Scaling Terabit Networks: Breaking Through Capacity Barriers and Lowering Cost with New Architectures and Technologies

Report based on findings from a workshop that took place on 19-20 September 2013 at OSA Headquarters, Washington DC, USA

Lead authors:
Keren Bergman, *Columbia University;* **Vincent Chan**, *Massachusetts Institute of Technology*
Daniel Kilper, *CIAN, College of Optical Sciences, University of Arizona;* **Inder Monga**, *ESnet, Energy Science Network*
George Porter, *University of California, San Diego;* **Kristin Rauschenbach**, *Raytheon BBN Technologies*

Acknowledgements

Support for this workshop was provided by the National Science Foundation. The Center for Integrated Access Networks, an NSF Engineering Research Center, and the Optical Society served as co-hosts of the workshop, part of the OSA Incubator Program, and special thanks to the Optical Society for the dedication of their staff and the use of their facilities. The workshop was support through NSF grant CNS-1346666 and CIAN acknowledges support under NSF grant EEC-0812072.

Workshop Participants

Brandt-Pearce, Maïté
Chan, Vincent
Chang, Gee-Kung
Dagenais, Dominique
Dexheimer, John
Dutta, Rudra
Farber, Dave
Fisher, Darleen
Frankel, Michael
Fox, Geoffrey
Fumagalli, Andrea
Gerstel, Ori
Kilper, Daniel
Koley, Bikash
Liou, Chris
Lyles, Bryan
Magill, Pete
Marzullo, Keith
Monga, Inder
Papen, George
Partridge, Craig
Peleg, Avner
Ramamurthy, Byrav
Rauschenbach, Kristin
Rouskas, George
Saleh, Adel
Sethumadhavan, Chandrasekhar
Shah, Jag
Smith, Jonathan
Song, Houbing
Subramaniam, Suresh
Tierney, Brian
Thompson, Kevin
Veeraraghavan, Malathi
Vokkarane, Vinod
Xin, Yufeng
Yoo, Ben
Zussman, Gil

Executive Summary

The consolidation of massive amounts of computation, storage, and networking into online data centers and "cloud" computing deployments is fundamentally transforming our ubiquitous access to unprecedented amounts of information. The continued progress in societal benefits and our capabilities to unlock new or previously intractable applications in science, healthcare, media, and business are dependent more than ever on the scalability of this critical infrastructure. For these online services to scale to hundreds of billions of users and devices—an Internet of Things, the networks underpinning these services must likewise increase their scale by orders of magnitude. The needed exponential increase cannot be met with continued evolutionary advances in optical component technology and network architecture. Traditional networking methods are facing a combination of hard physical limits, bottlenecks in energy and footprint, and new requirements well beyond what they were originally envisioned to handle. A focused and concerted research effort is required, combining multiple facets of technology development from optical components through network protocols and applications, and including design purview from the edge through the core of the network. Such an effort would necessarily span cross-disciplinary research groups, and include multiple, interrelated research topics. Addressing the key "grand challenges" necessarily requires a synergistic, community-wide approach with a common long-term vision toward realizing the transformation of future networks.

Our current networking infrastructure is predominantly built upon a separated static optical layer for delivering transport capacity. This underlying structure is simply incapable of effectively scaling to meet the explosive bandwidth demands of future Terabit-scale networks. New networking paradigms that aggressively drive dynamic functions deeper into the optical layer offer the opportunity to create future networks with terabits of readily accessible bandwidth at reduced costs and power requirements. Fueling these architectural advances will be new ways of leveraging emerging photonic technologies for introducing dynamic programmability directly at the physical layer. The envisioned next-generation multi-Terabit networks will intelligently scale to address the growing demands in information processing, storage, and transport enabling new and previously unimaginable applications.

To address the challenge of scaling terabit networks, the Center for Integrated Access Networks, the Optical Society, and the National Science Foundation convened a workshop in September, 2013. Participants in the workshop were selected from academia, industry, and governmental organizations, and represent the entire network stack. The technical program of the workshop was organized into a series of sessions that built toward the development of a prioritized list of research challenges specifically identifying candidate topics to form 'grand research challenges'. Three plenary speakers gave their perspectives on key issues in technical areas related to scaling future networks.

Underpinning the discussion of future networks is the identification of several concrete trends, obstacles, and opportunities. The workshop participants identified the following key network evolution issues: (1) **Optical technology is penetrating further to the edge** - as network hosts scale from 10 Gbps to 100 Gbps and beyond, optics is already reaching the top-of-rack switch in datacenters and migrating further to end user and home networks. Advanced optical network technologies are needed for a wide range of new application spaces including data center, metro, aggregation, mobile backhaul, and access networks, each of which has unique requirements and design constraints that are often very different from those of traditional backbone networks upon which most fiber systems are based today. (2) **New applications and services are demanding new architectures**, specifically to address "Big Data" and large scientific deployments. The consolidation of computing into high-density data centers will result in dramatic changes to access and aggregation networks. (3) **The rise of data centers** will require new network topologies as well as novel distributed control paradigms. In this context, software-defined networking is creating a critical need for (4) **new end-to-end control planes and virtualization**. These control planes will present compute, storage, and networking as a seamless, reconfigurable resource, rather than as discrete components. (5) **The proliferation of mobile networks**, are driving increased capacity requirements with a greater need for optics, leading potentially to new coding, advanced modulation formats, high capacity small cells, and distributed and smart antenna techniques. (6) **Spectral efficiency is reaching the Shannon limits of optical transmission**, and is leading to (7) **increased parallelism in optical networks**, which is fueling the transition to photonic integration and scalable platforms such as CMOS compatible Silicon photonics. New integrated and parallel photonics, tightly embedded with the electronics, may enable greater computer system disaggregation. Thus, not only are optical networks being used in new ways and for new applications, but the optical technologies themselves are undergoing an evolution, opening the door for innovation and potentially accelerating the impact of research.

The workshop participants identified the key research priorities grouped around **three** grand challenges and the need for a national scale multi-user networking test-bed:

1. **Programmable, Virtualized and Intelligent Optical Networks for the Future Internet**

Research is required to meet the grand challenge of realizing an optical network with intelligent capabilities that can support direct programmability in the physical layer, cognitive and autonomic control, and dynamically adaptive security capabilities. Innovations are needed to break through the current barriers created by physical complexities in order to realize a truly programmable optical network. The intelligent networking paradigm will foster new levels of optimization to meet the accelerating needs of the Future Internet. Programmability in optical networks would open up fundamentally new applications and network capabilities while providing efficient scaling to large volumes of traffic at finer granularities. Due to their intimate relationship with the physical layer properties, programmable optical networks may need significant changes to the upper layer network protocols and architectures.

2. **Cross-Layer Optical Network Architectures for Datacenters & Cloud Computing**

A grand challenge for future datacenters and cloud computing is the creation of new architectures that can essentially remove the network bottleneck and approach optical transit time limited delays. Addressing this challenge will require research on application-aware control and cross-layer functionality to the optical layer. Increased functionality in the lower layers, particularly through the use of optical switching and aggregation require information and/or coordination across multiple layers. Research is needed on new layering paradigms with network architectures designed for the needs of future datacenters and cloud computing. This grand challenge stresses end-to-end, edge-through-core, solutions that can eliminate conventional capacity bottlenecks.

3. **Clean Slate Architectures and Component Technologies for Optical Networks**

Research is needed on clean-slate architectures that can better define viable solutions for future requirements and serve as a guide for standards and network evolution. A holistic network design approach is envisioned for rethinking the architecture in a manner that leverages optics as the end-to-end interconnection and switching platform. In the different application spaces of metropolitan, enterprise, and datacenter networks, the requirements in terms of switching, aggregation, and traffic engineering are vastly different from long haul systems. New clean slate architectures that are co-designed for wireless-optical networks would yield increased efficiency and future scalability. Under this grand challenge, research on sustainable optical networks to support ubiquitous and low power monitoring and control would be developed for varied applications including the power grid, home heating and cooling, automobile traffic management, lighting, and home appliances. Greater access to Silicon photonic foundries and maturation of photonic integrated circuit technologies will launch a new generation of optical components that will both drive new optical network innovation and benefit from knowledge derived from clean slate optical networking research. This grand challenge will address efforts to co-develop innovative optical components and architectures together to best meet the challenges facing future networks.

4. **Collaborative, Multi-User Test-beds for Optical Terabit Scale Experiments**

The workshop participants identified the important need for creating a collaborative multi-user national scale test-bed specifically for Terabit optical networking experiments. This capability will be vital to each of the research grand challenges described above and would provide an essential competitive advantage relative to other research worldwide. It provides a common environment for proving-out next generation technology, and will help drive the common goals of the grand challenges across multiple research endeavors. This capability would also provide an experimental platform to study networks built on emerging devices based on Silicon photonics and other photonic integration and would be synergistic with the development of new foundry facilities. The envisioned test-bed infrastructure would provide a unique capability with the potential to enable the US to leap-frog Europe and Asia to once again take the lead in optical networking research.

Contents

Background Statement

Commercial transmission and switching systems support 100 Gb/s, with edge rates routinely exceeding 10 Gb/s and efficient optical transport reaching to the enterprise and home. Terabit per second signal transmission has been demonstrated in labs, with commercial systems on the horizon. Emerging applications are expected to comprise bursty, large and unscheduled transactions, with resulting network traffic exhibiting heavy tail behavior. Component commoditization is changing the traditional balance of storage, processing and communications resources, making way for new architectures. Furthermore, disaggregated computer platforms will bring optics past the top-of-rack switch directly to the end host. These trends demand research on innovative new terabit-scalable network architectures designed for end-to-end optical and electronic technologies.

New terabit-scale architectures must consider changes at all layers of the stack. Optical devices are very different from their electronic counterparts, in some cases without an electronic analog, so it's likely the optimum optical network architectures will not be the same as today's Internet architecture. Scaling transmission to terabit rates has exceeded serial scaling methods, driving ever-more parallelism. Bridging parallel optical connectivity with multi-core processing will result in higher overall system utilization, but only if that connectivity can be used effectively. Researchers must rethink how photonics can improve network scaling, and determine high-layer architecture changes that better exploit optics.

Today, Software Defined Networking (SDN) focuses primarily at the electronic layer and is being explored for the transport layer. In this workshop we considered optical communication and networking hardware technologies together with higher layer protocols for the exploration of new agile, programmable and efficient optical network architectures. Software programmable optical devices can figure prominently in new virtualizable and flexible architectures that incorporate optical layer agility and scalability. Cross-disciplinary research is needed for "programmable optics" to reach full potential and to address the concomitant control and management challenges for both systems and networks.

Networks and data centers that support scientific research are experiencing exponential traffic growth, and so research backbones (ESNet, I2, NLR) and regional networks are anticipated early adopters for terabit-scale networks. On the commercial side, video production, digital healthcare, warehouse-scale-computing and latency-driven financial applications will benefit immediately from this new capability.

Methodology

The technical program of the workshop was organized into a series of sessions that built toward the development of a prioritized list of research challenges specifically identifying candidate topics to form 'grand research challenges'. Throughout each of the sessions the workshop background statement above was used to guide the scope and technical focus of the activities. Participants were encouraged to suggest additional areas for discussion during the course of the meeting and suggest these either directly to the chairs or through the discussion sessions. In developing ideas for new topics of research, the following metrics were applied.

Research challenge metrics:

Good research challenge problems have roughly the following properties…

1. Admit to multiple approaches to solving them (we don't know up front the best approach – pulls in multiple groups)
2. Have, or are amenable to the formation of, communities that will want to work on them
3. Will require and will generate progress in a wide variety of areas. For example, integration of optical devices, networking, network addressing, network control, network host interfaces and operating systems, application structures.
4. Allow the community to carefully manage the risk associated with required invention of new science, engineering or devices
5. Are not problems that industry will solve on its own (at least anytime soon)
6. If solved will be actually deployable in some important context: don't require you to "boil the ocean" in order to have impact. Moreover, impactful results likely throughout project life.
7. It is easy to state and understand both the problem and its importance. You can describe them to (mostly) non-technical senior leadership and their impact will be understood.

Following the welcome session, three plenary speakers gave their perspectives on key issues in technical areas related to scaling future networks. They provided background data and motivation for the discussion sessions to follow. The main workshop activities took place during three breakout sessions and three plenary discussion/prioritization sessions. There was also a plenary panel session on day one in order to further frame and stimulate discussions. The breakout sessions progressed systematically toward the main workshop objective of defining and prioritizing key research areas and grand challenges.

The first session focused on the motivation and background: the 'why' behind the workshop. The second breakout session considered research directions from a metric and measurement point of view: what are important quantities to improve or targets to achieve? For the different metrics, participants also considered the capabilities that are required to measure and address these metrics—can it be measured? Is it achievable? The material and findings of the first day sessions were reviewed at the opening of the second day. This material then served as input to the third breakout session, which focused on defining a set of research challenges or priorities and their associated requirements. The lists and definitions generated in the third breakout session were reviewed in a plenary session. Participants were then asked to score the list electronically during the plenary prioritization session. The final technical plenary session was an open discussion to further clarify the final top 10 research topics and to identify the main stakeholders and early adopters.

Key Trends, Obstacles, and Opportunities

Continued traffic growth is the underlying engine that is driving and will largely shape the evolution of communication technologies and optical networks in particular. While this has been the case since the start of optical communications and the Internet, there have been important changes in recent years to the details of this story [1]. It's not just more of the same. Services and applications such as cloud computing, big data, and mobile data admit new statistical behavior and performance requirements. Emerging trends are anticipated to change where traffic is generated geographically and how it is distributed among the different network domains: enterprise, access, mobile backhaul, metropolitan, long haul, and data center. At the same time that the nature of the traffic is changing, there are technology shifts underway that will impact how the traffic demands are met and play an important role in shaping the economics of how networks are used. Optical systems are reaching physical limits that are forcing a move to parallelism. The electronics in communication and computing systems are increasing in parallelism and continue to encounter challenges associated with further Moore's law scaling. In addition, there are constraints beyond traffic growth that have recently started to place increased pressure on network technology and will be expected to gain importance over time. These include factors such as energy efficiency, greenhouse gas impact, and system footprint.

Optics Penetrates Further to the Edge

Optical transmission and switching technologies are most cost effective for communication at high capacity and over distance. Recently the tipping point at which optical systems have found application is at roughly in the bandwidth distance product range 10-100 Gb/s-meters [2]. Thus, most applications with data rates of 10 Gb/s and above will be most cost effectively handled in the network by optical systems and devices. 10 Gb/s interfaces are now common on servers and optics has become ubiquitous throughout the network from servers in the data center to the home. With optical USB entering the market, optics is even reaching to the home computer.

There is much interest in the industry in what is referred to as 'embedded optics'. Today, a server or switch will use a pluggable network interface card or unit, which contains all of the communication optics. At high data rates, however, the losses in moving data from these interface cards across the board to the CPU electronics becomes problematic. The solution is to use 'embedded optics' which place the network interface optoelectronics on the board close to the processing electronics, replacing the network interface card with a simple passive fiber connector. The consensus of the workshop panel was that the move to embedded optics is inevitable, driven by technology requirements. This trend has the potential to change the supply chain as stand-alone interface cards become on board components in a server or switch platform. At the same time this opens up the potential for on-board optics and a possible evolution to optics on chip—with the network interfaces built into the processor chips. Progress in Silicon photonics provides a technological basis for this trend. An important new capability that results from embedded optics is an optical 'data bus' with the potential to extend off the board.

Today optical or wavelength signals are used throughout data networks from the Internet backbone to the servers in data centers to the home networks and are anticipated to extend onto the board in the near future with the long term potential to even reach onto the processor cores. This trend is a direct consequence of the demand for higher data rates at the edge of the network.

New Applications and Services Demanding New Architectures

A number of trends on the application side have changed the way that we use networks. Most traffic now goes through large data centers located in each major metropolitan area. This trend is resulting in metropolitan traffic growing at a faster rate than the long haul/core traffic. Recent reports place cloud computing as the main traffic load, generating 34% of the total data center traffic [1]. Big data is a major new application that involves both collecting a high volume of small data sets and then manipulating and processing the very large aggregate data sets. Machine to machine traffic, particularly within and between data centers is a source of high capacity data flows or so called elephant flows. In fact today small user requests can generate large elephant flows within data centers—for example, to extract a single statistic out of a data set for a large population.

In big science applications the data source is generally not in the same place as the computational infrastructure. Low latency and high bandwidth is needed to move this science big data to the computers. Delocalization and global collaborative science are generating distributed big data sources that require distributed compute resources of a similar scale. Optical connections with high bandwidth and low latency could enable distributed compute resources to perform as if they were at the same location. A similar application in the commercial space is visualization in geographically distributed sessions including medical consultations or surgery.

At the same time that new applications are creating a need for very high bandwidth in demanding scenarios such as low latency or end to end connectivity, there are examples of an increasing number of small flows in the network. The widespread use of sensors and sensor networks, the Internet of Things, and smart energy systems are all examples of emerging applications that will increase the number of small transactions across the Internet. Some applications such as smart grid power control systems may have unique security, reliability, and latency requirements that will benefit from intelligence in the network. Thus, service differentiation may be essential even for applications that involve very small data sets. Using the network as a large, dumb pipe that carries all services indiscriminately does not allow for such differentiation. Recent years have seen a proliferation of complex rules in the IP layer, requiring deep packet inspection of many small transactions, an approach that does not scale well.

The traditional assumption that what works for long haul will work for other networks no longer applies. As optical networks are more widely used they must handle applications with increasingly more divergent requirements. Data center, campus, mobile backhaul, and metropolitan networks have unique aggregation and traffic management needs expressed over distances and topologies not found in the long haul.

Applications today are exhibiting an extreme divergence in their use-cases and service requirements from the network. This represents a major new challenge for optical networks. How can these diverse requirements be supported over a common network infrastructure? Is it needed or even practical to create different paradigms or architectures to meet these diverging needs? These challenges are being dealt with today by breaking the traditional network layering to create new functionality and by increasing the dynamic capabilities of equipment that in the past has been static. In the case of optical networks, protocols such as GMPLS are being bootstrapped to include optical transport design constraints such as optical impairment or reach information. These trends provide motivation for revisiting network layering and considering new architectures to optimize dynamicity across layers—considered using a holistic perspective from the application layer all the way down to layer 0.

In 1997, Internet design was guided by a network of plenty, not scarcity and transport was guided by the needs of data systems, not network design assumptions. As data systems have increasingly asked more of networks, fiber capacity has become a scarce resource and network constraints have been bootstrapped onto data protocols and management systems to diminishing effect.

Rise of Data Centers and Revolution in Computing Systems

The networks in data centers are receiving increased attention as the east-west or server to server traffic continues to increase rapidly. High bisection bandwidth requires many layers of electronic switching and routing that do not scale well. Large data centers can have petabits/second in port capacity with point to point fiber links that may reach close to 1 km. Cable management is a challenging problem. High port count optical switches have the potential to flatten the network and provide better scalability by pushing the network interfaces to the edge and allowing for graceful capacity growth over a transparent core switching infrastructure. Higher capacity and larger networks will likely lead to a transition to optical networks including wavelength division multiplexing in the fibers. Many of the engineering rules required in wide area networks will not apply over the relatively short distances in data centers opening the door to more switching and greater agility in the optical systems. As a self-contained single operator network, many of the inter-operator, inter-domain issues that complicate cross layer networking solutions are also removed. Data centers have proven to be an effective entry point for software defined networking and therefore may serve as a convenient springboard for many optical networking technologies.

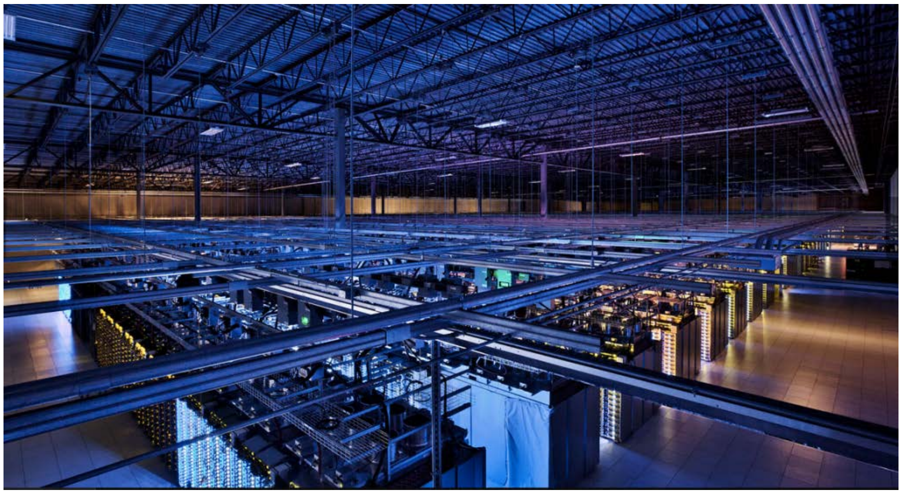

Figure 1. A large scale data center can include thousands of servers connected through point to point fiber connections covering hundreds of meters. Courtesy of Bikash Koley, Google.

Research on optical networks may provide opportunities to re-think next generation computing architectures. Maxims in the computing world that the network consists of fixed bandwidth pipes with perfect reliability could be removed. One example would be to make data the focus of the computer architecture and consider compute as free. This would elevate the data input/output interface and drive an integration of the computing systems with the network architecture and protocols.

Historical questions of whether computing should be centralized or distributed and whether the client should be thin or thick, continue to find new life as both the technologies evolve and new applications and services emerge. Increased use of cloud computing is creating a resurgence of the thin client paradigm in which the data center is now playing the role of the mainframe computational engine. The network is central in determining whether one approach is preferred over another. From a network perspective a key question may be: ‘where is the cache?’ Content delivery networks have been increasingly placing cache within the network in order to improve latency and other performance measures. There is much fine tuning of domain names and virtual topologies in order to further enhance performance, often in ways that are ad-hoc and not scalable over the long term. Future Internet architectures such as named data networking or information centric networking make use of an intelligent network with cache distributed throughout. Such in-network functionality can play an important role in how applications are implemented. With increasing traffic demands, optical networking needs to be considered in this context.

End-to-end control plane and virtualization for agility and efficiency

In data centers some vendors have moved away from the traditional server platform oriented architecture in which each server is running its own operating system and is assigned its own set of tasks as applications. The warehouse scale computing alternative is to introduce a unified cross-platform operating and file system that effectively uses the servers as pooled computational resources in which to most efficiently process applications (See Fig. 2). Similar opportunities may exist within networks. The parallel to the traditional server platform in networking is the OSI layer model in which the protocols in each layer serve as operating systems for the corresponding hardware: optical transport, label/circuit switching, and IP routing. Following the warehouse scale computing model, software defined networking (SDN) could be used to create a single, centralized or unified network operating system that efficiently utilizes the hardware resources across all layers [3].

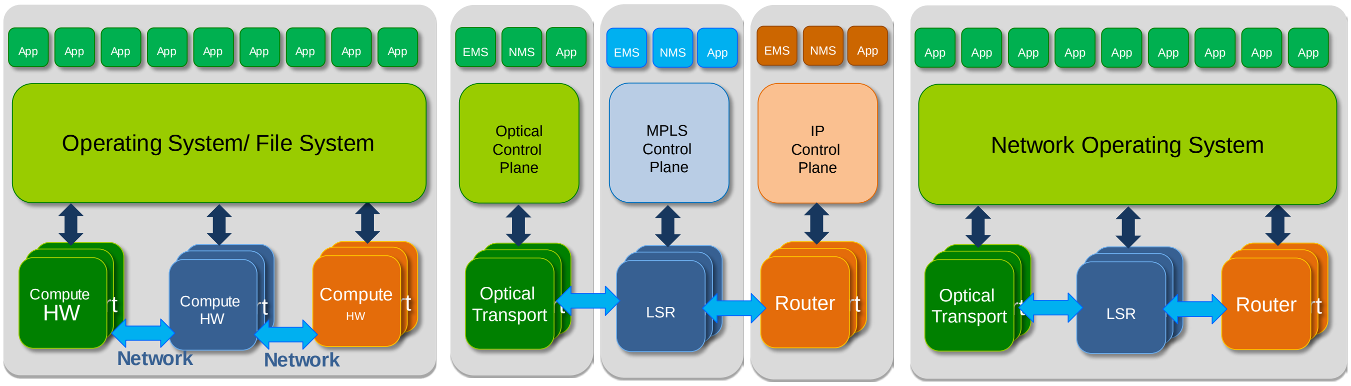

Figure 2. The warehouse scale computer data center model (left) can be applied to the traditional network layer model (middle) to create an SDN-based network operating system model (right). Courtesy of Bikash Koley, Google.

In some respects the traditional layering structure of networks has provided a form of virtualization of the optical infrastructure. A given wavelength channel is shared by 1000's of applications, and routers and switches efficiently groom signals into these wavelengths, to enable the most efficient sharing of optical bandwidth. With machine to machine, big data, or network as a service type applications, an individual application may require multiple wavelength channels. Furthermore, increasing bandwidth demand will likely lead to multiple parallel fiber systems, as is already occurring. This brings up the possibility of virtualization of the fibers or transmission systems/bands. Thinking of fibers in the same way as servers in a data center, network orchestration tools might dynamically and efficiently pack optical wavelengths into the minimum number of fibers—an ability that is likely to be critical from an efficiency standpoint when this parallel fiber networking regime is reached. While simple in concept, this would be extremely complex to realize with the current practice in optical transmission. Removing wavelength blocking and fragmentation may require multiple contingent operations distributed through a network. Transmission engineering rules and physical impairments will be difficult to calculate and manage, particularly on time scales that may be needed for efficient operation.

Mobile Networks: Increased capacity and novel Architectures

Through the use of novel antenna designs and small cells, mobile data has the potential to see another 100x or more increase in data rates to the end users. Already many mobile devices support functionality such as 4k high definition video that requires high data rates. Mobility has become the primary delivery portal for most services and relies heavily on the underlying wireline network infrastructure. This dependence was made clear during the early smartphone deployments in which the capacity limitations of copper backhaul networks, widely used at the time, proved to be a critical bottleneck. The subsequent rapid roll-out of optical backhaul alleviated the problem. As mobile data continues to increase, the solution for the underlying network may not be as straight forward. Optical networks today are not well suited to handle the inherent burstiness in both time and location of mobile data. The traditional optical system practice of provisioning one time for peak demand may not be feasible. Fiber infrastructure is expensive and available bandwidth within the fiber is becoming scarce. As coherent receivers and advanced modulation formats, including OFDM, are more widely used in optical systems these might be exploited in a similar way to how they are used in mobile systems.

In addition to increased application traffic, achieving high data rates in mobility places new requirements on the network in terms of coordination between cells for handover, particularly important in a small cell environment where it may happen frequently. Nanosecond scale timing markers may be needed, placing challenging requirements on the latency of mobile signals between multiple antenna sites and processing locations. Seamless hand-overs as users move will require transitions of session states that will amplify data bursting at the fiber edge.

Recently optics has been used in base stations to create ‘remote radio heads’ in which the antenna at the top of the tower is connected to the baseband processing unit through an optical fiber. This approach dramatically reduces the power requirements of the cabling to the antenna and has opened up the possibility for new architectures using highly distributed antennas and new multiple input multiple output (MIMO) configurations. An important question on the table today is: where should the baseband processing be placed? Is the traditional home in the cell tower baseband housing the right place or can it be moved to a data center for more efficient processing and better coordination amongst multiple cells? The optical network architecture will play a critical role, especially due to the difficult latency requirements for timing coordination.

Shannon Limits and Optical Transmission

As seen in Fig. 3, optical systems from the mid 1980’s through 2000 enjoyed a 1000-fold capacity increase, which was essential for the early growth of the Internet. Over the past decade, capacity growth slowed substantially as wavelength division multiplexing (WDM) technology matured. Continued growth is now following the serial interfaces trend, which is meeting new challenges as interfaces are nearing the fundamental Shannon channel capacity limit [4] [5]. Spectral efficiency in state of the art optical systems is within a factor of two of the Shannon Limit. At the same time installing new fiber and building out parallel systems is expensive and does not scale.

This situation is analogous to transitioning to smaller cells in mobile networks or moving to multi-core processors in computing systems. The key takeaway point is: there are no fundamental technological obstacles to increasing capacity in optical systems or networks today, the difficulty is that all of the known approaches for doing this involve adopting growth through parallelism. This represents an important change from data rate or serial system capacity scaling to parallel system scaling. Note that although WDM is a form of parallelism, its implementation was within a single system, fiber, and amplifier band. As a result, capacity could be grown by adding WDM interfaces at the edges of the network and used network wide within a single system—providing the scalability that was so important in the early days of the Internet. Moving to multiple parallel systems would require adding new amplifiers and networking gear at every node and hut in the network and possibly laying new fiber throughout the network. New optical network research is needed to develop techniques to create parallel systems that 'scale', meaning they can grow exponentially in dimensions such as capacity or number of nodes without incurring an exponential increase in cost, footprint or energy.

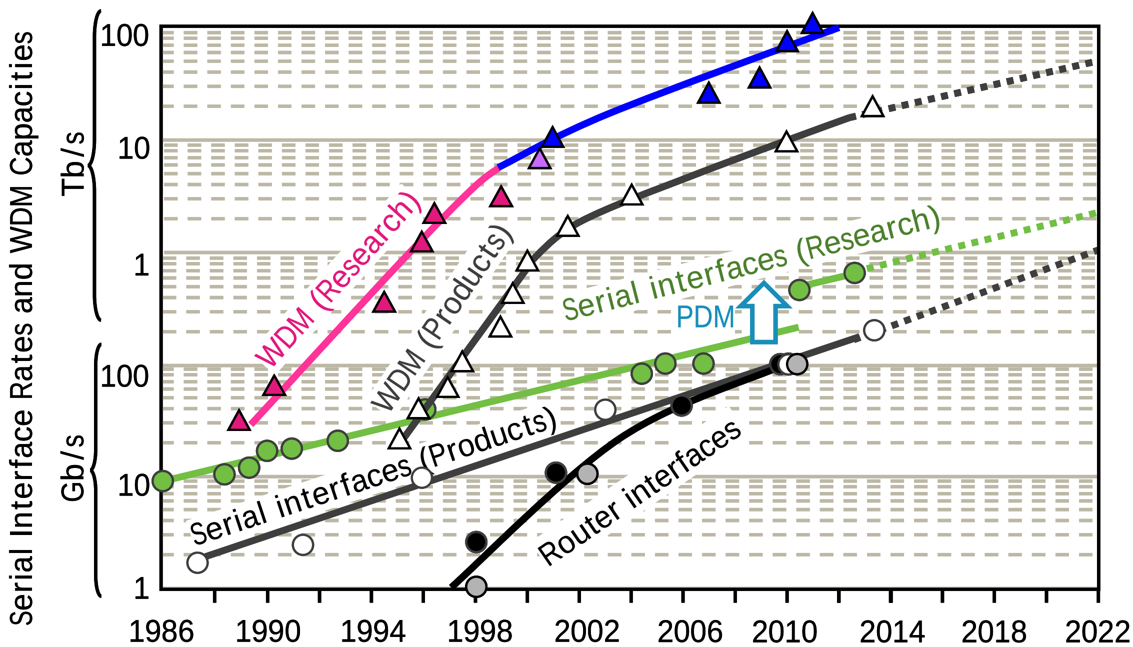

Figure 3. Trends in optical system and interface rates. Since 2005 both WDM system and Router (IP) interfaces have been following serial optical interface scaling, which in each case is much slower than historical growth rates. Courtesy of P. Winzer [6].

Moore's law is scaling faster than optical network capacity; traffic demands are growing faster than the data rate in a wavelength channel; metropolitan traffic is growing faster than core traffic; scaling up optical capacity is becoming a hard problem: if not addressed these factors will contribute to optical networks becoming a bottleneck to continued application growth and network expansion.

Parallelism in Optical Networks

Growth through increased parallelism brings greater focus to the networking side of optical systems, both because parallelism involves expansion through increasing the number of constituent elements and also because efficiency becomes a central strategy to achieve scalability. Systems with 10-100 optical fibers on a link and ~80 wavelength channels per fiber will have hundreds to thousands of signals to manage and switch. Bandwidth will become fragmented and wavelength blocking will limit network utilization. The practice of building optical networks for peak traffic has been effective because of steady exponential capacity growth. With this slowing, methods for adapting to traffic variations and resource sharing become important. If a network is provisioned for peak traffic then there is no value to turning off an under-utilized wavelength channel. In a dynamic system, however, such a signal would occupy spectrum thereby blocking that channel in that portion of the network from being used by other signals which need additional capacity. Increased network functionality in the optical systems would further open new possibilities for operations such as multi-casting, any-casting, and in-casting.

It is also the case that increases in data volume and traffic demands growing faster than the growth of data rate in one WDM channel is resulting in large relative granularity, more statistical fluctuation and shorter coherence times. High capacity services at the edge, requiring wavelength channel capacities will need dynamic and unscheduled session times for 10 milliseconds and longer. The protocols such as TCP/IP that would traditionally handle these sessions are not expected to scale to 100x greater speeds that today.

Similar to the case for electronics, parallel optical systems can benefit from increased integration. Photonic integration can take multiple discrete components and combine them into a single or multi chip device. Progress in this area will not only benefit the optical systems, but can also be exploited to better interface the optics with the electronics. Electronics are already well down the path of parallelism, particularly in communication systems for which high speed requires many parallel elements. In fact, parallelism in the electronics is running into complications due to long on chip and on board interconnects. Novel solutions that take advantage of parallelism in the optics and in the electronics may be possible if optical networks are extended to the board or chip level [7], for example through embedded optics. Serialization at the network interfaces could be removed enabling parallel optical systems to mate with parallel computing cores. Network addressable chips and memory would enable new levels of disaggregation opening up new possibilities both for the network and the computing systems. Emerging optical aggregation and modulation techniques such as orthogonal frequency division multiplexing (OFDM) might create interesting possibilities.

Networking and parallel system design in optics involves unique technical challenges that create both obstacles and opportunities. In short, for data center applications, when compared to electronic interconnects in computing systems, fiber links are essentially lossless and relatively independent of the data rate. This opens up the potential for new architectures using disaggregated or distributed resources and pooling of common elements for greater efficiency. At the same time, optical systems do not have a viable queueing mechanism for more than a few bits. In general electronics are more cost and energy efficient for computational tasks that involve multiple operations [8]. While there is still much opportunity for creative new optical devices in computation, the photon energy remains a hard physical limit that is orders of magnitude larger than the thermal and capacitive energies that limit electronic computing systems. This energy difference tends to drive cost and footprint scaling. For this reason, electronic devices are expected to continue to provide the intelligence and queueing functions in the network [9]. The challenge then is to find new ways to use these two technology platforms together to overcome the bottlenecks facing future networks. Optical networking will benefit from new or revisited approaches—short queues or circuits, separate control and data, analog processing—that are designed to best inter-operate with traditional packet networking and electronic processing and storage systems.

Metrics, Targets, & Capabilities

Capacity

Exponential scaling is essential in order for networks to keep pace with electronics, which continue to follow Moore's law trends. Network requirements are driven by increasing user traffic demands, with increasing numbers of users and devices worldwide, and with increasingly complex applications using the network. Ongoing research must be focused on solutions that will ultimately enable 100-1000 fold performance improvements over a 10 year time period. Already end devices and servers support up to 10 Gb/s, and thus edge networks will need to support 1-10 Tb/s and the long haul core would approach 1 Pb/s. Although the relative magnitude of long haul versus edge could change dramatically if the rate of edge traffic continues to increase faster than the core. As described above, emerging big data or elephant flow applications might be amenable to circuit-based network designs delivering Tb/s rates from end to end. On the opposite end of the spectrum, networks in the future will need to handle many millions of transactions at low latency. Today there is no direct path to achieving these levels of performance improvement in most cases.

Providing a Tb/s rate to the end user in a cost and energy effective manner on sub-millisecond time scales was postulated as a long term goal that, if achieved, would future-proof networks for the foreseeable future. However, such a network may not be well suited to handling millions of transactions at low latency. Thus while the ability to reposition the 'fat pipe' plumbing of the Internet all the way to the end user would be a major step forward, it is not clear that this would support the requirements of next-generation distributed systems. The brunt of the problem comes from the interplay between the network and the computing systems, which may differ widely among emerging applications.

Scalability

Scalability was an undercurrent throughout the discussion on metrics. While it is traditional to examine metrics in terms of the appropriate cost trade-offs, today it is equally important to include scaling properties. Performance against a particular metric must not only be considered against prevailing requirements, but also in light of increasing network dimensions such as traffic volume, the number of nodes or ports, and number of users and devices. Within the community, there value in better defining the sweet spots and capacity breakpoints for a given metric against its incumbent cost, and to understand how these may evolve over time and under different technology scenarios. For example, 10-20 Gbaud is a cost breakpoint for today's serial interfaces. Above 10-20 Gb/s modulation rates, it becomes more cost effective to move to advanced modulation formats or multi-carrier solutions.

Benchmarks

Another important issue for optical networking related metrics is that there are no convenient or commonly accepted benchmarks that can be used to compare different architectures or technologies. Progress in computing systems has historically benefited from benchmarks, such as LINPACK [10]. The internal complexity of computers makes benchmarks critical for obtaining quantitative comparisons on different implementations. Similarly, metrics such as bits/sec/km have been valuable in promoting research on long haul optical transmission systems. This metric provided a simple, clear goal that could be achieved through many different technical solutions, and was applicable to systems far below the Shannon limit. This metric was also well connected to the market as longer reach and higher capacity systems provided lower cost by reducing reliance on regeneration. As the systems approach the Shannon limit or optical add/drop and other networking functions were added to these systems, the situation became complicated as the number of node hops and other networking based performance measures become relevant. In metropolitan and smaller networks, transmission distance is no longer as essential and instead the node hops or amplifier hops becomes important. The bandwidth and distance related metrics also do not capture energy, time and utilization metrics, which are central to switched systems. A set of network performance benchmarks might provide a means to obtain quantitative performance comparisons for these metrics. Other additional reference architecture dimensions such as number of nodes and network diameter would likewise be needed since time and energy depend on these factors as well. Today, certain reference networks such as the NSFNET and European COST networks serve this function, often using canonical traffic patterns such as uniform loading.

Quantitative Metrics

Quantitative metrics that will be important for optical networks include:

- Network utilization
- Goodput and goodput variance
- Power and energy efficiency (joule/bit/network utilization factor)
- Reliability
- Availability (including. end to end)
- Latency and provisioning time
- Efficiency: spectral for optical systems, capacity and energy for networks, fewest bits sitting idle
- Agility

Timing Related Metrics

Agility, latency, and provisioning time are all closely related in optical networks and deserve special consideration as they can also strongly impact many of the other metrics. The physical complexities of optical transmission will often require significant system setup and tuning or provisioning of a channel when any switching or adaptation is applied to a wavelength signal. Thus latency in optical networks is more of a tuning or setup time delay rather than a jitter or congestion related phenomena. Important time scales for such delays were identified as:

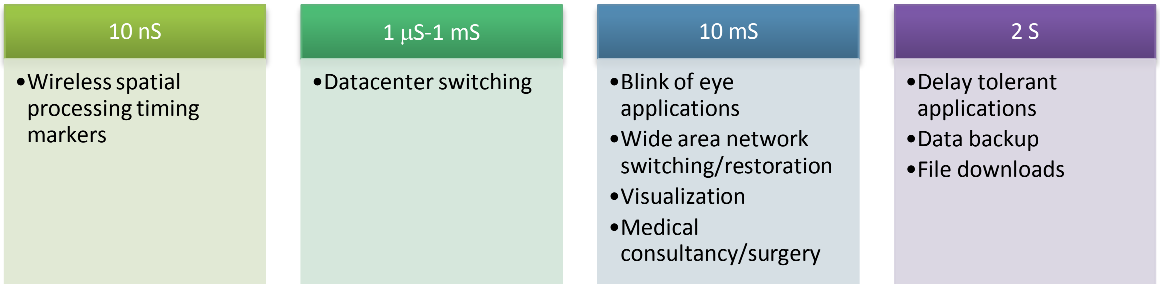

Figure 4. Timescales for dynamic service management

Spatial processing in mobile systems such as network MIMO or other distributed antenna systems have the tightest timing requirements. Optical network switching speeds in data centers have different timing requirements depending on the granularity or level of aggregation of the signals at which the optical switching is applied, and also on nature of the processing jobs. As optics moves closer to the servers and processing cores, faster switching is needed as the traffic becomes more bursty. Often the control plane, both for the optical systems and the higher layer processing, will be the limiting factor. Opportunities may exist to consider cross-layer or hybrid optical and electronic solutions to achieve the lowest latency and greatest agility in service delivery. In wide area networks, time of flight fiber transmission times can be on the order of milliseconds, and thus 10 ms can be considered to be on the order of time of flight. This time scale more than satisfies the traditional 50 ms protection switching requirement from telephony systems. It is also considered a good target for 'blink of the eye' applications such as remote surgery or data visualization across geographically distributed locations. For delay tolerant applications such as data backup or diurnal traffic management, switching delays can be as large as several seconds. Longer time delays would exhibit diminishing benefit as the setup delay should be much shorter than the service transaction time.

The time scales in Fig. 4 should be viewed as service delivery and management targets rather than optical switching targets. Optical switches with sub-nanosecond transition times have been commercially available for at least two decades. The challenge is to create a network that can deliver and manage services at these speeds. Higher speed optical switches might still be needed, but only for switches that provide the functionality required by the network. For example, high port count wavelength selective switches are currently limited to millisecond switching speeds. Recently a world record was reported for turning up capacity on a long haul link from Amsterdam to Hamburg: 8 Tb/s of capacity was set up in 19 minutes. This record established 16 optical signals at a rate of more than 1 minute per signal across a predetermined optical route. More complex switching or provisioning patterns over longer distances or larger networks would take longer. This highlights a critical aspect of optical switching and agility in networks: the speed depends on network details such as distances, number of nodes, and the amount of capacity involved in the operation. In applications such as data center networks, there is a similar dependence on a different set of factors such as the end-to-end loss, number of switching elements, the switching capacity, and the number of amplifiers in the case of amplified systems.

If we can shorten the network transport times between racks in data centers and among data/cloud centers in the WAN to only propagation times and a little more how is computing transformed? What happens if we are able to effectively extend computer data buses across networks?

Reliability and Resilience

Reliability in telecom systems is typically quantified through metrics such as the availability, outage probability and mean time between failures. Reliability of a network or system is determined through a probability analysis based on the failure in time (FIT) rates of the components. Carrier systems commonly apply a target of 'five-nines' or 99.999% availability measured as the fraction of down time in a network, for example, over a year. As the applications for optical networks have broadened with the convergence of transport and data centers, there may be opportunities to consider different requirements, definitions, and approaches to reliability in optical networks. In a future highly heterogeneous environment, the network resilience to failures rather than strict reliability metrics may be more relevant. In many data networking applications lower levels of availability may be acceptable. Indeed, dropped packets, congestion, and other forms of system outages are often accepted in order to enable the desired features at the given cost point. The reliability of applications and services across data and mobile networks are included in the overall cost-performance trade-offs and optimization. Of course some applications require very high reliability. An open question is whether optical networks should apply very high reliability as strict requirement across all applications as is done today or if its reliability should be part of the design and market trade-off. Today this is done to some extent through the use of protection. The gold standard of 1+1 protection (one back up connection for each service connection) is not universally applied and is sometimes sold as a service feature.

In computing systems redundancy is often used as a means to achieve reliability and different cost points. Using large numbers of lower reliability and therefore lower cost, commodity parts can achieve a high level of reliability through redundancy. Often this approach can achieve a target reliability at a lower cost than a system that is engineered from the bottom up with high reliability components. It also allows for service differentiation. As the economic and energy costs of redundancy become prohibitive, advanced high-performance systems are embracing the concept of resilience where software-hardware co-design is being used to create systems that are resilient to failures. This approach to reliability also relies on the ability for the overall system to efficiently adapt and repair from the more frequent failures in the individual sub-systems or components. While this may be problematic for traditional 'static' optical systems, new dynamic configuration and switching capability in optical systems may be a solution. Introducing shared protection and agility to optical networks could enable new methods of designing for reliability—with the potential to be more scalable and cost effective.

Qualitative Metrics

Important qualitative metrics include:

- Security
- Usability: from an operator perspective
- Reconfigurability
- Programmability
- Interoperability across multiple layers and domains/platforms

Security

As optical networks take on increasing functionality and reach broadened usage, security will become ever more important. Opportunities may arise to use optics in new ways to achieve security. In general optical networking has much the same security issues as electronic based networking. The fact that optical signals require electronics for any digital processing means that there are fewer opportunities to access the signal through conventional channels—fewer touch points. At the same time, this makes optics more vulnerable because it is more difficult to sense, authenticate and manage the security of optical signals. Furthermore, accessing the signal in the physical layer provides an entry point to all layers. If optical control planes are penetrated, particularly if they include advanced and automated functionality, they may trigger attacks across multiple layers. Centralized SDN control planes might provide localized points for attack. An open question is whether circuit based optical networks are more or less secure and what level or forms of authentication are required. Thus, optical networks or hybrid electronic and optical networks may have unique security requirements or implementations due to differences in how the data is bundled and handled. Both issues related to securing data and related to securing the network itself from attack need to be considered.

Adaptive Management and Control

The metrics for adaptive management and control in response to dynamic traffic or services will vary by application and network platform. In general for inter-domain or cross-layer protocols there is a need to understand the minimum set of information and methods for efficient operation—both in terms of the application data and of the management and control itself. Metrics should therefore incorporate the protocol and cross-layer coordination as part of addressing and assessing the overall network efficiency. For example, for the virtualization of network resources important metrics will include: the time to migrate services across the virtualized infrastructure, the isolation between virtual systems or networks, and the level of utilization achieved through the virtualization. A global target in this respect is to meet variable and increasing customer demand for a reasonable or scalable network infrastructure investment through better sharing of resources and increased utilization.

Mobility Related Metrics

For mobile systems the backhaul capacity will continue to be an important parameter and will drive greater reliance on optical networks. For small cell networks, optical systems will not only backhaul data, but also move large quantities of data between cell sites for spatial processing, hand-off coordination, and interference mitigation. Optical backhauling of sampled or native RF signals will introduce unique transmission performance requirements in terms of spectral efficiency, distortion, and reach. The inter-cell 'east-west' traffic may be similar to the situation in data center networks—short distances but high port counts and fast switching and management requirements. As indicated in Fig. 4, mobile systems may in fact have the tightest delay and delay jitter requirements. The reliability and speed of the handover of network and session states across all layers during movement between RF access points will be important performance parameters.

Research Priorities and Grand Challenges

Below is a prioritized list of the research areas, in the form of priorities and/or grand challenges, identified through the voting of the workshop participants. Note that no attempt was made to make these topics mutually exclusive. Therefore several identify similar or overlapping research areas. The voting results and the full list of topics is also provided in Appendix A and B, respectively.

1. Programmable Optical Networks with Intelligent Control Planes

Programmability in optical networks holds the potential to open up fundamentally new applications and network capabilities while providing the capability to efficiently scale systems to higher granularity and volumes of traffic. The potential exists to realize supercomputing on demand across the network, extend the computer data bus network wide, and unlock scalable and distributed data center functionality. A small user interaction with the cloud can create elephants in data and computing centers (on the scale of Terabytes). Currently, the optical physical layer network is a complex distributed system that utilizes cumbersome offline tools for management and provisioning, and although many functions are automated they often rely on an element of manual intervention. New applications will be unscheduled and bursty and can require setup times on the order of 10s of miliseconds, not minutes. This research topic calls for innovations to break through the barrier created by these physical complexities in order to realize a truly programmable optical network. Recent progress in the area of software defined networks provides a further basis for this functionality, but programmability should not be confined to SDN extensions, which is largely the direction pursued by industry and thus far has not adequately addressed the serious problems of rate adaptation and admission control during congestion and fast restoration to link and path failures. In order to be transformative, programmability is needed over a wide range of capabilities and requirements that include:

- Optical processes in response to dynamic unscheduled sessions that often have short delay service requirements
- Network wide allocation of resources with stability and fairness under network congestion
- Flexible and adaptive bandwidth provisioning/switching/adaptation, especially in the presence of network congestion

– Dynamically separating and optically differentiating services such as elephant and mice data flows
– Provide end-to-end data transport reliability and rate adaptability beyond TCP.

Forward looking programmability should be focused on creating new functionality and fast rate allocation adaptability to congestion. An essential part of this will be to expose more control knobs to the higher layers. Security, reliability, and stability must be maintained while opening up these additional degrees of freedom. In particular security should be designed into systems from day one, otherwise it will be difficult to retrofit and plug holes at a later time. Monitoring and control capabilities should be designed not for the first generation of programmable networks, but for future networks scaled up across multiple dimensions and delivering the full efficiency and functionality that will be needed. In particular, scalable network state sensing and reporting (as in traditional link state protocols) should be redesigned to deal with fast dynamic sessions and "elephants' that exacerbate widely varying traffic fluctuations. Solutions will be needed across multiple optical network domains or platforms, spanning data centers to metropolitan networks to long haul networks. However, such capability will be most needed in the data centers and near the edge of the network where there exist more service dependent demands and more bursty and changing traffic. Furthermore, the optical transmission engineering requirements are relaxed for these networks, making them more amenable to physical layer programmability.

A dynamically adaptive control plane of the type that deals with bursty elephant computer traffic will be a huge challenge. Conventional control planes are quasi-static and thus can be inefficient and slow and still satisfy service expectations. Future optical networks at many orders of magnitude increase in data rates will need fast adaptation. The amount of per session level network state reporting and the complexity of computation for the optimum (or sub-optimum) schedules for a general optical network with long reach represent a large amount of information being passed and large computation loads. The optical network architect must pursue designs that are computationally tractable and with a low burden of network control traffic.

Stake Holders and Early Adopters: Programmability will need to be implemented in the optical system hardware. Given the broadening range of network domains in which optics is now finding application, there are many entry points for new systems. The green field domains such as data centers or enterprise networks will likely be a source of early adopters. Similar to the case for SDN, the service providers and network operators stand the most to gain and will provide market pull, particularly if SDN finds more widespread success and can therefore provide a basis for optical network programmability to build upon.

2. Clean Slate Optical Network Architectures

Expected future ultra-high capacities and unprecedented traffic volatility will not be adequately addressed by conventional, incremental and backward-compatible approaches to optical network technology and architecture. What is needed is bold, and long-term research that seeks to avoid technology lock-in and encourages new components co-developed with novel architectures, designed with new services and operational paradigms in mind. Given the past 20+ years of optical technology development, if we were to start over with a clean slate, how would we design optical infrastructure?

Research on clean-slate architectures can better define viable solutions for future requirements and serve as a guide for standards and network evolution. Optical network technologies may also be achieving a level of maturity that can open new possibilities when considered in cooperation with the prevailing electronic technologies. As optics becomes viable at the compute board and chip level, new scaling, economics, and system architectures come into play. New directions for clean slate research include:

Holistic network design making use of optics as the end-to-end interconnection and switching platform. This includes rethinking the design of compute, storage, and networking in light of scaling and bandwidth afforded by optical technologies, cross-layer optimizations and re-design of the network layers, and opportunities to better integrate networks across technologies and layers. An important element of such clean-slate network research would be the development of new network validation or performance verification methods at scale—akin to a network benchmark such as the LINPACK benchmarks used for computational systems.

We can also re-design fiber highways, taking into account the different application spaces of metropolitan, enterprise, and data center networks where. the requirements in terms of switching, aggregation, and traffic engineering are vastly different from long haul systems New clean-slate architectures, targeted to unique application spaces, are needed, including issues such as optical switching, parallel fiber infrastructure, and new approaches to multiplexing and aggregation. We can consider opening up the fiber transparency window using un-amplified or wide-band amplification. The new architectures would be driven by the applications and therefore should be designed holistically together with the protocols and network layering and tier structures, much like dynamic spectrum allocation in mobile systems including development of cognitive network techniques to manage such systems.

Stake Holders and Early Adopters: This research area represents a long term investment in the overall network. Network operators and service providers will ultimately need to adopt the technologies that emerge from this program. Similar programs in the past such as the DARPA MONET program infused the community with new technologies and served to develop experts to fuel the transfer of those technologies to the marketplace, e.g. through the development of WDM systems.

3. New Optical Network Architectures for Datacenters & Cloud Computing

The high volume of east-west or intra-data center traffic and the emergence of cloud computing as a central application within data centers have brought focus to scalable high capacity networking within the data center. Recent studies have indicated that optical switching on time scales of 1 μs to 1 ms, accessible with many advanced optical technologies, may be suitable for many high bandwidth data center applications. Furthermore, these data center networks have unique requirements in terms of network scale, cost, footprint, and energy. Research on optical networking and hybrid optical and electronic networking is needed to address the emerging requirements for data centers. This includes optical architectures that can meet the highly dynamic demands of data center applications while scaling in port count and capacity well beyond current commercial systems. Optical systems will need to support data center functions such as point-to-point, narrowcasting for MapReduce type applications, multi-tenancy and virtualization, including the ability to form pools or puddles of network resources on demand. High rate optical interfaces are an important potential opportunity to achieve scalability, particularly considering optics that reaches onto the board and to the chip, exploiting multi-core processor architectures. A central goal for this research will be to realize new network architectures that become invisible to the computational infrastructure—removing the network as a bottleneck and approaching optical transit time limited delays.

Stake Holders and Early Adopters: Data center operators and cloud computing service providers will be the main stakeholders. Large 'mega' data center operators in particular will be the likely early adopters. Recent reports indicate that the network infrastructure in data centers is turned over every 3-5 years due to higher speed and energy efficiency requirements over time. Therefore innovation in this area will likely find rapid adoption and strong market pull.

4. Adaptive & Cognitive Optical Network Architectures

Networks today are supporting an ever wider range of applications and services many of which have diverse requirements. At the same time, the bandwidth demands of these applications are rapidly approaching the levels that are best handled using optical signals end-to-end without intermediate nodes converting optical signals to electronics for processing. The traditional approach of treating optical systems as fat and static pipes is not suited to these emerging trends. New optical network methods are needed that enable robust, rapid adaptation of the optical infrastructure. The ultimate per-transaction allocation of network resources requires almost real time sensing of network states (both occupancy and transmission impairments). Concepts such as elastic bandwidth using flexible data rate and modulation format optical transceivers and so-called flex-grid systems that remove the fixed channel spacing and grid are being applied today in commercial systems to improve the provisioning efficiency, dynamic load balancing and congestion control. This trend opens up the possibility to go much farther and realize real-time adaptation of spectrum utilization for even greater efficiency and flexibility. These capabilities would enable more efficient protection mechanisms including restoration through re-provisioning in an agile network. Networks could be made more resilient to disruptions, attacks, and technology evolution, while achieving higher levels of utilization. Adaptation also extends to impairment compensation and mitigation. Similar to mobile systems, such flexible bandwidth allocation would be driven by innovation in cognitive techniques. New on-line, in-network sensing and intelligence would be able to negotiate the diverse application requirements against the complex transmission engineering rules. Sophisticated offline system tools and manual controls and decision making must be replaced by distributed physical layer and cross-layer control algorithms. Solutions would need to balance global, centralized intelligence with control speed and scalability. The challenges are:

1. What is the abstraction of the Physical Layer and Data Link Control Layer properties that need to be passed up to the Routing and Transport Layers so network performance can be close to optimum?
2. Passing all link state information upwards may introduce too much complexity for the upper layers to perform fast resource allocation. Passing too little information may prevent the network from operating at high efficiency. What should be the sensible network states to report so the network control system decision algorithms can be manageable and scalable (see figure below)?

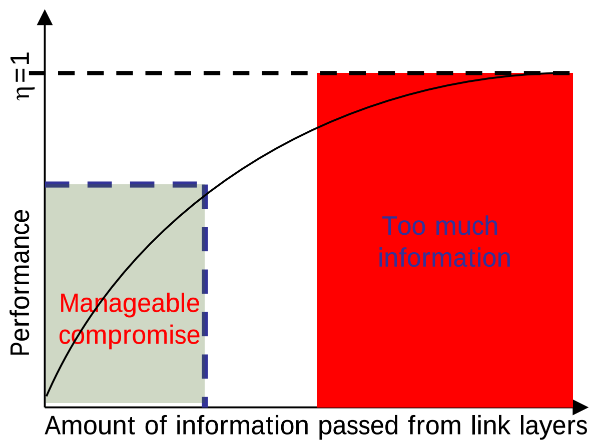

Stake Holders and Early Adopters: This topic shares the same considerations as that of topic #1. Cognitive networking capabilities will need to be implemented in the optical system hardware. The complexity of cognitive algorithms may help create a market for optical control system software as part of the supply chain to the optical system vendors.

5. Optical Network Architectures Across Layers - Cross Layer Design

The current layered network structure has been essential for the rapid growth and success of the Internet. Layering enables different network functions and operations to be designed and optimized separately. Problems that otherwise would be overly complex can be readily understood, formulated, and solved. Layering was also critical for allowing systems that were not designed for data communications to be used to support the Internet, e.g. building a packet network on top of a circuit switched telephone network. The dramatic convergence merging among today's networks is removing some of the historical drivers for layering. Many functions such as network survivability are implemented on multiple layers in ways that can be redundant or even conflicting. Increased functionality in the lower layers, particularly through the use of optical switching and aggregation require information and/or coordination across multiple layers. Research on new layering paradigms, building on top of an optical network, is needed to address these issues. Algorithmic solutions that work across the current layers 1-3 (physical layer to IP layer) will enable the automatic setup of optical signals.

This includes identifying minimal sets of parameters to share/communicate across layers for efficiency and simplicity. Redundant functions should be collapsed into coordinated, cross-layer operations. New cross-layer controls and algorithms should be designed to work seamlessly with emerging software defined networking standards and protocols.

Stake Holders and Early Adopters: This topic shares the same considerations as that of topic #1; however, the cross layer aspect necessarily brings in equipment vendors at different layers. Many of the larger equipment vendors offer a full line of transmission and switching equipment and therefore would be the most likely to use technologies developed within this program.

6. Application-Aware Optical Network Architectures

Many applications including big data, network as a service, and other cloud based processing require network resources at the capacities only afforded by optical systems. However, these applications have varied and unique network infrastructure requirements that demand a responsive network. What is needed is optical networks that are application-aware. Application-awareness relates to the capacity of an intelligent network to maintain current information about connecting applications and respond by optimizing both network functions and the other applications or systems that they control. By maintaining application state and resource requirements, resources are allocate on-demand, to best effect. A variety of features and functions can be considered including dynamic topology and path selection, latency, data rate, capacity, and protocols.

Current networks are not architected to support application specific functionality in the optical transmission systems. By implementing application-aware control down to the optical layer, new capabilities on the scale of a 'supercomputer on demand' would become possible in large part owing to the high capacities and low latencies afforded by optical signals. As an example, research could address the general problem of joint scheduling of compute, storage, and networking resources optimized for different applications or services, which would enable network-wide data center disaggregation. Methods to facilitate peer-to-peer parallelism at the application layer would enable end to end high bandwidth connections running many parallel applications, including big data processing. This would include scheduling and sharing of optical resources across multiple applications.

Stake Holders and Early Adopters: While this topic shares the same considerations as that of topic #1, application awareness will be of particular interest to large enterprises that wish to deliver differentiated services over the network. Network operators may be resistant to such capabilities unless a marketplace for such services is created or they are allowed to provide differentiation for their own services.

7. Collaborative, Multi-User Test-beds for Optical Physical Layer Experiments

Experimental progress on optical networks is challenging but essential due to the complex and distributed nature of these systems. Optical systems are fundamentally analog in nature, requiring detailed physical simulations encompassing fiber nonlinear effects and amplifier dynamics. Numerous approximations and abstractions are required in order to simulate even basic network behavior. At the same time, continental scale fiber communication networks can cost in the range of $100 million in equipment and require considerable support staff and expertise. Even small network systems are beyond the resources of an individual PI. Commercial systems in the field or in industrial research and development labs are highly proprietary and when access is granted for external research, limitations are placed on what can be published.

Research center level funding is needed to realize optical physical layer testbeds and field systems for optical network research and development. A key attribute of these systems would be that they include ‘dark fiber’ which can be lit with novel amplifiers, optical transceivers, and other experimental equipment. One goal would be to create an ‘optical sandbox’ to experiment with different networking innovations. Such testbeds would be synergistic with foundry capabilities in the US for integrated Photonics, particularly Silicon photonics. Phyiscal layer testbeds will be needed to study and evaluate the new generation of devices that would emerge from such a foundry. New understanding generated through experiments using these testbeds can motivative the development of novel devices or key performance targets. At the same time, in order to support research on topics such as programmable networks, application awareness, cross-layer networking, intelligent and cognitive networks, it is essential that the testbeds also provide infrastructure to support experiments that examine all layers of the protocol stack and even novel architectures that break away from layering. Exploring a federated model where currently funded application testbeds can be incorporated with an ‘optical sandbox’ would be of merit and will leverage the currently funded infrastructure. Initiatives could range from individual research centers considering different network platforms: datacenter, metropolitan, regional, or enterprise with a potential approach to federating them together to construct an end-to-end research platform, if needed by the researcher. Research should include a combination of field deployments, testbed implementations of field systems similar in scale to a commercial network test facility, and novel testbed implementations that emulate larger systems, for example the way a recirculating loop apparatus is used to emulate long haul optical transmission. Focus should be given to research to identify and to address bottleneck issues. Research community input would be needed in order to identify the right network dimensions, scope, and capabilities to cover a broad range of investigation. Facilities would also play a valuable role in education and training that is not available at universities today.

Stake Holders and Early Adopters: Similar to topic #2, this area is a long term investment that will promote research across many optical networking areas. However, such facilities would likely benefit optical system vendors, particularly smaller or startup organizations that would lack resources to build and maintain their own infrastructure. Testbed facilities may even be able to become self-supporting by offering services for a fee and provide a mechanism to ease the transition of research ideas into pre-production by early adopters like the R&E community. However, industry support should be balanced against basic research needs.

8. Co-Architected Optical and Wireless Networks

Today's fiber networks are essential for providing backhaul capacity to cellular data infrastructure. This convergence will be important for the migration to small cells for increased efficiency and higher data rates. Optics has recently found application within the base stations, providing an efficient high bandwidth connection between the top of the tower antennas and the baseband equipment enabling remote radio heads. A key question is how much inter-domain functionality for optical and wireless should be used? More specifically, when and where should baseband, digitized RF, or analog over fiber signals be used? New methods and devices may facilitate the aggregation and switching of these unique optical signals. Answers to these questions will depend on the mobile architecture and extent of small cells—nano, pico, femto—and distributed or array antennas. Optics relaxes delay and distance requirements opening the potential to place the digital signal processing at different locations from the antenna to the baseband unit or even in a central office or data center. New optical components designed for mobile systems may create further possibilities by achieving lower power, near transit time limited delay, and more efficient bandwidth utilization. Optical technologies may further advance the performance for RF over fiber solutions. Innovations in both the control plane and data plane can improve the coordination of mobile and optical systems. New research has the potential to take systems beyond just a simple bolting together of a picocell to an Ethernet switch, instead co-architecting the mobile and optical systems to create a more efficient and scalable solution.

Stake Holders and Early Adopters: Cellular network operators will be the primary customer for technologies in this area. Early adoption could come from operators who don't have their own fiber network and are looking to invest in the direction; however, the more likely case would be an operator that has both wireline and wireless networks, particularly those with optical access networks so that right of way and other issues are not a barrier. The market pull for increasing mobile data traffic will likely fuel rapid adoption of associated technologies, particularly where new standards are not required.

9. Network Optimization

As network traffic continues to increase and networks continue to expand in their utility and application space, efficiency and coordination become critical across multiple dimensions including cost, energy, footprint, and spectrum. This brings more value to network optimization. At the same time, increased intelligence and functionality in the optical systems introduces new problems that demand novel approaches to optimization. Optical systems have traditionally been quasi-static and optimization is carried out on network snapshots in time considering a set of resources and configuration objectives. Non-realtime or off-line algorithms with large complexity that needs minutes to hours to compute can be used. As networks become dynamic, optimized algorithms are needed for the time varying control and management systems. Not only are optimal configurations sought, but also optimal pathways between configurations and optimal control methods for reaching those configurations. Due to the complexity and scale of optical networks, optimization methods will require new heuristics, and real time controls will need new on-line methods for fast and reliable estimation. New opportunities exist for multi-layer and multi-domain scenarios in light of emerging optical capabilities and agility. Pricing across network resources along multiple dimensions can likewise be highly dynamic and time varying, requiring novel fast optimization strategies. The new network algorithms must work with the new paradigm of operating sometimes without detailed network states or stale network states due to fast dynamics in the network. Network optimization should be an integral part of the new control plane of the optical network that we have alluded to before and should be designed with agility and scalability in mind.

Stake Holders and Early Adopters: Research in this area will benefit both service/network providers and the optical system vendors. As the control plane for optical systems becomes more complex, there might develop a market for control plane software. System vendors and network operators pursuing more advanced optical network functionality would likely be the early adopters.

10. Economic, Reliable Networking over Unreliable Optical Sub-Systems

Reliability and scaling of large systems can be realized through a variety of different strategies. Communication networks have typically followed the telephony model of developing highly reliable parts in order to build a large, reliable system. One approach that has found success in the 'big data' data center community is to build data center systems from low reliability elements and achieve high system reliability through redundancy and higher layer recovery. Reducing the reliability requirements on the constituent parts in this case enables the use of lower cost, commodity parts and ultimately speeds up the rate at which innovative technologies are adopted. Economies of scale result in a potentially lower cost solution that achieves high system reliability. Also important in these systems is an architecture that is designed anticipating a high failure rate and therefore is able to seamlessly adapt and support low cost maintenance practices. The current high reliability constraint on optical network equipment sets a high bar that impacts most aspects of the system architecture and optical research and innovation pipeline. Research into the cost reliability trade-off for optical networks may unlock new performance and scaling that could not be achieved in systems that must meet high reliability requirements down to the component level. Progress in this area would require new control plane and management strategies for less reliable components, and may even require co-design with specific applications. Protection can be traded off with agility to explore new protection and restoration architectures to accomplish system-wide resilience.

Stake Holders and Early Adopters: Optical network/service providers would be the primary customers of technologies developing from this area. Newer service providers and data center vendors and 'big data' providers that already utilize these methods for their compute hardware would likely be the early adopters.

Other Noteworthy Topics

In voting on different topics there are a number of risks that may cause important areas to be overlooked. First, cross-disciplinary topics may not fare well depending on the mix of workshop participants. Although a wide range of expertise was represented at the meeting, the majority of participants work in the area of optical networking. Another problem is that by going through two rounds of voting, topics that are similar may divide votes in the first round and therefore neither would make it through to the final round. With this in mind, here we examine and look for themes in the first round topics that did not make it through to the second round, with special emphasis on cross-disciplinary topics. Indeed, cross-disciplinary subjects did not do well in this process, but because of their high impact potential deserve particular attention. These topics might benefit from more focused workshops or progress within the community in order to develop the research concepts further.

Security in Optical Networks

Security was a topic of much discussion during the workshop and was divided into two main subtopics: (1) security of the network and (2) security of the data. With regard to the network, there are many questions: who gets to turn a particular knob, and when? who gets access? and what third-party support is provided? Can optical network control flexibility be exposed to operators and users without compromising security? If information about network control is publicly known it might open the possibility for attacks that are designed to exploit the existing control cycles, e.g., through timed attacks. New optical circuits will need to be authenticated and secured. Methods for physical layer fingerprinting might be developed as well as ways to secure and identify 'optical trustable components.' Additional monitoring and controls may be needed throughout the network to detect attacks or intrusion. This includes not only the control plane, but the optical network elements themselves—optical amplifiers, fiber plant, and transceivers. From a data security point of view, there is the question of optical tapping. In the past, methods such as optical chaos and quantum cryptography have been studied as ways to provide data security for the native optical signal. These and other techniques might be applicable to security in optical networks.

Co-Design of Optical and Electronic Networks

Several interesting directions for research in this area were introduced during the workshop. The first was the concept of designing optical networks based on 'embedded optics' at the edge. This unifies the network and computing systems to create a new paradigm worthy of exploration. Embedded optics was considered by some as an inevitable technology path. Another topic considered research into new electronic techniques that are optimized for integrating optics economically across all layers and platforms. Coherent optical receivers might benefit from analog signal processing rather than doing everything purely digital. Research in this area would also consider how best to use parallelism in the electronic cores together with parallelism in the optical systems. Going beyond the hardware integration, another research topic called for the co-design of computer and network architectures requiring a network aware operating system, software, and applications. Research could embrace a data-centric view in which bandwidth and network connectivity is the expensive resource to be most efficiently utilized. Research in the area would benefit from greater access to foundry capabilities for integrated photonics and help to accelerate developments coming from such facilities.

Sustainable Optical Networks

Networks are frequently relied upon for so-called smart energy solutions, in which greater energy efficiency is achieved through better monitoring and control of the energy use. Applications include the power grid itself, home heating and cooling, automobile traffic management, lighting, and home appliances [11]. Designing networks to support these types of services can be considered a form of application-aware networking. However, future applications will likely need to consider the energy efficiency and sustainability of the networks that are delivering the smart services. This would include optical network elements designed to operate off renewable energy sources, which may deliver limited and highly variable power. Data centers today are being located near large sources of lower carbon energy to increase efficiency and carbon impact. Networks are by nature distributed and ubiquitous connectivity will require other approaches, most likely leveraging distributed micro-grids of renewable energy. Designing for a limited, distributed, and highly variable power source may dramatically alter the network design choices and optimizations. Sustainability should also be considered from a full lifecycle perspective including material sustainability. Many materials common in network devices such as Indium and Tantalum have questionable supply levels. Heavy reliance on rare materials may make Internet growth dependent on global mining markets.

Challenge Statements

Statements of potential research challenges emerged throughout the workshop and are reflected in many of the research topics described above. In this section we highlight a few statements that might serve as alternative ways of describing the different research topics or better bringing attention to the topics.

"Super computer on demand", "Optical network as a computer bus"

Realizing "super computer" computational capabilities across a network or essentially turning the wide area networks into an extension of the computer bus would require a high level of intelligence and agility for large granularity data flows. Progress would be required across the topics of application awareness, programmable optical networking, cross-layer networking, and cognitive networking. It would also address cooperation between computing and optical network systems.

"Fiber-Optic Information Mainstreet" or "Mainstreet Optical Networks"

Shifting focus from the long haul or backbone networks—the information super highways—to the metropolitan, enterprise, mobile, and data center networks was a pervasive theme at this workshop. The concept of optics for the information mainstreet is an intriguing way to express this that broadly encapsulates many of the research areas identified here.

"Any Data, Any Compute (Any Where?)"

This topic is similar to the notion of a super computer on demand, but is more general and allows for the consideration of small as well as large or big data flows. It also brings in the concept of cooperation with mobile networks if 'any where' is included. Currently there is a large hadron collider physics project under this name. Applying this to optical networks opens up new dimensions in terms of flexible capacity.

Grand Challenges

In this section we group the different research topics into prominent themes that address major research challenges related to network scaling into the Terabit regime.

A. Programmable, Virtualized and Intelligent Optical Networks for the Future Internet

This grand challenge is based on research priorities: 1, 4, and 9, along with the other noteworthy topic of Security in Optical Networks

Networks of the future are expected to support a wide range of applications and services with very diverse requirements. At the same time, the bandwidth demands of these applications will be at levels that are best handled using optical signals end-to-end without intermediate nodes converting optical signals to electronics for processing. The traditional approach of treating optical systems as fat and static pipes is not suited to the Internet of the Future. Research is required to meet the grand challenge of realizing an optical network with intelligent capabilities for the needs of the Future Internet. Such a network would support features such as programmability in the physical layer, cognitive and autonomic control capabilities, and the capabilities to secure both the data and the network itself from attack, intrusion, and theft. A more intelligent network will also allow new levels of optimization and require new approaches to network optimization.

Programmability in optical networks might open up fundamentally new applications and network capabilities while providing the capability to efficiently scale systems to higher granularity and volumes of traffic. The potential exists to realize supercomputing on demand across the network, extend the computer data bus network wide, and unlock scalable and distributed data center functionality. A small user interaction with the cloud can create elephants in data and computing centers (~TB). Currently, the optical physical layer network today is a complex distributed system that utilizes cumbersome offline tools for management and provisioning, and although many functions are automated they often rely on an element of manual intervention. New applications will be unscheduled and bursty that sometimes require setup times of the order of 10s of mS not minutes. This grand challenge calls for innovations to break through the barrier created by these physical complexities in order to realize a truly programmable optical network. Recent progress in the area of software defined networks provides a further basis for this functionality, but programmability should not be confined to SDN extensions, which is largely the direction pursued by industry but thus far have not adequately addressed the serious problems of rate adaptation and admission control during congestion and fast restoration to link and path failures. In order to be transformative, programmability is needed over a wide range of capabilities and requirements that include:

- Optical processes in response to dynamic unscheduled sessions that often have short delay service requirements
- Network wide allocation of resources with stability and fairness under network congestion
- Flexible and adaptive bandwidth provisioning/switching/adaptation especially at network congestions

- Dynamically separating and optically differentiating services such as elephant and mice data flows
- Provide end-to-end data transport reliability and rate adaptability beyond TCP.

Forward looking programmability should be focused on creating new functionality and fast rate allocation adaptability to congestion. An essential part of this will be to expose more control knobs to the higher layers. Security, reliability, and stability must be maintained while opening up these additional degrees of freedom. In particular security should be designed into systems from day one, otherwise it will be difficult to retrofit and plug holes at a later time. Monitoring and control capabilities should be designed not for the first generation of programmable networks, but for future networks scaled up across multiple dimensions and delivering the full efficiency and functionality that will be needed. In particular, scalable network state sensing and reporting (as in traditional link state protocols) should be redesigned to deal with fast dynamic sessions and "elephants' that exacerbate widely varying traffic fluctuations. Solutions will be needed across multiple optical network domains or platforms, spanning data centers to metropolitan networks to long haul networks. However, such capability will be most needed in the data centers and near the edge of the network where there exist more service dependent demands and more bursty and changing traffic. Furthermore, the optical transmission engineering requirements are relaxed for these networks, making them more amenable to physical layer programmability.

A dynamically adaptive control plane of the type that deals with bursty elephant computer traffic will be a huge challenge. Conventional control planes are quasi-static and thus can be inefficient and slow and still satisfy service expectations. Future optical networks at many orders of magnitude higher data rates will need fast adaptation. The amount of per session level network state reporting and the complexity of computation for the optimum (or sub-optimum) schedules for a general optical network with long reach represents a large amount of information being passed and large computation loads. The optical network architect must pursue designs that are computationally tractable and with a low burden of network control traffic.

The ultimate per-transaction allocation of network resources requires almost real time sensing of network states (both occupancy and transmission impairments). Concepts such as elastic bandwidth using flexible data rate and modulation format optical transceivers and so-called flex-grid systems that remove the fixed channel spacing and grid are being applied today in commercial systems to improve the provisioning efficiency, dynamic load balancing and congestion control. This trend opens up the possibility to go much farther and realize real-time adaptation of spectrum utilization for even greater efficiency and flexibility. These capabilities would enable more efficient protection mechanisms including restoration through re-provisioning in an agile network. Networks could be made more resilient to disruptions, attacks, and technology evolution, while achieving higher levels of utilization. Adaptation also extends to impairment compensation and mitigation. Similar to mobile systems, such flexible bandwidth allocation would be driven by innovation in cognitive techniques. New on-line, in-network sensing and intelligence would be able to negotiate the diverse application requirements against the complex transmission engineering rules. Sophisticated offline system tools and manual controls and decision making must be replaced by distributed physical layer and cross-layer control algorithms. Solutions would need to balance global, centralized intelligence with control speed and scalability. The challenges are:

1. What abstractions of the Physical Layer and Data Link Control Layer properties are needed to be passed to the Routing and Transport Layers so network performance can be close to optimum?
2. Passing all link state information upwards may introduce too much complexity for the upper layers to perform fast resource allocation. Passing too little information may prevent the network from operating at high efficiency. What should be the sensible network states to report so the network control system decision algorithms can be manageable and scalable (see figure below)?

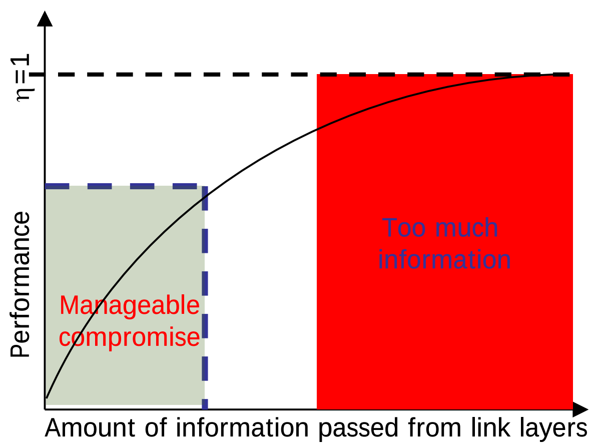

Increased intelligence and functionality in the optical systems introduces new problems that demand novel approaches to optimization. Optical systems have traditionally been quasi-static and optimization is carried out on network snapshots in time considering a set of resources and configuration objectives. Non-real-time or off-line algorithms with large complexity that needs minutes to hours to compute can be used. As networks become dynamic, optimized algorithms are needed for the time varying control and management systems. Not only are optimal configurations sought, but also optimal pathways between configurations and optimal control methods for reaching those configurations. Due to the complexity and scale of optical networks, optimization methods will require new heuristics, and real time controls will need new on-line methods for fast and reliable estimation. New opportunities exist for multi-layer and multi-domain scenarios in light of emerging optical capabilities and agility. Pricing across network resources along multiple dimensions can likewise be highly dynamic and time varying, requiring novel fast optimization strategies. The new network algorithms must work with the new paradigm of operating sometimes without detailed network states or stale network states due to fast dynamics in the network. Network optimization should be an integral part of the new control plane of the optical network that we have alluded to before and should be designed with agility and scalability in mind.

With intelligence comes a need for greater security so that more advanced functionality is not misused and it also creates the potential for new and more sophisticated methods to ensure security. Security can be divided into two main subtopics: (1) security of the network and (2) security of the data. With regard to the network, there are many questions: who gets to turn a particular knob, and when? who gets access? and what third-party support is provided? Can optical network control flexibility be exposed to operators and users without compromising security? If information about network control is publicly known it might open the possibility for attacks that are designed to exploit the existing control cycles, e.g., through timed attacks. New optical circuits will need to be authenticated and secured. Methods for physical layer fingerprinting might be developed as well as ways to secure and identify ‘optical trustable components.’ Additional monitoring and controls may be needed throughout the network to detect attacks or intrusion. From a data security point of view, there is the question of optical tapping. In the past, methods such as optical chaos and quantum cryptography have been studied as ways to provide data security for the native optical signal. These and other techniques might be applicable to security in optical networks.

B. Cross-Layer Optical Network Architectures for Datacenters & Cloud Computing

This grand challenge is based on research priorities: 3, 5, and 6.

Cloud computing and datacenters have become an essential part of the Internet and the backbone of our cyber economies and cyber social infrastructure. Applications such as Big Data, Social Networking, and Online Customer Relations Management are dependent on datacenters and cloud accessability. Big Science, involving the processing and management of large data sets, and other scientific computing is now frequently handled through datacenters in the cloud. With Zettabytes of data moving through datacenters, networks both inside and between data centers are now becoming a bottleneck for many applications. With 30% traffic growth expected per year for the foreseeable future, the networks in datacenter and for cloud computing applications will need to scale up by 10-100x to go beyond the next decade. Realizing optical networks to meet this demand for cloud computing and intra-data center applications is a grand research challenge. A key part of this challenge is to enable high bandwidth cloud applications to efficiently communicate through the layers to access optical or wavelength scale performance, thus creating an application aware physical layer.

The high volume of east-west or intra-data center traffic and the emergence of cloud computing as a central application within data centers have brought focus to creating scalable high capacity networking within the data center. Recent studies have indicated that optical switching on time scales of 1 μs to 1 ms, accessible with many advanced optical technologies, may be suitable for high bandwidth data center applications. Furthermore, these data center networks have unique requirements in terms of network scale, cost, footprint, and energy. Research on optical networking and hybrid optical and electronic networking is needed to address the emerging requirements for data centers. This includes optical architectures that can meet the highly dynamic demands of data center applications while scaling in port count and capacity well beyond current commercial systems. Optical systems will need to support data center functions such as point-to-point, narrowcasting for MapReduce type applications, multi-tenancy and virtualization, including the ability to form pools or puddles of network resources on demand. High rate optical interfaces are an important potential opportunity to achieve scalability, particularly considering optics that reaches onto the board and to the chip, exploiting multi-core processor architectures. A central goal for this research will be to realize new network architectures that become invisible to the computational infrastructure—removing the network as a bottleneck and approaching optical transit time limited delays.

Many cloud and datacenter applications including big data, network as a service, and other cloud based processing require network resources at the capacities only afforded by optical systems. However, these applications have varied and unique network infrastructure requirements that demand a responsive network. What is needed is optical networks that are application-aware. Application-awareness relates to the capacity of an intelligent network to maintain current information about connecting applications and respond by optimizing both network functions and the other applications or systems that they control. By maintaining application state and resource requirements, resources are allocated on-demand, to best effect. A variety of features and functions can be considered including dynamic topology and path selection, latency, data rate, capacity, and protocols.

Current networks are not architected to support application specific functionality in the optical transmission systems. By implementing application-aware control down to the optical layer, new capabilities on the scale of a 'supercomputer on demand' would become possible in large part owing to the high capacities and low latencies afforded by optical signals. As an example, research could address the general problem of joint scheduling of compute, storage, and networking resources optimized for different applications or services, which would enable network-wide data center disaggregation. Methods to facilitate peer-to-peer parallelism at the application layer would enable end to end high bandwidth connections running many parallel applications, including big data processing. This would include scheduling and sharing of optical resources across multiple applications.

The current layered network structure has been essential for the rapid growth and success of the Internet. Layering enables different network functions and operations to be designed and optimized separately. Problems that otherwise would be overly complex can be readily understood, formulated, and solved. Layering was also critical for allowing systems that were not designed for data communications to be used to support the Internet, e.g. building a packet network on top of a circuit switched telephone network. The dramatic convergence merging among today's networks is removing some of the historical drivers for layering and this is particularly true for datacenters and cloud computing, where many of the traditional rules do not apply. Many functions such as network survivability are implemented on multiple layers in ways that can be redundant or even conflicting. Application awareness requires cross-layer functionality. Increased functionality in the lower layers, particularly through the use of optical switching and aggregation require information and/or coordination across multiple layers. Research on new layering paradigms, building on top of an optical network and designed for the needs of future datacenters and cloud computing, is needed to address these issues. Algorithmic solutions that work across the current layers 1-3 (physical layer to IP layer) will enable the automatic setup of optical signals for cloud or datacenter applications. This includes identifying minimal sets of parameters to share/communicate across layers for efficiency and simplicity. Redundant functions should be collapsed into coordinated, cross-layer operations in support of cloud and datacenter functions. New cross-layer controls and algorithms should be designed to work seamlessly with emerging software defined networking standards and protocols.

C. Clean Slate Architectures and Component Technologies for Optical Networks

This challenge combines research priorities 2, 8 and 10, along with the other noteworthy topics of co-design of optical and electronic networks and sustainable optical networks.

Expected future ultra-high capacities and unprecedented traffic volatility will not be adequately addressed by conventional, incremental and backward-compatible approaches to optical network technology and architecture. What is needed is bold, and long-term research that seeks to avoid technology lock-in and encourages new components co-developed with novel architectures, designed with new services and operational paradigms in mind. Given the past 20+ years of optical technology development, if we were to start over with a clean slate, how would we design optical infrastructure?

Research on clean-slate architectures can better define viable solutions for future requirements and serve as a guide for standards and network evolution. Optical network technologies may also be achieving a level of maturity that can open new possibilities when considered in cooperation with the prevailing electronic technologies. As optics becomes viable at the compute board and chip level, new scaling, economics, and system architectures come into play. New directions for clean slate research include:

Holistic network design making use of optics as the end-to-end interconnection and switching platform. This approach includes rethinking the design of compute, storage, and networking in light of scaling and bandwidth afforded by optical technologies, cross-layer optimizations and re-design of the network layers, and opportunities to better integrate networks across technologies and layers. An important element of such clean-slate network research would be the development of new network validation or performance verification methods at scale—akin to a network benchmark such as the LINPACK benchmarks used for computational systems.

Re-designing the fiber highways. In the different application spaces of metropolitan, enterprise, and datacenter networks, the requirements in terms of switching, aggregation, and traffic engineering are vastly different from long haul systems. New clean-slate architectures, targeted to these unique application spaces, are needed, including issues such as optical switching, parallel fiber infrastructure, and new approaches to multiplexing and aggregation. We can consider opening up the fiber transparency window using un-amplified or wide-band amplification. The new architectures would be driven by the applications and therefore should be designed holistically together with the protocols and network layering and tier structures, much like dynamic spectrum allocation in mobile systems including development of cognitive network techniques to manage such systems.

Economic, reliable optical networking over lower reliability sub-systems. Reliability and scaling of large systems can be realized through a variety of different strategies. Communication networks have typically followed the telephony model of developing highly reliable parts in order to build a large, reliable system. One approach that has found success in the 'big data' data center community is to build data center systems from low reliability elements and achieve high system reliability through redundancy and higher layer recovery. Reducing the reliability requirements on the constituent parts in this case enables the use of lower cost, commodity parts and ultimately speeds up the rate at which innovative technologies are adopted. Economies of scale result in a potentially lower cost solution that achieves high system reliability. Also important in these systems is an architecture that is designed anticipating a high failure rate and therefore is able to seamlessly adapt and support low cost maintenance practices. The current high reliability constraint on optical network equipment sets a high bar that impacts most aspects of the system architecture and optical research and innovation pipeline. Research into the cost reliability trade-off for optical networks may unlock new performance and scaling that could not be achieved in systems that must meet high reliability requirements down to the component level. Progress in this area would require new control plane and management strategies for less reliable components, and may even require co-design with specific applications. Protection can be traded off with agility to explore new protection and restoration architectures to accomplish system-wide resilience.

Co-Architected Optical and Wireless Networks. Today's fiber networks are essential for providing backhaul capacity to cellular data infrastructure. This convergence will be important for the migration to small cells for increased efficiency and higher data rates. Optics has recently found application within the base stations, providing an efficient high bandwidth connection between the top of the tower antennas and the baseband equipment enabling remote radio heads. A key question is how much inter-domain functionality for optical and wireless should be used? More specifically, when and where should baseband, digitized RF, or analog over fiber signals be used? New methods and devices may facilitate the aggregation and switching of these unique optical signals. Answers to these questions will depend on the mobile architecture and extent of small cells—nano, pico, femto—and distributed or array antennas. Optics relaxes delay and distance requirements opening the potential to place the digital signal processing at different locations from the antenna to the baseband unit or even in a central office or data center. New optical components designed for mobile systems may create further possibilities by achieving lower power, near transit time limited delay, and more efficient bandwidth utilization. Optical technologies may further advance the performance for RF over fiber solutions. Innovations in both the control plane and data plane can improve the coordination of mobile and optical systems. New research has the potential to take systems beyond just a simple bolting together of a picocell to an Ethernet switch, instead co-architecting the mobile and optical systems to create a more efficient and scalable solution.

Co-design of optical and electronic networks. This approach can include several different aspects of the design and operation of optical and electronic based networking technologies. One strategy is to design optical networks based on 'embedded optics' at the edge. This unifies the transmission systems with the computing and networking systems to create a new paradigm. Another approach is to consider research into new electronic techniques that are optimized for integrating optics economically across all layers and platforms. Coherent optical receivers might benefit from analog signal processing rather than doing everything purely digital. Research in this area would also consider how best to use parallelism in the electronic cores together with parallelism in the optical systems. Going beyond the hardware integration, another research approach is the co-design of computer and network architectures requiring a network aware operating system, software, and applications. Research could embrace a data-centric view in which bandwidth and network connectivity is the expensive resource to be most efficiently utilized, rather than the computing cores. Research in the area would benefit from greater access to foundry capabilities for integrated photonics and help to accelerate developments coming from such facilities.

Sustainable Optical Networks: Networks are frequently relied upon for so-called smart energy solutions, in which greater energy efficiency is achieved through better monitoring and control of the energy use. Applications include the power grid itself, home heating and cooling, automobile traffic management, lighting, and home appliances. Designing networks to support these types of services can be considered a form of application-aware networking. However, future applications will likely need to consider the energy efficiency and sustainability of the networks that are delivering the smart services. This would include optical network elements designed to operate off renewable energy sources, which may deliver limited and highly variable power. Data centers today are being located near large sources of lower carbon energy to increase efficiency and carbon impact. Networks are by nature distributed and ubiquitous connectivity will require other approaches, most likely leveraging distributed micro-grids of renewable energy. Designing for a limited, distributed, and highly variable power source may dramatically alter the network design choices and optimizations. Sustainability should also be considered from a full lifecycle perspective including material sustainability. Many materials common in network devices such as Indium and Tantalum have questionable supply levels. Heavy reliance on rare materials may make Internet growth dependent on global mining markets. How would optical networks be engineered if sustainability is the primary design criteria?

D. Collaborative, Multi-User Test-beds for Optical Physical Layer Experiments

This topic is precisely research priority 7 as described above. This capability will be vital to each of the grand challenges described above and would provide an essential competitive advantage relative to other research worldwide. Europe has gone through several rounds of framework programs (FP) with strong support for research in optical networking [12]*. Japan has seen similar support. Although much progress was made, neither had the benefit of large scale physical layer testbeds and simulation capabilities, which hold promise to open up new and transformative advances. Meeting this grand challenge of realizing large scale, multi-user collaborative testbeds for optical networking, built using a physical layer that is fully open to experimentation, would provide a unique capability with the potential to enable the US to leap-frog Europe and Asia to once again take the lead in optical networking research.*

Works Cited

[1] Cisco, "The Zettabyte Era--Trends and Analysis," 2013.

[2] D. Neilson, "Photonics for Switching and Routing," *IEEE J. Sel. Top. Quantum Electron.,* vol. 12, pp. 669-678, 2006.

[3] X. Zhao, V. Vusirikala, B. Koley, V. Kamalov and T. Hofmeister, "The Prospect of Inter-Data-Center Optical Networks," *IEEE Communication Magazine,* vol. 51, pp. 32-38, 2013.

[4] A. R. Chraplyvy, "The coming capacity crunch," in *European Conference on Optical Communications*, Vienna, 2009.

[5] R.-J. Essiambre, G. J. Foschini, G. Kramer and P. J. Winzer, "Capacity Limits of Information Transport in Fiber-Optic Networks," *Phys. Rev. Lett.,* vol. 101, p. 163901, 2008.

[6] P. J. Winzer, *IEEE Com. Mag.,* 2010.

[7] D. A. B. Miller, "Device requirements for Optical Interconnects to Silicon Chips," *Proceedings of the IEEE,* vol. 97, no. 7, pp. 1166-1185, 2009.

[8] K. Hinton and a. et., "Switching Energy and Device Size Limits on Digital Photonic Signal Processing Technologies," *IEEE J. Sel. Top. Quantum. Electron.,* vol. 14, no. 3, p. 938, 2008.

[9] R. Tucker, "Optical Packet Switching: A Reality Check," *Optical Switching and Networking,* vol. 5, p. 2, 2008.

[10] J. J. Dongarra, C. B. Moler, J. R. Bunch and G. W. Stewart, "LINPACK: users' guide," Society for Industrial and Applied Mathematics, 1979.

[11] The Climate Group, "SMART 2020: Enabling the low carbon economy in the information age," 2008.

[12] S. Figuerola, D. Simeonidou, J. F. Palacios, A. Di Giglio, N. Ciulli, J. A. Garcia, R. Nejabati, X. Masip, R. Munoz, G. Landi, M. Yannuzzi and R. Casellas, "Research Trends on ICT Convergence from the CaON Cluster," in *International Conference on Transparent Optical Networks (ICTON)*, Coventry, 2012.

Appendix A

Below are the voting results for the top 10 research areas. Vote counts are based on a prioritization made by each participant present at the time of voting, using an anonymous electronic voting system provide by the Optical Society of America.

	Responses (percent)	(count)
Clean Slate Optical Network Architecture	12.56%	194
Adaptive & Cognitive Optical Network Architecture	11.20%	173
Programmable Optical Network with Intelligent Control Plane	13.01%	201
Optical network architecture across layers - Cross Layer Design,	10.23%	158
Application aware optical network architectures	8.93%	138
Economic, reliable networking on unreliable optical networks	7.64%	118
Test-bed for optical physical layer for multiple users to collaboratively experiment	8.74%	135
Co-architecture of optical and wireless networks	8.35%	129
Network optimization	7.70%	119
New Optical Network Architectures for Datacenters & Cloud Computing	11.65%	180
Totals	100%	1545

Appendix B

This is the full set of research topics identified during the workshop, listed in no particular order.

1. Clean Slate Optical Network Architecture
2. Adaptive & Cognitive Optical Network Architecture
3. Compute Network Architecture Using Embedded Optics
4. Programmable Optical Network with Intelligent Control Plane
5. Programmable optics and emerging technologies
6. Optical network architecture across layers - Cross Layer Design,
7. Management and operations
8. Application aware optical network architectures
9. Co-design computer/network architecture require network aware OS/Software/Apps/APIs
10. Economic, reliable networking on unreliable optical networks
11. Sustainable optical networks: Energy, materials.
12. Test-bed for optical physical layer for multiple users to collaboratively experiment
13. How to build network when statistical multiplexing is not valid
14. Electronics optimized for optics integration economically across platforms
15. Cooperative optics and wireless
16. Rethinking the end-point of the fiber
17. Security of the network
18. Co-architecture of optical and wireless networks
19. Network optimization
20. Security of the data
21. System stability in terms of optical devices – stability of optical network layer
22. New Optical Network Architectures for Datacenters & Cloud Computing
23. Capacity Exhaustion
24. Optical MIMO Networking
25. Switching technologies – optics and electronics