Making spatial multiplexing a reality

Peter J. Winzer

To avoid a ‘capacity crunch’, future optical networks will need to simultaneously transmit multiple spatial channels. For spatial multiplexing to be practical, the upgrade path from legacy wavelength-division multiplexed systems needs to be smooth and to consider integration-induced crosstalk from the outset.

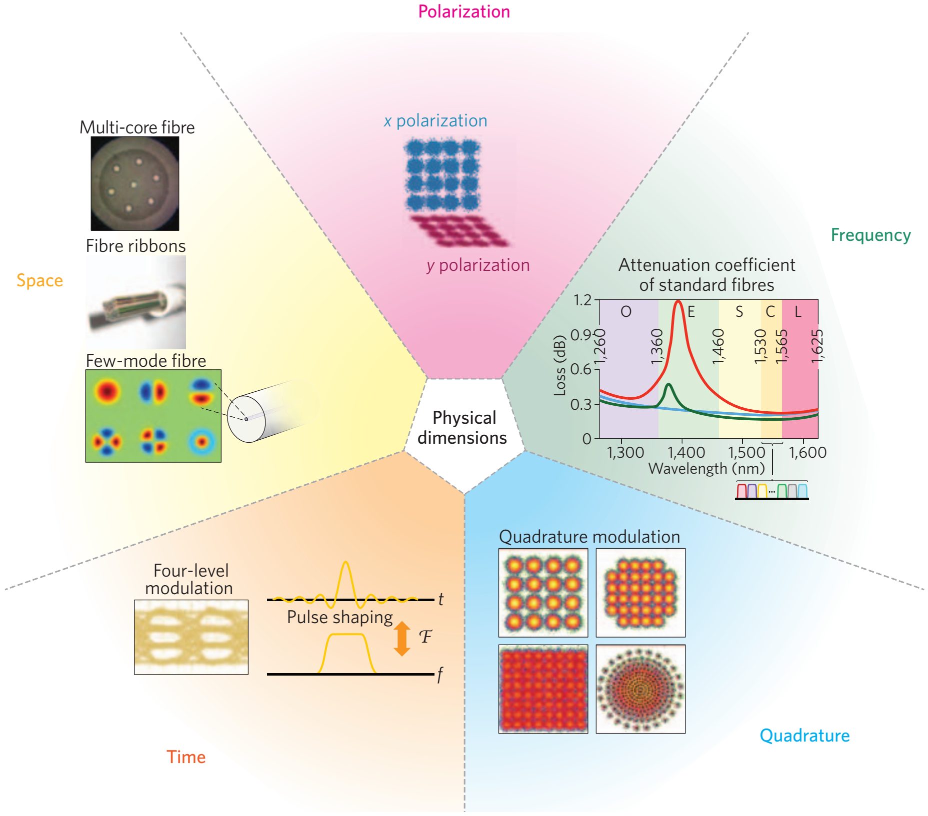

Figure 1 | Five physical dimensions (polarization, frequency, quadrature, time and space) form the basis of all electromagnetic communication techniques. Specific examples pertaining to optical communications are shown.

In addition to being dependent on conventional global resources such as fresh water, fossil fuels and electricity, modern society is becoming increasingly reliant on another critical resource — the speed (bit rate) at which digital data can be transmitted around the world. Yet, as is the case with many resources, the general public simply takes an adequate supply for granted. The problem is that the demand for data is increasing exponentially with annual growth rates between 30% and 90% (refs 1,2). These enormous growth rates apply to all segments of the network — from mobile wireless and fixed access to supercomputer and data-centre interconnects and to long-haul transport. As a result, a global community of researchers and engineers is relentlessly striving to design network infrastructure that can carry more data, more efficiently than ever before.

Five physical dimensions can be employed to carry optical data (Fig. 1): time, frequency, space, polarization and quadrature[3,4]. These dimensions can be simultaneously used to greatly increase the bit rate of a communication system.

The *time dimension* is exploited by sending communication symbols in temporal succession — just like assembling words and sentences in written text by concatenating characters from a predefined alphabet. Pulse shaping may be used to compress the spectrum of communication pulses subject to fundamental time–frequency constraints, and multilevel modulation may be employed to increase the number of information bits carried per pulse[4].

For communication channels that modulate pulses onto a carrier frequency much higher than the symbol rate (such as a microwave or optical carrier), both sine and cosine (or real and imaginary) components of the carrier wave may be exploited; these components are referred to as the two *quadrature dimensions*. This results in two-dimensional symbol alphabets, such as the examples of quadrature amplitude modulation shown in Fig. 1 (ref. 5).

Using the *frequency dimension*, one may transmit multiple communication signals in parallel on distinct carrier frequencies over the same transmission medium. This technique is known as wavelength-division multiplexing (WDM) in optical communications. The scalability limits of frequency multiplexing may be determined by regulatory bandwidth constraints on an inherently shared medium (for example, in mobile wireless) or by fundamental physical or engineering limitations on waveguides (for example, for coaxial, twisted-pair or fibre cables).

In some cases, such as in coherent optical communications, the *polarization dimension* may additionally be exploited to realize simultaneous transmission of multiple information streams.

Finally, making use of the *spatial dimension* entails a wide variety of techniques ranging from parallel tracks (data buses) on printed circuit boards, ribbon cables in a personal computer, multiple parallel twisted wire pairs in Ethernet cables, fibre ribbons, and spatial re-use and multi-antenna techniques in cellular wireless.

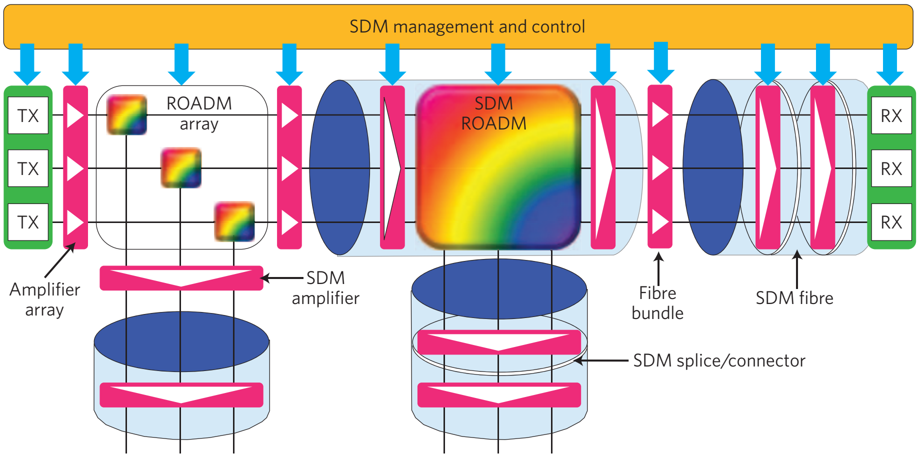

Figure 2 | An SDM network will have to operate across a diverse infrastructure, making use of the installed WDM infrastructure to the fullest possible extent. (Black lines indicate parallel spatial paths; TX, transmitter; RX, receiver; ROADM, reconfigurable optical add/drop multiplexer.)

Multiplexing order matters

Although the above-mentioned five physical dimensions form the common toolbox for all known communication techniques involving electromagnetic radiation, the preference and order in which they are adopted to increase the transmission capacity is highly application specific and economically sensitive. As a general starting point, it is usually preferable to modulate as fast as economically feasible in the time dimension before resorting to the use of any other physical dimension. In the context of commercial optical communication systems, economics currently limits channel modulation speeds to the range 10–50 Gbaud.

One next needs to consider the specifics of a system. For example, if a system permits many parallel optical fibres to be used, the space dimension can be easily exploited, making spatial multiplexing an attractive solution; this is often the case for board-to-board or rack-to-rack interconnects that have reaches of up to ~100 m. Longer-reach systems are usually expected to operate over single strands of transmission fibre, and consequently space becomes a design constraint. It is then most cost effective to first tap into the frequency dimension by deploying multiple carrier wavelengths (that is, wavelength-division multiplexing (WDM)). The bandwidth over which WDM signals may be deployed is limited by the low-loss window of optical fibres (~1,260 nm to 1,625 nm, corresponding to ~50 THz for standard fibre; see Fig. 1) and, in systems extending beyond several tens of kilometres, by the amplification bandwidth of optical amplifiers (typically ~5 THz for standard C-band amplifiers; see Fig. 1). Hence, increasing system capacity requires squeezing as much information as possible into a limited optical amplification bandwidth, making spectral efficiency (the ratio of the total information bit rate to the total system bandwidth) a key system parameter to be optimized. Techniques that can increase spectral efficiency[4] include the use of higher-order symbol constellations and polarization-division multiplexing (PDM); both of these techniques are based on digital coherent detection.

By exhausting all the physical dimensions except space, long-haul optical transmission research has now reached spectral efficiencies that approach (within a factor of two in the few-thousand-kilometre transmission regime) a fundamental limit called the nonlinear Shannon limit[6].

Indeed, commercial systems introduced in 2013 operate at factors of only four to seven from the nonlinear Shannon limit, with C-band system capacities of ~20 Tb s^{-1} (ref. 7). At an annual traffic growth rate of 30% (a conservative forecast), we anticipate that commercial systems will need to support in excess of 80 Tb s^{-1} by 2018.

The associated C-band spectral efficiencies are fundamentally impossible to achieve over required transmission distances, a situation that has become known as the 'capacity crunch' in the optical communications community[8]. As this conclusion is fairly insensitive to even substantial variations in fibre parameters[9], a parallel approach, either in frequency (entering a new spectral window of fibre transmission) or space, must be adopted. Exploiting multiple optical amplification bands across the low-loss window of a deployed fibre could increase the bandwidth by a factor of about ten relative to that available in the C-band, resulting in an approximately fivefold increase in system capacity when effects such as higher amplifier noise figures, increased span losses, spectral overlap between signals and Raman pumps, and potential 'fibre fuse' problems[10] are considered. The spatial dimension is the only remaining degree of freedom capable of offering the multiple orders of magnitude of capacity scalability that will ultimately be required. Hence, it is not a question of whether space-division multiplexing (SDM) will be adopted in long-haul transmission systems, but rather when it will be.

A smooth transition to SDM

When discussing the need for SDM to scale capacities beyond those of existing WDM systems, it is highly instructive to consider the big technology changes in long-haul transmission that occurred in the late 1970s, when fibre optics started to replace widely deployed coaxial cables for long-haul transport. Back then, massive coaxial cables with 3-inch outer diameters were exhausting duct space on many routes. These cables were deployed in 400-m sections with repeater spacings of about 3 km (ref. 11). Field trials and commercial deployments conducted between 1976[12] and 1986[13] demonstrated that fibre cables could carry over two orders of magnitude more traffic than coaxial cables. In addition, fibre cables were almost an order of magnitude thinner and almost two orders of magnitude lighter than coaxial cables; consequently, they could be installed much more easily and in much longer sections. Importantly, fibre cables also allowed for an order of magnitude increase in repeater spacings. Finally, their cost was significantly lower[14]. These multiple-orders-of-magnitude improvements quickly led to the widespread adoption of fibre optics in long-haul transmission, and fibres rapidly

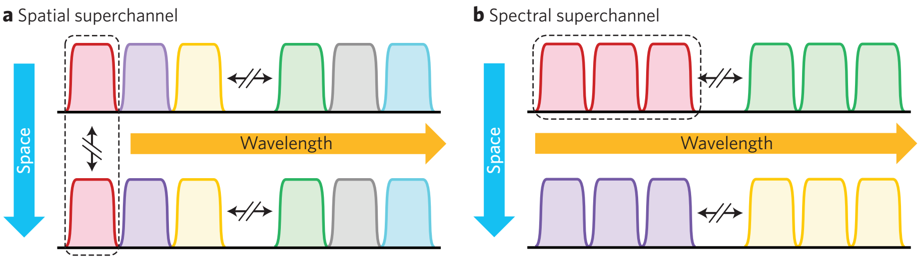

Figure 3 | Optical superchannels achieve a desired interface rate through the use of parallel data streams. **a**,**b** Parallelization may take place in space, leading to spatial superchannels (**a**), or in frequency, leading to spectral superchannels (**b**). In the presence of crosstalk among SDM paths, the need for MIMO processing forces network architectures using spatial superchannels (**a**) as opposed to spectral superchannels on independent spatial paths (**b**).

replaced all previously used technologies as well as contemporary contenders (such as millimetre-wave hollow waveguides[15]). These figures (particularly those for capacities and repeater spacings) illustrate that a revolutionary transmission medium similar to what fibre provided over coaxial copper is clearly not even theoretically in sight today.

In the absence of a new breakthrough transmission medium that could warrant a radical technology displacement similar to the copper-to-fibre transition in long-haul transport some 30 years ago, it is critically important to provide a smooth yet long-term viable upgrade path from existing WDM systems to future SDM networks. Several important aspects need to be considered in this context.

Compatibility: As long as parallel fibre strands are available within a deployed cable, operators will want to make use of those, based on existing and deployed wavelength bands for which mature components are available. Hence, hybrid network architectures that use parallel fibre strands on some spans and possibly new SDM-specific fibre on other spans must be supported (see Fig. 2). A transition to new wavelength bands seems unlikely unless it results in a substantial (orders of magnitude) increase in system capacity and/or repeater spacing.

Integration: To follow historic trends and reduce cost and energy consumption per transmitted bit by about 20% per year[16], integration and unification of equipment in parallel WDM systems is essential. As shown in Fig. 2, integration may take place on a systems level (including network management and control), on a network element level (including reconfigurable optical add/drop multiplexers (ROADMs)[17–19]), on a transponder level (similar to what is currently being done for spectral superchannel interfaces, where multiple signals at adjacent wavelengths are cohesively bonded to form a single architectural entity[20]), on an optical amplifier level, and on a fibre and splice/connection level to reduce installation costs. Various integration efforts in the context of recent SDM research are reviewed in refs 7, 21 and 22. Importantly, when proposing integrated SDM solutions, researchers will need to keep in mind the cost and/or energy benefit of the respective integration solution within the bigger picture. As Tucker[16] pointed out, a significant fraction of a subsystem's energy consumption is absorbed by an overhead that is inessential to the subsystem's key functionality. As a result, a reduction in the cost or energy consumption of the key functionality may have little impact on the efficiency of the entire system. Furthermore, although the integration of multiple fibres into multicore or few-mode structures has so far received the most attention in SDM research[23–26], it is unlikely to be the first aspect of commercial interest. The deployment of SDM-specific optical waveguides will probably be considered only if they provide beneficial interfacing (lower cost and more efficient optical coupling) to integrated SDM components[27]. The physical size reductions associated with integrated SDM waveguides do not seem to be of great relevance unless these new SDM waveguides, once cabled and installed, are much cheaper than a cable with an equivalent number of single-mode fibre strands[28].

Crosstalk: Hardware integration at various levels will inevitably result in crosstalk among parallel SDM paths and consequently in transmission penalties[4]. If the end-to-end crosstalk from integrated transponders, multiple ROADMs, optical amplifiers and transmission fibre with splices every few kilometres (that is, on the order of 1,000 splices for long-haul links) exceeds tolerable limits, crosstalk mitigation will need to be implemented at the receiver. This can be done by adapting multiple-input–multiple-output (MIMO) techniques developed for multi-antenna wireless systems[29] to optical transmission[30]. In essence, MIMO processing aims to restore the set of transmitted signal streams from a random superposition of received signal streams. In its simplest form, a MIMO processing receiver estimates the channel matrix and subjects the received signals to the inverse channel matrix to 'undo' the transmission channel. This is analogous to a PDM (= 2 × 2 MIMO) receiver, which estimates and inverts the fibre's Jones matrix to separate the two signal polarizations[4]. When MIMO processing is used, the amount of accumulated crosstalk becomes irrelevant. However, the need to use MIMO processing has huge consequences for network architectures. For a viable and smooth long-term upgrade path, the potential long-term need for MIMO should be built into SDM architectures from the very beginning, even if initial SDM systems are not limited by crosstalk. In particular, the need to use MIMO processing forces the adoption of spatial superchannels, where multiple spatial paths at a given wavelength must originate and terminate on the same line card so that they can be MIMO coprocessed. Figure 3a depicts the spatial superchannel concept, where an end-to-end signal occupies all spatial paths within a given wavelength slot. In contrast, the use of independent spectral superchannels and the addition of SDM paths as needed (Fig. 3b) — a direct extension of today's single-mode fibre WDM systems — would rule out the option of mitigating spatial crosstalk through MIMO coprocessing and potentially hamper the long-term evolution of SDM networks. The notion of adding, dropping and re-routing individual SDM tributaries is as incompatible with MIMO as the adding, dropping or re-routing of individual polarizations in today's PDM systems. Unless a clear value proposition for spatial switching in SDM networks is established that outweighs component cost and energy reductions through crosstalk-inducing integration, the potential need for MIMO lets the wavelength remain the unit of switching in an optical SDM network.

Conclusions

Although short-term stopgap solutions to an imminent capacity crunch in optical networks may use multiple optical amplification bands on legacy fibre, long-term capacity scalability can be guaranteed only through exploiting the spatial dimension. However, owing to its relatively modest value proposition compared to that of the copper-to-fibre transition that occurred some 30 years ago, practical SDM solutions must offer compatibility with and a smooth upgrade path from legacy WDM systems. Hardware integration across system components will be essential to provide the necessary cost and energy reductions compared to individually deployed parallel WDM systems, and SDM waveguides will likely prove more valuable due to improved interfacing than through saving physical space. Furthermore, hardware integration may introduce unacceptable amounts of spatial crosstalk, which can be mitigated through introducing MIMO signal processing at the receiver. As an architectural consequence, MIMO forces the use of spatial superchannels, leaving 'wavelength' as the optical routing dimension. The above considerations should be kept in mind in a holistic systems context when turning SDM research into practical solutions for scaling optical network capacity. ❐

Peter J. Winzer is at Bell Laboratories, Alcatel-Lucent, 791 Holmdel-Keyport Road, Holmdel, New Jersey 07733, USA.
e-mail: peter.winzer@alcatel-lucent.com

References

1. Tkach, R. W. *Bell Labs Tech. J.* **14,** 3–9 (2010).
2. IEEE 802.3 industry connections Ethernet bandwidth assessment (IEEE 802.3 BWA ad hoc report) http://www.ieee802.org/3/ad_hoc/bwa/BWA_Report.pdf (2012).
3. Winzer, P. J. *IEEE/LEOS Newsletter* http://photonicssociety.org/newsletters/feb09/modulation.pdf (2009).
4. Winzer, P. J. *J. Lightwave Technol.* **30,** 3824–3835 (2012).
5. Liu, X., Chandrasekhar, S. & Winzer, P. J. *IEEE Sig. Process. Mag.* **31,** 16–24 (2014).
6. Essiambre, R.-J., Kramer, G., Winzer, P. J., Foschini, G. J. & Goebel B. *J. Lightwave Technol.* **28,** 662–701 (2010).
7. Winzer, P. J. *Euro. Conf. Opt. Commun.* paper We.1.D.1 (2013).
8. Chraplyvy, A. in *Proc. Euro. Conf. Opt. Commun.* (2009).
9. Essiambre, R.-J. & Tkach, R. W. *Proc. IEEE* **100,** 1035–1055 (2012).
10. Morioka, T. *Proc. Optoelectron. Commun. Conf.* paper FT4 (2009).
11. Duvall, G. H. & Rackson, L. M. *Bell Syst. Tech. J.* **48,** 1065–1093 (1969).
12. Schwartz, M. I., Reenstra, W. A., Mullins, J. H. & Cook, J. S. *Bell Syst. Tech. J.* **57,** 1881–1888 (1978).
13. Fishman, D. A. *et al. Proc. Opt. Fib. Commun. Conf.* paper PD11 (1986).
14. Porter, M. E. *Cases in Competitive Strategy* (Macmillan, 1983).
15. Warters, W. D. *Bell Syst. Tech. J.* **56,** 1825–1827 (1977).
16. Tucker, R. S. *IEEE J. Sel. Top. Quantum Electron.* **17,** 245–260 (2011).
17. Feuer, M. D. *et al. Proc. Opt. Fib. Commun. Conf.* paper PDP5B.8 (2013).
18. Chen, X. *et al. J. Lightwave Technol.* **31,** 641–647 (2013).
19. Fontaine, N. *et al. Proc. Opt. Fib. Commun. Conf.* paper OTh1B.1 (2013).
20. Chandrasekhar, S. & Liu, X. *J. Lightwave Technol.* **30,** 3816–3823 (2012).
21. Morioka, T. *et al. IEEE Commun. Mag.* **50,** S31–S42 (2012).
22. Richardson, D. J., Fini, J. M. & Nelson, L. E. *Nature Photon.* **7,** 354–362 (2013).
23. Sano, A., Hidehiko, T., Kobayashi, T. & Miyamoto, Y. *Proc. Opt. Fib. Commun. Conf.* paper Tu2J.1 (2014).
24. Ryf, R. *et al. Proc. Opt. Fib. Commun. Conf.* paper PDP5A.1 (2013).
25. Ip, E. *et al. Proc. Opt. Fib. Commun. Conf.* paper PDP5A.2 (2013).
26. Sleiffer, V. A. J. M. *et al. Opt. Express* **22,** 749–755 (2014).
27. Taubenblatt, M. A. *J. Lightwave Technol.* **30,** 448–457 (2012).
28. Korotky, S. K. *J. Opt. Commun. Netw.* **4,** 426–435 (2012).
29. Foschini, G. J. *Bell Labs Tech. J.* **1,** 41–59 (1996).
30. Winzer, P. J. & Foschini, G. J. *Opt. Express* **19,** 16680–16696 (2011).