# Proliferation of the Optical Transport Network: A Use Case Based Study

*Ashwin Gumaste, Indian Institute of Technology*
*Nalini Krishnaswamy, Ciena Corporation*

## ABSTRACT

The ITU-T G.709 standard is perhaps the most underrated and yet widely used of all telecommunication standards, especially pertaining to the transport layer. In this article we focus on OTN as a delivery platform for emerging services. We will understand why OTN was originally proposed and how it can be a critical enabler for emerging services. To this end, we build use cases of OTN especially for Ethernet transport, mobile backhaul, IP router interconnection, sub-wavelength support, and support of overlaid services, especially across multiple domains. The use cases will bring to the fore key advantages of OTN and how this technology solution is ideal for emerging applications. We also discuss future movements in this technology and how these can affect the transport network.

## INTRODUCTION

International Telecommunication Union — Telecommunication Standardization Sector (ITU-T) G.709- or ITU-T G.872-based optical transport network (OTN) [1, 2] is perhaps one of the most important standard technologies that govern the transport of data through the Internet. OTN has evolved as a telecommunication class transport mechanism that does away with some of the non-data-centric issues posed by synchronous optical network/digital hierarchy (SONET/SDH). It paves a way to meet the future of transport technology, for high line rates, flexible granularity support, and doing away with the basic time-division multiplex (TDM) hierarchy, while providing for excellent operations, administration, maintenance, and provisioning (OAM&P) capabilities. The OTN philosophy is well adapted to meet the paradigm requirements of IP over wavelength-division multiplex (WDM), enabling IP routers to have OTN-compatible interfaces and facilitating IP data being mapped to wavelengths in a WDM network. OTN is a true telecom-class transport protocol that is designed as a mechanism to transport contemporary TDM services as well as serve as an enabler technology for emerging services such as carrier Ethernet and cloud services. OTN has evolved as an impressive transport technology ideal for replacing SONET/SDH — taking into consideration the requirements of *both* the overlay and the underlay network.

The overlay of IP, fiber channel, and Ethernet services are characterized by bursty, packet-oriented, or user-defined granular requirements. Mapping the plethora of overlay protocols onto SONET/SDH is inefficient and unaffordable. SONET/SDH was designed as a hierarchy to transport TDM signals whose base line-rate was a DS0. However, with the increasing dominance of IP-centric communication, the reliance on TDM services is fast vanishing. Mobile backhaul, especially in the third-/fourth-generation (3G/4G) variants, also has traffic that is bursty, yet requires latency guarantees, implying a need for a stable transport layer.

The underlying network is primarily based on WDM technologies. Mapping user-defined signals onto WDM lambdas is possible using a multitude of protocols, such as Ethernet, SONET/SDH, and OTN. OTN, however, is especially suited for the WDM layer — it was designed keeping in consideration the fact that WDM would be its primary physical layer component. OTN provides a mechanism to transport signals directly over lambdas, a unique feature considering the fact that variable rate client signals can now be transported and effectively recovered through a WDM network. OTN provides for digital monitoring techniques and forward error correction (FEC) — both of which act as force multipliers, leading to support of higher line rates and better reach [3].

The future of the data layer seems to be moving toward Ethernet, with OTN having a strong role to play; specific OTN mechanisms are available to transport 10 Gb/s, 40 Gb/s, and 100 Gb/s Ethernet through the WAN. OTN also works well with reconfigurable optical add/drop multiplexers (ROADMs), especially when we consider multidegree ROADM hubs that allow seamless movement of wavelengths across any two ports. The per-channel monitoring capabilities induced by the OTN layer enable us to provision end-to-end all-optical circuits over a ROADM core network, relegating the circuit management to the OTN layer.

The aforementioned discussion makes a compelling case for OTN as the transport technology of the future, perfectly mapping to both WDM below it and the vast number of overlaid proto-


0163-6804/10/$25.00 © 2010 IEEE

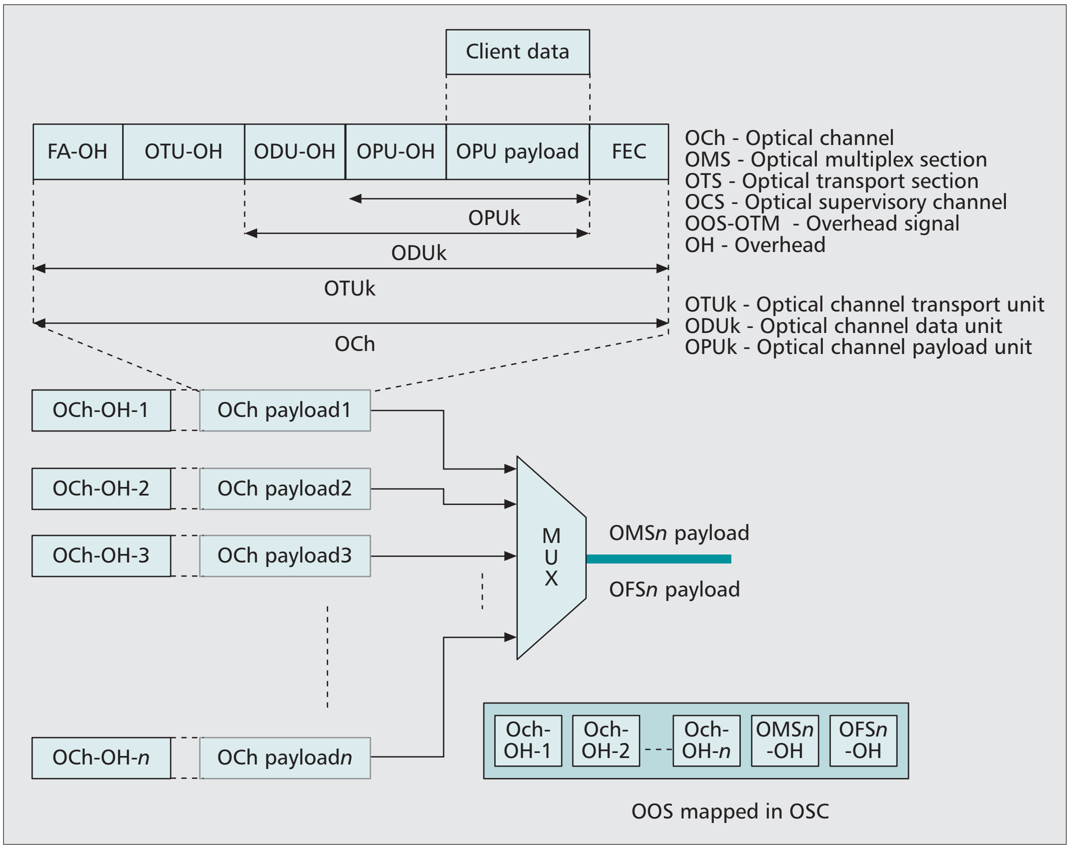


**Figure 1.** *OTN hierarchy.*

*The contemporary solution of using SONET/SDH as the transport layer has limitations in addressing many of these requirements. Besides SONET/SDH was not designed for the kind of emerging services that are seen on the horizon today.*

cols. In this article our aim is to understand how OTN can proliferate itself in different user scenarios. Our approach is to investigate existing and hypothetical scenarios of pragmatic interest and, while doing so, bring to the forefront how OTN can become a smarter and more effective transport technology.

In the next section we briefly recap some of the key findings of ITU-T Recommendations G.709 and G.872 [1, 2], and discuss how OTN works. We skip the *whys* of OTN in favor of *what* the specifics imply. We then detail a number of use cases of OTN. Some of the use cases also describe a comparison to allied technology. We then include thoughts on future work in this area, and the final section summarizes this article.

## OTN Primer

The emergence of WDM technology has allowed good utilization of the installed fiber plant, enabling a provider to add new services on existing fibers. Critical components like ROADMs based on wavelength selectable switching (WSS) technology furthered the use of WDM networks with the flexibility service providers needed in designing the network to accommodate emerging demand-based services. Service providers are now subject to the need for a new transport layer infrastructure that:

- Supports transport of a diverse range of client signals (arising from the need to support services such as video, data center, and storage networking)
- Is flexible to accommodate the varying bandwidth granularities
- Is agnostic to client signal types — thus enabling the use of a single transport layer for all service types
- Enables internetworking of equipment from different vendors and interworking of networks from different carriers [1, 2, 4, 5]

The contemporary solution of using SONET/SDH as the transport layer has limitations in addressing many of these requirements. Besides, SONET/SDH was not designed for the kind of emerging services that are seen on the horizon today. Primary limitations include its rigidity to support transparent transport of LAN type signals and limitations in supporting network-wide provider operations.

The optical transport network architecture defined in ITU-T G.872 consists of the following three layers: optical transport layer (OTS), optical multiplex layer (OMS), and optical channel layer (OCh). The optical signals at each of these layers are terminated at different points in the network. The OTS is terminated at every regenerator/reshaper/retimer (3R) regenerator. The OTS is the physical interface that contains an OMS with a one-to-one mapping. The OTS is terminated at an optical element, like a section layer in SONET or regenerator section in SDH. The OMS is terminated at each point where changes to the wavelengths are made (e.g., at a WDM ROADM). The OCh layer maps onto a wavelength, and its termination is at the point where the OMS is terminated for adding or dropping the client signals. Overhead added at

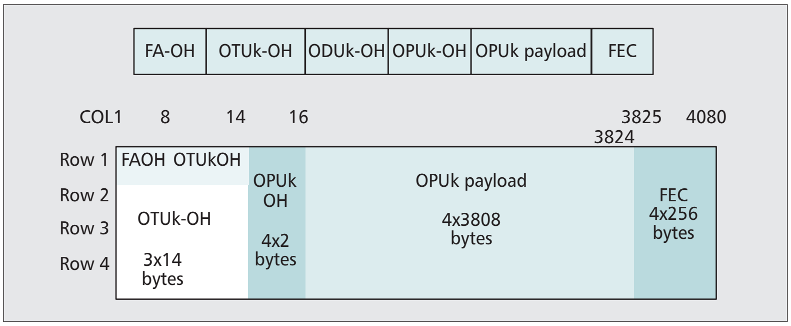


**Figure 2.** *OTUk frame structure*

each of these layers for network maintenance and management purposes are transported through the optical supervisory channel (OSC) in the WDM network. While client signals of any type — SONET/SDH, IP, Ethernet, and so on — could be directly transported over a wavelength (OCh), service providers needed support for individual wavelength monitoring capability and OAM&P. The much needed SONET/SDH-like OAM&P capabilities were not achievable at the optical layer due to limitations in contemporary optical technologies as well as cost.

ITU-T G.709 defines interfaces to the OTN that expand the OCh to add per-wavelength OAM&P support. It makes use of the OEO conversion that is currently required at the 3R regenerator points in the network to provide OAM&P capabilities in the electrical domain. G.709, in addition to adding OAM&P related overheads to the client signal at the wavelength level, specifies FEC to help increase the span between 3R regeneration points, resulting in capital expenditure (CAPEX) savings. It also provides a layered structure comprising the optical channel transport unit (OTU), optical channel data unit (ODU), and optical channel payload unit (OPU) layers for mapping client signal payload. Together, the following factors make OTN a perfect transport layer for next-generation networks that require a converged transport platform that exploits the bandwidth potential offered by WDM technology [6–8]:

- Overhead at each of these layers
- Support for mapping a wide variety of client signal types — legacy TDM and emerging packet services
- Ability to support any client signal rate — rigid SONET rates and packet data at non-SONET rates
- Per wavelength OAM&P capability that is agnostic to client signal type

The OTU layer defined in OTN is the electrical content of the OCh. Supported OTU*k* line rates are at 2.7 Gb/s, 10.7 Gb/s, 43.018 Gb/s, and 111.8 Gb/s. These rates are specifically selected to accommodate SONET/SDH client payloads at 2.5G, 10G, and 40G rates; and Ethernet at 100G. The OTU layer signal corresponding to these rates are OTU1, OTU2, OTU3, and OTU4, respectively, and the common representation is OTU*k* where $k = 1, 2, 3$, or 4.

The OTU*k* layer is the section/regenerator equivalent in SONET/SDH, and is terminated along with the OCh. The client payload is mapped into the OPU*k* payload, and the information about payload type and rate justification is added in the OPU*k* overhead. The OPU*k* is comparable to the SONET/SDH path layer.

The ODU*k* layer is equivalent to the line/multiplex section in SONET/SDH, and it encapsulates the client data in the OPU*k* layer with overheads that provide necessary information for end-to-end supervision.

ODUk path layer data rates defined in ITU.T Recommendation G.709 include only the SONET/SDH rates of 2.5 Gb/s, 10 Gb/s, and 40 Gb/s originally. The second revision of the Recommendation extended these to include:

- A higher 100 Gb/s rate (ODU4) client payload
- A lower 1.25 Gb/s (ODU0) rate payload
- A flexible container (ODUFlex) for variable and constant bit rate (CBR) payloads at any rate $n \times 1.25$ Gb/s ($1 \leftarrow n \leftarrow 80$) and GFP-F mapped packet clients
- 10.4 Gb/s rate (ODU2e) payload support

The ODU0 rate is added to support bandwidth-efficient transport of Gigabit Ethernet signal and any other client signal of rate less than 1.25 Gb/s (the ODU0 rate). The ODU2e rate has been included to transparently transport the 10GBASE R LAN signal. ODUFlex was added to handle the data traffic requirements arising from services with payloads whose rates do not match the SONET/SDH or Ethernet rates. The addition of these rates make OTN a converged transport layer. *Note that ODU0, ODU2e, and ODUFlex signals do not have a corresponding OTU layer signal.* These are mapped into the payloads of a higher-order ODU*k* where $k = 1$, 2, 3, or 4, and transported over the corresponding OTU*k* layer signal.

The G.709 standard also supports multiplexing of lower-rate ODU*j* signals into higher-rate ODU*k* signals where $k > j$. This provides the sub-wavelength networking capability shown in one of our use cases. For this purpose, the first revision of the standard divided the higher-rate ODU*k* signal into multiple ODU1-sized tributary slots. For example, ODU2 is divided into four tributary slots and ODU3 into eight tributary slots. The second revision of the standard supports dividing all higher-rate ODU*k* signals into 1.25 Gb/s tributary slots, corresponding to the ODU0 rate.

## OTN Use Cases

In this section we describe use cases for OTN that are instructive in defining the role of OTN for the future of the transport network. A total of seven use cases are discussed in this section. The first use case, apart from being popular, is also the one that has contemporary impact. It involves provisioning Ethernet services over OTN. The second use case is that of provisioning services that require sub-wavelength granularity over a metro network. This use case is especially important in the metro access/collector network, where the traffic demands are dynamic and do not require full-wavelength granularity, implying that by provisioning sub-wavelength granular services we are able to achieve lower CAPEX. The third use case we consider is perhaps the most important from a

CAPEX perspective: using OTN as a transport technology for an IP-network overlay. In this use case two aspects are considered — provisioning IP over OTN and hence achieving the paradigm of IP over WDM, as well as using OTN networks as a bypass to IP routers. The final four use cases focus on the use of OTN as wavelength services, OTN as an enabler for multidomain communication, OTN as a transport solution in the emerging context of flexible channel spacing, and, finally, using OTN for mobile backhaul, especially in the context of 3G/4G services.

## USE CASE 1: PROVISIONING GIGABIT ETHERNET OVER OTN

Ethernet LAN services dominate much of enterprise and residential traffic, especially due to the cost advantage Ethernet brings along with it. However, transporting Ethernet directly over fiber in the WAN environment has challenges. While colored pluggable optics are an option to send Gigabit Ethernet and 10 Gigabit Ethernet directly over the fiber as lambdas, this option is not very scalable when we consider large distances or across disparate networks and topologies. Provisioning of such Ethernet services across a metropolitan network core is traditionally undertaken through packet-over-SONET/SDH mechanisms whereby Ethernet is mapped into a SONET/SDH payload and transported across the network. The reason for having SONET/SDH as the transport solution is to protect the health of the signal as it moves across the network as well as to provide for restoration in case of fiber cuts and equipment failures. However, one must ask whether SONET/SDH is a necessary technology for transporting of Ethernet signals, given the high cost involved in SONET/SDH interfaces and the associated rigidity in bandwidth provisioning. Synchronization, protocol conversion, and stringent mapping are some of the factors that contribute to making SONET/SDH interfaces expensive. A plausible method of transporting Ethernet seamlessly across disparate networks or across large distances over a WDM core is by the use of OTN, as shown in Fig. 3. In the figure Gigabit Ethernet is mapped into OTN. The OTN framing technology ensures that non-standard-rate signals (10Gig LAN PHY cannot be mapped into existing SONET/SDH without losing transparency) are also mapped into a pragmatic ODU*k*, which is then provisioned as an OCh. The OMS hierarchy can now be established over any state-of-the-art WDM system such as a ROADM or an optical cross-connect that efficiently transports the Ethernet-over-OTN-over-WDM signal across the network.

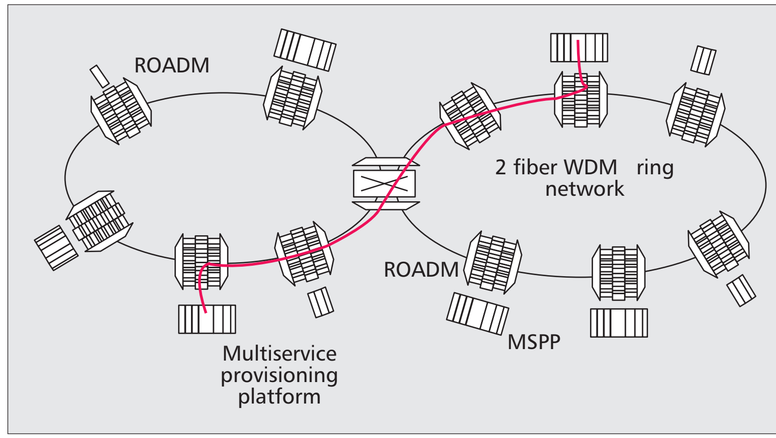


**Figure 3.** *Use case for provisioning Gigabit Ethernet over OTN networks.*

## USE CASE 2: PROVISIONING SUB-WAVELENGTH GRANULAR DEMANDS THROUGH AN OTN INFRASTRUCTURE

In this use case we consider the situation when multiple nodes have traffic demands to each other, and these demands cumulatively do not amount to a full wavelength granularity. Consider Fig. 4 where nodes A, B, C, D, E, and F have traffic to each other, such that the total traffic is less than a wavelength granularity, say 2.5 Gb/s. In this case we can have two provisioning situations:

- Each node has a dedicated lambda provisioned to its peer, leading to $N^2 - N$ connections and hence wavelengths, with each connection being suboptimally used.
- Deploy SONET/SDH solution at each node and run a 2.5 Gb/s SONET/SDH slot train across nodes A to F (and of course a corresponding reverse solution from F to A with concatenation of time slots at nodes, such that the nodes share the bandwidth efficiently).

The second solution leads to a situation wherein each node gets sub-wavelength granularity, and all the nodes effectively and efficiently time-share the wavelength through SONET/SDH time slots. The problem with this solution is the cost involved and the rigidity to scale the same were the demands at the individual nodes to change and increase. An option to do away with expensive SONET/SDH gear and yet be protocol-friendly would be to offer an OTN solution that could achieve a similar performance metric to SONET/SDH: sub-wavelength granular provisioning to discrete nodes while removing the constraint of rigid bandwidth support. To provision such a solution using OTN would require bundling multiple ODUs together within a single OCh. This would require smart traffic engineering of the network. Traffic from A would now go to B, but could either cut through B, using the OCh-based OTN switching fabric, or be dropped entirely at B were B the destination node. Similarly, B could now add its own data into the OCh. However, for traffic from both A and B to coexist would require the data plane to be modified to ensure that the client signals from both A and B get a *fair* share of the OTN data plane bandwidth. This implies multiple ODU*k* sections mapped into a higher-rate ODU*j*. This modification is important — it involves mapping two slower rate ODUa and ODUb signals into an ODUk signal, such that $a + b <= k$. Another mechanism of achieving sub-wavelength granular communication is by incorporating and mapping the discrete sub-wavelength signals into ODUflex and then com-

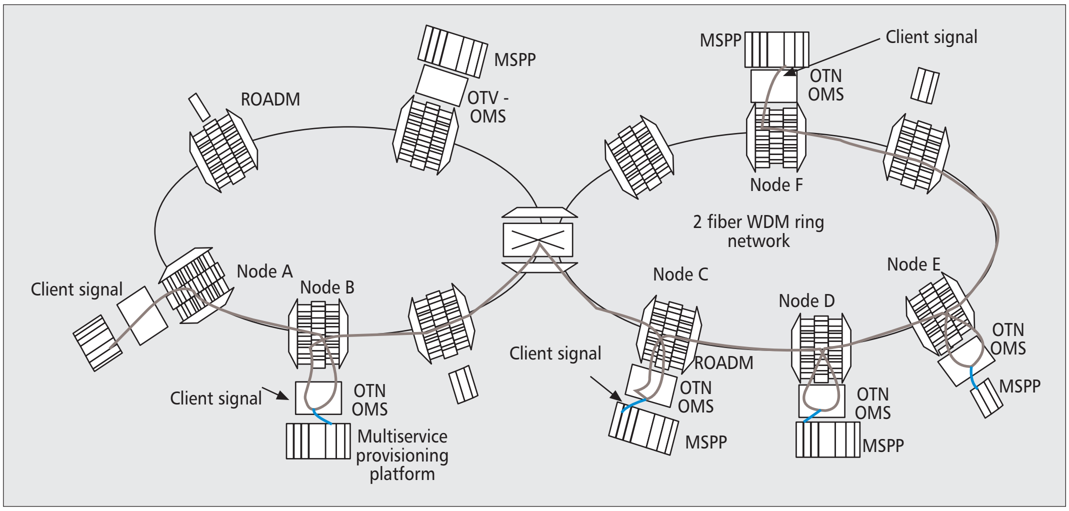


**Figure 4.** *Sub-wavelength bandwidth provisioning over OTN.*

bining these into a higher-rate ODUflex signal. The case of two signals (just described) is now extended to *n* signals leading to the network hierarchy shown in Fig. 4 and hence obtaining sub-wavelength granularity to each constituent node.

## Use Case 3: IP over WDM and the Value of OTN Bypass

In this use case we discuss how to provision IP over WDM using OTN as a technology enabler. A second use case arising out of this solution is to create an OTN underlay that can act as an IP bypass, thereby exploiting the cost advantages of the optical layer. The popularity of IP as a network layer protocol and that of WDM as an efficient physical layer transport solution make the marriage of IP over WDM legitimate and a necessity in an industry where CAPEX demands are primary business drivers. The infancy of optical packet technology combined with the necessity for a framing protocol makes OTN a good choice for provisioning IP directly over WDM. Router interfaces can now have colored optics and support an underlay of OTN transport, mapping IP traffic directly over an ODU*k* interface. The good aspect of having an OTN interface on an IP router is that the signal can now be mapped directly into a colorless port of a ROADM, thereby achieving a full dynamic wavelength routed network. Referring to Fig. 5, we observe that we are making use of provisioned colored optics with OTN interfaces at routers to support an IP overlay (see the lower figure in Fig. 5). By deploying OTN interfaces to IP routers we are now able to directly send data (as OCh) through the ROADM interfaces, thereby achieving the desired IP-over-WDM functionality, critical for CAPEX reduction. This kind of a solution does away with unnecessary digital cross-connect equipment. The advantage of this solution, as compared to IP over Ethernet over WDM, is the ability to monitor each OTN connection through the rich OTN OAM&P functionality.

A by-product use case arising from IP over WDM is that of providing the value of IP bypass using OTN technology. Optical bypass of IP routers is important when we need direct connections between routers. This absolves the need for data to be sent through an intermediate router. Such requirements are particularly needed in a multiprotocol label switching (MPLS) environment where the routing overlay network is traffic engineered into a virtual topology. Even though a classic SONET/SDH network can be configured to provide optical bypass function, this would be restrictive at the optical layer, since there would be an additional need for a SONET/SDH cross-connect that would pass through the data as it moves through a node.

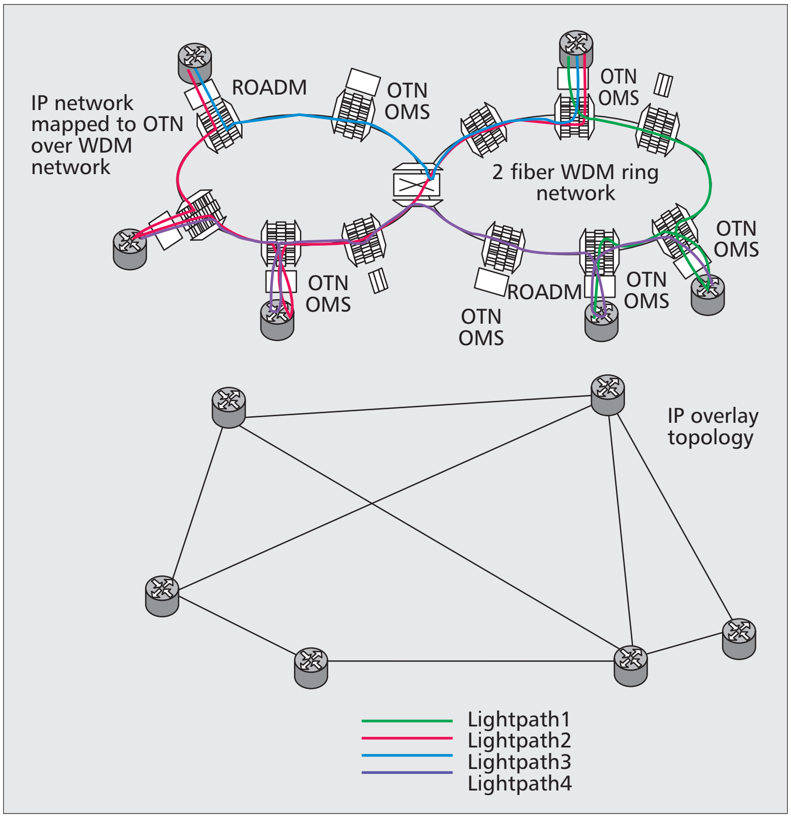


**Figure 5.** *IP over WDM using OTN interfaces and provisioning.*

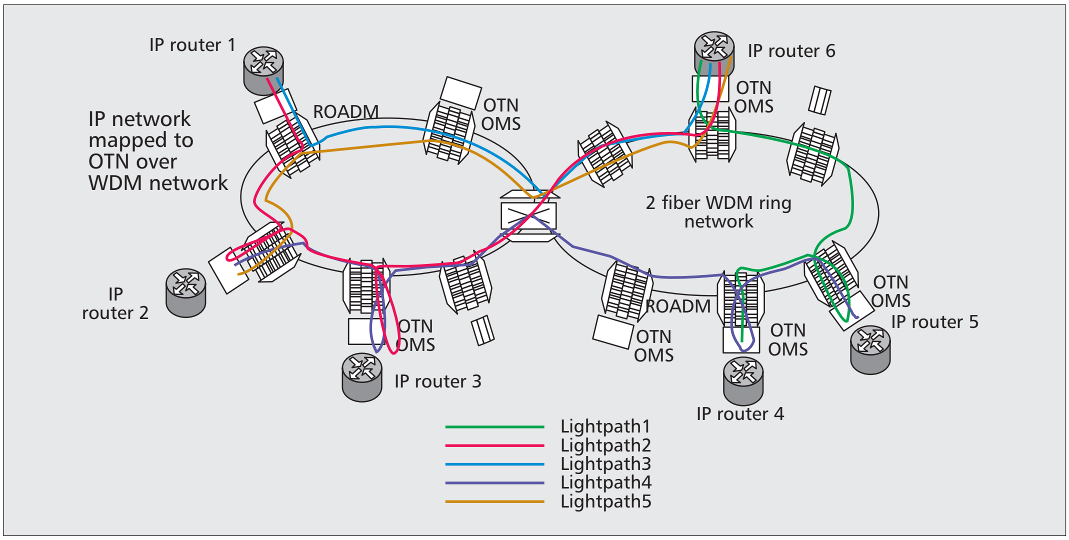


**Figure 6.** *IP over WDM: the value of optical bypass.*

OTN, on the other hand, can allow all-optical movement through a ROADM while still providing end-to-end management and monitoring capabilities.

## USE CASE 4: MULTIDOMAIN NETWORKING USING OTN

The proliferation of data services and applications such as IT virtualization have meant that requirements today often span across several domains/autonomous systems. While in the case of conventional Internet traffic this requirement of moving data across multiple domains has existed and is well provisioned through protocols such as BGP, this problem is severe when we consider the provisioning of managed services across multiple domains. In such a situation, each autonomous system or domain could possibly support its own layer 2 protocol manifestation. In some cases there might also be wavelength translation as we traverse from the ingress domain to the egress domain. Monitoring such a network in the presence of layer 2 and layer 1 anomalies is a challenge, especially if the provisioned service is managed, implying the acute need to enforce strong physical layer service level agreements (SLAs). Our approach to such managed service provisioning is to deploy OTN as the end-to-end transport protocol. The end-to-end circuit is provisioned as an OCh. Segments of contiguous wavelengths are provisioned as OMS, while spans between nodes would be provisioned as OTS. To verify digital signal quality, the OTU could be monitored at domain boundaries if necessary. In this way, the egress node in the last domain has access to information about the entire signal health as the signal traverses across multiple domains. No other technology can seamlessly provision a connection, as shown in Fig. 7, across multiple domains — some manual information exchanger would be necessary to meet the demands of such a system. The OTN hierarchy of engulfing a digital signal in a wrapper and ensuring delivery at the egress with necessary transport information needed for monitoring the health of the signal is able to meet the needs of such multidomain all-optical communication.

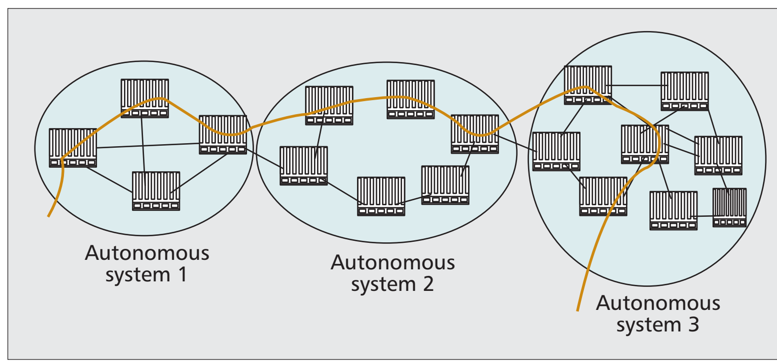


**Figure 7.** *Multi-domain communication using OTN.*

## USE CASE 5: FLEXIBLE WAVELENGTH SPACING AND USE OF OTN AS A TRANSPORT ENABLER

The exhaustion of fiber bandwidth has forced researchers to think of ways to exploit smart schemes to increase fiber utilization. One popular approach in the research community is to introduce non-ITU grid optics and hence flexible lambda spacing. The inherent idea in [9, 10] is to send data on wavelengths with non-congruent spacing. Higher-line-rate channels or those traversing a larger distance are given a wider passband than those traversing a shorter distance or having a lesser line rate. Such flexible channel spacing has significant CAPEX and operating expenditure (OPEX) benefits as shown in [9]. However, none of the approaches has considered the monitoring and signal integrity aspect of provisioning data over such flexible channels. The primary benefit of flexible wavelength spacing is the betterment of physical layer performance. However, without an effective monitoring technique, this benefit is limited. By using OTN as an enabler for provisioning any protocol traffic over the optical layer, we are able to provide the in-line monitoring function.

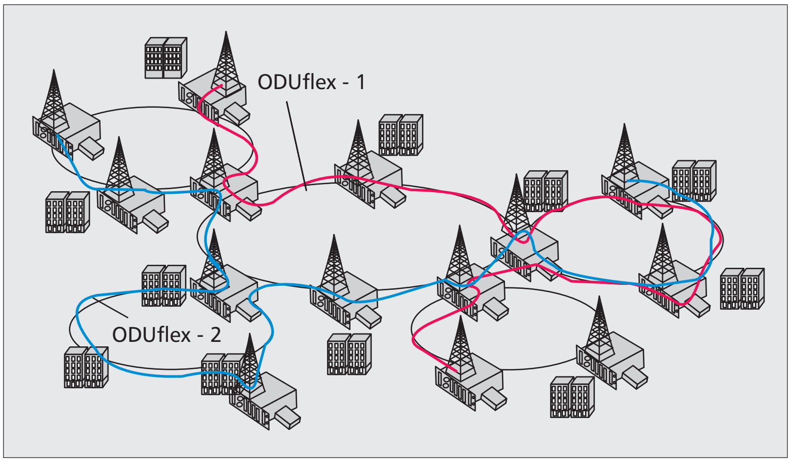


**Figure 8.** *Mobile backhaul network with ODUflex support.*

Especially with the use of ODUflex, variable bit rate signals can now directly be provisioned over the WDM layer, thus taking good advantage of the flexible channel spacing available.

### USE CASE 6: OTN FOR WAVELENGTH SERVICES

A byproduct use case of the flexible wavelength spacing is when we provision wavelength services. Such wavelength services require independent monitoring, which can at times be customer-driven, implying that the network operator is agnostic to the monitoring technique. It is our understanding that for deployment of wavelength services, OTN provides the perfect transport technology — enabling each channel to have its own OAM&P suite associated with it. The added advantage is that the OAM&P features of one channel are not dependent or restrictive on the OAM&P features of another channel in the same fiber, thus making independent wavelength services a reality — much needed for growing enterprise leased wavelengths.

### USE CASE 7: MOBILE BACKHAUL

As we move from 2G to 3G and toward 4G mobile technology, the backhaul becomes an important aspect of the mobile network. The mobile backhaul network has characteristics that are defined by a need for dynamic bandwidth allocation (between base stations) and the ability to provision sub-wavelength granular traffic — even with 3G technology, the bandwidth requirement at a base station is less than 10 Gb/s. OTN, with its ability to provision quality bandwidth (good OAM&P) and at sub-wavelength needs (e.g., using ODUflex), is an ideal technology for the backhaul. In fact, no single technology can meet the requirements of the mobile backhaul the way OTN can. Especially when we move to 4G technology, the backhaul may primarily be based on IP — and IP, to be cost efficient, will be transported on lambdas. OTN comes in handy as a transport mechanism moving bandwidth chunks between base stations as well as provisioning itself as an excellent switching core, using OTN cross-connects whenever there is a need for IP router bypass. Thus, mobile backhaul represents a peculiar use case of OTN, encompassing the features of OTN necessary for sub-wavelength provisioning, IP over WDM, and IP bypass. Shown in Fig. 8 is a mobile backhaul network with ODUflex support. Note that two circuits have been provisioned, each of which is on a different wavelength, and with multiple base-stations sharing the wavelength channel. Furthermore, with the automatic switched optical network (ASON) control plane, it is possible to provision OTN circuits on a dynamic basis, meeting the needs of the mobile backhaul.

## FUTURE WORK

One of the issues with OTN is the generation of OSC from an OTN multiplexed signal. However, when we consider a classical WDM system, not all the channels are OTN compatible, with a likelihood that some channels might be SONET/SDH and some Gigabit/10 Gigabit Ethernet. While for SONET/SDH OSC has limited meaning, for Ethernet directly mapped to WDM, OSC has strong meaning, in the sense that all the provisioning information is carried in the generalized MPLS (GMPLS) control plane sent through the OSC. An important area of investigation is to examine how the OTN generated information transported through OSC works in conjunction with a GMPLS control plane of other layer networks.

One of the reasons Ethernet has been successful is because of its ability to provide backward compatibility with its existing or earlier manifestations. In the case of OTN, this kind of backward compatibility is perhaps more important. It is hence possible that in the future we will have both kinds of networks: where OTN is mapped into a SONET/SDH hierarchy, and where SONET/SDH can be mapped into OTN tunnels, replicating some of the pseudo wire functions known in the data world.

### ACKNOWLEDGMENTS

Ashwin Gumaste was funded by the Ministry of Communications and Information Technology, Government of India.

## CONCLUSION

In this article we have showcased the reasons for the proliferation of OTN. Our study is primarily through the understanding of ITU-T G.709 and G.872, and then applying them as use cases to problems of critical interest. We studied a total of seven use cases and how OTN alleviates some of the sour points of existing networks. These use cases cover the breadth of problems that affect the metropolitan and core networking space. In particular, our focus is on absolving the transport issues in Ethernet, IP, and managed service networks. The use cases, perhaps for the first time, document the breadth of solutions that can be accomplished by the use of OTN technology. We also examine future uses of this fast moving and popular technology solution.

### REFERENCES

[1] ITU-T Rec. G.872, "Architecture of Optical Transport Networks."
[2] ITU-T Rec. G.709, "Network Node Interface for the Optical Transport Network (OTN)."

[3] G. P. Agrawal, *Fiber Optic Communication Systems*, 3rd ed., Wiley.
[4] P. Bonenfant and A. Rodriguez-Moral, "Optical Data Networking," *IEEE Commun. Mag.*, Mar. 2000, pp. 63–70.
[5] TPACK Application Note A/S, "ODU0 and ODUflex: A Future-Proof Solution for OTN Client Mapping," Feb. 2010; http://www.tpack.com
[6] E. L. Varma *et al.*, "Architecting the Services Optical Network," *IEEE Commun. Mag.*, Sept. 2001, pp. 80–87.
[7] EXFO Application Note, "The G.709 Optical Transport Network — An Overview"; http://www.exfo.com.
[8] J, Roese *et al.*, "Optical Transport Network Evolving with 100 Gigabit Ethernet", *IEEE Commun. Mag.*, Mar. 2010, pp. S28–S34.
[9] A. Gumaste, "Reach Optimized Architecture for Multi-Rate Transport System (ROAMTS): One Size Does Not Fit All," *25th IEEE/OSA OFC*, San Diego, 2009.
[10] M. Jinno *et al.*, "Spectrum-Efficient and Scalable Elastic Optical Path Network: Architecture, Benefits, and Enabling Technologies," *IEEE Commun. Mag.*, Nov. 2009, pp. 67–76.

## BIOGRAPHIES

ASHWIN GUMASTE (ashwing@ieee.org) iis currently a faculty member in the Department of Computer Science and Engineering at the Indian Institute of Technology Bombay, where he holds the James R. Isaac Chair. He is also the DAE-SRC Outstanding Research Investigator. He was a visiting scientist and visiting scholar at the Massachusetts Institute of Technology (MIT), Cambridge, in the Research Laboratory for Electronics — Claude-E-Shannon group from 2007 to 2009. He was previously with Fujitsu Laboratories (USA) Inc. as a member of research staff in the PHotonics Networking Laboratory (2001–2005). He has also worked in Fujitsu Network Communications R&D and prior to that with Cisco Systems in the Optical Networking Group. His work on light-trails has been widely referenced and recognized by both industry and academia. He has 18 granted U.S. patents and over 30 pending patent applications. He has published about 120 papers in refereed conferences and journals. He has also authored three books in broadband networks: *DWDM Network Designs and Engineering Solutions* (a networking bestseller), *First-Mile Access Networks and Enabling Technologies*, and *Broadband Services: User Needs, Business Models and Technologies* for Wiley. He is also a guest editor for *IEEE Communications Magazine* and *IEEE Network*, and the founding editor of the IEEE ComSoc ONTC's newsletter *Prism*. He has been with IIT Bombay since 2005, where he convenes the Gigabit Networking Laboratory: http://www.cse.iitb.ac.in/gnl.

NALINI KRISHNASWAMY (nalini.krishnaswamy@gmail.com) is a lead engineer at Ciena focusing on embedded software driver architecture and development for the core switch products. Prior to this she was with Fujitsu Network Communications for seven years in the R&D group focusing on embedded software for various SONET MSPP and DWDM optical transport products. Prior to that she was with Blue Wave Systems. She has been associated with IIT Madras and has research interests in DSPs and embedded software systems, as well as optical transport networks.