Evolution of Packet-Optical Integration in Backbone and Metropolitan High-Speed Networks: A Standards Perspective

Ashwin Gumaste[1] and Shamim Akhtar[2]
[1]Indian Institute of Technology, Bombay
[2]Comcast Cable LLP, Reston, VA, USA.
Email: ashwing@ieee.org

Abstract: Network operators are currently in the midst of a situation whereby there is doubling of bandwidth requirement every two years, though the growth in revenue is nearly flat. In this situation providers must plan their networks to optimize for future growth and take into consideration flat-to-declining operating margins. The OPEX benefits of keeping the data in the lower layers of the network stack can be leveraged to plan future networks. We focus on the impact of evolving standards in the lower layers of the telecommunications stack, namely through the contributions of the IEEE, ITU and the IETF. Specifically, we focus on WDM, Carrier Ethernet, OTN, MPLS and MPLS-TP set of transport standards. These standards have evolved and are being adapted to meet the requirements of next generation networks and being currently considered in the standard bodies. We show through a simulation study that by adopting the packet-optical integrated standard families, providers would benefit because of CAPEX reduction and thereby enhance margins – benefiting the e-community through better services.

I. INTRODUCTION

With the explosion of handheld devices and wired networks the *compounded average growth rate* (CAGR) of bandwidth is almost 63% [1]. Further analysis of the CAGR concludes that the increase in the number of users is not proportional to the increase in bandwidth consumption – implying that users are indulging in more bandwidth-heavy applications as well as using a larger-number of bandwidth consuming devices. As a result, the stress on existing networks continues to grow, while due to the stagnant number of users the revenue continues to be flat. This implies that for a provider to hold its market-share, a recurring investment into the network is mandatory without substantial returns [1]. This problem is particularly pronounced in the metro and core of the network, whereby ports need to be added and line-rate continues to increase with the explosion of bandwidth usage. Ways to compensate include: (1) usage based billing – a strategy that has inertia among users, and (2) smart network planning – with enough foresight into bandwidth growth to be able to meet future needs with existing infrastructure. As part of the continued efforts to plan networks in the metro and regional segments, providers are adopting an initiative to create a multi-vendor supported network infrastructure that facilitates better provisioning and the ability to scale in terms of number of users and bandwidth. While initial CAPEX optimization is critical, it must also be noted that OPEX minimization becomes a key objective, given its more potent business impact. A key aspect of this minimization is to enable data to be kept in the lower layers of the stack, resulting in minimal processing (and hence lower energy consumption by avoiding the use of expensive network processor based routing elements). To be able to architect a network that is scalable and future-proof, an operator has to rely and influence standards that would support multi-vendor environments as well as facilitate growth that is not dependent on frequent network upgrades.

To this end, is the vision of packet-optical integration – a process that has already begun, and would potentially lead to the best of both the IP and the WDM worlds, using enablers such as Carrier Ethernet and the *Optical Transport Network* (OTN).

In this article, we capture the essence of the work being done in the IEEE, ITU, IETF standard bodies from the perspective of next generation network planning in the optical, data and network layers. In Section II we discuss a vision of a managed transport infrastructure based on low-CAPEX-and-OPEX equipment. We discuss how standards will drive the bigger picture of a low-cost network that can be soft-upgraded to meet new service and user needs. Specifics of contributions in the IEEE, ITU and IETF are discussed in Section III-VI across the three layers of the stack. Section VII discusses enablers to standard development. Section VIII showcases numerical results that justify the inclusion of standards in the next generation networks.

II. THE PACKET-OPTICAL NETWORK: THE BIG PICTURE

The packet-optical network is optimized for both transport as well as user services. In parallel such a network achieves scalability at low price-points using platforms that support low-energy-consumption and adaptable policies for optimized switching and routing. Providers currently have networks that are not optimized for efficiency and use multi-layer equipment that only adds redundancy across the layers. What is desired is a network that is designed to scale and optimized for performance. In quest of such a design, standards play a critical role – as markers that guide deployment as well as facilitate *multi-source agreements* (MSAs). We desire to be able to achieve the following network hierarchy through the evolution of packet-optical standards.

1. *Scalable, fault-tolerant, topology aware routing*: The network layer is dominated by the IP/MPLS overlay. IP/MPLS allows operators good *control* over traffic. The problem is the inherent high-cost-per-port. The problem is particularly pronounced when we consider the scalability of IP/MPLS routers. As we populate more ports to existing routers, these saturate the chassis, and there is in fact a degradation of throughput due to lack of scalability of the non-blocking fabric. It is hence desired to be able to scale the port-count of IP/MPLS routers to meet future service needs. A second issue is that of line-rate. As we move to 40/100Gbps and beyond, the ports must not only support the existing line-rate, but in many cases be able to seamlessly migrate to new ones. Due to the higher degree of functionality of IP/MPLS routers, these are desired to be minimized across a network to conserve both CAPEX and OPEX. Currently, there is no direct established relationship between traffic (and traffic churn) and the amount of IP/MPLS core routers required. Standardization of such a relationship can lead to an optimized network and good quality-of-experience to the end-user. This is perhaps a challenge that is difficult to address.

2. *Carrier-class, service-aware forwarding plane*: Layer-2 services and devices have efficiently migrated to 10Gbps and in many instances to 40Gbps as well as 100Gbps per-port. Layer-2 technologies, especially Ethernet-based are carrier-class, implying good *operations, administration, maintenance and provisioning* (OAM&P) capabilities, in addition to efficient sub-50 millisecond restoration using *connectivity-fault-management* (CFM) standards. The functionality of *Carrier Ethernet* (CE) is not fully utilized. The performance of CE is such that it is possible to get circuit-like-performance, at fractional CAPEX of SONET/SDH systems. However, the replacement of SONET/SDH gear with Carrier Ethernet is slow. This is because of the lack of insights into the full capabilities of Carrier Ethernet network architecture. While the standards in this area such as the 802.1Qay for *provider backbone bridging-traffic engineering* (PBB-TE) and RFC 5317 (and beyond) for *multi-protocol label switching-transport profile* (MPLS-TP) have matured, these do not dwell on network level issues. There is a gap in how CE standards can be applied to within a network and how efficient mapping of services to CE occurs. From the work in both IETF and the IEEE the new drafts in this area focus on aspects such as segment-protection (802.1Qbf), data-center bridging, shortest-path bridging, Generalized Multi-Protocol Label Switching (GMPLS) control of CE networks etc. – all of which point to the lack of a relationship between the technology (CE) and the service.

3. *Transport-friendly optics, intelligent, dynamic optical cross-connects*: Since the photonic bubble-burst in the early 2000s, there has been a new effort in the photonic networking domain. We have witnessed the development of the *Reconfigurable Optical Add-Drop Multicast* (ROADM) as an infrastructure to support multi-wavelength add-drops and cross-connect functionality leading to the concept of a Wavelength Switched Optical Network (WSON) [2]. In parallel, we also have experienced the advent of new photonic technology for coherent optics – to support 40Gbps, and 100Gbps systems in metro and regional networks. The deployment of photonic networks as an efficient underlay to IP/MPLS infrastructures has implied that there is a need for the optical network to understand, follow, and perhaps even optimize the overlay. The power of the optical bypass (as is popularly known) is a tool that can withstand the lack of scalability of IP/MPLS routers (as will be seen through our simulations study). There is the discussion around flexible grid spacing [3] (grid-less optics) that aim to pack more channels in the fiber by flexibly using the spectrum and doing away with rigid channel-spacing, thus resulting in more bandwidth in the network. Both optical line-rate and wavelength-count do not increase in proportion to the rate of data-growth. Lack of MSA for new optical components and standards that can modularize the different aspects of higher-line-rate transponders means that there are several non-technical bottlenecks that impact faster production of higher-speed systems. The concept of software-defined optics, and pluggable modules that can be replaced for achieving higher-rate are instructive for future standards development.

4. *Malleable, service-centric control plane*: With disparate layers, it is of interest to control the network using a unified control plane. The discussions around the creation of a unified control plane are at least a decade old. However, despite much effort, not much of deployment of a unified control plane has happened in the field. Scalability, inter-domain issues, multi-layer equipment, upgrade of network gear, service-rollouts are some of the issues that impact control plane deployment. There is a significant agreement between vendors of layer 1-3 gear about the use of GMPLS as an over-the-top standard. Despite this agreement, the problem is that most vendors deploy control through a proprietary protocol. The functionality of the control plane is restricted to only service provisioning or support for low-impact equipment. This implies that between the control plane and the *network management system* (NMS), there is a clear distinction preventing the same control plane to be used across equipment from different vendors. There is a serious lack of effort (and perhaps business-will) in being control-plane-wise interoperable.

To be able to meet the above features, we need specific technology changes, primarily brought about by standards that are quintessential to the next generation network. These are captured as below:

1. *Ability to support Flexible Spacing Optical Networks*: With the massive growth of traffic, there is a need to maximally utilize the fiber. Advance modulation

formats such as QPSK and OFDM are enablers at high line-rates (100Gbps and beyond). With the advent of coherent optics, it is possible to provision a large number of channels at narrow spectral spacing. Such ultra-dense multiplexing techniques would not only facilitate the optimization by an optical bypass, but may also facilitate the long term vision of a wavelength to an enterprise. Flexible spacing capable optical network proliferation (also called grid-less networks) would require a significant cost reduction of coherent receivers by means of photonic integration technologies.

2. *Ability to support carrier-class packet transport*: Most services today are in packet format. For both application and network flexibility, packet-mode communication is a must. This means that a provider must be able to provision a packetized service that is billable and can be administered and maintained. Carrier-class packet transport technologies are maturing and their adoption would be critical to support next generation services, particularly those within a cloud environment or for mobile backhaul.

3. *Ability to support mass network migration through dynamic networking*: Contemporary providers have to facilitate a delicate provisioning-balance between *content distribution network* (CDN) traffic and user-emanated traffic. At times, CDN traffic can lead to uncontrolled surges, especially video traffic that is voluminous and delay sensitive. For example, a webcast of a popular sport-event has the capability of disrupting network operations and such spurting demands must be well planned or accounted for while designing the network. There is not sufficient effort in the CDN standards community to interact with providers leading to mutually beneficial provisioning and peering norms.

4. *Multi-domain network analytics and monitoring*: With the growth and diversity of services – mobile rollouts, data-center/cloud providers etc. make the Internet a collection of specialized entities. Tier-1 providers have the onus of interacting with multiple tier-2 and tier-3 entities. In such a multi-domain environment with diverse technology deployments, managing end-to-end services and maintaining network analytics is a challenge. Individual technology standards such as the Metro Ethernet Forum for Ethernet services exist. However, there is an absolute lack of standards that focus on multi-technologies across multi-domains for the same service. This creates *quality of experience* (QoE) discrepancy and leads to non-optimized service performance.

5. *Network scalability using soft-upgrades*: Lack of standardization specific to scalability has implied that it is difficult to scale existing equipment to newer technologies and line-rate. With the current line-rates, devices use specialized technologies for transmission, switching, control and protocol stack. There is a need for standards to be followed within the architecture of a platform – so that MSAs can be adhered to and upgrades of technology can be “soft” rather than “hard”. This would tremendously save CAPEX as well as work well for the vendor who would have a lock-in with the operator across generations of features and upgrades.

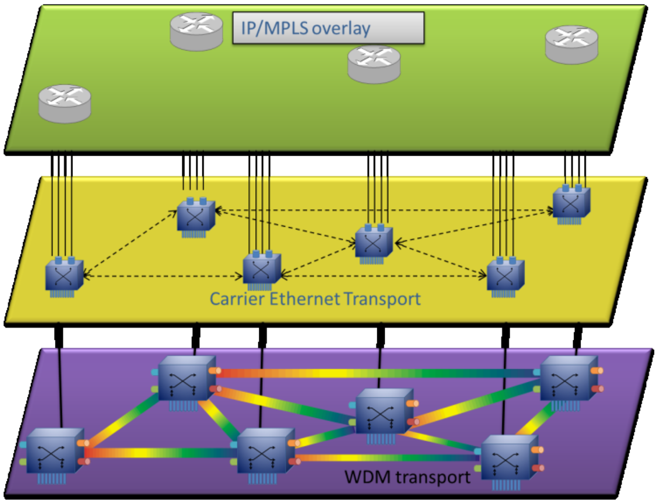

Fig. 1. Multi-layer optimized network.

III. SCALING AND EXPANDING IP/MPLS NETWORKS

The vast deployment of IP/MPLS technology is an efficient mechanism to route packets between source-destination pairs. This one-stop technology is critical in providing complete routing, management and networking functionality. It comes at a price – high initial CAPEX, difficult to scale without significant periodic CAPEX infusion and control plane scalability issues. A bigger problem with IP/MPLS systems is the difficulty in mapping revenue bearing services to the IP/MPLS network. Long Term Evolution (LTE) and LTE-Advanced backhaul services are optimized for IP/MPLS, and their coexistence with other best-effort services in the IP domain implies non-carrier-class support in the backhaul. In an era where uncontrolled CDN traffic can create havoc on a providers' infrastructure [4], the limitations of IP/MPLS technology specifically pertaining to support of carrier-class managed pipes becomes critical. While almost the entire IETF and large parts of the ITU are devoted to work on various aspects of IP/MPLS technology, these are areas where new work is required.

Specifically, scalability of IP/MPLS platforms is not well discussed in the standards. It can be argued that scalability of IP/MPLS platforms is a vendor specific issue. However, such argument does not stand when we consider a holistic network view with the need to support scalable platforms at diverse locations from different vendors. As line-rates move from 10Gbps to 40Gbps and to 100Gbps, there is a need to scale the cross-connect fabric to support larger platforms. In an ideal world, for different line-rates, line-cards and cross-connect fabrics should be interchangeable, given that the forwarding plane adheres to the IETF RFCs. One would expect that the degree of scalability across vendors be such that operators are able to design a network with enough dynamic features. In current settings, it is difficult to plan a network, and once provisioned, it is difficult to replace gear easily. For example, if we have three vendors X, Y, Z then, one problem is that not all of Y's gear can be placed anywhere. Similarly, Z's gear may not be able to be used in conjunction with that of X to overcome its scalability deficiencies with Y. The control plane of Z may not be interoperable with X and Y but its forwarding plane may be superb in a stand-alone network. The aforementioned issues are not just technology impairments but are due to unavailability of standards for vendors to design platforms.

IV. PACKET TRANSPORT: PBB-TE AND MPLS-TP

The high CAPEX and OPEX in SONET/SDH networks along with the inability to support statistical packet multiplexing has accelerated Carrier Ethernet standards – in the IETF, ITU with MPLS-TP and IEEE with PBB-TE. Using packets for transport seemed as the logical step in order to support the vast data-traffic needs at low price-points. The current deployment of CE platforms however is quite the opposite. Even though it is expected that the deployments would go up in the near future – these would be no where compared to what was expected. Partially this is due to lack of CAPEX budgets for operators. However, this is also due to a lack of clear visibility for the technology itself. We bring to the fore certain issues that are addressed only partially through CE standardization efforts.

The CE standards (both PBB-TE and MPLS-TP) do not consider the legacy network. For example, several million T1s and E1s exist across the globe, and would it not make sense to migrate these to an all-Ethernet infrastructure? While such Ethernet-centric migration has happened, the standards do not focus on these. As a result, almost every provider deploying CE for end-to-end connectivity has a unique solution to offer. Both flavors of CE use service identifiers to map payloads which forwarding elements can differentiate and process the data-flows. To set these identifiers (Backbone VLAN Identifier or BVIDs in PBB-TE and labels in MPLS-TP) requires a control plane. While both standards dwell upon the control aspects of the protocol, there is little or no mention on identifier selection – which is instructive for multi-vendor compatibility. As a result an operator deploying gear from a particular vendor can expect only forwarding-plane interoperability between vendors. When we come to the more popular flavor of CE, i.e. MPLS-TP, there is significant ambiguity in the relevant RFCs. RFC 5317 and 5860 discuss basic MPLS-TP and OAM aspects. The ambiguity for implementation and lack of completeness are causes of implementation concerns. The implementation ambiguity allows vendors to subscribe to their proprietary mechanisms. When compared to SONET/SDH, which achieved good interoperability, CE interoperability performs poorly. The architectural problems due to interoperability issues are immense: equipment from vendor A claiming to be MPLS-TP compliant is similar to vendor B's MPLS equipment with some minor changes. The tragedy of the standardization process is the conformance of equipment to conflicting standards. Vendor D's equipment for example conforms to a paradox of carrier-class as well as a best-effort packet standard (PBB-TE and PBB). Theoretically, it is possible to support both, but when implemented in a large network, it would create havoc – PBB's Multiple Spanning Tree Protocol (MSTP) convergence would destroy all notions of carrier-class behavior in the network leading to confusion, toggling and crashing of the entire network. While none of these have actually happened in those select networks that deploy vendor D's gear, it is just a question of time before some or all of the catastrophic effects occur.

V. THE OPTICAL TRANSPORT NETWORK

As Ethernet continued to reinvent itself from (CSMA/CD) based LANs to switched Ethernet, to Ethernet in the air (WiFi) and now Carrier Ethernet in WANs, there was a similar push in SONET/SDH towards being more packet-friendly. The first sign of such adoption was the facilitation of packet mapping using techniques such as Link Channel Adjustment Scheme (LCAS) and Virtual Concatenation (VCAT). Generic Framing Procedure or GFP (ITU.T 7041) was proposed as a better way to provide packet-over-SONET (POS) capabilities for variable data-rates. POS led to Next-Generation SONET/SDH (NG-SONET) that had rich packet transport features including efficient traffic grooming and control – but yet maintained the high price-points. Operators that adopted these standards often did so due to compatibility with legacy gear. In addition, the 3-decades operational expertise with SONET/SDH implied a trained workforce. To support SONET/SDH like performance while taking advantage of WDM technologies across multiple-provider domains, the ITU standardized the *Optical Transport Network*. Apart from being readily able to map and encapsulate packetized traffic and provide SONET/SDH like OAM environment, OTN has two other major advantages: (1) the ability to include Forward Error Correction (FEC) enhances reach to 40/100Gbps (2) providing Tandem Channel Monitoring (TCM) – enables a user/operator to track a signal through multiple domains. The OTN hierarchy and its use cases have been described in [5]. ODU-switching and ODU-flex are two popular techniques that facilitate enhanced usage of the OTN hierarchy to map services to the OTN layer as well as allow for efficient grooming. OTN standards have focused on the architecture (ITU.T. G.872), protection, signaling among other features. There are two aspects of the OTN standards portfolio that require more work for operator networks:

(1) Making good use of OTN switching. Given the redundancy of switching across multiple layers – IP,

Ethernet, OTN and lambda, it is not clear in most operator networks as to the role of OTN switching and how it would impact the services portfolio leading to business. Some questions that we attempt to answer in our simulations study includes: (a) At what service-granularity can ODU switching best work? (b) Would ODU switching work better than packet switching for certain services, and how do we define these services.

(2) Impact on network equipment and translation of OTN information across layers. Router vendors are using OTN chipsets to encapsulate Ethernet frames with OTN. The question operators ask is whether this utilizes the full spectrum of OTN capabilities? If the idea is only to utilize FEC, and hence enhance reach, then the "OTN-izing" of the router interface has limited value. From the perspective of standards, unless the aforementioned issues are standardized, it is difficult to interoperate OTN equipment, especially from the perspective of a large network. Vendor "Os" equipment may support a smaller ODU-cross-connect than vendor "Ps". What type of equipment and in what quantity must an operator deploy? Unless the switching architecture and guarantees are standardized, it is impossible for a designer to plan for future growth.

VI. TOWARDS SPECTRALLY EFFICIENT FLEXIBLE OPTICAL NETWORKING

With the commercialization of coherent optics as well as extremely tight-spacing multiplexer technology, flexible spacing oriented optical networking is becoming commercially feasible. The approaches by component vendors are quite diverse and there is a critical lack of standardization as to how metro or core equipment would be designed to conform to the super-tight specifications juxtaposed by the flexible/grid-less optical network. The goal of being able to provide a wavelength or a *spectral slice* [9] to a user can be achieved in multiple ways: In one embodiment, a single wavelength [6] consisting of many sub-wavelengths formed by OFDM techniques is sent to a group of users, each of which will extract their relevant data (on a sub-channel). In another embodiment, the spacing between channels is made extremely narrow, and an Arrayed Waveguide based multiplexing system is used to create the composite signal. Some of the spectrally efficient approaches consider spanning across access, metro and regional networks using coherent optics to achieve reach [7]. Such a network has the potential of giving a completely new meaning to metro and access networks with good reach and efficient spectral utilization. The differentiation between the access and the metro domains can potentially be negated. The network has the potential of reducing the amount of electronic equipment across the network. In this manner bypassing the cross-connect at the metro-access interconnection as well as providing for IP-bypass by using optical channels from the core to the edge of the network is possible. The flexible spacing oriented optical network can manifest in multiple architectures. Shown in Fig. 2 is the concept of a wavelength or slice-per-user architecture. The ROADM architecture required at the edge of the network is now replaced by a Flexible grid supporting ROADM that not just provides optical add-drop capabilities, but also absolves the need for edge routers/switches. Several questions emerge: the optical architecture that spans the unified network is undefined as to how to support such large number of channels. We must also standardize the equipment at which to terminate the channels in the metro. This termination equipment can be a collection of IP routers or CE switches or some combination of both.

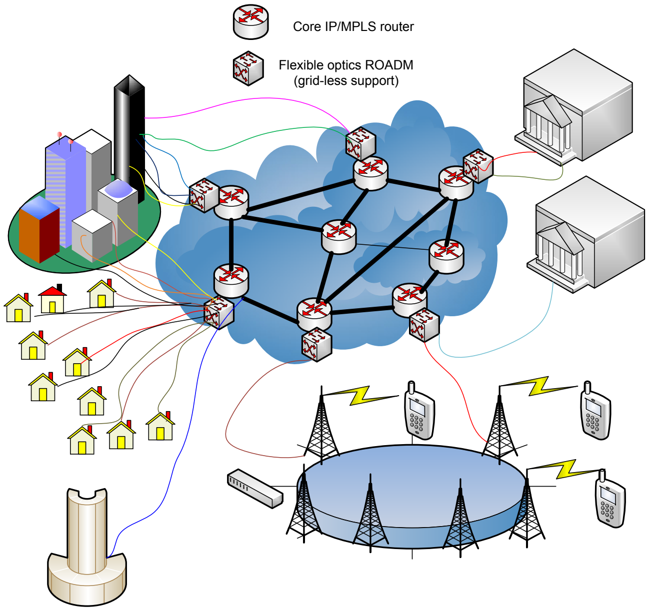

Fig. 2. Flexible (Grid-less) Optical Network Architecture.

VII. VERIFICATION

To better understand the merits of a packet-optical network we built a multi-layer simulations model in Matlab. The model assumes a 120-node network with 24 core-nodes and the remaining metro-nodes. Each node supports up to 100K broadband users (using a distribution network). Each core-node supports a cluster of metro nodes in an interconnected ring topology, while the core nodes form a mesh network. Our goal is to compute the CAPEX impact using network equipment as well as with optimized standards that may be available in the future. For the purpose of cost computation we assume the following table. Note that this table is derived from discussions with vendors in North America as well as Asia and costs are averaged and then normalized [10]. The costs for flexible optical technology are only target costs as such equipment is currently not available.

TABLE 1: Cost Computation in ‘000, USD

Equipment	Cost for 1Gbps	Cost for 10Gbps	Cost for 40Gbps	Cost for 100Gbps
IP/MPLS	4	10	18	36
OTN	1	2.5	5	10
CE	0.25	2	6	12
Gridless TX/RX	0.5	1.5	4	8
ROADM TX/RX	0.25	0.5	1.8	4
ROADM Cross Connect	4	4	8	16

Based on the table above, we performed simulations by assuming user demands to arrive according to Poisson distribution and with exponential holding times. 4-levels of QoS were assumed. 1/100th of traffic was assumed to be dedicated for network control.

In Fig. 3 we showcase the impact of a network with IP/MPLS, CE and optical bypass (ROADMs). We compare the use of optical bypass with select IP/MPLS nodes as well as optical bypass with CE switches and select IP/MPLS nodes. We observe that with optical bypass there is a direct cost reduction in the mesh network validating our argument of keeping the data in the lower layers of the stack. A further justification is seen when we add CE nodes to the simulation model. The cost further reduces, as those paths which could not be set up purely with the optical bypass can now be set up using CE. The role of IP/MPLS is restricted to the core of the network. In fact a side-result of this viewgraph is that the adoption of larger switch fabrics for IP/MPLS routers can be delayed by the optimized CE+ROADM network in conjunction with the IP/MPLS network.

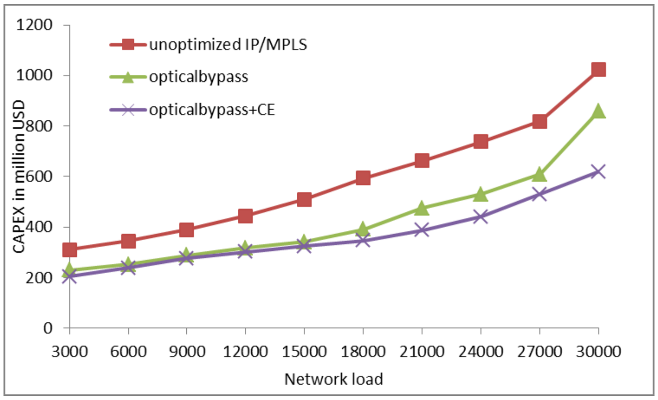

Fig. 3. Impact of IP/MPLS networks with standardized bypass and CE.

Shown in Fig. 4 is the impact of OTN switching. It is observed that OTN switching has no impact on a network that already deploys CE and IP/MPLS. However, when used individually, i.e. IP+OTN or CE+OTN, there is a cost advantage. While the latter (CE+OTN) solution is less-expensive, it is difficult to scale, especially when requests are dynamic and require rapid provisioning – which happens through an IP/MPLS network.

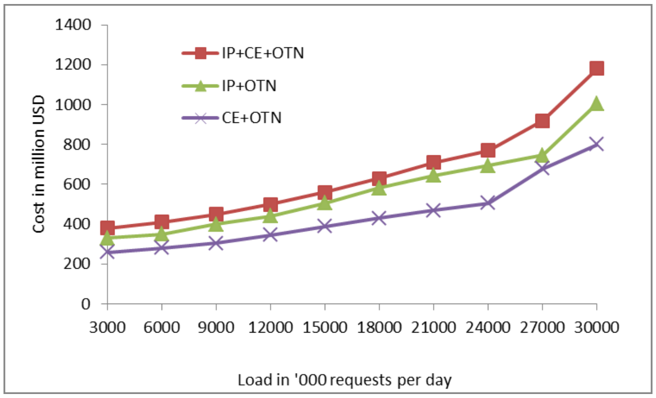

Fig. 4. Impact of OTN switching.

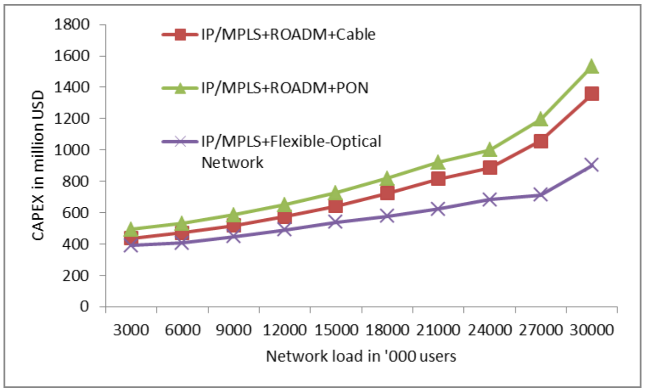

Fig. 5. Impact of flexible optical network on CAPEX.

Shown in Fig. 5 is the impact of flexible optical networking on CAPEX and comparison with cable and xGPON options. Flexible optics significantly reduces end-to-end CAPEX by eliminating the need for aggregation/edge routers and replacing these with spectrally efficient optical gear. In fact, a spectrally optimized optical network is lower in cost than a cable or xGPON network, because despite the higher target cost of optics, there is significant reduction of the core and metro equipment. What we have observed is that 37% of connections in the network can be provisioned point-to-point optical channels (as spectral slices) absolving the need for higher-layer (and more expensive) equipment. We further observe that 29% of connections are provisioned such that the aggregation happens only in the core, thus implying the metro part of the network is largely free of aggregation equipment.

IX. CONCLUSION

We have presented the vision of a packet-optical integrated network architecture. Specifically, we have focused on issues pertaining to network scalability and service delivery across IP/MPLS, Carrier Ethernet, OTN and optical domains. Qualitative and quantitative impact of the standards is also shown. Opportunities for new standardization approaches are discussed from the perspective of an operator. A key observation from this article is that standards must impact both network architecture and planning as well as product architecture and should perhaps be extended to service delivery. This would result in benefit towards the entire ecosystem.

References

[1] IEEE802.3 Industry Connections Ethernet Bandwidth Assessment Ad Hoc Group Report: http://www.ieee802.org/3/ad_hoc/bwa/BWA_Report.pdf

[2]"A Framework for the Control of Wavelength Switched Optical Networks (WSONs) with Impairments", draft-ietf-ccamp-wson-impairments-10 http://tools.ietf.org/html/draft-ietf-ccamp-wson-impairments-10

[3] "A Framework for control of Flex Grid Networks", draft-syed-ccamp-flexgrid-framework-ext-01.txt, http://tools.ietf.org/html/draft-syed-ccamp-flexgrid-framework-ext-01

[4]H. Chang, S. Jamin, and W. Willinger. To peer or not to peer: Modeling the evolution of the Internet's AS-level topology. In INFOCOM 2006.

[5] A. Gumaste and N. Krishnaswamy, "Proliferation of ITU 709 - Optical Transport Network (OTN) – A Use Case Based Study," IEEE Communications Magazine, September 2010.

[6] J. Armstrong, "OFDM for optical communications," J. Ligthw. Technol., vol. 27, no. 3, pp. 189–204, Feb. 2009.

[7] H. Rohde et al, "Coherent Optical Access Networks," Proc. of the Optic Fiber Conference (OFC) OTuB1, March 2011, San Diego.

[8]Content Delivery Networks Interconnection: http://datatracker.ietf.org/wg/cdni/

[9] M. Jinno et al, "Spectrum-efficient and scalable elastic optical path network: architecture, benefits, and enabling technologies," IEEE Communications Magazine, Vol.47. No 11. November 2009, Page(s):66 - 73

[10] A. Gumaste, et al, "Using Global Content Balancing to Solve the Broadband Penetration Problem in the Developing World: Case Study – India," IEEE Communications Magazine, May 2012. Volume: 50, Issue: 5, Page(s): 74-81

Bio: Ashwin Gumaste is currently the Institute Chair Associate Professor in the Department of Computer Science and Engineering at the Indian Institute of Technology (IIT) Bombay. From 2008-2012 he was also the J. R. Isaac Chair Assistant Professor. He was a Visiting Scientist with the Massachusetts Institute of Technology (MIT), Cambridge, USA in the Research Laboratory for Electronics from 2008 to 2010. He was previously with Fujitsu Laboratories (USA) Inc in the Photonics Networking Laboratory (2001-05). He has also worked in Fujitsu Network Communications R&D (in Richardson TX) and prior to that with Cisco Systems in the Optical Networking Group (ONG) and has been a consultant to Nokia Siemens Networks, Munich working on Next Gen access standards. His work on light-trails has been widely referred, deployed and recognized by both industry and academia. His recent work on Omnipresent Ethernet has been adopted by tier-1 service providers and also resulted in the largest ever acquisition between any IIT and the industry. This has led to a family of transport products. Ashwin has 20 granted US patents and over 30 pending patent applications. Ashwin has published about 150 papers in referred conferences and journals. He has also authored three books in broadband networks called 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 John Wiley. Owing to his many research achievements and contributions, Ashwin was awarded the Government of India's DAE-SRC Outstanding Research Investigator Award in 2010 as well as the Indian National Academy of Engineering's (INAE) Young Engineer Award (2010). He was also the recipient of the Vikram Sarabhai Research Award in 2012. Ashwin was the IBM Faculty Award winner for the year 2012. He has served Program Chair, Co-chair, Publicity chair and workshop chair for IEEE conferences and as Program Committee member for IEEE ICC, Globecom, OFC, ICCCN, Gridnets etc. Ashwin is also a guest editor for IEEE Communications Magazine, IEEE Network and the founding Editor of the IEEE ComSoc ONTC's newsletter Prism. He is the Chair of the IEEE Communication Society's Technical Committee on High Speed Networks (TCHSN) 2011-2013. He has been with IIT Bombay since 2005 where he convenes the Gigabit Networking Laboratory (GNL): www.cse.iitb.ac.in/gnl. The Gigabit Networking Laboratory has secured over 11 million USD in funding since its inception and has been involved in 4 major technology transfers to the industry. Ashwin can be reached through www.ashwin.name.