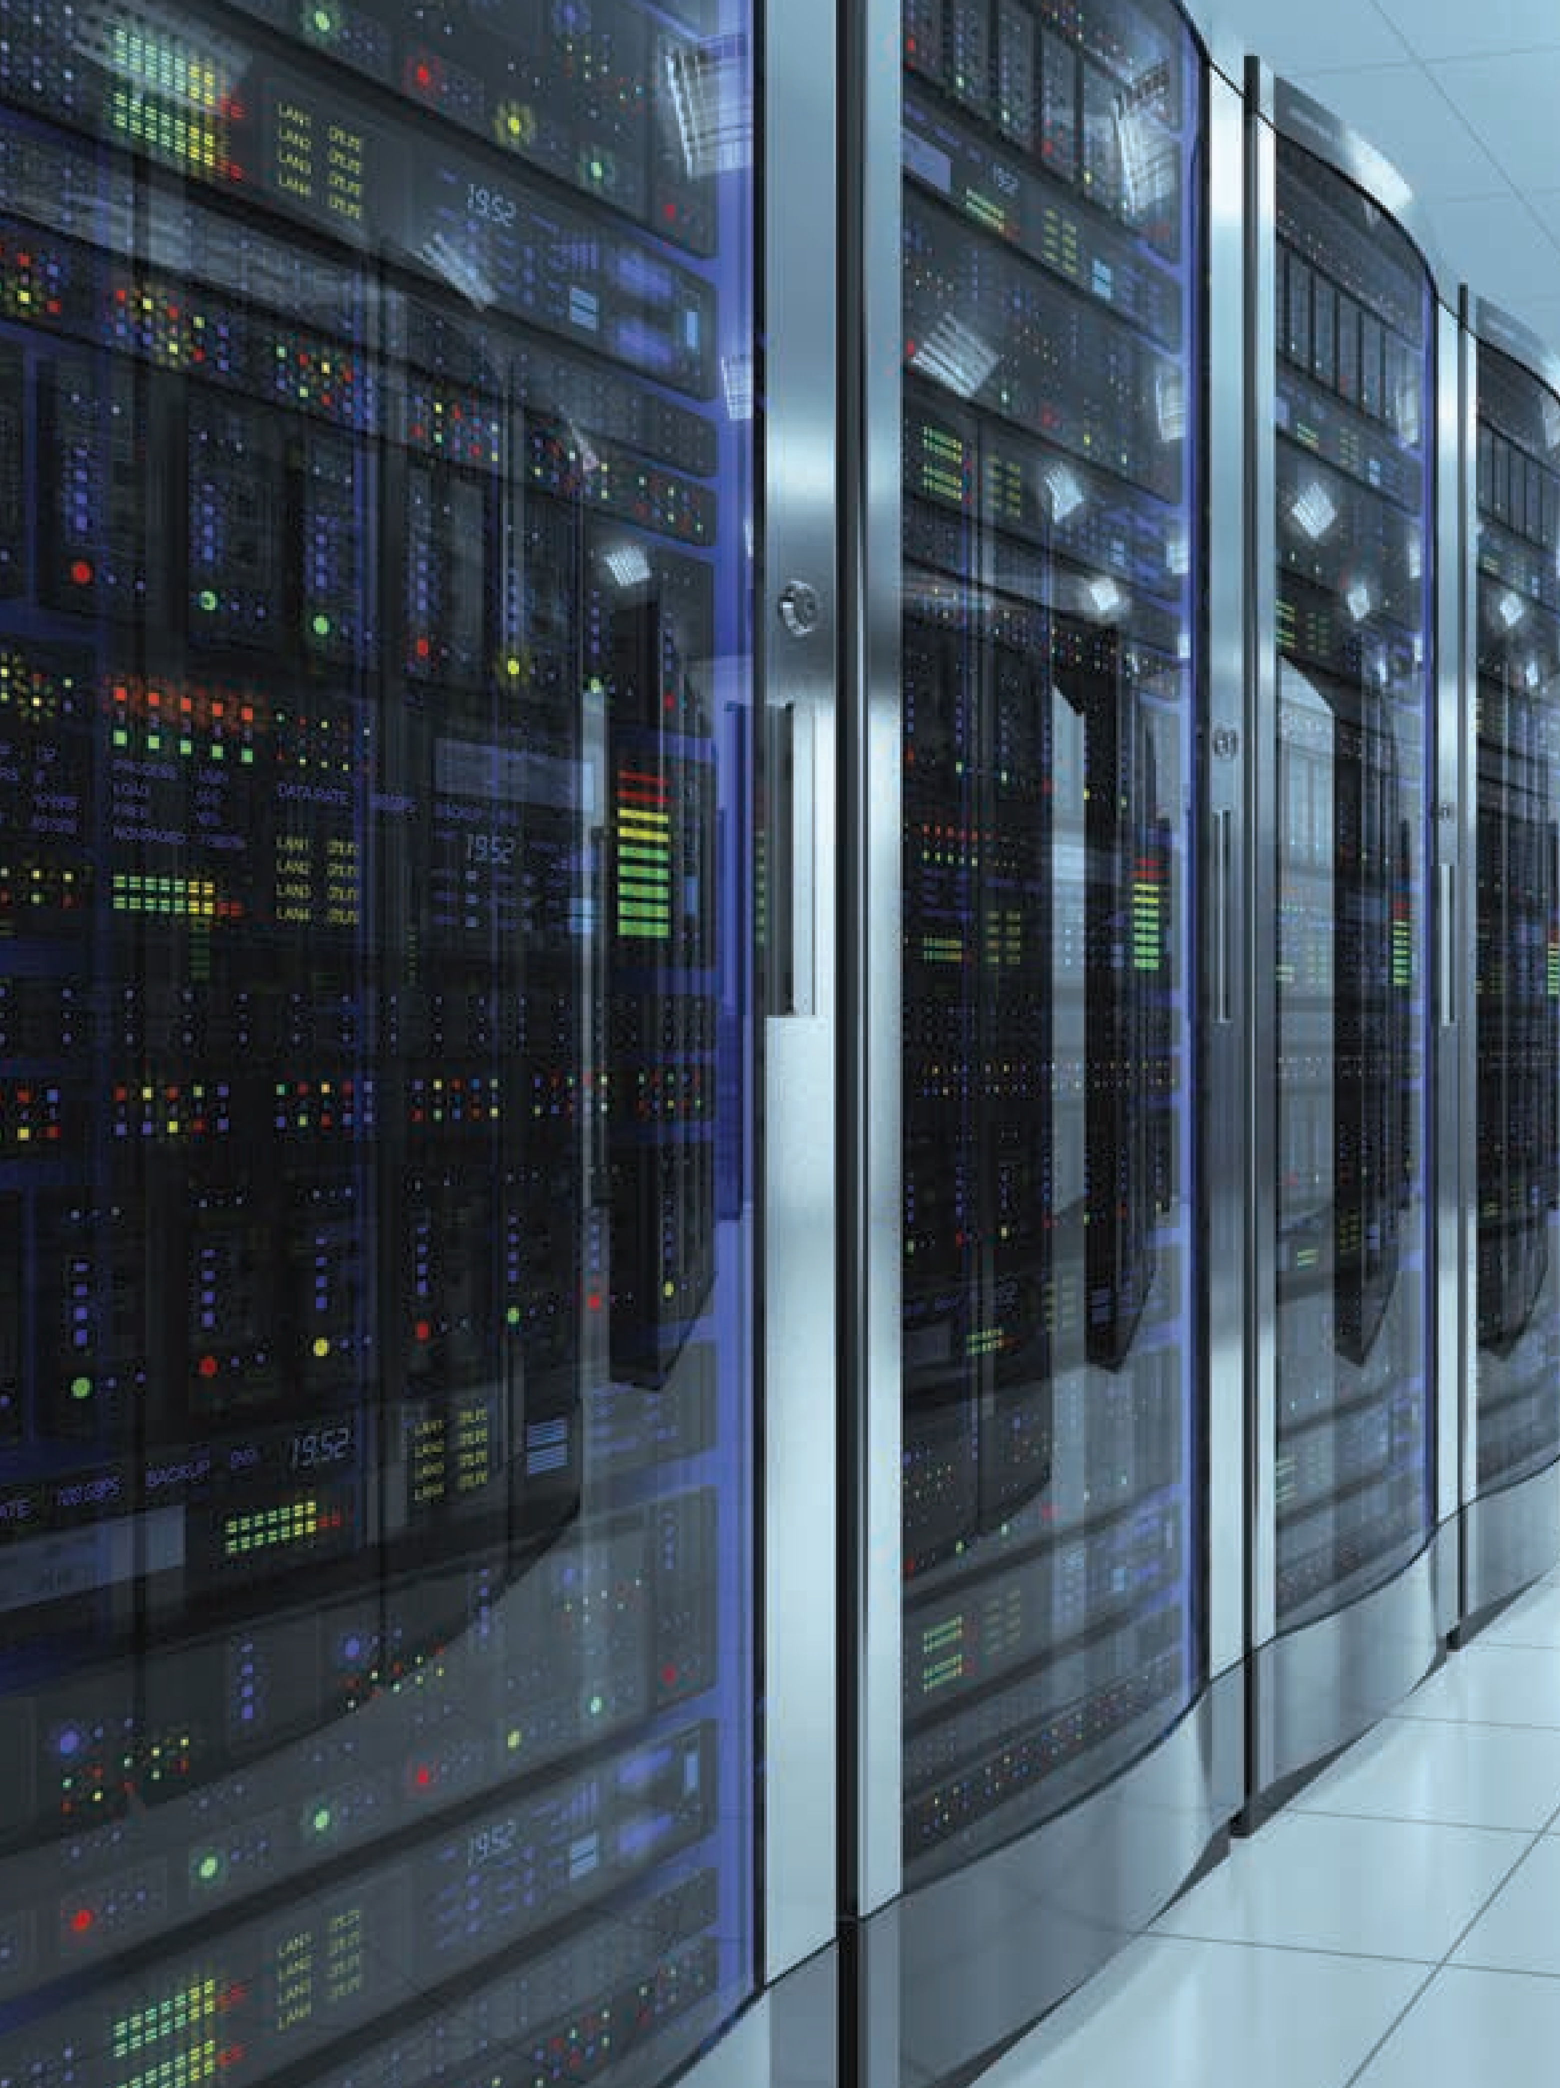

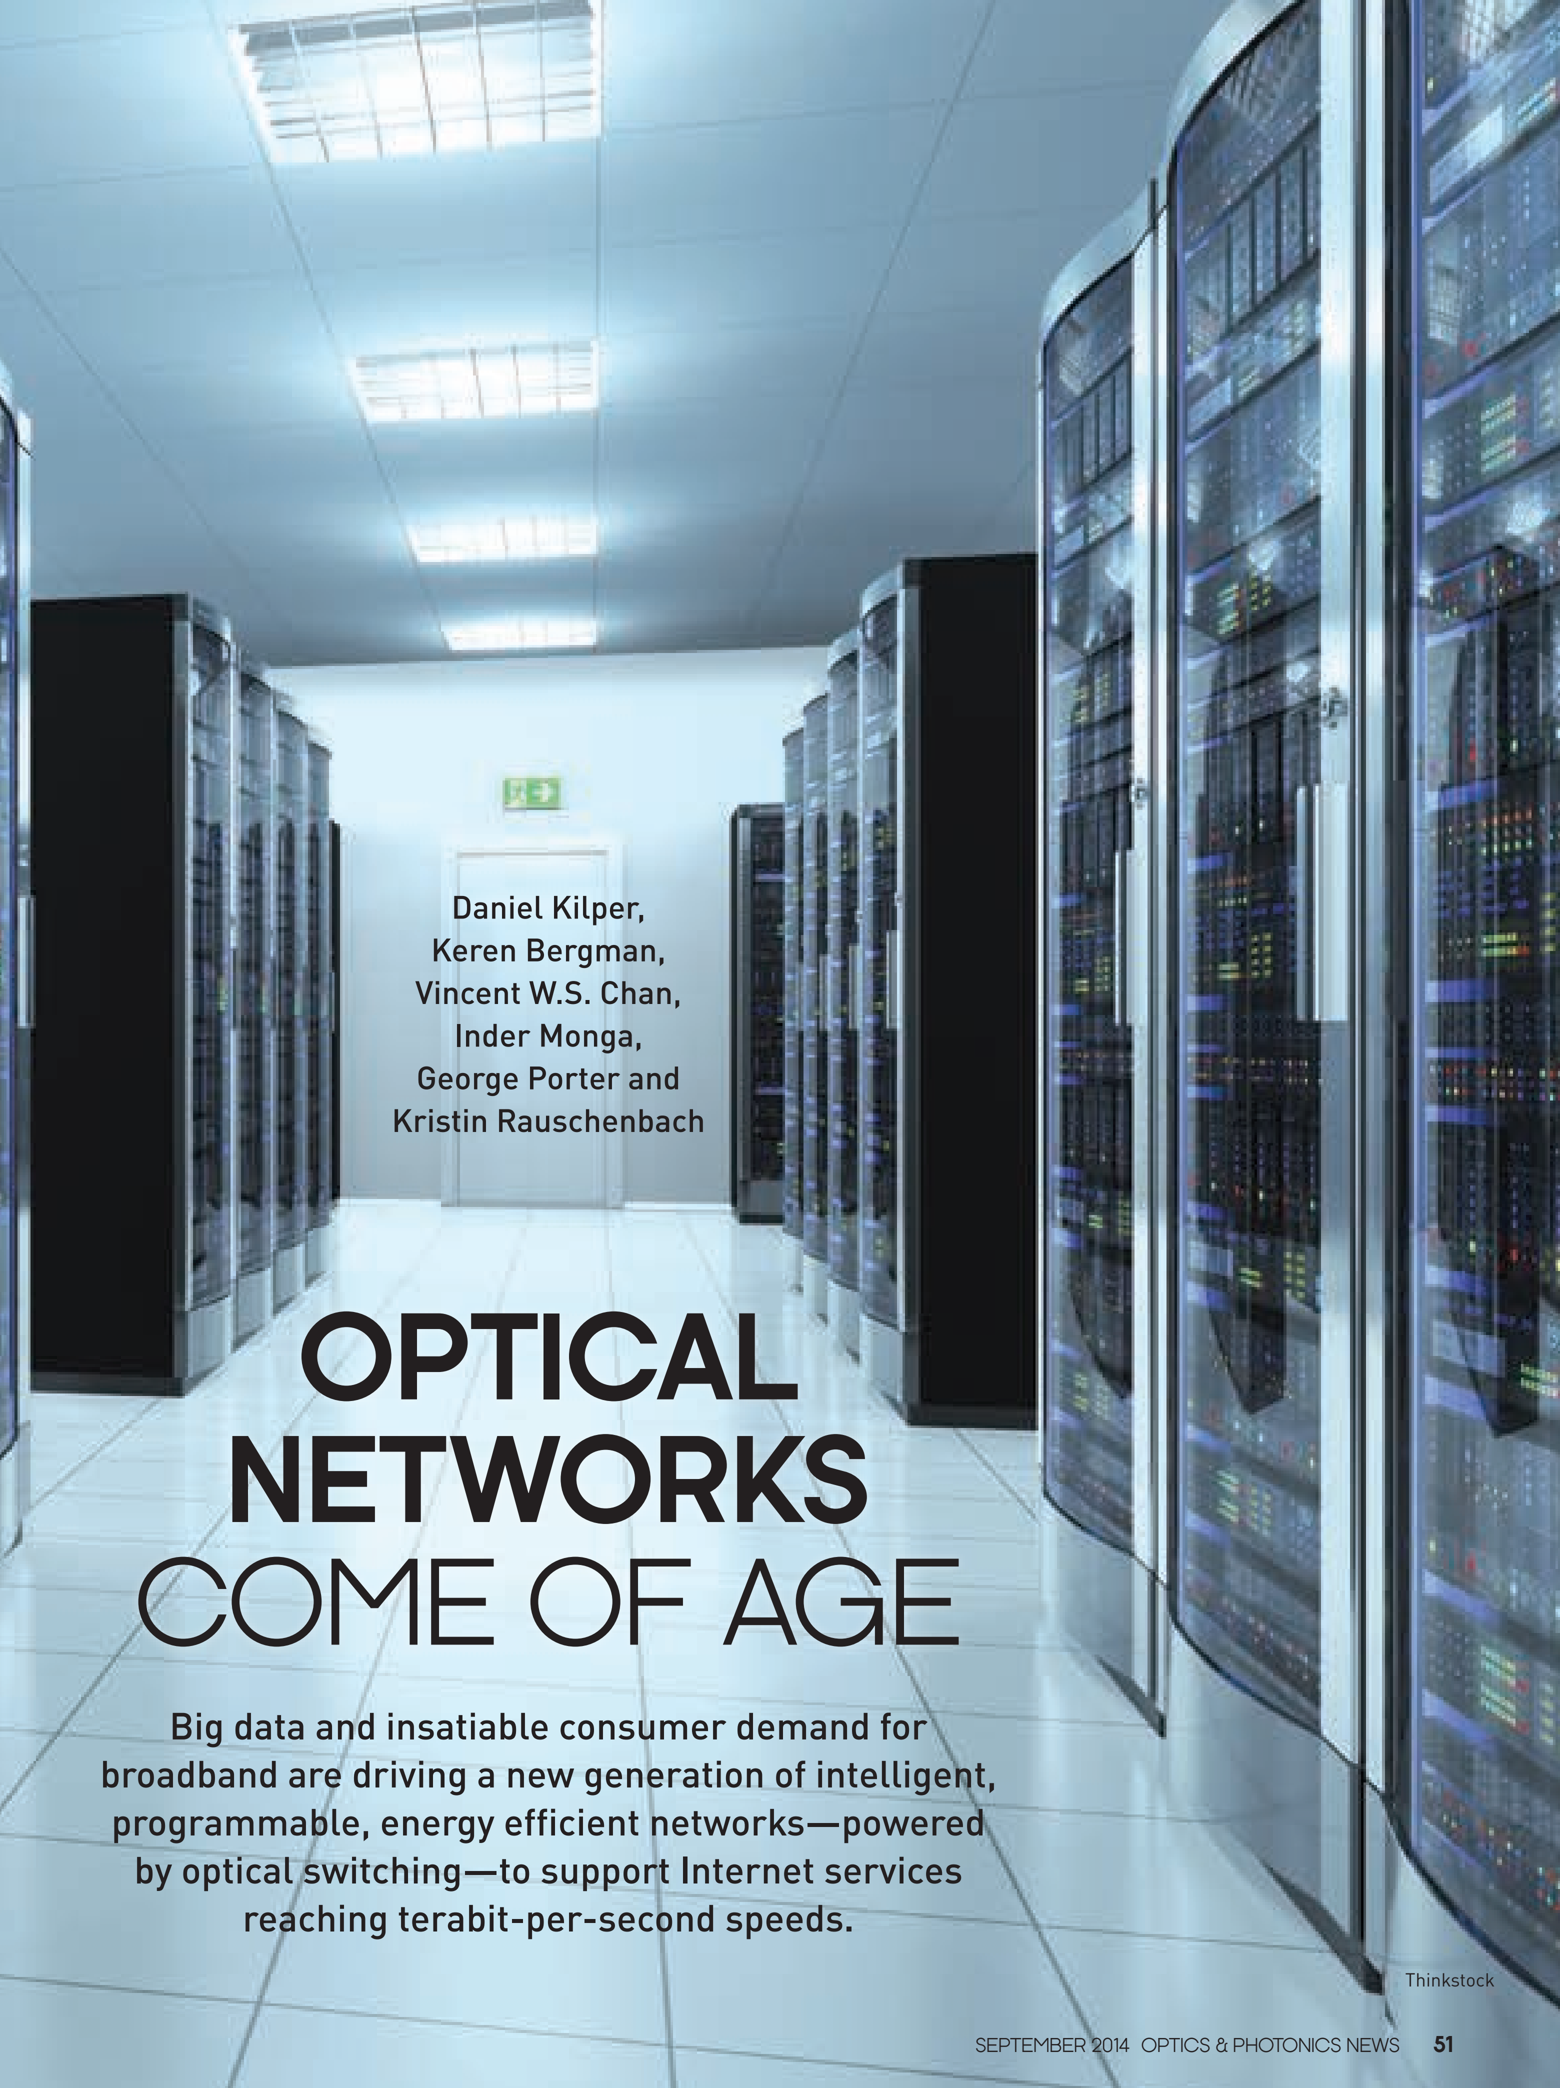

Thinkstock

Daniel Kilper,
Keren Bergman,
Vincent W.S. Chan,
Inder Monga,
George Porter and
Kristin Rauschenbach

OPTICAL NETWORKS COME OF AGE

Big data and insatiable consumer demand for broadband are driving a new generation of intelligent, programmable, energy efficient networks—powered by optical switching—to support Internet services reaching terabit-per-second speeds.

The story of fiber optic transmission capacity growth is well known: seven orders of magnitude in a short two decades. The telecom bubble of the late 1990s accelerated continent-scale wavelength division multiplexing (WDM) technology, which has steadily expanded from the Internet backbone into metropolitan regions. Fiber optic links are now making their way onto the edges of the network, into the home and the largest data centers, as photonic devices replace traditional electrical links. The smartphone and mobile-data explosion forced carriers to rapidly upgrade their cellular base-station Internet connections over to fiber optics as well.

What few may realize, however, is that the functionality of these fiber systems has been confined largely to static transmission links connecting large electronic circuit and packet switches. Optical systems act mainly as "fat pipes," the large-scale plumbing of the Internet.

A changing landscape in fiber optic communication technologies is stimulating a resurgence of interest in optical switching.

Greater use of optical networks—particularly in network edge applications that carry less aggregated, more "bursty," service traffic—and continued traffic growth will soon revise this picture. A changing landscape in fiber optic communication technologies is stimulating a resurgence of interest in optical switching.

These trends are coming together in ways that hold promise for the long-anticipated widespread deployment of optically switched fiber networks that respond in real time to changing traffic and operator requirements. The ultimate mission is to enable the next-generation Internet—one that can support terabit-per-second speeds, but that remains economical and energy efficient.

In late 2013, an OSA Incubator meeting on "Scaling Terabit Networks" investigated the opportunities and challenges of building optical networks to support the next phase of massive Internet growth. This article reviews some of those challenges, and the potential solutions that lie in the development of a new generation of intelligent optical networks.

Scoping the challenge

The enormous increase in Internet traffic over the past two decades could not have happened without the deployment of a worldwide fiber optic network. Long-haul fiber networks, previously simple, sparse backbones connecting large nodes of the early Internet, have evolved into a complex web, which, in fractal fashion, reveals an additional zone of mesh connectivity in metropolitan areas as fiber continues to push out into the network's edges. In many areas this continues one step further, forming a fiber-to-the-home/premises/curb network. In sum, fiber is no longer solely for intercity links, but increasingly also supplies metro and even last-mile connectivity.

The ability of static fiber networks to support the exponential traffic growth stems from fiber's characteristics as a fat pipe—a channel that, aided by WDM techniques, could support orders of magnitude of increased capacity as demand grew with the consumer appetite for on-demand videos, photo sharing, social media, and other bandwidth-hungry applications. That capacity growth has been compared with (and, in fact, has actually outpaced) Moore's Law in the semiconductor realm. And, as with Moore's Law, the question has increasingly arisen as to when the exponential growth will run up against hard physical limits.

For optical systems, those limits are embodied in the Shannon channel capacity limit as applied to a nonlinear fiber channel, which caps the spectral density the channel can support. In recent years, existing systems have started to close in on that limit, creating the so-called fiber capacity crunch in optical transmission systems. We are getting to the point where we are packing as much information into a single optical channel as is physically possible. And these limits may be starting to show up in a slowing of traffic growth, as the network bumps up against this physical reality.

1047-6938/14/09/50/8-$15.00 ©OSA

Challenges and Opportunities for the Future Internet

Tomorrow's networks face unprecedented and ever-growing demand for speed and bandwidth, but must overcome fundamental spectral and energy limitations. These problems are pushing optical solutions further into the network edges.

Traffic growth is slowing as optical transmission capacity starts to run up against spectral limits

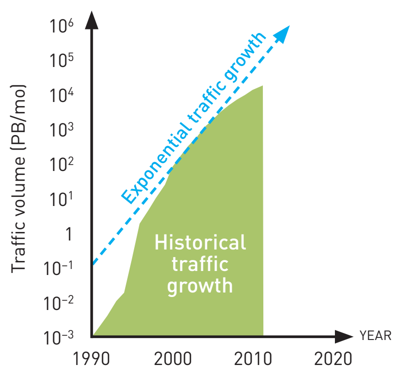

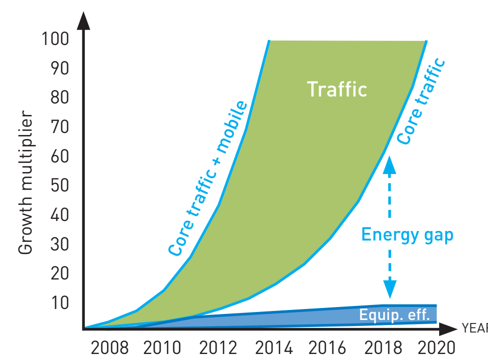

... yet new sources of demand continue to emerge **...**

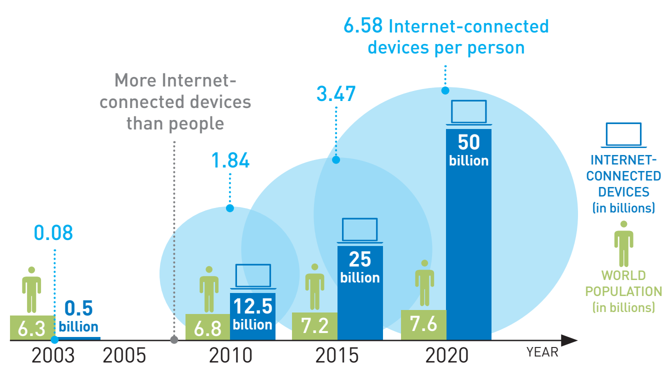

THE "INTERNET OF THINGS"

The number of data-hungry devices connected to the Internet is expected to be **50 billion**—more than **6.5 times** the world's population—**by 2020**.

... and energy efficiency isn't scaling with the growth in demand. While North American wireline and mobile traffic are growing at exponential rates, improvements in energy efficiency are barely linear, an unsustainable situation.

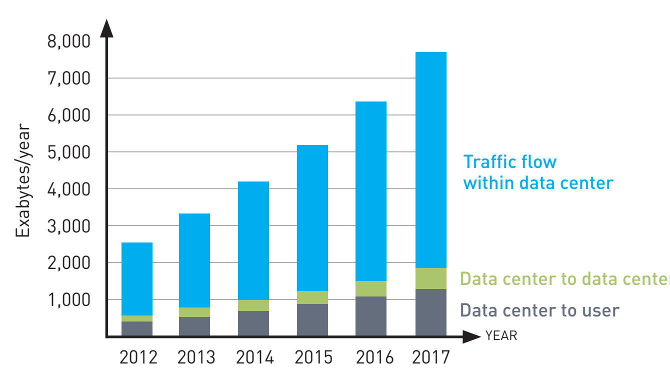

THE RISE OF THE DATA CENTERS

Driven by "**Big Data**," traffic flows between and especially within data centers are driving **new demands for optical switching** and other solutions.

Infographic by Stewart Wills and Alessia Kirkland
Source (clockwise from top left): Korotky, BLTJ **18**(3), 5 (2013)/Kilper et al., IEEE J. Sel. Top. Quantum Electron. **17**, 275 (2011)/Cisco Global Cloud Index (2013)/Cisco IBSG (2011).

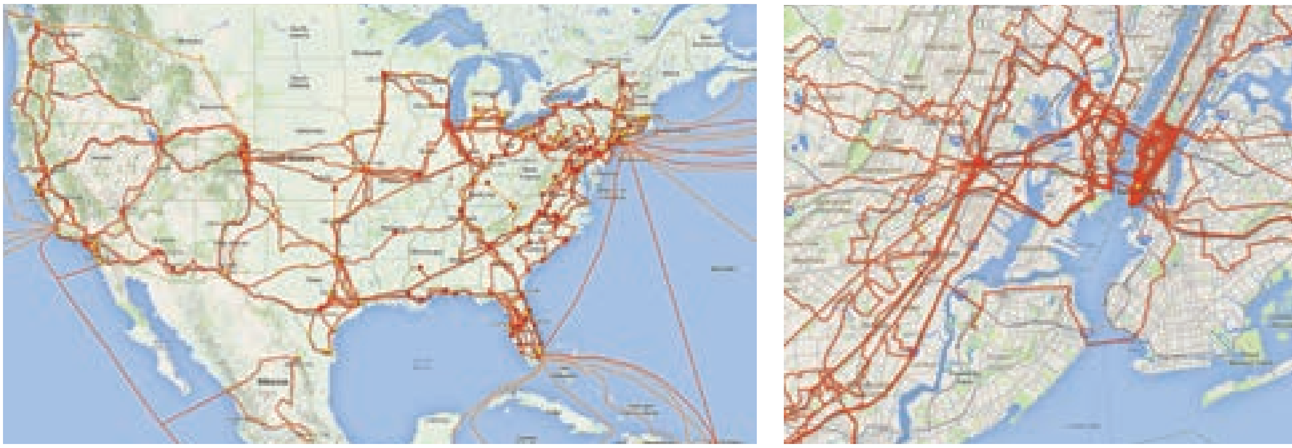

From long-haul to metropolitan

(Left) U.S. and Central American fiber network for Level(3); (right) greater New York City metropolitan area fiber network for Level(3).
Source: http://maps.level3.com/

As is discussed below, finding a solution to this quandary will require increased parallelism in optical systems—and that, in turn, will require new advances in optical networks and their architecture. Those advances are badly needed, as a variety of trends will require the networks underpinning the future Internet to increase their scale by orders of magnitude:

The "Internet of Things." In 2008, the number of devices connected to the Internet exceeded the world's population for the first time—and the trend to attach more "things" to the Internet, such as cameras, cars, and even household appliances such as refrigerators, is accelerating. Within six years, it's estimated that 50 billion devices will be connected to the Internet, well over six times the human population. That could expand the effective user or access base enormously, particularly if these things communicate with each other and act autonomously, without user intervention—the so-called Internet of Things.

Big data and internal data center flows. When data centers first emerged, their focus was on delivering content to end users—web pages, videos, music. This user-facing traffic is often referred to as "north-south" traffic in the data center. More recently, cloud-computing and big-data applications have spurred massive growth of "east-west" traffic within the data center itself. In an era of big data, a single user request from outside the data center can trigger a cluster of computers within a data center, all working together to generate the response, and sometimes processing very large data sets to extract, for example, statistical information on a single customer's preferences.

All of this requires parallel computing among, and traffic between, many hosts before the response to the user leaves the data center. Intra-data-center traffic is estimated to constitute 76 percent of the total data center traffic, far outstripping data-center-to-user and inter-data-center traffic.

An increasingly metropolitan focus. The increase in content-delivery networks and data centers is spurring rapid growth in metro-only Internet traffic, which is projected to surpass long-haul traffic for the first time this year and to exceed it by 40 percent by 2017.

These trends, which point to a more heterogeneous set of requirements, are forcing a shift in perspective on the future of optical networks. The historical focus on optics as a fat pipe envisioned a gradual evolution from fixed point-to-point optical systems to optical circuit switching and even optical packet switching. More recently, the consensus has shifted toward a hybrid network evolution, in which electronic packet switching and optical switching are used together. Under such a view, small, bursty data flows and data processing-intensive applications are best handled using electronics, while the larger "elephant" flow traffic is best handled using optics.

Parallelism and its problems

Given that optical channels are approaching their physical capacity limits, one consequence is that, to support the order-of-magnitude capacity increases of the

The challenge is not so much how to replace electronic systems with optics, but how to enable these systems to work together.

next-generation Internet, these systems will need to go parallel. Just as electronic processor clock speeds hit a ceiling and processors moved to multicore architectures, optical system channel capacities are hitting a ceiling and are beginning to move to parallel solutions.

Going parallel still offers tremendous room for growth. The number of optical channels per fiber can expand perhaps tenfold. Cables as large as 1,000 fibers are already commercialized, and 10,000 parallel fibers are conceivable. A host of other techniques are also on the horizon, including the use of multi-core fibers (a denser form of a fiber cable) and multiple spatial modes. All told, parallel growth could offer six to seven orders of magnitude of additional capacity.

But there's a problem with parallel systems: built today, their cost, energy and footprint would all increase almost linearly with the capacity. Energy constraints place limits on the density or footprint of the equipment and thereby impact cost as well. Recent studies suggest that anticipated equipment efficiency improvements lag far beyond expected traffic growth, opening a huge and unsustainable "energy gap" over the coming years.

The historical seven-orders-of-magnitude capacity increase was able to proceed without a commensurate increase in cost, power and footprint, because it involved increasing the data rate of the optics within a system chassis and increasing the number of WDM channels within a single amplification band—higher capacity in a fiber system. Going forward, with the data rate per channel near its limit, the challenge becomes how to support even more optical channels in the same system or chassis—ultimately moving to multiple fibers and other parallel solutions. The optics will need to become denser and lower power to scale in parallel. Integration through silicon photonics, which can take advantage of much of the technology and scalability of silicon, is one promising solution for increasing density and reducing power requirements that is receiving much attention today.

But photonic integration does not fully solve the challenges of continued Internet growth and higher capacity demands at the network edges. For this, we need optical networking solutions that can reduce the processing of high-bandwidth data flows in electronic switches and routers. Architectural innovations are needed for the network itself to implement a less passive, more dynamic role for optical switching and components.

That will require truly programmable physical-layer optical systems—and those, in turn, will demand cross-disciplinary advances in optical transmission and switching, cognitive and autonomic control, dynamically adaptive security capabilities, and real-time optimization. While optics is exceedingly efficient at transporting or switching large quantities of data, it cannot achieve the density or efficiency of electronics in computations requiring more than a few operations. The challenge, therefore, is not so much how to replace electronic systems with optics, but how to enable these systems to work together with increasing reliance on optical functionality.

Toward dynamic optical systems

Historically, the focus in networks has been on rapid growth through packet switching—with optics largely used as a (passive) fat pipe for data. Optical switching today is used to provide flexibility when signals are provisioned or restored; active switching of optical signals in response to changing traffic demands and network conditions is generally not done.

With such fixed optical channels, capacity is provisioned for peak traffic levels, often greater than a factor of two in capacity over average levels, with additional capacity for network resilience, so that a backup path is always available. Signal quality is further guaranteed by allocating performance margins for the many signal impairments, system configurations and aging.

Moving to more flexible optical systems would allow for tighter margins, since the system can adapt in response to the changing channel conditions. Furthermore, the capacity on a fiber might be allocated in proportion to the instantaneous traffic demand, yielding greater network efficiency. Such a dynamic optical system, however, would mark a paradigm shift away from how such systems are engineered and managed today, creating a large barrier to such new approaches.

In some respects, this vision of active switching using optical systems has been around almost since the first fiber communications, but the cost and complexity have always been prohibitive. The emerging need to eke ever-greater efficiency out of optical systems could prove the motivation to overcome these challenges. Other

The Incubator meeting

On 19-20 September 2013, an OSA Incubator meeting, "Scaling Terabit Networks: Breaking Through Capacity Barriers and Lowering Cost with New Architectures and Technologies," examined many of the issues raised in this feature. The workshop, held at OSA's Washington, D.C., headquarters and hosted by the NSF's Center for Integrated Access Networks (CIAN) with funds from the Computer and Information Systems Engineering directorate of the NSF (grant #CNS-1346666), framed the grand challenges for research priorities in optical networks, and also focused on the need for a reliable testbed for terabit-scale experiments in optical networking. The full whitepaper from the workshop is available to OSA members at www.osa.org/members/membership/member_documents/stnfinalreport/ .

factors have also improved the picture for active optical switching: optical communication technologies have matured; coherent modulation techniques are now cost-effective; coherent reception can compensate for many physical phenomena such as chromatic dispersion and polarization mode dispersion; "elastic networking" allows a range of quadrature and amplitude modulation constellations to be realized depending on channel conditions.

Along with these transceiver advances have come optical switch developments. Today, optical switches not only switch individual wavelength channels between fibers, but also adapt the channels' bandwidth and center frequency. This flexible grid capability opens new degrees of freedom that can be exploited when flexibly switching channels within a network.

To fully exploit the new environment for dynamic optical switching, optical systems must be programmable and virtualizable, so that they can scale to a large number of processes and entities. Can an optical system support multiple network control systems to meet different performance objectives, in the same way that a computer might run multiple operating systems or virtual hosts? Dynamic adaptation via new fast algorithms offers a mechanism toward direct control and use of optical systems.

Tomorrow's optical networks: programmable and virtualized

Traditionally, optical transmission system control has involved proprietary, highly optimized control systems unique to each company. Enabling this physical-layer control to interact with the higher-layer packet or circuit switch control is a long-sought-after, and elusive, capability. Controls and protocols for the digital packet systems can be very different from those used for optical systems, which are still very analog in nature. This technology disconnect has been difficult to overcome.

One recent incarnation of a network control plane achieves a more unified, open-access structure via software application interfaces. Referred to as software-defined networking (SDN), the framework has taken hold particularly in electronic packet systems, because it enables the network user or operator to directly configure the routing tables via third-party software applications—for example, to optimize the network in a data center for a particular set of applications. The wave of success in data centers has carried the SDN framework to other systems such as optical transmission and switching, as a cross-layer controller.

One aspect of SDN that makes it attractive for optical systems is that it uses system abstractions to simplify control across layers. These abstract models provide the higher control layers with only the essential details that are needed to make reasonable decisions regarding the optics. Thus the complex details are handled through proprietary methods within the optical system itself, allowing simplified control interfaces to be provided to the software applications. This approach avoids technology lock-in at the optical component and subsystem layer. As more efficient, compact, capable optical devices emerge, they integrate easily into existing control systems. Such capabilities do not depend on SDN and may be used in other cross-layer control approaches as well.

A key question is what level of detail the abstractions for optical systems should use. More detail will provide the potential for greater optimization—just the right equipment, capacity and performance for the application needs. However, due to the complexity of optical system engineering, too much detail may be cumbersome and difficult for a cross-layer controller to manage, potentially resulting in poor performance. From an optics perspective, this opens up an interesting cross-disciplinary problem of control for the complex, distributed physical systems in optical transmission. A trade-off can also be made between the use of hardware controls or

To fully exploit the new environment for dynamic optical switching, optical systems will need to be programmable and virtualizable.

more optical signal regeneration versus sophisticated software controls. Increased hardware drives up cost and energy, but simplifies the control problems.

Other research priorities from the Incubator

At the OSA Incubator meeting, workshop participants identified programmable, virtualized and intelligent optical networks, described earlier, as the highest-priority research area for addressing the order-of-magnitude increases in traffic and complexity envisioned for the future Internet. Participants also focused on several other areas to guide and advance future research.

Cross-layer architectures for data centers and cloud computing. Future data centers and cloud computing will benefit from new architectures that can remove the network bottleneck and approach delays limited only by optical transit time. Addressing this challenge will require research on application-aware control and cross-layer functionality to the optical layer, and on new layering paradigms with network architectures designed for the needs of future data centers and cloud computing.

High-port-count optical space switches, long of interest in long-haul systems, have found renewed interest in data centers to provide high bandwidth connections on demand. Making this work, however, will require a new hybrid switching control plane that can seamlessly direct data flows to either an optical switch or an Ethernet switch, depending on application requirements or performance objectives, as well as, more generally, integration of new optical technologies and control capabilities into existing data centers and emerging cloud infrastructures. A new generation of low-cost optical space switches, together with hybrid switching control through methods such as SDN, may hasten commercial adoption.

Clean-slate architectures and technologies. Given the complex constraints and entrenched technology evolution and standardization cycles in the networking area, clean-slate research is crucial for breaking through the status quo and finding novel solutions for addressing long-term scalability needs. Rethinking the network to leverage optics as the end-to-end transmission and switching platform may be a fruitful approach given the diversity of optical technologies available today together with the heterogeneity of high-capacity services, addressable through optics.

A testbed for terabit-scale experiments. The workshop participants identified the vital need for a collaborative, multiuser, national-scale testbed specifically for terabit-scale optical networking experiments. Today's optical-networking experiments are largely confined to small three-node networks that do not capture the performance of larger systems or allow for the combined study of physical-layer optics and cross-layer control systems. The optical infrastructure needs to be research grade, not fully commercial, with dark-fiber infrastructure configured to accept new optical networking hardware and for studying control dynamics and higher-layer interactions.

Each research challenge described above depends on having such a testbed, to provide a common environment for proving-out next-generation technology and for studying networks built on emerging devices based on photonic integration. Such an infrastructure would offer a unique capability globally, providing international leadership in optical networking research. OPN

Daniel Kilper (dkilper@optics.arizona.edu) is at the University of Arizona, U.S.A.; Keren Bergman is at Columbia University, U.S.A.; Vincent W.S. Chan is at the Massachusetts Institute of Technology, U.S.A.; Inder Monga is with the Energy Science Network (ESnet), Lawrence Berkeley National Laboratory, U.S.A.; George Porter is with the University of California, San Diego, U.S.A., and Kristin Rauschenbach is with Notchway Solutions LLC, U.S.A.

References and Resources

- D.C. Kilper et al. IEEE J. Sel. Top. Quantum Electron. **17**, 275 (2011).
- G. Porter et al. "Integrating microsecond circuit switching into the data center," SIGCOMM'13, ACM 978-1-4503-2056, 2013.
- Cisco Visual Networking Index (VNI), www.cisco.com/c/en/us/solutions/service-provider/visual-networking-index-vni/index.html .
- Cisco Global Cloud Index: Forecast and Methodology, 2012-2017, www.cisco.com/c/en/us/solutions/collateral/service-provider/global-cloud-index-gci/Cloud_Index_White_Paper.html .
- S.K. Korotky. "Semi-empirical description and projections of Internet traffic trends using a hyperbolic compound annual growth rate," Bell Labs Tech. J. **18**(3), 5, 2013.
- V.W.S. Chan. "Optical flow switching networks," invited paper, special issue, Proc. IEEE, March 2012.
- R.S. Tucker and K. Hinton. Photonics J. IEEE **5**, 821 (2011).
- C. Kachris et al., Eds. *Optical Interconnects for Future Data Center Networks*, Springer Science + Business Media, New York (2013).