

Scaling



The background of the cover is a dark, almost black, space filled with a complex network of glowing lines. These lines, representing fiber optic paths, are colored in a gradient from deep blue and purple on the left to bright yellow and orange on the right. The lines are not straight but curve and branch out, creating a sense of dynamic movement and connectivity. In the lower right quadrant, there are several lines of binary code (0s and 1s) in a light yellow color, appearing to be part of the data flow.

g Optical Fiber Networks

Challenges and Solutions

Peter J. Winzer


Increases in consumer demand and machine-to-machine network traffic are creating big challenges for letting optical communications continue to scale cost-effectively. Meeting those demands will require new forms of optical parallelism.

Virtually every phone call we make today, every text message we send, every movie we download, every Internet-based application and service we use is at some point converted to photons that travel down a vast network of optical fibers. More than two billion kilometers of optical fibers have been deployed, a string of glass that could be wrapped around the globe more than 50,000 times. Well over 100 million people now enjoy fiber optic connections directly to their homes. Optical fibers also link up the majority of cell towers, where the radio frequency photons picked up from billions of mobile phone users are immediately converted to infrared photons for efficient fiber optic backhaul into all-fiber metropolitan, regional,

long-haul and submarine networks that connect cities, span countries and bridge continents.

Driven by a continual succession of largely unanticipated emerging applications and technologies, network traffic has grown exponentially over decades. More recently, it is no longer just the human ability to consume information that may ultimately set limits to required network bandwidth, but the by now dominant amount of machine-to-machine traffic arising from data-centric applications, sensor networks and the growing penetration of the Internet of Things, whose limits are primarily based on the economic value that these services can provide to society. While historical data and forecasts of network

Key components of an optical transport network



Phil Saunders/spacechannel.org

traffic vary widely among service providers, geographic regions and application spaces, annual traffic growth rates between 20 percent and 90 percent are frequently reported, with close to 60 percent being a typical value observed for the expansion of North American Internet traffic since the late 1990s.

The role of optical fiber communication technologies is to ensure that cost-effective network traffic scaling can continue to enable future communications services as an underpinning of today's digital information society. This article overviews the scaling of optical fiber communications, highlights practical as well as fundamental problems in network scalability, and points to some solutions currently being explored by the global fiber optic communications community.

The modern transport network

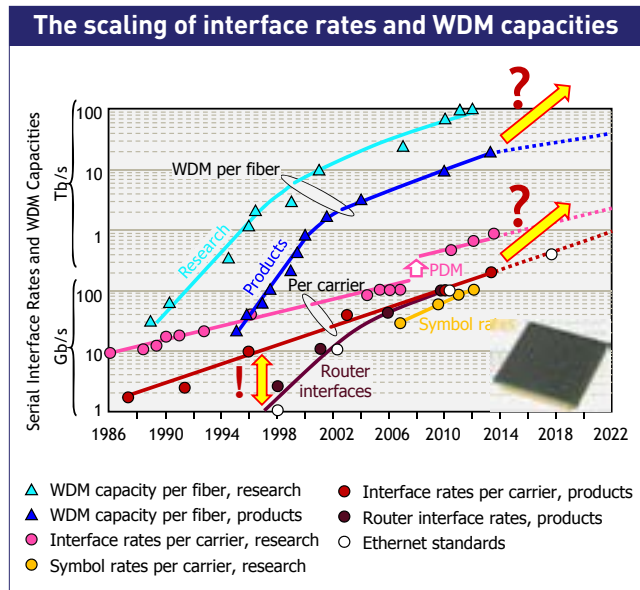
An optical transport network (see figure on facing page) interconnects Internet Protocol (IP) packet routers that pass data packets from a data source to the intended recipient, preferably along minimum-hop transmission paths. These routers are connected through optical client interfaces, which today offer connections of up to 100 Gbit/s over distances of around 40 km. Compact and low-cost client interfaces can directly tie a router to other nearby routers or connect a router to an optical transport system that in turn establishes a connection to distant routers.

At its client-facing side, the optical transport system terminates one or more short-reach client interface signals and converts them into long-reach signals that it subsequently transmits on its line interfaces. These signals can traverse thousands of kilometers of fiber without any intermediate electronic processing, passing only through optical amplifiers and optical filter components that can be dynamically reconfigured to add and drop signals or to switch them to different parts of the network, through reconfigurable optical add-drop multiplexers (see "ROADMs," p. 36).

In contrast to optical client signals, optical line signals are designed with spectral stacking in mind. Modern wavelength division multiplexed (WDM) optical transport systems carry about 100 optical signals at up to 200 Gbit/s each, on a 50 GHz optical frequency grid, for an overall capacity of about 20 Tbit/s on a single optical fiber.

The need for optical parallelism


Historically, commercial line interface rates—the bit rates carried on a single wavelength in an optical



transport system—have increased steadily at around 20 percent per year, a growth rate sufficient to scale voice-dominated networks. The rise of data traffic beginning in the 1990s, however, changed this picture. Router capacities have shown much faster growth rates of around 40 to 60 percent, as have router interface rates. Those rates are coupled to the evolution of computer interface speeds, which themselves are driven by the evolution of microprocessor computational capabilities, ultimately rooted in Moore's law.

As a result of these scaling differences between packet-switching and optical transport, the capabilities of optical line interfaces began to limit the growth of router interface rates by around 2005 (see figure above). The standardization of 100G Ethernet and the 100 Gbit/s Optical Transport Network (OTN) as a common packet-optical interface rate in 2010, and the anticipated standardization of 400G Ethernet in late 2017, underline this evolution. Another direct consequence of the packet-optical scaling disparity is the reduced aggregation potential of multiple router interfaces into a single optical channel, which has implications for overall network design, in that it deemphasizes the need for subwavelength add/drop capabilities in core networks.

Interestingly, keeping up with the past decade's 20 percent annual scaling in commercial optical interface rates could not have happened by simply increasing modulation speeds from 10 Gbit/s to 40 Gbit/s to 100 Gbit/s. Rather, systems had to rely on optical parallelism, and thereby exploit other physical dimensions in addition to modulating only the intensity of optical



The five physical dimensions and their use in optical modulation and multiplexing

pulses, the predominant technique used at rates up to 10 Gbit/s.

As seen in the figure above, optical parallelism was introduced by independently modulating the complex optical field's real and imaginary parts—or, in engineering terms, its in-phase and quadrature components—as well as by modulating both polarizations with individual signal streams (polarization division multiplexing, or PDM). Most 100 Gbit/s optical line interfaces modulate four parallel electrical signals at a more comfortable electronics speed of around 30 Gbit/s (25 Gbit/s, plus overhead for forward error correction). To make this happen required extracting the full optical field information at the receiver—which, in turn, meant that systems had to transition from direct detection of the optical pulse intensity to coherent detection of the optical field.

Coherent detection and optical superchannels

Although heavily researched in the 1980s, optical coherent detection was abandoned with the advent of erbium-doped optical amplifiers in the early 1990s. The rebirth of coherent detection in the 2000s was technologically enabled by the capabilities of digital electronic

signal processing (DSP), including the necessary digital-to-analog converters (DACs) and analog-to-digital converters (ADCs).

Today, CMOS technology can provide converter speeds of up to 90 GSamples/s, integrated with DSP engines containing ~100 million gates. The use of DACs enables the generation of Nyquist-shaped and magnitude/phase-predistorted optical pulses at the transmitter, while ADCs allow the faithful conversion of the full optical field of high-speed received signals into the digital electronic domain for further digital processing. In research experiments, leading records are rapidly approaching 1 Tbit/s per optical wavelength, with symbol rates of about 100 GBaud, carrying higher-level quadrature amplitude modulation (QAM) formats with bit rates of up to 864 Gbit/s today.

To increase the scaling of optical interface rates even further, optical parallelism in more physical dimensions has to be exploited. With the time, quadrature, and polarization dimensions already taken, optical interface rate scaling relies on the frequency dimension to overcome its scalability bottleneck. This is done by grouping multiple carriers to form a single logical interface called an optical superchannel, increasing efficiency through dense signal

To increase the scaling of optical interface rates even further, optical parallelism in more physical dimensions has to be exploited.

frequency packing and increasing economics through shared or integrated transponder componentry. Using superchannel technologies, optical interface rates of terabits per second and beyond are feasible today.

Spectral efficiency limits

To most efficiently use the embedded, costly optical fiber infrastructure, WDM systems try to pack as many optical signals as possible into an optical fiber—or, more specifically, into the limited bandwidth of the optical amplifiers placed periodically along the transmission link. For example, erbium-doped fiber amplifiers typically cover the C-band between 1,530 nm and 1,565 nm.

The spectral information density (spectral efficiency) that can be transmitted over a fiber of a given length, however, faces some hard limits—fundamental limits tied to amplification noise and Kerr nonlinearities that lead to various types of signal distortions, and practical limits stemming from technological imperfections of transponders and optical amplifiers as well as from the onset of catastrophic damage through fiber fuse.

The former, fundamental limit is called the nonlinear Shannon limit, depicted on the chart below. The logarithmic scaling of transmission distance versus the linear scaling of spectral efficiency lets the fiber capacity limit depend only moderately on transmission distance. Not surprisingly, the trade-off between transmission distance and spectral efficiency of record WDM experiments follows a slope similar to the fundamental capacity limit. There is a range of available and roadmapped commercial WDM products (yellow ellipse in the chart), many of which, called “flexponders,” offer dynamic modulation format adaptability, allowing systems to trade off capacity for transmission reach in a software-defined manner.

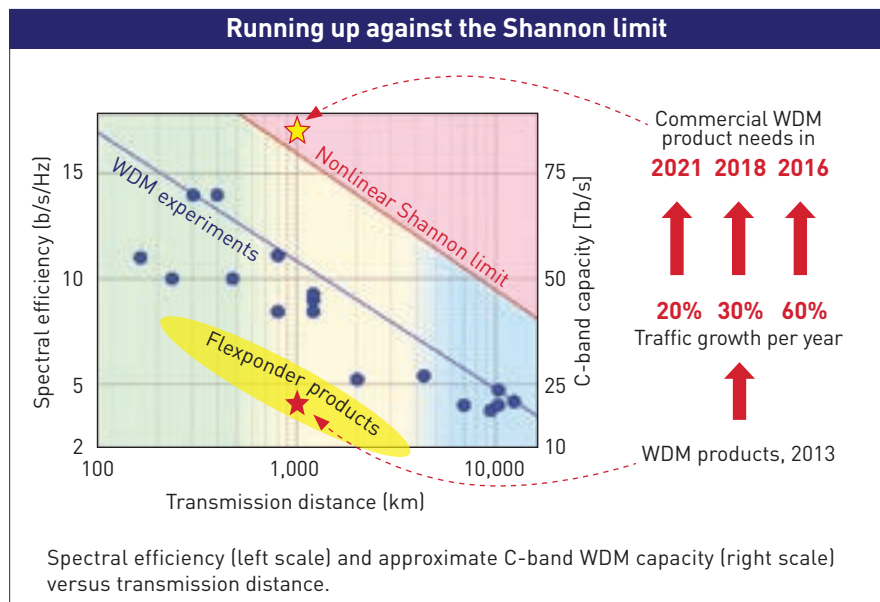
In 2013, leading system integrators started to offer, and leading service providers started to deploy, WDM products that, once fully populated with WDM signals, would support close to 20 Tbit/s over a distance of around 1,000 km. Assuming annual traffic growth rates of 20, 30 or 60 percent, leading-edge service providers will likely


demand systems that can scale beyond 75 Tbit/s over the same distance by 2021, 2018 or 2016.

Parallelism’s next frontier: spatial multiplexing

The nonlinear Shannon limit, of course, makes such systems fundamentally impossible to build. Yet the demand for these systems is not that far in the future, and will require development of entirely new technologies to accommodate these capacity needs. This conundrum has become known as the “capacity crunch” in optical networking.

The nonlinear Shannon limit is fairly insensitive to even heroic changes in fiber parameters such as fiber loss or nonlinearity coefficient. Going parallel in the frequency dimension by using more optical amplification bands across the low-loss window of deployed optical fiber may result in a capacity increase of only about a factor of five. Undoubtedly, using both better fibers and wider amplification bands will provide critical stopgap solutions in the near future. None of these techniques, however, can provide a sustainable path to overcome the optical network capacity crunch. To scale optical networks into the next several decades, optical parallelism must extend to yet another physical dimension—and the only dimension not already exploited is space. That means that spatial multiplexing, or space division multiplexing (SDM), is not just an attractive long-term solution, but the only viable solution on the horizon.





Some fibers for SDM transmission demonstrated over the past few years


Two key factors will drive the adoption of SDM in future optical networks:

Integration. Deploying individual optical transport systems in parallel, the simplest form of SDM, increases capacity, but also keeps constant the cost and energy consumption per bit. Hence, this form of SDM does not provide the exponential cost and energy reduction that has historically been expected of optical networking hardware and that has economically enabled the Internet, with its associated network traffic growth. Integration, then, constitutes an essential ingredient for sustained economics. Examples include the integration of transponders to form a spatial superchannel transponder, integration of optical amplifiers into optical amplifier arrays that share common hardware and control elements, or integration of multiple optical switching elements into switching arrays that handle multiple spatial paths at a reduced cost per path.


Another aspect of integration lies in new transmission waveguides such as multicore or few-mode optical fibers. These may yield cost savings both in capital and operating expenditure, for example, if the associated connectors, splices or fiber-to-chip connections can be made cheaper with new SDM fiber compared with an array of individual single-mode fibers. Research labs worldwide have made tremendous progress over the past few years designing new fiber structures and the associated mode- and core-coupling elements, with fascinating interdisciplinary synergies to other fields such as multi-mode astronomy and endoscopy. Impressive record optical transmission results over such fibers have been demonstrated on a systems and network level.

Frequently in engineering practice, integration comes at the expense of crosstalk. Fortunately, with the help of powerful DSP engines, crosstalk can be handled using multiple-input/multiple-output (MIMO) techniques originally developed for cellular wireless and digital


To scale optical networks into the next several decades, optical parallelism must extend to yet another physical dimension—and the only dimension not already exploited is space.




Sakaguchi, OFC 2012




Kobayashi, ECOC 2013



Qian, FIO 2012



Cia, IPS SumTop 2012



Mizuno, OFC 2014

subscriber lines (DSL). In fact, polarization-multiplexed coherent systems on the market today already employ 2x2 MIMO processing (required owing to the inherent polarization rotations occurring along the single-mode transmission fiber). As DSP chips are commonly low cost but high power, while optoelectronic transponder components are typically high cost but low power, developing successful SDM systems will critically hinge on finding the best balance between optoelectronic integration and DSP for SDM applications.

A smooth upgrade path. Whatever technology is being used, ensuring a smooth upgrade path from existing fiber optic networks will prove essential. Operators will not accept systems that fundamentally require the deployment of new transmission fiber, unless these new waveguides offered as revolutionary an advantage as fiber did when it started to replace coaxial cables and microwave relays in the late 1970s to mid-1980s. At that time, fiber cables could carry 100 times more traffic than coaxial cables, with the potential to scale capacity by another five orders of magnitude; they were also 10 times thinner and 100 times lighter, and allowed for 10 times longer repeater spacings, all of which made them easier to install. In today's terms, to yield a similar disruption, a new waveguide would have to support speeds of several petabits (thousands of terabits) per second without amplification over a 1,000 km span, with the ability to scale to support several hundred exabits (thousands of petabits) per second.

For now, at least, such waveguides clearly belong to the realm of fiction. Consequently, SDM systems must reuse the existing fiber infrastructure and available optical system components to the maximum possible extent. SDM networks will have to operate over a mixed infrastructure of parallel deployed single-mode fiber—which, when exhausted on certain spans, may gradually be replaced by the new waveguide technologies on

which optical communications research is so intensely working today. [OPN](#)

I gratefully acknowledge valuable discussions with S. Chandrasekhar, A. Chraplyvy, N. Fontaine, R.-J. Essiambre, S. Korotky, X. Liu, D. Neilson., S. Randel, G. Raybon, R. Ryf, R. Tkach and S. Trowbridge. Fiber cross-section photographs were kindly provided by R. Lingle and B. Zhu (OFS), T. Sasaki (Sumitomo), K. Imamura (Furukawa), Y. Miyamoto (NTT), S. Matsuo (Fujikura), M.-J. Li (Corning), D. Richardson (ORC Southampton), G. Li (CREOL), and R. Ryf and N. Fontaine (Alcatel-Lucent Bell Labs).

Peter J. Winzer (peter.winzer@bell-labs.com) is with Bell Labs, Alcatel-Lucent, Holmdel, N.J., USA.

References and Resources

- ▶ A. Chraplyvy. "The coming capacity crunch," plenary talk, European Conference on Optical Communications (ECOC), Vienna, Austria (2009).
- ▶ T. Morioka. "New generation optical infrastructure technologies: EXAT initiative towards 2020 and beyond," Proc. Optoelectron. Commun. Conf. (OECC), FT4 (2009).
- ▶ R.W. Tkach. "Scaling optical communications for the next decade and beyond," Bell Labs Tech. J. **14**, 3 (2010).
- ▶ R.S. Tucker. "Green optical communications—part I: Energy limitations in transport," IEEE J. Sel. Topics Quantum Electron. **17**(2), 245 (2011).
- ▶ R.-J. Essiambre and R.W. Tkach. "Capacity trends and limits of optical communication networks," Proc. IEEE **100**(5), 1035 (2012).
- ▶ D.J. Richardson et al. "Space-division multiplexing in optical fibres," Nat. Photon. **7**, 354 (2013).
- ▶ C. Laperle and M. O'Sullivan. "Advances in high-speed DACs, ADCs, and DSP for Optical Coherent Transceivers," J. Lightwave Tech. **32**(4), 629 (2014).
- ▶ S.G. Leon-Saval. "Multimode photonics, optical transition devices for multimode control," Proc. OECC, 2014.
- ▶ X. Liu et al. "Digital signal processing techniques enabling multi-Tb/s superchannel transmission," IEEE Signal Processing Magazine **31**(2), 16 (2014).
- ▶ Y. Miyamoto and H. Takenouchi. "Dense space-division-multiplexing optical communications technology for petabit-per-second class transmission," NTT Tech. Rev. **12**(12), 1 (2014).
- ▶ P.J. Winzer. "Spatial multiplexing in fiber optics: The 10x scaling of metro/core capacities," Bell Labs Tech. J. **19**, 22 (2014).
- ▶ K. Kikuchi. "Coherent optical communication technology," Proc. Opt. Fiber Commun. Conf. (OFC), paper Th4F.4 (2015).
- ▶ G. Raybon. "High symbol rate transmission systems for data rates from 400 Gb/s to 1Tb/s," OFC, paper M3G.1 (2015).