of all failed non-protected VCG members. The results are shown in Fig. 6 for minimum distance routing. As expected, the restoration rate decreases with load and increases with larger protection factors. Here, full protection via legacy 1+1 SONET or NGS with $K = 4/\rho = 1$ provides 100 percent recovery, but carried loads are much lower. Conversely, pure restoration ($K = 4, \rho = 0$) gives the highest load carried load — almost five times above legacy



■ Figure 6. Fractional EPL services recovery under high loading (up to 20 percent BBR).

1+1 SONET and two times above NGS dedicated protection — but notably lower recovery rates, 75–80 percent. These values translate to an availability of 98.56 percent across all connections (link failures only). Meanwhile, tiered/partial NGS protection delivers a very good trade-off between these two extremes. For example, 50 percent protection ($\rho = 0.50$) gives over 97.8 percent recovery for both MTTR values — roughly 20 percent higher than regular restoration — and approximately 50 percent higher load versus full VCAT protection. The corresponding availability here is 99.91 percent, even at high 20 percent BBR. Such notable increases in restoration rates/availability (10–20 percent) represent very tangible values in terms of service pricing and will enable carriers to broaden their offerings. These gains are a direct result of increased multipath diversity as connections are less susceptible to single VCG failures. Note that restoration times average in the 100s of milliseconds (assuming 10–20 ms node processing delays for recomputation and signaling).

Overall, the above findings confirm that NGS can greatly reduce resource overbuild and circumvent the all-or-nothing protection of legacy SONET/SDH. This is particularly attractive when provisioning EPL services with generally lower cost-per-bit pricings. Namely, carriers can deploy tiered NGS survivability to recoup their margins, thereby achieving high CAPEX and OPEX reductions.

CONCLUSIONS

The bandwidth services market continues to evolve, and carriers are actively deploying new technologies with improved provisioning capabilities. In particular, next-generation SONET/SDH NGS offers much promise as a genuine multi-tiered enabler with high scalability, multi-protocol transparency, and service flexibility. More importantly, NGS is fully compatible with entrenched infrastructures, enabling carriers to rapidly expand new service footprints with minimal infrastructure outlays. Results from sample

Results from sample NGS studies indicate that indeed it can provision a wide range of service types with high service efficiency and availability. These saliciencies provide many new avenues for improved carrier revenue-generation.

NGS studies indicate that indeed it can provision a wide range of service types with high service efficiency and availability. These saliciencies provide many new avenues for improved carrier revenue-generation.

ACKNOWLEDGMENTS

This research has been supported in part by the United States National Science Foundation (NSF) Computer and Information Sciences Division (CISE) under Award 0806637.

REFERENCES

- [1] N. Ghani *et al.*, "On IP-WDM Integration: A Retrospective," *IEEE Commun. Mag.*, vol. 41, no. 9, 2003, pp. 42–45.
- [2] A. Kasim, Ed., *Delivering Carrier Ethernet: Extending Ethernet Beyond the LAN*, McGraw Hill Publishers, Nov. 2007.
- [3] E. Hernandez-Valencia, "Hybrid Transport Solutions for TDM/Data Networking Solutions," *IEEE Commun. Mag.*, vol. 40, no. 5, May 2002, pp. 104–12.
- [4] G. Bernstein *et al.*, "VCAT/LCAS in a Clamshell," *IEEE Commun. Mag.* (standards report), vol. 44, no. 5, May 2006, pp. 34–36.
- [5] D. Zhou and S. Subramaniam, "Survivability in Optical Networks," *IEEE Network Mag.*, vol. 14, no. 6, Nov.–Dec. 2000, pp. 16–23.
- [6] R. Srinivasan and A. Somani, "Dynamic Routing in WDM Grooming Networks," *Photonic Network Commun.*, Mar. 2003.
- [7] K. Zhu *et al.*, "Ethernet-over-SONET (EoS) over WDM in Optical Wide-Area Networks (WANs): Benefits and Challenges," *Photonic Network Commun.*, vol. 10, no. 1, Jan. 2005, pp. 107–8.
- [8] A. Srivastava, "Flow Aware Differential Delay Routing for Next-Generation Ethernet over SONET/SDH," *IEEE ICC 2006*, Istanbul, Turkey, June 2006.
- [9] S. Acharya *et al.*, "PESO: Low Overhead Protection for Ethernet Over SONET Transport," *IEEE Infocom 2004*, Hong Kong, Mar. 2004.
- [10] R. Roy and B. Mukherjee, "Degraded-Service-Aware Multipath Provisioning in Telecom Mesh Networks," *IEEE/OSA OFC 2008*, San Diego, CA, Feb. 2008.
- [11] C. Ou *et al.*, "Survivable Virtual Concatenation for Data over SONET/SDH in Optical Transport Networks," *IEEE/ACM Trans. Networking*, vol. 14, no. 1, Feb. 2006, pp. 218–31.
- [12] S. Rai, *et al.*, "Reliable Multipath Provisioning for High-Capacity Backbone Mesh Networks," *IEEE/ACM Trans. Networking*, vol. 15, no. 4, August 2007, pp. 803–12.
- [13] J. Jones *et al.*, "Interoperability Update: Dynamic Ethernet Services via Intelligent Optical Networks," *IEEE Commun. Mag.*, vol. 43, no. 11, Nov. 2005, pp. S4–S10.
- [14] N. Rao *et al.*, "Ultra Science Net: Network Testbed for Large-Scale Science Applications," *IEEE Commun. Mag.* (OCS), vol. 43, no. 11, Nov. 2005, pp. S12–S17.
- [15] N. Ghani and S. Park, "Multi-Tiered Service Survivability in Next-Generation SONET/SDH Networks," *Photonic Network Commun.*, vol. 13, no. 1, Jan. 2007, pp. 79–92.

BIOGRAPHIES

NASIR GHANI [S'95, M'97, SM'01] (nghani@ece.unm.edu) received a B.S. in computer engineering from the University of Waterloo, Canada, in 1991, an M.S. in electrical engineering from McMaster University, Canada, in 1992, and a Ph.D. in electrical and computer engineering from the University of Waterloo, Canada, in 1997. Currently, he is an associate professor in the Department of Electrical and Computer Engineering at the University of New Mexico, where he is actively involved in a wide range of funded research projects. He has gained a wide range of industrial and academic experience in the networking area and has held senior positions at Nokia, IBM, Motorola, Sorrento Networks, and Tennessee Tech University. He has authored over 100 publications (including journal and conference papers, book chapters, industry magazine articles, and standardization drafts) and has two patents granted. He is currently a co-chair of the IEEE INFOCOM High-Speed Network (HSN) series and has served as a co-chair for the optical networking symposia for IEEE ICC 2006 and IEEE GLOBECOM 2007. He is also a program committee mem-

ber for IEE/OSA OFC and numerous other conferences and has served regularly on NSF, DOE, and other international panels. He is an associate editor of IEEE Communications Letters and has guest edited special issues of IEEE Network, IEEE Communications Magazine, and Cluster Computing. He is a recipient of the NSF CAREER Award.

QING LIU received a B.E. degree in telecom engineering and an M.E. degree in computer science from the Nanjing University of Post and Telecommunications, China, in 2001 and 2004, respectively. He is currently completing a Ph.D. degree in the Department of Electrical and Computer Engineering at the University of New Mexico, Albuquerque, USA. In the past, he has worked as a summer intern at Motorola China and Acadia Optronics LLC (USA), as well as a software engineer at Bell Labs/Lucent Technologies China. His research interests include optical networks, multi-domain networks, traffic engineering, multi-media communications, and wireless networks. He has authored over 15 publications (including six journal articles and one book chapter).

ASHWIN GUMASTE is currently a faculty member in the Department of Computer Science and Engineering at the Indian Institute of Technology, Bombay, and a visiting scholar at the Massachusetts Institute of Technology. He was previously with Fujitsu Laboratories (USA) Inc. as a member of research staff in the Photonics Research Laboratory (2001–2005). Prior to this he worked at Fujitsu Network Communications R&D, and prior to that with Cisco Systems in the Optical Networking Group. He has over 40 pending U.S. and EU patents, and has published close to 60 articles in refereed conferences and journals. He has authored three books on broadband networks: *DWDM Network Designs and Engineering Solutions* (a networking bestseller), *First-Mile Access Networks and Enabling Technologies* (Pearson Education/Cisco Press), and *Broadband Services: User Needs, Business Models and Technologies* (Wiley). He is also an active consultant to industry, and has worked with both service providers and vendors. In addition, he has served as Program Chair, Co-Chair, Publicity Chair, and Workshop Chair for various IEEE conferences, and has been a Technical Program Committee member for IEEE ICC, IEEE GLOBECOM, IEEE Broadnets, IEEE ICCCN, Gridnets, and so on. He has also been a Guest Editor for *IEEE Communications Magazine* and was General Chair of the First International Symposium on Advanced Networks and Telecommunication Systems, held in Bombay, India.

JOHN LANKFORD has over 15 years' experience working with IP, Ethernet, SONET, and DWDM technologies in the networking industry. Currently he is a senior systems engineer with Ciena. Prior to this he was with Force10 Networks, and prior to that he was the senior network architect at the University of Tennessee. He is also the chief engineer of the NSF funded International Research Network Project Global Ring Network for Advanced Application Development (GLORIAD). Since 1998 his primary research interest has been the design of large-scale multitiered IP over optical networks. Currently his focus is on programmable service delivery architectures in multiservice multilayer gigabit-edge terabit-core networks.

ABDALLAH SHAMI [S'01, M'03] received a B.E. degree in electrical and computer engineering from the Lebanese University, Beirut, Lebanon, in 1997, and a Ph.D. degree in electrical engineering from the Graduate School and University Center, City University of New York, New York, in February 2003. Since July 2004, he has been with the University of Western Ontario, London, Ontario, where he is currently an assistant professor in the Department of Electrical and Computer Engineering. From 2002 to 2004, he was an assistant professor in the Department of Electrical Engineering at Lakehead University, Ontario, Canada. His current research interests are in the areas of network protocol design and analysis for access wireless and optical networks. In 2002, he received the Irving Hochberg Dissertation Fellowship Award at the City University of New York. He serves as an associate editor for IEEE Communications Letters and as an associate editor for the Journal of Computer Systems, Networks, and Communications. He served as a panel co-chair for AccessNets 2007 and as a symposium co-chair for ChinaCom 2008.

CHADI M. ASSI [S'01, M'03, SM'08] received his B.Eng. degree from the Lebanese University, Beirut, Lebanon, in 1997 and a Ph.D. degree from the Graduate Center, City

University of New York, New York, in April 2003. He is an associate professor with the Concordia Institute for Information Systems Engineering at Concordia University, Montreal, Canada. His current research interests are in the areas of optical networks, multihop wireless and ad hoc networks, and security. Before joining the Institute in August 2003 as an assistant professor, he was a visiting scientist for one year at Nokia Research Center, Boston, working on quality-of-service in optical access networks. He received the prestigious Mina Rees Dissertation Award from the City University of New York in August 2002 for his research on wavelength-division-multiplexing optical networks. He is on the editorial board of the IEEE Communications Surveys and Tutorials, serves as an associate editor for the IEEE Communications Letters and also an associate editor for Wiley's Wireless Communications and Mobile Computing.

AHMAD KHALIL received a B.E. degree in mechanical engineering (Lebanese University, 1997), an M.S. degree in mechanical engineering (Ohio University, 2000), an M. Phil. degree in electrical engineering (The Graduate Center/CUNY, 2002), and a Ph.D. degree in electrical engineering (The Graduate Center/CUNY, 2005). He is an associate professor in the Mathematics, Engineering, and Computer Science department at LaGuardia College (CUNY). In July 2008, he joined the Ethernet Engineering team at Tata Communications. He is currently leading the Ethernet Network Architecture group, and his main respon-

sibility is to design, architect, and optimize the L2 network. His current research interests are in the areas of data/optical networking, traffic grooming and multicast in optical networks, EPON, and ad hoc networks. During his studies at the CUNY Graduate Center, he was awarded the Robert E. Gilleece Fellowship from 2000 to 2004, one of the most prestigious fellowships at CUNY.

DRISS BENHADDOU received a B.S. in physics (electronics) from University Abdelmalek Esaadi, Morocco. He received an M.S. in optoelectronics from the University of Montpellier II, France, in 1991. He holds doctoral degrees in optoelectronics and telecommunication engineering, from the University of Montpellier II in 1995, and from the University of Missouri in 2002, respectively. He is an assistant professor at the University of Houston (UH), Texas, where he is actively involved in optical networking and sensor network research activities. Prior to joining UH, he was a senior technical staff member at Lambda Optical Systems Inc., where he played a key role in protocol development and systems integration activities. In particular, he led system test and integration activities for the Advanced Technology Demonstration Network (ATDNet) testbed project and worked closely with the Naval Research Laboratory (NRL) and the Laboratory for Telecommunication Sciences (LTS). During his earlier tenure at Sprint, he also implemented an extensive broadband testbed for vendor equipment certification and research/development activities.