



Fiber-optic transmission and networking: the previous 20 and the next 20 years [Invited]

PETER J. WINZER,* DAVID T. NEILSON, AND ANDREW R. CHRAPLYVY

Nokia Bell Labs, 791 Holmdel Road, Holmdel, NJ 07733, USA

*peter.winzer@nokia-bell-labs.com

Abstract: Celebrating the 20th anniversary of Optics Express, this paper reviews the evolution of optical fiber communication systems, and through a look at the previous 20 years attempts to extrapolate fiber-optic technology needs and potential solution paths over the coming 20 years. Well aware that 20-year extrapolations are inherently associated with great uncertainties, we still hope that taking a significantly longer-term view than most texts in this field will provide the reader with a broader perspective and will encourage the much needed out-of-the-box thinking to solve the very significant technology scaling problems ahead of us. Focusing on the optical transport and switching layer, we cover aspects of large-scale spatial multiplexing, massive opto-electronic arrays and holistic optics-electronics-DSP integration, as well as optical node architectures for switching and multiplexing of spatial and spectral superchannels.

© 2018 Optical Society of America under the terms of the [OSA Open Access Publishing Agreement](#)

OCIS codes: (060.0060) Fiber optics and optical communications (200.4650) Optical interconnects.

References and links

1. P. J. Winzer and D. T. Neilson, "From scaling disparities to integrated parallelism: A decathlon for a decade," *J. Lightwave Technol.* **35**(5), 1099–1115 (2017).
2. R. S. Tucker, "Green optical communications—Part I: Energy limitations in transport," *IEEE J. Sel. Top. Quantum Electron.* **17**(2), 245–260 (2011).
3. S. H. Williamson, "Annualized growth rate of various historical economic series," On-line: <https://www.measuringworth.com/graphs/> (2018)
4. "Nortel launches first 10 Gbit/s transmission system in Asia," On-line: <https://www.hpcwire.com/1997/06/13/nortel-launches-first-10-gbits-transmission-system-in-asia/> (1997).
5. P. Trischitta, M. Colas, M. Green, G. Wuzniak, and J. Arena, "The TAT-12/13 cable network," *IEEE Commun. Mag.* **34**(2), 24–28 (1996).
6. A. H. Gnauck, R. W. Tkach, A. R. Chraplyvy, and T. Li, "High-capacity optical transmission systems," *J. Lightwave Technol.* **26**(9), 1032–1045 (2008).
7. H. Kogelnik, "On optical communication: reflections and perspectives," *Proc. European Conf. on Optical Comm. (ECOC), Mo1.1.1* (2004).
8. J. Cho, X. Chen, S. Chandrasekhar, G. Raybon, R. Dar, L. Schmalen, E. Burrows, A. Adamiecki, S. Corteselli, Y. Pan, D. Correa, B. McKay, S. Zsigmond, P. Winzer, and S. Grubb, "Trans-Atlantic field trial using high spectral efficiency probabilistically shaped 64-QAM and single-carrier real-time 250-Gb/s 16-QAM," *J. Lightwave Technol.* **36**(1), 103–113 (2018).
9. T. Zami, B. Lavigne, O. B. Pardo, S. Weisser, J. David, M. Le Monnier, and J. Faure, "31.2-Tb/s real time bidirectional transmission of 78x400 Gb/s interleaved channels over C band of one 90-km SMF span," *Proc. Optical Fiber Comm. Conf. (OFC), W1B.5* (2018).
10. G. Raybon, A. Adamiecki, J. Cho, P. J. Winzer, A. Konczykowska, F. Jorge, J.-Y. Dupuy, M. Riet, B. Duval, K. Kim, S. Randel, D. Piori, B. Guan, N. Fontaine, and E. C. Burrows, "Single-carrier all-ETDM 1.08-Terabit/s line rate PDM-64-QAM transmitter using a high-speed 3-bit multiplexing DAC," *Proc. IEEE Photonics Conf. (IPC'15), post-deadline session* (2015).
11. K. Schuh, F. Buchali, W. Idler, T. A. Eriksson, L. Schmalen, W. Templ, L. Altenhain, U. Dümmler, R. Schmid, M. Möller, and K. Engenhardt, "Single carrier 1.2 Tbit/s transmission over 300 km with PM-64 QAM at 100 GBaud," *Proc. Optical Fiber Comm. Conf. (OFC), Th5B.5* (2017). R. Maher, K. Croussore, M. Lauermaun, R. Going, X. Xu, and J. Rahn "Constellation shaped 66 GBd DP-1024QAM transceiver with 400 km transmission over standard SMF," *Proc. European Conf. on Optical Comm. (ECOC), Th.PDP.B.2* (2017).
12. X. Chen, S. Chandrasekhar, G. Raybon, S. Olsson, J. Cho, A. Adamiecki, and P. Winzer, "Generation and intradyne detection of single-wavelength 1.61-Tb/s using an all-electronic digital band interleaved transmitter," *Proc. Optical Fiber Comm. Conf. (OFC), Th4C.1* (2018).

13. D. Qian, M. Huang, E. Ip, Y. Huang, Y. Shao, J. Hu, and T. Wang, "101.7-Tb/s (370x294-Gb/s) PDM-128QAM-OFDM transmission over 3x55-km SSMF using pilot-based phase noise mitigation," Proc. Optical Fiber Comm. Conf. PDPB5 (2011).
14. A. Sano, T. Kobayashi, S. Yamanaka, A. Matsuura, H. Kawakami, Y. Miyamoto, K. Ishihara, and H. Masuda, "102.3-Tb/s (224 x 548-Gb/s) C- and extended L-band all-Raman transmission over 240 km using PDM-64QAM single carrier FDM with digital pilot tone," Proc. Optical Fiber Comm. Conf. (OFC), PDP5C.3 (2012).
15. J. Renaudier, A. C. Meseguer, A. Ghazisaeidi, P. Tran, R. R. Muller, R. Brenot, A. Verdier, F. Blache, K. Mekhazni, B. Duval, H. Debregeas, M. Achouche, A. Boutin, F. Morin, L. Letteron, N. Fontaine, Y. Frignac, and G. Charlet, "First 100-nm continuous-band WDM transmission system with 115Tb/s transport over 100km using novel ultra-wideband semiconductor optical amplifiers," Proc. European Conf. on Optical Comm. (ECOC), Th.PDP.A.3 (2017).
16. D. Soma, Y. Wakayama, S. Beppu, S. Sumita, T. Tsuritani, T. Hayashi, T. Nagashima, M. Suzuki, H. Takahashi, K. Igarashi, I. Morita, and M. Suzuki, "10.16 Peta-bit/s dense SDM/WDM transmission over low-DMD 6-mode 19-core fibre across C+L Band," Proc. European Conf. on Optical Comm. (ECOC), Th.PDP.A1 (2017).
17. S. L. I. Olsson, J. Cho, S. Chandrasekhar, X. Chen, E. C. Burrows, and P. J. Winzer, "Record-high 17.3-bit/s/Hz spectral efficiency transmission over 50 km using probabilistically shaped PDM 4096-QAM," Proc. Optical Fiber Comm. Conf. (OFC), Th4C.5 (2018).
18. J.-X. Cai, H. G. Batshon, M. V. Mazurczyk, O. V. Sinkin, D. Wang, M. Paskov, C. R. Davidson, W. W. Patterson, M. A. Bolshtyansky, and D. G. Foursa, "51.5 Tb/s capacity over 17,107 km in C+L bandwidth using single mode fibers and nonlinearity compensation," Proc. European Conf. on Optical Comm. (ECOC), Th.PDP.A.2 (2017).
19. A. V. Turukhin, O. V. Sinkin, H. G. Batshon, Y. Sun, M. Mazurczyk, C. R. Davidson, J.-X. Cai, M. A. Bolshtyansky, D. G. Foursa, and A. N. Pilipetskii, "High capacity ultralong-haul power efficient transmission using 12-core fiber," J. Lightwave Technol. **35**(4), 1028–1032 (2017).
20. R. S. Kerdock and D. H. Wolaver, "Atlanta fiber system experiment: results of the Atlanta experiment," Bell Syst. Tech. J. **57**(6), 1857–1879 (1978).
21. J. Hecht, *City of Light: The Story of Fiber Optics*, Sloan Technology (2004).
22. M. I. Schwartz, W. A. Reenstra, J. H. Mullins, and J. S. Cook, "The Chicago lightwave communications project," Bell Syst. Tech. J. **57**(6), 1881–1888 (1978).
23. R. W. Berry, D. J. Brace, and I. A. Ravenscroft, "Optical fiber system trials at 8 Mbits/s and 140 Mbit/s," IEEE Trans. Commun. **26**(7), 1020–1027 (1978).
24. A. Moncalvo and F. Tosco, "European field trials and early applications in telephony," in IEEE J., Select. Areas Comm. **1**(3), 398–403 (1983).
25. A. R. Chraplyvy, "The coming capacity crunch," European Conf. on Optical Comm. (ECOC), Vienna, Austria, Plenary Session (2009).
26. E. Agrell, M. Karlsson, A. R. Chraplyvy, D. J. Richardson, P. M. Krummrich, P. J. Winzer, K. Roberts, J. K. Fischer, S. J. Savory, B. J. Eggleton, M. Secondini, F. R. Kschischang, A. Lord, J. Prat, I. Tomkos, J. E. Bowers, S. Srinivasan, M. Brandt-Pearce, and N. Gisin, "Roadmap of optical communications," J. Opt. **18**(6), 063002 (2016).
27. P. J. Winzer, "Scaling optical fiber networks: Challenges and solutions," Opt. Photonics News **26**(3), 28–35 (2015).
28. R. J. Mears, L. Reekie, I. M. Jauncey, and D. N. Payne, "Low-noise erbium-doped fibre amplifier operating at 1.54 μ m," Electron. Lett. **23**(19), 1026–1028 (1987).
29. E. Desurvire, J. R. Simpson, and P. C. Becker, "High-gain erbium-doped traveling-wave fiber amplifier," Opt. Lett. **12**(11), 888–890 (1987).
30. G. Agrawal, *Nonlinear Fiber Optics*, 5th edition, Academic Press (2012).
31. F. Forghieri, R. W. Tkach, and A. R. Chraplyvy, "Fiber nonlinearities and their impact on transmission systems," in *Optical Fiber Telecommunications IIIA*, I. Kaminow and T. Koch (eds.), Academic Press, ch. 7, 196–264 (1997).
32. A. R. Chraplyvy, R. W. Tkach, and K. L. Walker, "Optical fiber for wavelength division multiplexing," U. S. Patent 5,327,516 (1994).
33. A. R. Chraplyvy, A. H. Gnauck, R. W. Tkach, and R. M. Derosier, "8 x 10 Gb/s transmission through 280 km of dispersion-managed fiber," IEEE Photonics Technol. Lett. **5**(10), 1233–1235 (1993).
34. R. W. Tkach, R. M. Derosier, A. H. Gnauck, A. M. Vengsarkar, D. W. Peckham, J. L. Zyskind, J. W. Sulhoff, and A. R. Chraplyvy, "Transmission of eight 20-Gb/s channels over 232 km of conventional singlemode fiber," IEEE Photonics Technol. Lett. **7**(11), 1369–1371 (1995).
35. R.-J. Essiambre, P. J. Winzer, and D. F. Grosz, "Impact of DCF properties in system design," in *Fiber Based Dispersion Compensation*, S. Ramachandran (ed.), Springer, ch. 12, 425–496 (2007).
36. A. M. Vengsarkar, A. E. Miler, and W. A. Reed, "Highly efficient single-mode fiber for broadband dispersion compensation," Proc. Optical Fiber Comm. Conf. (OFC), PD-13 (1993).
37. A. M. Vengsarkar, A. E. Miller, M. Haner, A. H. Gnauck, W. A. Reed, and K. L. Walker, "Fundamental-mode dispersion-compensating fibers: Design considerations and experiments," Proc. Optical Fiber Comm. Conf. (OFC), ThK2 (1994).
38. G. J. Foschini and C. D. Poole, "Statistical theory of polarization dispersion in single mode fibers," J. Lightwave Technol. **9**(11), 1439–1456 (1991).

39. J. P. Gordon and H. Kogelnik, "PMD fundamentals: Polarization mode dispersion in optical fibers," *Proc. Natl. Acad. Sci. U.S.A.* **97**(9), 4541–4550 (2000).
40. H. Kogelnik, R. M. Jopson, and L. E. Nelson, "Polarization-mode dispersion," in *Optical Fiber Telecommunications IVB*, I. Kaminow and T. Li (eds.), Academic Press, ch. **15**, 735–86 (2002).
41. M. Brodsky, N. J. Frigo, M. Boroditsky, and M. Tur, "Polarization mode dispersion of installed fibers," *J. Lightwave Technol.* **24**(12), 4584–4599 (2006).
42. A. C. Hart, Jr., R. G. Huff, and K. L. Walker, "Method of making a fiber having low polarization mode dispersion due to a permanent spin," U. S. Patent 5,298,047 (1994).
43. A. R. Chraplyvy, J. A. Nagel, and R. W. Tkach, "Equalization in amplified WDM lightwave transmission systems," *IEEE Photonics Technol. Lett.* **4**(8), 920–922 (1992).
44. A. R. Chraplyvy, R. W. Tkach, K. C. Reichmann, P. D. Magill, and J. A. Nagel, "End-to-end equalization experiments in amplified WDM lightwave systems," *IEEE Photonics Technol. Lett.* **5**(4), 428–429 (1993).
45. A. H. Gnauck, R. M. Derosier, A. R. Chraplyvy, and R. W. Tkach, "160 Gbit/s (8 x 20 Gbit/s WDM) 300 km transmission with 50 km amplifier spacing and span-by-span dispersion reversal," *Electron. Lett.* **30**(15), 1241–1243 (1994).
46. N. S. Bergano, J. Aspell, C. R. Davidson, P. R. Trischitta, B. M. Nyman, and F. W. Kerfoot, "Bit error rate measurements of a 14 000 km 5 Gb/s fiber-amplifier transmission system using a circulating loop," *Electron. Lett.* **27**(21), 1889–1890 (1991).
47. H. Toba, K. Oda, K. Nakanishi, N. Shibata, K. Nosu, N. Takato, and M. Fukuda, "A 100-channel optical FDM transmission/distribution at 622 Mb/s over 50 km," *J. Lightwave Technol.* **8**(9), 1396–1401 (1990).
48. A. H. Gnauck, P. J. Winzer, S. Chandrasekhar, X. Liu, B. Zhu, and D. W. Peckham, "10 × 224-Gb/s WDM transmission of 28-Gbaud PDM 16-QAM on a 50-GHz grid over 1,200 km of fiber," *Proc. Optical Fiber Comm. Conf. (OFC), PDPB8* (2010).
49. M. E. McCarthy, N. M. Suibhne, S. T. Le, P. Harper, and A. D. Ellis, "High spectral efficiency transmission emulation for non-linear transmission performance estimation for high order modulation formats," *Proc. European Conf. on Optical Comm. (ECOC), P.5.7* (2014).
50. N. S. Bergano and C. R. Davidson, "Circulating loop transmission experiments for the study of long-haul transmission systems using erbium-doped fiber amplifiers," *J. Lightwave Technol.* **13**(5), 879–888 (1995).
51. S. Lee, Q. Yu, L. S. Yan, Y. Xie, O. H. Adamczyk, and A. E. Willner, "A short recirculating fiber loop testbed with accurate reproduction of Maxwellian PMD statistics," *Proc. Optical Fiber Comm. Conf. (OFC), WT2* (2001).
52. M. Nissov, J. Cai, A. N. Pilipetskii, Y. Cai, C. R. Davidson, A. Lucero, and N. S. Bergano, "Q-factor fluctuations in long distance circulating loop transmission experiments," *Proc. Optical Fiber Comm. Conf. (OFC), FM6*(2004).
53. A. H. Gnauck, S. Chandrasekhar, and A. R. Chraplyvy, "Stroboscopic BER effects in recirculating-loop optical transmission experiments," *IEEE Photonics Technol. Lett.* **17**(9), 1974–1976 (2005).
54. T. Li, Private communication (1995).
55. B. Gowan, "The story behind the founding of Ciena," On-line at http://www.ciena.com/insights/articles/Ciena-20-The-Founding-of-Ciena_prx.html
56. R. E. Wagner, R. C. Alferness, A. A. M. Saleh, and M. S. Goodman, "MONET: Multi-wavelength optical networking," *J. Lightwave Technol.* **14**(6), 1349–1355 (1996).
57. W. T. Anderson, J. Jackel, G.-K. Chang, Hongxing Dai, Wei Xin, M. Goodman, C. Allyn, M. Alvarez, O. Clarke, A. Gottlieb, F. Kleytman, J. Morreale, V. Nichols, A. Tzathas, R. Vora, L. Mercer, H. Dardy, E. Renaud, L. Williard, J. Perreault, R. McFarland, and T. Gibbons, "The MONET project - A final report," *J. Lightwave Technol.* **18**(12), 1988–2009 (2000).
58. N. S. Bergano, "Wavelength division multiplexing in long-haul transoceanic transmission systems," *J. Lightwave Technol.* **23**(12), 4125–4139 (2005).
59. N. Bergano, "Undersea communication systems," in *Optical Fiber Telecommunications IVB*, I. Kaminow and T. Li (eds.), Academic Press, ch. **4**, 154–197 (2002).
60. S. Bigo, "Technologies for global telecommunications using undersea cables," in *Optical Fiber Telecommunications VB*, I. Kaminow, T. Li, and A. E. Willner (eds.), Academic Press, ch. **14**, 561–610 (2008).
61. K. Takehira, "Submarine system powering," in *Undersea Fiber Communication Systems*, 2nd ed., J. Chesnoy (ed.), Academic Press (2016).
62. A. Pilipetskii, "High capacity submarine transmission systems," *Proc. Optical Fiber Comm. Conf. (OFC), W3G.5* (2015).
63. A. Pilipetskii, "The role of SDM in future transoceanic transmission systems," *Proc. European Conf. on Optical Comm. (ECOC), Tu.1.E.1* (2017).
64. R. Dar, P. J. Winzer, A. R. Chraplyvy, S. Zsigmond, K.-Y. Huang, H. Fevrier, and S. Grubb, "Cost-optimized submarine cables using massive spatial parallelism," *J. Lightwave Technol.* **36**(18), 3855–3865 (2018).
65. A. Hasegawa and F. Tappert, "Transmission of stationary nonlinear optical pulses in dispersive dielectric fibers. I. Anomalous dispersion," *Appl. Phys. Lett.* **23**(3), 142–144 (1973).
66. L. F. Mollenauer, R. H. Stolen, and J. P. Gordon, "Experimental observation of picosecond pulse narrowing and solitons in optical fibers," *Phys. Rev. Lett.* **45**(13), 1095–1098 (1980).
67. S. K. Turitsyn, B. G. Bale, and M. P. Fedoruk, "Dispersion-managed solitons in fibre systems and laser," *Phys. Rep.* **521**(4), 135–203 (2012).

68. L. F. Mollenauer, J. P. Gordon, and P. V. Mamyshev, "Solitons in high-bit-rate, long-distance transmission," in *Optical Fiber Telecommunications IIIA*, I. Kaminow and T. Koch (eds.), Academic Press, ch. 12, 373–460 (1997).
69. J. P. Gordon and H. A. Haus, "Random walk of coherently amplified solitons in optical fiber transmission," *Opt. Lett.* **11**(10), 665–667 (1986).
70. L. F. Mollenauer, J. P. Gordon, and S. G. Evangelides, "The sliding-frequency guiding filter: An improved form of soliton jitter control," *Opt. Lett.* **17**(22), 1575–1577 (1992).
71. S. Yamamoto, H. Taga, I. Morita, N. Edagawa, S. Akiba, and M. Suzuki, "Reduction of Gordon-Haus timing jitter by periodic dispersion compensation in soliton transmission," *Electron. Lett.* **31**(23), 2027–2029 (1995).
72. N. S. Bergano, C. R. Davidson, M. Ma, A. N. Pilipetskii, S. G. Evangelides, H. D. Kidorf, J. M. Darcie, E. A. Golovchenko, K. Rotwitt, P. C. Corbett, R. Menges, M. A. Mills, B. Pedersen, D. Peckham, A. A. Abramov, and A. M. Vengsarkar, "320 Gb/s WDM transmission (64 x 5 Gb/s) over 7200 km using large mode fiber spans and chirped return-to-zero signals," *Proc. Optical Fiber Comm. Conf. (OFC)*, PD2 (1999).
73. C. R. Menyuk, G. M. Carter, W. K. Kath, and R.-M. Mu, "Fiber nonlinearities and their impact on transmission systems," in *Optical Fiber Telecommunications IVB*, I. Kaminow and T. Li (eds.), Academic Press, ch. 7, 305–328 (2002).
74. D. A. Fishman, W. A. Thompson, and L. Vallone, "LambdaXtreme® transport system: R&D of a high capacity system for low cost, ultra long haul DWDM transport," *Bell Labs Tech. J.* **11**(2), 27–53 (2006).
75. P. J. Winzer, G. Raybon, H. Song, A. Adamiecki, S. Corteselli, A. H. Gnauck, D. A. Fishman, C. R. Doerr, S. Chandrasekhar, L. L. Buhl, T. J. Xia, G. Wellbrock, W. Lee, B. Basch, T. Kawanishi, K. Higuma, and Y. Painchaud, "100-Gb/s DQPSK transmission: From laboratory experiments to field trials," *J. Lightwave Technol.* **26**(20), 3388–3402 (2008).
76. A. R. Pratt, P. Harper, S. B. Alleston, P. Bontemps, B. Charbonnier, W. Forsysiak, L. Gleeson, D. S. Govan, G. L. Jones, D. Nessel, J. H. Nijhof, I. D. Phillips, M. F. Stephens, A. P. Walsh, T. Widdowson, and N. J. Doran, "5,745 km DWDM transcontinental field trial using 10 Gbit/s dispersion managed solitons and dynamic gain equalization," *Proc. Optical Fiber Comm. Conf. (OFC)*, PD26 (2003).
77. A. Hasegawa and Y. Kodama, *Solitons in Optical Communications*, Clarendon, Oxford, UK (1995).
78. R. Dar, M. Feder, A. Mecozzi, and M. Shtaif, "Pulse collision picture of inter-channel nonlinear interference in fiber-optic communications," *J. Lightwave Technol.* **34**(2), 593–607 (2016).
79. S. K. Turitsyn, J. E. Prilepsky, S. T. Le, S. Wahls, L. L. Frumin, M. Kamalian, and S. A. Derevyanko, "Nonlinear Fourier transform for optical data processing and transmission: advances and perspectives," *Optica* **4**(3), 307–322 (2017).
80. J. E. Ford, V. Aksyuk, J. A. Walker, and D. J. Bishop, "Wavelength-selectable add/drop with tilting micro mirrors," *Proc. IEEE LEOS Annual Meeting*, PD2.3 (1997).
81. D. A. Fishman, J. Ying, X. Liu, S. Chandrasekhar, and A. H. Gnauck, "Optical add/drop multiplexer with asymmetric bandwidth allocation and dispersion compensation for hybrid 10-Gb/s and 40-Gb/s DWDM transmission," *Proc. Optical Fiber Comm. Conf. (OFC)*, OW164 (2006).
82. H. Onaka, H. Miyata, G. Ishikawa, K. Otsuka, and H. Ooi, Y. Kai, S. Kinoshita, M. Seino, H. Nishimoto, and T. Chikama, "1.1 Tb/s WDM transmission over a 150 km 1.3 μm zero-dispersion single-mode fiber," *Proc. Optical Fiber Comm. Conf. (OFC)*, PD-19 (1996).
83. A. H. Gnauck, A. R. Chraplyvy, R. W. Tkach, J. L. Zyskind, J. W. Sulhoff, A. J. Lucero, Y. Sun, R. M. Jopson, F. Forghieri, R. M. Derosier, C. Wolf, and A. R. McCormick, "One terabit/s transmission experiment," *Proc. Optical Fiber Comm. Conf. (OFC)*, PD-20 (1996).
84. T. Morioka, H. Takara, S. Kawanishi, O. Kamatani, K. Takiguchi, K. Uchiyama, M. Saruwatari, H. Takahashi, M. Yamada, T. Kanamori, and H. Ono, "100 Gbit/s x 10 channel OTDM/WDM transmission using a single supercontinuum WDM source," *Proc. Optical Fiber Comm. Conf. (OFC)*, PD-21 (1996).
85. K. Fukuchi, T. Kasamatsu, M. Morie, R. Ohhira, T. Ito, K. Sekiya, D. Ogasahara, and T. Ono, "10.92-Tb/s (273 x 40 Gb/s) triple-band/ultra-dense WDM optical-repeated transmission experiment," *Proc. Optical Fiber Comm. Conf. (OFC)*, PD24 (2001).
86. S. Bigo, Y. Frignac, G. Charlet, W. Idler, S. Borne, H. Gross, R. Dischler, W. Poehlmann, P. Tran, C. Simonneau, D. Bayart, G. Veith, A. Jourdan, and J.-P. Hamaide, "10.2 Tbit/s (256 42.7 Gbit/s PDM/ WDM) transmission over 100 km TeraLight™ fiber with 1.28 bit/s/Hz spectral efficiency," *Proc. Optical Fiber Comm. Conf. (OFC)*, PD25 (2001).
87. A. H. Gnauck, G. Charlet, P. Tran, P. J. Winzer, C. R. Doerr, J. Centanni, E. Burrows, T. Kawanishi, T. Sakamoto, K. Higuma, *et al.*, "25.6-Tb/s C+L-band transmission of polarization-multiplexed RZ-DQPSK signals," *Proc. Optical Fiber Comm. Conf. (OFC)*, PDP19 (2007).
88. P. J. Winzer and R.-J. Essiambre, "Advanced optical modulation formats," *Proc. IEEE* **94**(5), 952–985 (2006).
89. A. H. Gnauck and P. J. Winzer, "Optical phase-shift-keyed transmission," *J. Lightwave Technol.* **23**(1), 115–130 (2005).
90. S. Perrin, "Deployment and service activation of 100G and beyond," *Heavy Reading*, 1–9, March (2015).
91. T. Okoshi and K. Kikuchi, *Coherent Optical Fiber Communications*, Springer (1988).
92. R. A. Linke and A. H. Gnauck, "High-capacity coherent lightwave systems," *J. Lightwave Technol.* **6**(11), 1750–1769 (1988).
93. F. Derr, "Optical QPSK transmission system with novel digital receiver concept," *Electron. Lett.* **27**(23), 2177–2179 (1991).

94. F. Derr, "Coherent optical QPSK intradyne system: concept and digital receiver realization," *J. Lightwave Technol.* **10**(9), 1290–1296 (1992).
95. N. Alić, G. Papen, R. Saperstein, L. Milstein, and Y. Fainman, "Signal Statistics and Maximum Likelihood Sequence Estimation in Intensity Modulated Fiber Optic Links Containing a Single Optical Pre-amplifier," *Opt. Express* **13**(12), 4568–4579 (2005).
96. A. Farbert, S. Langenbach, N. Stojanovic, C. Dorschky, T. Kupfer, and C. Schulien, "Performance of a 10.7-Gb/s receiver with digital equalizer using maximum likelihood sequence estimation," *Proc. European Conf. on Optical Comm. (ECOC), Th4.1.5* (2004).
97. D. McGhan, C. Laperle, A. Savehenko, Chuandong Li, G. Mak, and M. O'Sullivan, "5120 km RZ-DPSK transmission over G652 fiber at 10 Gb/s with no optical dispersion compensation," *Proc. Optical Fiber Comm. Conf. (OFC), PDP27* (2005).
98. Bo Gowan, "Coherent optical turns 10: Here's how it was made," On-line:
<http://www.ciena.com/insights/articles/Coherent-optical-turns-10-Heres-how-it-was-made-prx.html>
99. A. Fercher, C. Hitzberger, M. Sticker, R. Zawadzki, B. Karamata, and T. Lasser, "Numerical dispersion compensation for partial coherence interferometry and optical coherence tomography," *Opt. Express* **9**(12), 610–615 (2001).
100. Z. Wang, B. Potsaid, L. Chen, C. Doerr, H.-C. Lee, T. Nielson, V. Jayaraman, A. E. Cable, E. Swanson, and J. G. Fujimoto, "Cubic meter volume optical coherence tomography," *Optica* **3**(12), 1496–1503 (2016).
101. M. G. Taylor, "Coherent detection method using DSP for demodulation of signal and subsequent equalization of propagation impairments," *IEEE Photonics Technol. Lett.* **16**(2), 674–676 (2004).
102. R. Noe, "PLL-free synchronous QPSK polarization multiplex/diversity receiver concept with digital I&Q baseband processing," *IEEE Photonics Technol. Lett.* **17**(4), 887–889 (2005).
103. D.-S. Ly-Gagnon, S. Tsukamoto, K. Katoh, and K. Kikuchi, "Coherent detection of optical quadrature phase-shift keying signals with carrier phase estimation," *J. Lightwave Technol.* **24**(1), 12–21 (2006).
104. S. J. Savory, A. D. Stewart, S. Wood, G. Gavioli, M. G. Taylor, R. I. Killey, and P. Bayvel, "Digital equalization of 40 Gbit/s per wavelength transmission over 2480 km of standard fibre without optical dispersion compensation," *Proc. European Conf. on Optical Comm. (ECOC), Th2.5.5* (2006).
105. A. Leven, N. Kaneda, A. Klein, U.-V. Koc, and Y.-K. Chen, "Real-time implementation of 4.4 Gbit/s QPSK intradyne receiver using field programmable gate array," *Electron. Lett.* **42**(24), 1421–1422 (2006).
106. C. R. S. Fludger, T. Duthel, T. Wuth, and C. Schulien, "Uncompensated transmission of 86 Gbit/s polarization multiplexed RZ-QPSK over 100 km of NDSF employing coherent equalization," *Proc. European Conf. on Optical Comm. (ECOC), Th4.3.3* (2006).
107. G. Charlet, N. Maaref, J. Renaudier, and H. Mardoyan, P/ Tran, and S. Bigo, "Transmission of 40 Gb/s QPSK with coherent detection over ultra-long distance improved by nonlinearity mitigation," *Proc. European Conf. on Optical Comm. (ECOC), Th4.3.4* (2006).
108. H. Sun, K.-T. Wu, and K. Roberts, "Real-time measurements of a 40 Gb/s coherent system," *Opt. Express* **16**(2), 873–879 (2008).
109. N. Eiselt, J. Wei, H. Griesser, A. Dochhan, M. Eiselt, J.-P. Elbers, J. J. Vegas Olmos, and I. Tafur Monroy, "First real-time 400G PAM-4 demonstration for inter-datacenter transmission over 100 km of SSMF at 1550 nm," *Proc. Optical Fiber Comm. Conf., W1K.5D* (2016).
110. D. Che, A. Li, X. Chen, Q. Hu, and W. Shieh, "Rejuvenating direct modulation and direct detection for modern optical communications," *Opt. Commun.* **409**, 86–93 (2018).
111. A. Mecozzi, C. Antonelli, and M. Shtaif, "Kramers–Kronig coherent receiver," *Optica* **3**(11), 1220–1227 (2016).
112. X. Chen, C. Antonelli, S. Chandrasekhar, G. Raybon, A. Mecozzi, M. Shtaif, and P. J. Winzer, "Kramers-Kronig Receivers for 100-km Datacenter Interconnects," *J. Lightwave Technol.* **36**(1), 79–89 (2018).
113. K. Kikuchi, "Coherent optical communication systems," in *Optical Fiber Telecommunications V/B*, I. Kaminow, T. Li, and A. Willner (eds.), Academic Press, ch. 3, 95–129 (2008).
114. K. Kikuchi, "Fundamentals of coherent optical fiber communications," *J. Lightwave Technol.* **34**(1), 157–179 (2016).
115. S. J. Savory, "Digital filters for coherent optical receivers," *Opt. Express* **16**(2), 804–817 (2008).
116. S. J. Savory, "Digital coherent optical receivers: algorithms and subsystems," *IEEE J. Sel. Top. Quantum Electron.* **16**(5), 1164–1179 (2010).
117. A. Leven, N. Kaneda, U.-V. Koc, and Y.-K. Chen, "Coherent receiver for practical optical communication systems," *Proc. Optical Fiber Comm. Conf. (OFC), OThK4* (2007).
118. P. J. Winzer, "High-spectral-efficiency optical modulation formats," *J. Lightwave Technol.* **30**(24), 3824–3835 (2012).
119. R.-J. Essiambre, P. J. Winzer, W. Lee, C. A. White, and E. C. Burrows; Xun Qing Wang, "Electronic predistortion and fiber nonlinearity," *IEEE Photonics Technol. Lett.* **18**(17), 1804–1806 (2006).
120. G. Charlet, J. Renaudier, O. Bertran Pardo, P. Tran, H. Mardoyan, and S. Bigo, "Performance comparison of singly-polarized and polarization-multiplexed at 10 Gbaud under nonlinear impairments," *Proc. Optical Fiber Comm. Conf. (OFC), OThU8* (2008).
121. V. Curri, P. Poggiolini, A. Carena, and F. Forghieri, "Dispersion compensation and mitigation of nonlinear effects in 111-Gb/s WDM coherent PM-QPSK systems," *IEEE Photonics Technol. Lett.* **20**(17), 1473–1475 (2008).

122. R.-J. Essiambre, G. Kramer, P. J. Winzer, G. J. Foschini, and B. Goebel, "Capacity limits of optical fiber networks," *J. Lightwave Technol.* **28**(4), 662–701 (2010).
123. R.-J. Essiambre and R. W. Tkach, "Capacity Trends and Limits of Optical Communication Networks," *Proc. IEEE* **100**(5), 1035–1055 (2012).
124. S. Ten, "Ultra low-loss optical fiber technology," *Proc. Optical Fiber Comm. Conf. (OFC), Th4E.5* (2016).
125. Y. Tamura, H. Sakuma, K. Morita, M. Suzuki, Y. Yamamoto, K. Shimada, Y. Honma, K. Sohma, T. Fujii, and T. Hasegawa, "The first 0.14-dB/km loss optical fiber and its impact on submarine transmission," *J. Lightwave Technol.* **36**(1), 44–49 (2018).
126. T. Hasegawa, Y. Yamamoto, and M. Hirano, "Optimal fiber design for large capacity long haul coherent transmission [Invited]," *Opt. Express* **25**(2), 706–712 (2017).
127. A. Splett, C. Kurtzke, and K. Petermann, "Ultimate transmission capacity of amplified optical fiber communication systems taking into account fiber nonlinearities," *Proc. European Conf. on Optical Comm. (ECOC), MoC2.4* (1993).
128. P. Poggiolini, G. Bosco, A. Carena, V. Curri, Y. Jiang, and F. Forghieri, "The GN model of fiber non-linear propagation and its applications," *J. Lightwave Technol.* **32**(4), 694–721 (2014).
129. A. Carena, G. Bosco, V. Curri, Y. Jiang, P. Poggiolini, and F. Forghieri, "EGN model of non-linear fiber propagation," *Opt. Express* **22**(13), 16335–16362 (2014).
130. R. Dar, M. Feder, A. Mecozzi, and M. Shtaif, "Accumulation of nonlinear interference noise in fiber-optic systems," *Opt. Express* **22**(12), 14199–14211 (2014).
131. R. Dar, M. Feder, A. Mecozzi, and M. Shtaif, *Nonlinear Interference Noise Wizard*, online: <http://nlinwizard.eng.tau.ac.il>
132. P. J. Winzer, "From first fibers to mode-division multiplexing," *Chin. Opt. Lett.* **14**(12), 120002 (2016).
133. I. Burrington, "A journey into the heart of Facebook: How the "sharing" company's data centers reveal its values," *The Atlantic* (2015); on-line: <https://www.theatlantic.com/technology/archive/2015/12/facebook-data-center-tk/418683/>
134. J. Cho, C. Xie, and P. J. Winzer, "Analysis of soft-decision FEC on non-AWGN channels," *Opt. Express* **20**(7), 7915–7928 (2012).
135. A. Alvarado, E. Agrell, D. Lavery, R. Maher, and P. Bayvel, "Replacing the soft-decision FEC limit paradigm in the design of optical communication systems," *J. Lightwave Technol.* **33**(20), 4338–4352 (2015).
136. A. Alvarado, T. Fehenberger, B. Chen, and F. M. J. Willems, "Achievable information rates for fiber optics: Applications and computations," *J. Lightwave Technol.* **36**(2), 424–439 (2018).
137. J. Cho, L. Schmalen, and P. J. Winzer, "Normalized generalized mutual information as a forward error correction threshold for probabilistically shaped QAM," *Proc. European Conf. on Optical Comm. (ECOC), M.2.D* (2017).
138. D. Charlton, S. Clarke, D. Doucet, M. O'Sullivan, D. L. Peterson, D. Wilson, G. Wellbrock, and M. Bélanger, "Field measurements of SOP transients in OPGW, with time and location correlation to lightning strikes," *Opt. Express* **25**(9), 9689–9696 (2017).
139. M. Birk, P. Gerard, R. Curto, L. Nelson, Xiang Zhou, P. Magill, T. J. Schmidt, C. Malouin, B. Zhang, E. Ibragimov, S. Khatana, M. Glavanovic, R. Lofland, R. Marcocchia, R. Saunders, G. Nicholl, M. Nowell, and F. Forghieri, "Coherent 100 Gb/s PM-QPSK field trial," *IEEE Commun. Mag.* **48**(7), 52–60 (2010).
140. P. J. Winzer, A. H. Gnauck, G. Raybon, M. Schnecker, and P. J. Pupalakis, "56-Gbaud PDM-QPSK: Coherent detection and 2,500-km transmission," *Proc. European Conf. on Optical Comm. (ECOC), PD2.7* (2009).
141. G. Raybon, A. L. Adamiecki, P. J. Winzer, C. Xie, A. Konczykowska, F. Jorge, J.-Y. Dupuy, L. L. Buhl, S. Chandrasekhar, S. Draving, M. Grove, and K. Rush, "Single-carrier 400G interface and 10-channel WDM transmission over 4,800 km using all-ETDM 107-Gbaud PDM-QPSK," *Proc. Optical Fiber Comm. Conf. (OFC), PDP5A.5* (2013).
142. P. J. Winzer, G. Raybon, and M. Duelk, "107-Gb/s optical ETDM transmitter for 100G Ethernet transport," *Proc. European Conf. on Optical Comm. (ECOC), Th4.1.1* (2005).
143. M. Nakazawa, M. Yoshida, K. Kasai, and J. Hongou, "20 Msymbol/s, 64 and 128 QAM coherent optical transmission over 525 km using heterodyne detection with frequency-stabilized laser," *Electron. Lett.* **42**(12), 710–712 (2006).
144. L. B. Du and A. J. Lowery, "Optimizing the subcarrier granularity of coherent optical communications systems," *Opt. Express* **19**(9), 8079–8084 (2011).
145. J.-X. Cai, M. Mazurczyk, O. V. Sinkin, M. Bolshtyansky, D. G. Foursa, and A. Pilipetskii, "Experimental study of subcarrier multiplexing benefit in 74 nm bandwidth transmission up to 20,450 km," *Proc. European Conf. on Optical Comm. (ECOC), W.3.D.4* (2016).
146. Y. Ma, Q. Yang, Y. Tang, S. Chen, and W. Shieh, "1-Tb/s per channel coherent optical OFDM transmission with subwavelength bandwidth access," *Proc. Optical Fiber Comm. Conf. (OFC), PDPC1* (2009).
147. R. Dischler and F. Buchali, "Transmission of 1.2 Tb/s continuous waveband PDM-OFDM-FDM signal with spectral efficiency of 3.3 bit/s/Hz over 400 km of SSMF," *Proc. Optical Fiber Comm. Conf. (OFC), PDPC2* (2009).
148. S. Chandrasekhar, X. Liu, B. Zhu, and D. W. Peckham, "Transmission of a 1.2-Tb/s 24-carrier no-guard-interval CO-OFDM superchannel over 7200-km of ultra-large-area fiber," *Proc. European Conf. on Optical Comm. (ECOC), PD2.6* (2009).

149. G. Gavioli, E. Torrenco, G. Bosco, A. Carena, V. Curri, V. Miot, P. Poggiolini, M. Belmonte, F. Forghieri, C. Muzio, S. Piciaccia, A. Brinciotti, A. La Porta, C. Lezzi, S. Savory, and S. Abrate, "Investigation of the impact of ultra-narrow carrier spacing on the transmission of a 10-carrier 1Tb/s superchannel," Proc. Optical Fiber Comm. Conf. (OFC), OThD3 (2010).
150. G. Bosco, V. Curri, A. Carena, P. Poggiolini, and F. Forghieri, "On the performance of Nyquist-WDM terabit superchannels based on PM-BPSK, PM-QPSK, PM-8QAM or PM-16QAM subcarriers," J. Lightwave Technol. **29**(1), 53–61 (2011).
151. H. Goto, K. Kasai, M. Yoshida, and M. Nakazawa, "Polarization-multiplexed 1 Gsymbol/s, 128 QAM (14 Gbit/s) coherent optical transmission over 160 km using a 1.4 GHz Nyquist filter," Proc. Optical Fiber Comm. Conf. (OFC), JThA45 (2008).
152. X. Zhou, L. E. Nelson, P. Magill, R. Isaac, B. Zhu, D. W. Peckham, P. I. Borel, and K. Carlson, "PDM-Nyquist-32QAM for 450-Gb/s per-channel WDM transmission on the 50 GHz ITU-T grid," J. Lightwave Technol. **30**(4), 553–559 (2012).
153. F. Buchali, F. Steiner, G. Böcherer, L. Schmalen, P. Schulte, and W. Idler, "Rate adaptation and reach increase by probabilistically shaped 64-QAM: An experimental demonstration," J. Lightwave Technol. **34**(7), 1599–1609 (2016).
154. S. Okamoto, M. Terayama, M. Yoshida, K. Kasai, T. Hirooka, and M. Nakazawa, "Experimental and numerical comparison of probabilistically shaped 4096 QAM and a uniformly shaped 1024 QAM in all-Raman amplified 160 km transmission," Opt. Express **26**(3), 3535–3543 (2018).
155. S. L. I. Olsson, J. Cho, S. Chandrasekhar, X. Chen, P. J. Winzer, and S. Makovejs, "Probabilistically shaped PDM 4096-QAM transmission over up to 200 km of fiber using standard intradyne detection," Opt. Express **26**(4), 4522–4530 (2018).
156. S. Chandrasekhar, B. Li, J. Cho, X. Chen, E. C. Burrows, G. Raybon, and P. J. Winzer, "High-spectral-efficiency transmission of PDM 256-QAM with parallel probabilistic shaping at record rate-reach trade-offs," Proc. European Conf. on Optical Comm. (ECOC), Th.3.C.1 (2016).
157. P. J. Winzer, "Spatial multiplexing in fiber optics: The 10x scaling of metro/core capacities," Bell Labs Tech. J. **19**, 22–30 (2014).
158. P. P. Mitra and J. B. Stark, "Nonlinear limits to the information capacity of optical fibre communications," Nature **411**(6841), 1027–1030 (2001).
159. R.-J. Essiambre, G. J. Foschini, G. Kramer, and P. J. Winzer, "Capacity limits of information transport in fiber-optic networks," Phys. Rev. Lett. **101**(16), 163901 (2008).
160. R. Dar, M. Shtaf, and M. Feder, "New bounds on the capacity of the nonlinear fiber-optic channel," Opt. Lett. **39**(2), 398–401 (2014).
161. T. Morioka, "New Generation Optical Infrastructure Technologies: EXAT initiative towards 2020 and beyond," Proc. Opto-Electronics and Comm. Conf., FT4 (2009).
162. Audouin, P. Bonno, C. Drion, E. Grard, O. Rofidal, C. M. Martin, E. Ringoot, I. van de Voorde, G. J. Eilenberger, F. Raeymaekers, J. Chauvin, D. H. B. Hoa, M. Morin, Xing-Zhi Qiu, J. Vandewege, T. Daeleman, J. Angelopoulos, and C. Taille, "Design of a cross-border optical core and access networking field trial: First outcomes of the ACTS-PELICAN project," J. Lightwave Technol. **18**(12), 1939–1954 (2000).
163. B. Collings, "New devices enabling software-defined optical networks," IEEE Commun. Mag. **51**(3), 66–71 (2013).
164. D. M. Marom, D. T. Neilson, J. Leuthold, M. A. Gibbons, and C. R. Giles, "64 channel 4x4 wavelength-selective cross-connect for 40 Gb/s channel rates with 10 Tb/s throughput capacity," Proc. European Conf. on Optical Comm. (ECOC), We4.P.130 (2003).
165. D. T. Neilson, C. R. Doerr, D. M. Marom, R. Ryf, and M. P. Earnshaw, "Wavelength selective switching for optical bandwidth management," Bell Labs Tech. J. **11**(2), 105–128 (2006).
166. J. K. Rhee, F. Garcia, A. Ellis, B. Hallock, T. Kennedy, T. Lackey, R. G. Lindquist, J. P. Kondis, B. A. Scott, J. M. Harris, D. Wolf, and M. Dugan, "Variable passband optical add-drop multiplexer using wavelength selective switch," Proc. European Conf. on Optical Comm. (ECOC), 550–551 (2001).
167. D. M. Marom, D. T. Neilson, D. S. Greywall, N. R. Basavanhally, P. R. Kolodner, Y. L. Low, F. Pardo, C. A. Bolle, S. Chandrasekhar, L. Buhl, C. R. Giles, S.-H. Oh, C. S. Pai, K. Werder, H. T. Soh, G. R. Bogart, E. Ferry, F. P. Klemens, K. Teffeu, J. F. Miner, S. Rogers, J. E. Bower, R. C. Keller, and W. Mansfield, "Wavelength selective 1x4 switch for 128 WDM channels at 50 GHz spacing," Proc. Optical Fiber Comm. Conf. (OFC), **857** (2002).
168. T. Ducellier, J. Bismuth, S. F. Roux, A. Gillet, C. Merchant, M. Miller, M. Mala, Y. Ma, L. Tay, J. Sibille, M. Alavanja, A. Deren, M. Cugalj, D. Ivancevic, V. Dhuler, E. Hill, A. Cowen, B. Shen, and R. Wood, "The MWS 1X4: A high performance wavelength switching building block," Proc. European Conf. on Optical Comm. (ECOC), 2.3.1 (2002).
169. A. N. Patel, P. N. Ji, J. P. Jue, and T. Wang, "Defragmentation of Transparent Flexible Optical WDM (FWDM) Networks," Proc. Optical Fiber Comm. Conf. (OFC), OTu8 (2011).
170. A. Morea, J. Renaudier, T. Zami, A. Ghazisaeidi, and O. Bertran-Pardo, "Throughput Comparison Between 50-GHz and 375-GHz Grid Transparent Networks [Invited]," J. Opt. Commun. Netw. **7**(2), A293–A300 (2015).
171. G. Baxter, S. Frisken, D. Abakoumov, Hao Zhou, I. Clarke, A. Bartos, and S. Poole, "Highly programmable wavelength selective switch based on liquid crystal on silicon switching elements," Proc. Optical Fiber Comm. Conf. (OFC), OTuF2 (2006).

172. L. Pascar, R. Karubi, B. Frenkel, and D. M. Marom, "Port-reconfigurable, wavelength-selective switch array for colorless/ directionless/contentionless optical add/drop multiplexing," Proc. International Conf. on Photonics in Switching (PS), PDP.2 (2015).
173. Y. Ikuma, K. Suzuki, N. Nemoto, E. Hashimoto, O. Moriwaki, and T. Takahashi, "8 × 24 wavelength selective switch for low-loss transponder aggregator," Proc. Optical Fiber Comm. Conf. (OFC), Th5A.4 (2015).
174. J. Hecht, "All-optical converters promise improved networks," Laser Focus World **4**(4), 159–164 (2001).
175. E. Ciaramella, "Wavelength conversion and all-optical regeneration: Achievements and open issues," J. Lightwave Technol. **30**(4), 572–582 (2012).
176. J. G. Webster, I. Lazarou, and H. Avramopoulos, "All-optical wavelength conversion," in *Wiley Encyclopedia of Electrical and Electronics Engineering*, J. G. Webster (ed.) (2016).
177. L. E. Nelson, S. L. Woodward, S. Foo, M. Moyer, D. J. S. Beckett, M. O'Sullivan, and P. D. Magill, "Detection of a single 40 Gb/s polarization-multiplexed QPSK channel with a real-time intradyne receiver in the presence of multiple coincident WDM channels," J. Lightwave Technol. **28**(20), 2933–2943 (2010).
178. R. Alvizu, G. Maier, N. Kukreja, A. Pattavina, R. Morro, A. Capello, and C. Cavazzoni, "Comprehensive survey on T-SDN: software-defined networking for transport networks," IEEE Comm. Surv. and Tutor. **19**(4), 2232–2283 (2017).
179. E. B. Desurvire, "Capacity demand and technology challenges for lightwave systems in the next two decades," J. Lightwave Technol. **24**(12), 4697–4710 (2006).
180. R. W. Tkach, "Scaling optical communications for the next decade and beyond," Bell Labs Tech. J. **14**(4), 3–9 (2010).
181. C. I. S. C. O. Visual Networking Index, Forecast and Methodology, 2013–2018 (2016). [Online]. Available: www.cisco.com/c/en/us/solutions/collateral/service-provider/ip-ngn-ip-next-generation-network/white_paper_c11-481360.pdf
182. G. Wellbrock and T. J. Xia, "How will optical transport deal with future network traffic growth?" Proc. European Conf. on Optical Comm. (ECOC), Th.1.2.1 (2014).
183. A. Gerber and R. Doverspike, "Traffic types and growth in backbone networks," Proc. Optical Fiber Comm. Conf. (OFC), OTuR1 (2011).
184. A. Singh, P. Germano, A. Kanagala, J. Provost, J. Simmons, E. Tanda, J. Wanderer, U. Hölzle, S. Stuart, A. Vahdat, J. Ong, A. Agarwal, G. Anderson, A. Armistead, R. Bannon, S. Boving, G. Desai, and B. Felderman, "Jupiter rising: A decade of Clos topologies and centralized control in Google's datacenter network," Proc. Sigcomm, 183–197 (2015).
185. A. W. S. Global Infrastructure, on-line: <https://aws.amazon.com/about-aws/global-infrastructure/>
186. Data center locations, on-line: <https://www.google.com/about/datacenters/inside/locations/index.html>
187. Azure regions, on-line: <https://azure.microsoft.com/en-us/regions/>
188. D. Mortenson, "2017 Year in review: Data centers," on-line: <https://code.facebook.com/posts/392743124493876/>
189. Gary Cook, "Greenpeace's investigation of Amazon's rapidly growing cloud," on-line; <http://www.greenpeace.org/usa/greenpeace-investigation-aws/>
190. R. Miller, "Scaling up: Google building four-story data centers," on-line: <https://datacenterfrontier.com/google-building-four-story-data-centers/>
191. James Hamilton, "Tuesday Night Live with James Hamilton", AWS re:Invent 2016, Las Vegas, NV, on-line: https://mvdirona.com/jrh/talksandpapers/ReInvent2016_James%20Hamilton.pdf
192. R. J. Stone, "Use of embedded optics to decrease power consumption in IO dense systems," Proc. Optical Fiber Comm. Conf. (OFC), Th3G.5 (2017).
193. Z. Tong, C. Lundström, P. A. Andrekson, C. J. McKinstrie, M. Karlsson, D. J. Blessing, E. Tipsuwannakul, B. J. Puttnam, H. Toda, and L. Grüner-Nielsen, "Towards ultrasensitive optical links enabled by low-noise phase-sensitive amplifiers," Nat. Photonics **5**(7), 430–436 (2011).
194. Y. Yamamoto, Y. Kawaguchi, and M. Hirano, "Low-loss and low nonlinearity pure-silica-core fiber for C- and L-band broadband transmission," J. Lightwave Technol. **34**(2), 321–326 (2016).
195. E. Ip and J. M. Kahn, "Compensation of dispersion and nonlinear impairments using digital backpropagation," J. Lightwave Technol. **26**(20), 3416–3425 (2008).
196. R. Dar and P. J. Winzer, "Nonlinear interference mitigation: Methods and potential gain," J. Lightwave Technol. **35**, 903–930 (2017).
197. P. J. Winzer, "Energy-efficient optical transport capacity scaling through spatial multiplexing," IEEE Photonics Technol. Lett. **23**(13), 851–853 (2011).
198. P. J. Winzer, "Making spatial multiplexing a reality," Nat. Photonics **8**(5), 345–348 (2014).
199. M. N. Petrovich, F. Poletti, J. P. Wooler, A. M. Heidt, N. K. Baddela, Z. Li, D. R. Gray, R. Slavik, F. Parmigiani, N. V. Wheeler, J. R. Hayes, E. Numkam, L. Grüner-Nielsen, B. Pálsdóttir, R. Phelan, B. Kelly, J. O'Carroll, M. Becker, N. MacSuihbne, J. Zhao, F. C. G. Gunning, A. D. Ellis, P. Petropoulos, S. U. Alam, and D. J. Richardson, "Demonstration of amplified data transmission at 2 μm in a low-loss wide bandwidth hollow core photonic bandgap fiber," Opt. Express **21**(23), 28559–28569 (2013).
200. D. J. Richardson, "Hollow core fibres and their applications," Proc. Optical Fiber Comm. Conf. (OFC), Tu3H.1 (2017).
201. Parallel single mode four lane per direction 100 Gbit/s optical interface, On-line: <http://psm4.org/>

202. B. Zhu, T. F. Taunay, M. F. Yan, J. M. Fini, M. Fishteyn, E. M. Monberg, and F. V. Dimarcello, "Seven-core multicore fiber transmissions for passive optical network," *Opt. Express* **18**(11), 11117–11122 (2010).
203. K. Saitoh and S. Matsuo, "Multicore fiber technology," *J. Lightwave Technol.* **34**(1), 55–66 (2016).
204. R. Ryf, S. Randel, A. H. Gnauck, C. Bolle, A. Sierra, S. Mumtaz, M. Esmaelpour, E. C. Burrows, R.-J. Essiambre, P. J. Winzer, D. W. Peckham, A. McCurdy, and R. Lingle, "Mode-division multiplexing over 96 km of few-mode fiber using coherent 6x6 MIMO processing," *J. Lightwave Technol.* **30**(4), 521–531 (2012).
205. R. Ryf and N. Fontaine, "Space-division multiplexing and MIMO processing", in *Enabling Technologies for High Spectral-efficiency Coherent Optical Communication Networks*, Wiley, ch. **16**, 547–607 (2016).
206. M. Bigot, D. Molin, K. de Jongh, D. Van Ras, F. Achten, and P. Sillard, "Next-generation multimode fibers for space division multiplexing," *Proc. Advanced Photonics Congress (IPR), NeM3B.4* (2017).
207. M. D. Feuer, L. E. Nelson, X. Zhou, S. L. Woodward, R. Isaac, Benyuan Zhu, T. F. Taunay, M. Fishteyn, J. M. Fini, and M. F. Yan, "Joint digital signal processing receivers for spatial superchannels," *IEEE Photonics Technol. Lett.* **24**(21), 1957–1960 (2012).
208. P. J. Winzer and G. J. Foschini, "MIMO capacities and outage probabilities in spatially multiplexed optical transport systems," *Opt. Express* **19**(17), 16680–16696 (2011).
209. S. Randel, R. Ryf, A. Sierra, P. J. Winzer, A. H. Gnauck, C. A. Bolle, R.-J. Essiambre, D. W. Peckham, A. McCurdy, and R. Lingle, Jr., "6x56-Gb/s mode-division multiplexed transmission over 33-km few-mode fiber enabled by 6x6 MIMO equalization," *Opt. Express* **19**(17), 16697–16707 (2011).
210. X. Chen, P. Dong, S. Chandrasekhar, K. Kim, B. Li, H. Chen, A. Adamiecki, A. Gnauck, and P. J. Winzer, "Characterization and digital pre-compensation of electro-optic crosstalk in silicon photonics I/Q modulators," *Proc. European Conf. on Optical Comm. (ECOC), Tu3.A.5* (2016).
211. G. Li and X. Liu, "Focus issue: space multiplexed optical transmission," *Opt. Express* **19**(17), 16574–16575 (2011).
212. K.-P. Ho and J. M. Kahn, "Statistics of group delays in multimode fiber with strong mode coupling," *J. Lightwave Technol.* **29**(21), 3119–3128 (2011).
213. C. Antonelli, A. Mecozzi, M. Shtaif, and P. J. Winzer, "Stokes-space analysis of modal dispersion in fibers with multiple mode transmission," *Opt. Express* **20**(11), 11718–11733 (2012).
214. L. Palmieri and A. Galtarossa, "Intramodal dispersion properties of step-index few-mode spun fibers," *J. Lightwave Technol.* **34**(2), 303–313 (2016).
215. S. Mumtaz, R.-J. Essiambre, and G. P. Agrawal, "Reduction of nonlinear penalties due to linear coupling in multicore optical fibers," *IEEE Photonics Technol. Lett.* **24**(18), 1574–1576 (2012).
216. A. Mecozzi, C. Antonelli, and M. Shtaif, "Coupled Manakov equations in multimode fibers with strongly coupled groups of modes," *Opt. Express* **20**(21), 23436–23441 (2012).
217. W. Klaus, B. J. Puttnam, R. S. Luis, J. Sakaguchi, J.-M. D. Mendinueta, Y. Awaji, and N. Wada, "Advanced space division multiplexing technologies for optical networks," *J. Opt. Commun. Netw.* **9**(4), C1 (2017).
218. C. Antonelli, M. Shtaif, and A. Mecozzi, "Modeling of nonlinear propagation in space-division multiplexed fiber-optic transmission," *J. Lightwave Technol.* **34**(1), 36–54 (2016).
219. R. Ryf, J. C. Alvarado, B. Huang, J. Antonio-Lopez, S. H. Chang, N. K. Fontaine, H. Chen, R.-J. Essiambre, E. Burrows, R. Amezcua-Correa, T. Hayashi, Y. Tamura, T. Hasegawa, and T. Taru, "Long-Distance Transmission over Coupled-Core Multicore Fiber," *Proc. European Conf. on Optical Comm. (ECOC), Th.3.C.3* (2016).
220. D. J. Richardson, J. M. Fini, and L. E. Nelson, "Space-division multiplexing in optical fibres," *Nat. Photonics* **7**(5), 354–362 (2013).
221. P. J. Winzer, M. Montoliu, R. Ryf, and S. Randel, "Complexity analysis of adaptive frequency-domain equalization for MIMO-SDM transmission," *Proc. European Conf. on Optical Comm. (ECOC), Th.2.C.4* (2013).
222. C. Xia, R. Amezcua-Correa, N. Bai, E. Antonio-Lopez, D. May-Arrijo, A. Schulzgen, M. Richardson, J. Liñares, C. Montero, E. Mateo, X. Zhou, and G. Li, "Low-crosstalk few-mode multi-core fiber for high-mode-density space-division multiplexing," *Proc. European Conf. on Optical Comm., Mo.1.F.5* (2012).
223. S. Randel, A. Sierra, R. Ryf, and P. J. Winzer, "Crosstalk tolerance of spatially multiplexed MIMO systems," *Proc. European Conf. on Optical Comm. (ECOC), P4.08* (2012).
224. S. Randel, S. Corteselli, D. Badini, D. Pileri, S. Caelles, S. Chandrasekhar, J. Gripp, H. Chen, N. K. Fontaine, R. Ryf, and P. J. Winzer, "First real-time coherent MIMO-DSP for six coupled mode transmission," *Proc. IEEE Photonics Conf. (IPC'15)* (2015).
225. D. Marom, P. D. Colbourne, A. D'Errico, N. K. Fontaine, Y. Ikuma, R. Proietti, L. Zong, J. M. Rivas-Moscoso, and I. Tomkos, "Survey of photonic switching architectures and technologies in support of spatially and spectrally flexible optical networking," *J. Opt. Commun. Netw.* **9**(1), 1–26 (2017).
226. E. Pincemin, M. Song, B. Baeuerle, A. Josten, D. Hillerkuss, J. Leuthold, R. Rudnick, D. M. Marom, S. Ben-Ezra, J. F. Ferran, D. Klionidis, and I. Tomkos, "Cascaded all-optical sub-channel add/drop multiplexing from a 1-Tb/s MB-OFDM or N-WDM super-channel with ultra-low guard-bands," *Proc. International Conf. on Transparent Optical Networks (ICTON), Tu.D1.3* (2017).
227. P. J. Winzer, "An opto-electronic interferometer and its use in subcarrier add/drop multiplexing," *J. Lightwave Technol.* **31**(11), 1775–1782 (2013).
228. S. J. Fabbri, S. Sygletos, A. Perentos, E. Pincemin, K. Sugden, and A. D. Ellis, "Experimental implementation of an all-optical interferometric drop, add, and extract multiplexer for superchannels," *J. Lightwave Technol.* **33**(7), 1351–1357 (2015).
229. C. Clos, "A study of non-blocking switching networks," *Bell Syst. Tech. J.* **32**(2), 406–424 (1953).

230. J. Pestic, E. Le Rouzic, N. Brochier, and L. Dupon, "Proactive restoration of optical links based on the classification of events," Proc. International Conf. on Optical Network Design and Modeling (ONDM), (2011).
231. F. N. Hauske, M. Kuschnerov, B. Spinnler, and B. Lankl, "Optical performance monitoring in digital coherent receivers," J. Lightwave Technol. **27**(16), 3623–3631 (2009).
232. J. Simsarian and P. J. Winzer, "Shake before break: per-span fiber sensing with in-line polarization monitoring," Proc. Optical Fiber Comm. Conf. (OFC), M2E.6 (2017).
233. Y.-J. Kim, J. E. Simsarian, N. Choi, N. N. Mohanasamy, and M. Thottan, "Cross-layer aware packet-optical link management in software-defined network operating system," Proc. Optical Fiber Comm. Conf. (OFC), Th2A.31 (2018).
234. N. Sambo, I. Tomkos, A. Shaikh, S. Bigo, M. Suzuki, and H.-J. Schmidtke, "Guest editorial: Optical networks supporting interoperability and white boxes," J. Lightwave Technol. **36**(15), 3058–3061 (2018).

1. Introduction: taking a 20-year horizon

Over the previous 20 years, since Optics Express was created in 1997, Internet Protocol (IP) traffic in North America has grown by a factor of 10,000; the capacity of IP router blades, which make sure that these packets reach their destination host, has grown by a factor of 1,000; the capacity of wavelength-division multiplexed (WDM) fiber-optic communication systems transporting the IP traffic across the globe has grown by the same factor of 1,000; and per-wavelength transponder interface rates have grown by a factor of between 10 and 100. These numbers reflect enormous growth in the demand for data traffic and its supply through information and communications technologies, on well understood long-term exponential scaling trends, as we shall discuss in this paper. Fitting for the occasion of the 20-year anniversary of Optics Express, we note that the number of pages published in this on-line only journal has grown by an equally impressive factor of 37, with ~900 pages published per year in its beginnings, and ~34,000 pages published annually today.

Adopting a 20-year view, we will review the evolution of optical fiber communication systems in this paper, and through a look at the previous 20 years attempt to extrapolate fiber-optic technology needs and potential solution paths for the coming 20 years. Well aware that 20-year technology extrapolations are inherently noisy, we hope that taking a significantly longer-term view than most texts in this field will provide the reader with a broader perspective and will encourage the much needed out-of-the-box thinking to solve the very significant technology scaling problems ahead of us. Similar to our companion paper [1], which "only" looks a decade into the future, and which may serve as both a complementary and an introductory source to the material discussed here, we base our 20-year extrapolations on well-established long-term historic traffic and technology growth trends, which we assume will continue going forward. Obviously, all exponential growth will eventually saturate, but many examples related to information and communications technologies have shown that exponential growth (as well as exponential energy reduction) can be maintained for decades [1], and even for a century [2]; examples of sustained exponential growth over multiple centuries can also be found, e.g., in economics [3]. While human perception is quick to accept long-term exponential scaling as a historic fact, there is often significant hesitation associated with accepting continued exponential scaling as a likely evolution path into the future. These conceptual difficulties likely arise because people tend to focus on the scaling of a given *technology*, while exponential scaling should rather be associated with *functional* scaling, whereby the scaling of a certain technology used to implement a certain function may saturate and be replaced by a new technology to continue the scaling of the considered function. Examples are the functional scaling of microprocessors, which first involved clock speed scaling and then scaling through multi-core architectures; the functional scaling of storage, which involved various generations of different magnetic technologies before transitioning to semiconductor technologies; and the functional scaling of long-haul transport, which first resorted to per-span regenerated time-division multiplexing using a succession of optical wavelength ranges over multi-mode, then single-mode waveguides, then incorporated optically amplified WDM, and is now shifting to parallel spatial paths. Individual applications driving traffic growth also may exhibit saturation, only to be replaced by new applications

that continue the scaling of traffic. Should established long-term scaling trends experience a pronounced change that prevents their underlying *functional* scaling at the exponential pace that our heavily information and communications centric society has gotten used to, the reduced growth is likely to have significant social and/or economic repercussions. These, we hope, will be avoided through the inventive spirit of the scientific and engineering communities addressed in this paper. In fact, while we expect there to be challenges in scaling networks over the next 20 years, we do not foresee any fundamental roadblocks, as we shall discuss in this paper.

When Optics Express was created in July of 1997, the era of WDM systems had only just begun in terms of significant commercial deployments. State-of-the-art systems featured 16 channels at 2.5 Gb/s, at a spectral efficiency (SE) of 0.0125 b/s/Hz. A key performance parameter of WDM systems, the SE is defined as the ratio of aggregate system capacity C_{Sys} to system bandwidth B_{Sys} , which for a homogeneous WDM system is equivalent to the ratio of the per-channel bit rate R_{Ch} to the frequency spacing F_{Ch} between WDM channels,

$$SE = \frac{C_{\text{Sys}}}{B_{\text{Sys}}} = \frac{R_{\text{Ch}}}{F_{\text{Ch}}}. \quad (1)$$

As an alternative to AT&T's and Ciena's WDM approach, Nortel introduced a single-wavelength system at 10 Gb/s as early as 1997 [4].

Regarding subsea transmission, the TAT12/13 transatlantic cables had just been commissioned. Each cable had 2 fiber pairs and transported a single wavelength at 5 Gb/s per fiber over 5,913 km using erbium-doped fiber amplifiers (EDFAs) as repeaters, for an aggregate cable capacity of 20 Gb/s; the cable's aggregate capacity x distance product was 118 Tb/s.km [5].

In record research experiments, an aggregate WDM capacity of 1 Tb/s (55 x 20 Gb/s and 25 x 20 Gb/s x 2 polarizations) had just been established in 1996, more than 10x the SE used in commercial products, cf. Tab.1. (A detailed tabulated account of the progress of multi-terabit research results through 2007, in terrestrial as well as subsea experiments, is given in [6]. Details of commercial lightwave systems up to 2004 are tabulated in [7].)

Today, 20 years later, commercial long-haul C + L-band lightwave systems carry up to 192 channels at up to 250 Gb/s on a 50-GHz grid, sporting an aggregate long-haul capacity of ~48 Tb/s at an SE of 5 b/s/Hz, and for short-reach applications up to 400 Gb/s at 8 b/s/Hz, for a total capacity of up to 76 Tb/s [8,9]. The highest-capacity submarine cable, the Pacific Light Cable Network (PLCN) carries an aggregate bidirectional 144 Tb/s across the Pacific, with an overall capacity x distance product across 6 fiber pairs of 3,686 Pb/s.km (counting both transmission directions separately).

Research records today (cf. Tab. 1) achieve a net per-carrier interface rate in excess of 1 Tb/s [10–12], and aggregate WDM capacities in single-mode fiber of up to 115 Tb/s have been reported [13–15]. Using multi-core fiber, system capacities of up to 10 Pb/s have been achieved [16]. Short-reach systems have demonstrated a record-SE of 17.3 b/s/Hz on a single-mode fiber [17]; capacity x distance products of up to 881 Pb/s.km on a single-core, single-mode fiber [18] and up to 1,508 Pb/s.km on a multi-core fiber [19] have been demonstrated. Looking 20 years into the future, and assuming that computing technologies continue their (functional) scaling at ~40%, a conservative assumption based on historic scaling trends [1], we extrapolate capacity needs to 2037 as shown in Fig. 1 and summarized in Tab 1. We will discuss the detailed technological rationale behind these extrapolations in the remainder of this paper, and in particular the role that space-division multiplexing (SDM) is bound to play in addition to WDM.




Fig. 1. Evolution of commercial optical transmission systems over the past 30 years and extrapolations for the coming 20 years (after [1]).

Table 1. The previous 20 years and an extrapolation into the next 20 years of fiber optic systems. (*Record numbers in independent experiments; †Submarine cable capacities, counting both directions)

	1997		2017		2037	
	Products	Research	Products	Research*	Products	
Interface rate R_{Ch}	2.5 – 10 Gb/s	20 Gb/s 40 Gb/s pol mux	100 – 400 Gb/s	1.3 Tb/s [12]	2 – 6 Tb/s	
Wavelengths	16	25 – 55	96 – 192	several hundred	16 – 66	
Single mode	C_{Sys}	40 Gb/s	1 Tb/s	10 – 76 Tb/s	115 Tb/s [15]	32 – 400 Tb/s
	$C_{Sys} \times L$	14.4 Tb/s.km	150 Tb/s.km	up to 240 Pb/s.km	881 Pb/s.km [18]	500 Pb/s.km
	SE	0.0125 b/s/Hz	0.27 – 0.4 b/s/Hz	2 – 8 b/s/Hz (flex.)	17.3 b/s/Hz [17]	7 – 20
Cable [†] /SDM	C_{Sys}	20 Gb/s	-	288 Tb/s	10 Pb/s [16]	5 – 100 Pb/s
	$C_{Sys} \times L$	118 Tb/s.km	-	3,686 Pb/s.km	1,508 Pb/s.km [19]	30,000 Pb/s.km
Router blade	2 x 2.5 Gb/s	-	6 x 400 Gb/s	-	128 x 6 Tb/s	
Router capacity	24 x 2.5 Gb/s	-	120 x 400 Gb/s	-	5,120 x 6 Tb/s	

2. The evolution of fiber-optic transport networks

Signals were first sent through an optical fiber field test system in 1977 [20]. Within months, live telephone traffic was transmitted through multimode fibers (MMF) by General Telephone and Electronics (GTE, at 6 Mb/s [21]), American Telephone & Telegraph (AT&T, at 45 Mb/s [21,22]), and the United Kingdom Post Office (at 8.4 Mb/s [21,23]). Early field trials and production traffic carrying fiber-optic systems in other European countries are summarized in [24]. Fiber-optic systems developments over the past 40 years can be neatly compartmentalized into four major eras of transmission [25,26]:

- (1) the era of regeneration (1977 ~1995),
- (2) the era of amplified dispersion-managed systems (1995 ~2008),
- (3) the era of amplified coherent systems (2008 ~present), and
- (4) the era of space division multiplexing (actively researched since ~2008)

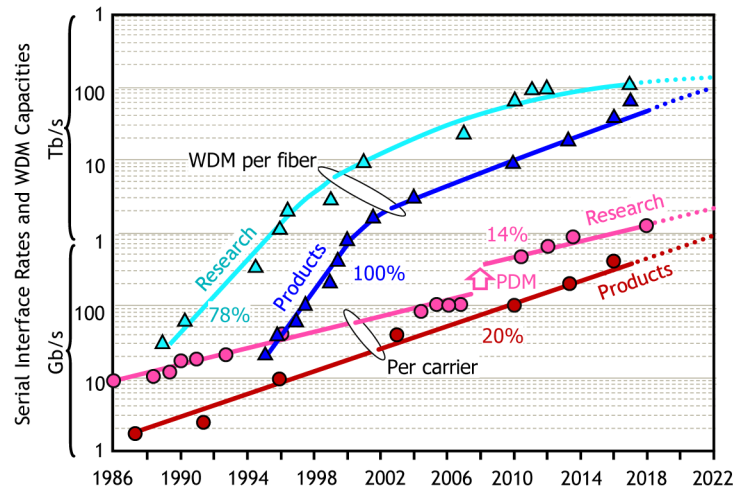


Fig. 2. Comparison of products and research records in terms of per-carrier interface rates and WDM capacities. The discontinuity in single carrier interface records around 2007 is due to the introduction of coherent detection using polarization division multiplexing (PDM). (Figure after [27].)

2.1 The era of regeneration

The capacity of the early span-by-span regenerated fiber-optic transmission systems depended on the transceivers' *interface rates*, i.e., the net bit rate that transceivers were able to support. Progress in interface rates was painfully slow, not only in commercial systems but also in research experiments. Figure 2 shows the roughly 20%/year (or $10 \log_{10}(1.2) = 0.8$ dB/year) improvement in commercial interface rates, which quite remarkably has persisted over more than three decades. Research results have been scaling at a slower pace, around 14% per year. In 1989 the first commercial opto-electronically *regenerated* 2-wavelength WDM system (each wavelength at 1.7 Gb/s) was introduced, but further progress was limited to increasing the data rate from 1.7 Gb/s to 2.5 Gb/s.

The first optical *undersea* transmission systems deployed across the Atlantic Ocean (TAT8) and Pacific Ocean (TPC3) in the late 1980s were regenerated systems operating at 1.3 μm and carrying 280 Mb/s on each of its three fiber pairs. TAT9 and TPC4 were the first transoceanic systems operating at 1.55 μm with double the capacity of the previous systems, but they were also still opto-electronically regenerated.

2.2 The era of amplified dispersion-managed systems

Cost-efficient WDM systems clearly would have been impossible without the invention of practical optical amplifiers, and in particular of the EDFA as the game-changing enabling technology that started this era of optical transmission [28,29]. Note, though, that the invention of the EDFA by itself could not have enabled the explosive exponential capacity growth of both research and commercial systems depicted in Figs. 1 and 2. Indeed, it was the combination of a handful of key inventions in the early 1990s, related to managing fiber nonlinearities across the bandwidth provided by EDFAs, that enabled the phenomenal growth of 100%/year (3 dB/year) of commercial WDM capacities from the mid-1990s to the early 2000s, again with research scaling slower than commercial systems, at 78% per year.

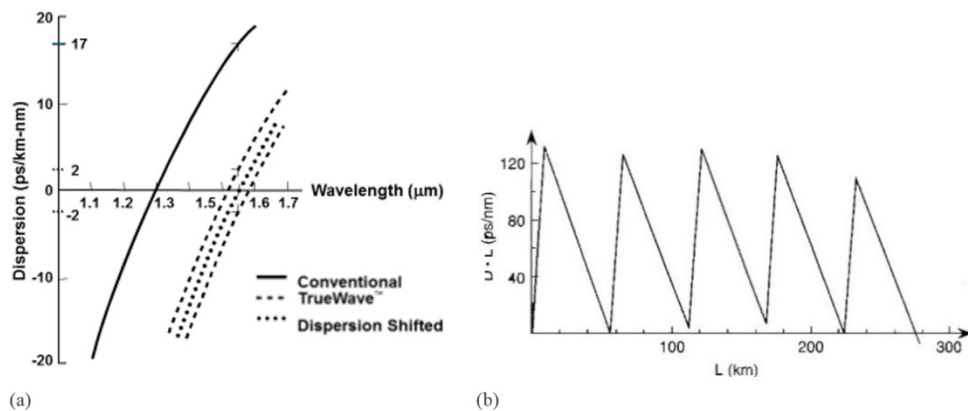


Fig. 3. (a) Dispersion versus wavelength of SSMF, DSF, and two flavors of TrueWave fiber [6,32]; (b) First implementation of a dispersion map [33].

2.2.1 Nonlinearity management in amplified WDM systems

In the early 1990s there were only two commercial transmission fibers available: *Standard single-mode fiber* (SSMF) had a chromatic dispersion (CD) of 17 ps/nm/km at a wavelength of 1.55 μm , allowing for 2.5-Gb/s regeneration distances of about 800-1000 km (using chirp-free sources) and about 100 km (directly modulated lasers), both larger than the attenuation-limited reach at the time. However, since the dispersion-limited reach shrinks quadratically with symbol rate, even chirp-free sources could only support ~ 60 -km SSMF transmission at the anticipated 10-Gb/s wavelengths that were needed to support the required capacity growth rates. (In the early 1990s, there were no practical CD compensation techniques available, especially no broadband ones.) This led to large-scale deployments of *dispersion-shifted fiber* (DSF, nominal zero-dispersion wavelength at 1.55 μm) around the world, led by widescale deployments in Japan. However, it was not widely appreciated at the time that the low dispersion inherent to DSFs would be fatal in WDM scenarios with closely spaced wavelengths, which makes WDM deployments in networks that still include DSF a complicated task to this day. Four-photon mixing (FPM), one of a handful of nonlinearities present in silica glass [30], thrives in low-CD environments. As a consequence, many closely spaced wavelengths produce new wavelengths, called FPM products. In fact, hundreds of such FPM products could be produced by less than 10 signal wavelengths. Not only do these FPM products rob the original signals of their power, but more importantly, in an evenly spaced wavelength scenario, these newly generated FPM products fall at precisely the wavelengths of the signals, thus coherently mixing with the signals. This means that FPM products that contain just 1% of the power of the signals would produce 1-dB power fluctuations in the signals [31]. Researchers at AT&T Bell Laboratories realized that simple modifications to the fiber manufacturing process could readily produce fibers with low, but non-zero CD [32]. The resulting fiber was known as TrueWave fiber, later generically standardized as *nonzero dispersion fiber*, *NZDF*, ITU-T Standard G.655. Interestingly, TrueWave fiber could readily be produced in “two flavors”, fibers that had either slightly positive or slightly negative CD at 1.55 μm . The dispersion curves of SSMF, DSF, and the two flavors of TrueWave fiber are shown in Fig. 3(a). This inspired what became known as *dispersion management* (DM) [33], a powerful technique that concatenates transmission fibers with opposite signs of dispersion so that the overall CD of a link is nearly zero but the “local” CD everywhere along the link is high enough to suppress FPM. This way, DM mitigates the effects of both FPM and CD. Since its invention in 1993, DM was universally used in all high-speed dense WDM systems, both research and commercial, until the introduction of digital coherent systems in 2008 (cf. Sec. 2.3). The first demonstration of DM

used the simple dispersion map of Fig. 3(b), a mixture of SSMF and negative-dispersion TW fiber [33].

Initially fairly simple dispersion maps quickly became insufficient for more aggressive, higher-bit-rate systems at 20 Gb/s and beyond, which led to the discovery of dispersion pre-compensation [34]. As the throughput of systems increased, first in research and then commercially, ever more sophisticated dispersion maps were invented to support more wavelengths at higher line rates [35]. The widespread introduction of dispersion management in commercial optically amplified lightwave systems was gated by two developments: the introduction of 10-Gb/s line rates (as the reach of 2.5-Gb/s transmitters was not limited by CD for the terrestrial-system needs of the mid-1990s), and the introduction of practical dispersion-compensating fiber (DCF) in the mid-1990s [36,37], and later slope-matched DCF. DCF has a much higher negative dispersion than NZDF and could therefore be packaged as a spool in conjunction with EDFAs, thus avoiding the need to deploy two different kinds of transmission fiber in the line system. On the other hand, subsea systems continued to use a concatenation of positive and negative dispersion transmission fiber as opposed to localized DCFs until coherent transponders started to ask for dispersion uncompensated link designs (cf. Sec. 2.3).

In terms of the then emerging field of optically switched networking (cf. Sec. 2.2.5), DM added additional complexity, since optimal dispersion maps did not natively provide dispersion compensated signals at *all* add/drop nodes throughout the network, which required, typically on a per-wavelength basis, additional DCFs (in mesh networks even tunable dispersion compensators) at the add/drop ports, with the associated cost, complexity, and additional insertion losses affecting the system's link budget. Networking flexibility and transmission performance started to ask for different trade-offs in terms of system design.

2.2.2 Polarization-mode dispersion

There were other “irritants” that could degrade system performance (but none was as serious as the combination of CD and optical nonlinearities). One such irritant was polarization-mode dispersion (PMD), whereby randomly varying optical birefringence in fibers causes unpredictable relative delays between the two polarizations carrying an optical signal. For polarization-ignorant direct-detection receivers, this manifested itself in pulse broadening, leading to inter-symbol interference (ISI) penalties. As PMD becomes more severe at higher bit rates, this appeared to be a daunting problem for bit rate scaling because of the relatively high values of PMD in fibers produced before 1992. What makes the problem even worse is its slowly varying statistical nature [38–41], which results in low but finite PMD-induced system *outage probabilities*. For example, the “five nines” reliability typical of telecom systems corresponds to an allowed outage probability of 10^{-5} , or 5 minutes of outage per year. System design must ensure that PMD-induced penalties only exceed the margin allocated for PMD with a probability sufficiently below the specified values, and system verification and testing must rely on the validity of the theoretically established [38] and practically refined [41] statistical models.

A key invention (initially referred to as the “GULP method”, now called “spinning”) employed a twisting motion on the fiber during the fiber draw process [42]. This introduced birefringence in fibers on a length scale short compared to the correlation or coupling length of the two polarizations. This invention significantly postponed the need for transceivers with PMD compensation capabilities. With the introduction of 40-Gb/s line rates, equipment vendors developed transceiver packs with optical PMD compensation (PMDC). Unfortunately, as the random effects of PMD are only correlated across a limited bandwidth [40], each wavelength in a WDM system typically requires its own PMDC. As circuit packs including PMDC were more expensive and larger than conventional transceivers, PMDCs were implemented with reluctance in commercial systems. In the cases of cables with large fiber counts, it was easier and less expensive to select fibers with sufficiently low PMD to

carry 40-Gb/s services. For the next increase in bit rates to 100 Gb/s, PMD essentially became a non-issue, because with the introduction of 100 Gb/s came coherent systems and digital signal processing, which effectively compensates both chromatic and polarization-mode dispersion.

2.2.3 Increasing the useful WDM system bandwidth

Increasing WDM capacity by populating a system with more and more signal wavelengths (while at the same time increasing per-wavelength interface rates) was initially a matter of increasing the useable EDFA bandwidth. EDFAs inherently have gain/saturation properties that vary with wavelength. These variations are compounded in multi-span systems. In the early 1990s, practical gain flattening filter technology was unavailable, thus limiting the transmission distance of wide-band WDM systems. A simple algorithm was devised in 1992 [43,44] that prescribed the appropriate launch powers of the various input wavelengths in a WDM system so that the optical signal-to-noise ratio (OSNR) of all the signals at the output of the link would be equal. This algorithm, now referred to as optical pre-emphasis, greatly extended the reach and/or the channel count of WDM systems. Although optical filter technology has progressed to a point where gain equalization can be accomplished in-line, optical pre-emphasis is still used in many subsea systems.

2.2.4 Two powerful laboratory tricks

As part of the beginning of optically amplified WDM systems in the early 1990s, two exceptionally powerful experimental laboratory techniques, were invented that to this day have enabled experiments, design, and development of large-scale WDM systems: *Multi-wavelength modulation* using a single external modulator [33,45] and the *optical recirculating loop* [46]. Both approaches addressed two major stumbling blocks to WDM experiments in the early days, which were economic ones, namely a massive duplication of equipment. In terms of WDM transmitters, one of the last large-scale system experiments accomplished by brute force was a demonstration involving 100 distributed feedback (DFB) lasers, each one individually frequency-shift keyed (to minimize chirp) by direct modulation at 622 Mb/s [47]. As bit rates were pushed to 10 Gb/s and channel spacing reduced to less than 200 GHz, direct modulation of DFB lasers was no longer possible due to laser chirp. Outfitting each wavelength with its own external lithium-niobate modulator plus driver plus pattern generator was a non-starter. The new experimental technique (initially called the “single-modulator trick” at Bell Labs) used a single modulator simultaneously for all wavelengths (see Fig. 4(a)), which had been first wavelength-multiplexed onto a single fiber, followed by a temporally dispersive element to delay-decorrelate the various modulated wavelength channels. At 10 Gb/s, the CD of a short segment of SSMF provided sufficient decorrelation delays. At higher bit rates, this no longer worked without introducing significant dispersive pulse broadening during the decorrelation step, and the transition was made to split wavelength channels into “odd” and “even” groups and use a separate modulator and data source for each of the two groups (this of course assumes that most degradations due to linear crosstalk or nonlinear interactions occur between neighboring channels). Alternatively, in some cases signals can be split into odd and even sets by a steep de-interleaver following joint modulation of all signals by a single modulator [48]. The savings in equipment using these techniques has enabled large-scale WDM experiments to this day. Only most recently an even simpler technique was introduced, valid for high-SE, dispersion-unmanaged coherent systems. This technique uses amplified spontaneous emission (ASE) of EDFAs, whose Gaussian optical field statistics accurately emulates a tightly packed broadband multiplex of spectrally and constellation shaped WDM channels [8,49].

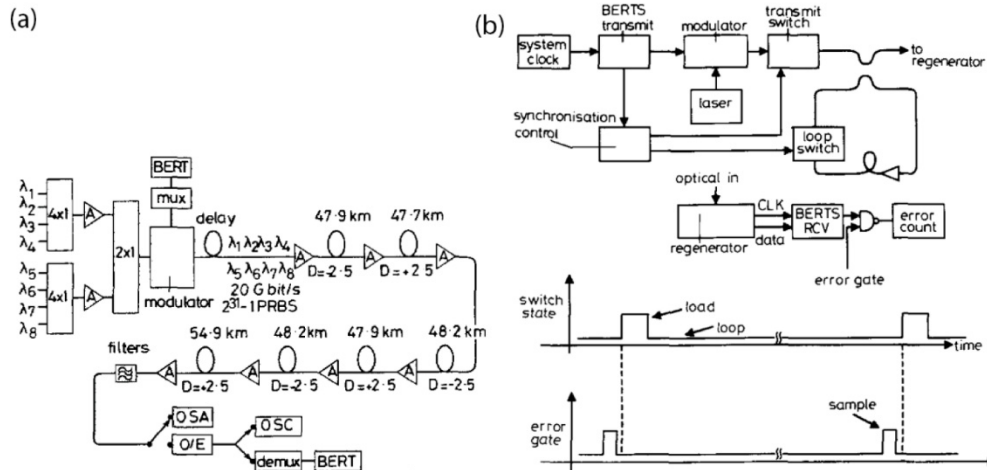


Fig. 4. (a) The “single modulator trick” enabling massive WDM experiments [45], and (b) the re-circulating loop enabling long-haul fiber transmission experiments [46].

As for the recirculating loop, experiments to study undersea and long-haul terrestrial systems were limited by the amount of fiber, EDFAs, and other system components available to experimentalists. The concept of recirculating loop experiments was first devised by Bergano et al. at Bell Labs [46], cf. Fig. 4(b). A loop setup involves precise timings to load, through an optical load switch (transmit switch), a typically several-hundred-km-long fiber loop with a bit stream whose length is no longer than the fiber making up the loop. The load switch is then opened and the loop switch is closed, which allows the bit stream to circulate continuously within the loop. Bit error ratios (BERs) are then measured by tapping off the signal after a number of full (or partial) circulations corresponding to the desired system length that is to be emulated. Up to this day, concerns about the validity of measurements using recirculating loops have been countered by continuously improving loop designs to more accurately reflect long-haul transmission conditions on a straight line [50–53]. Despite its shortcomings, there has not been any other way to practically do long-reach system experiments without resorting to full-length system testbeds whose exceedingly high costs make them only suitable for product verification testbeds. Recirculating loops can therefore be found in most optical transmission systems laboratories worldwide.

2.2.5 Terrestrial and subsea WDM deployments

The first commercial terrestrial WDM systems (8 wavelengths at 2.5 Gb/s on a 200-GHz grid) were deployed in the AT&T network in 1995 [7], without any public announcement, for the simple reason that these systems could not be produced fast enough to satisfy internal AT&T network demand, let alone to be supplied to outside customers [54]. In 1996 Ciena announced its first WDM product, comprising 16 wavelengths carrying 2.5 Gb/s per wavelength on a 100-GHz grid [55]. It is perhaps no coincidence that the first AT&T and Ciena commercial WDM systems were very similar to the WDM research being concurrently conducted under the MONET program, which several of the principal investigators at the two companies were heavily contributing to [56,57].

The inaugural subsea systems employing EDFAs were TAT12/13 and TPC5, deployed in 1995 with only a single 5-Gb/s wavelength per fiber (2 fiber pairs per cable) [5]. Since then, the evolution of *cable capacities* (given in terms of an aggregate bidirectional capacity, typically on 8 fiber pairs) of commercial intercontinental subsea systems has tracked the evolution of *per-fiber capacities* of commercial terrestrial systems (cf. Fig. 1). Optically amplified undersea systems, while qualitatively similar to terrestrial systems, are generally designed quite differently from terrestrial systems [58–60]. The obvious key difference

between the two types of systems is their overall reach. Initially, undersea systems were roughly an order of magnitude longer than terrestrial systems, although that ratio has been reduced with the deployment of transcontinental-scale terrestrial systems. To achieve comparable noise performance, the EDFA spacing in undersea systems had to be considerably smaller (35-50 km) than the terrestrial amplifier spacing (80-100 km). Over time, as various constituent technologies of subsea systems improved, the amplifier spacing in undersea applications has gradually increased to over 70 km. On the other hand, undersea systems were not as much at the mercy of geographical constraints as terrestrial systems and could thus place amplifiers at strictly regular intervals. This allowed the deployment of “identical” amplified spans along the entire route, thereby allowing amplifier properties to be precisely tailored to the particular system. In terrestrial systems EDFAs have to operate over a wide range of span conditions on a mixed fiber infrastructure. In addition, because each new undersea cable is a one-off system, designers can optimize the system even if that involves using nonstandard components, bit rates, channel spacing, etc. Also, because the end terminals of undersea systems represent a much smaller fraction of the system capital and deployment costs than they do in terrestrial systems that are installed on already existing fiber infrastructure, more expensive transponder technologies could be used as long as these improved system performance. As mentioned previously, commercial terrestrial systems had no need for DM at 2.5 Gb/s and only embraced DM with the introduction of 10-Gb/s systems and broadband DCF that could be packaged inside the EDFAs. Such a mindset did not exist in undersea cable design. In fact, DCFs lumped at each EDFA increase the overall length of the system, which increases the accumulated nonlinearity and degrades the system noise figure due to the additional amplification required to compensate for DCF loss. In-line DM avoids these issues. The typical mixture of fibers in undersea cables using direct detection was negative-dispersion NZDF and SSMF. Another key difference between undersea and terrestrial systems is their electrical powering. In terrestrial systems, power for the repeaters is provided locally. In undersea systems, the power for all the submerged equipment is provided by DC voltages applied to the cable ends (e.g., +15 kV at one end, -15 kV at the other end, limited by the properties of the cable dielectric) [61]. This restriction has led to the recent realization that undersea systems might be on the verge of a fundamental paradigm shift [62–64]: Instead of pursuing the highest possible signal powers with concomitant nonlinearity mitigation and the lowest-noise-figure and widest-bandwidth amplifiers, the overall system cost per bit can actually be significantly reduced by using more spatial paths (more fiber pairs), lower signal powers (such that nonlinearity ceases to be of concern), and more energy-efficient optical amplifiers, even if these have higher noise figures and lower bandwidths (hence avoid gain flattening filters that dissipate optical energy).

2.2.6 Solitons

From the prediction [65] and first observation [66] of solitons in optical fibers, research of solitons has been a rich field of interesting physics [67], intriguing because of the fact that solitons retain their pulse shape upon nonlinear transmission due to a careful balance between the fiber’s chromatic dispersion and nonlinearity coefficient and the soliton pulse’s duration and power [30]. More importantly, solitons were initially the only fiber-optic transmission modality that promised high-speed information transmission over transoceanic distances [68]. However, real-world non-idealities such as fiber loss, bulk amplification, amplifier noise, (non-flat) fiber dispersion, and WDM nonlinearities exposed the fragility of solitons and made them less and less competitive with respect to more traditional modulation/transmission formats, despite heroic efforts by a number of scientists and engineers. Solitons interacting with noise or with other solitons can produce variations in the central frequencies of the solitons. Coupled with fiber chromatic dispersion, this can lead to jitter in the arrival time of the pulses at the receiver. Soliton-noise interactions produce random arrival times, known as the Gordon-Haus effect [69]. This effect can be reduced by a narrowband filter at every

amplifier. However, in dense WDM systems, much of the ASE noise would pass through the periodic closely-spaced filters for each wavelength channel. To combat this problem, the concept of “sliding-guiding” filters was invented [70]. The required collection of hundreds of filters for many-channel WDM long-haul systems makes this a costly proposition, however. Soliton-soliton collisions in amplified WDM systems are problematic as well. In an ideal case where the fiber is lossless and the chromatic dispersion is constant across the WDM band, the effect of soliton-soliton collisions can be minimal if not zero because the second half of a “complete” soliton-soliton collision undoes the effects of the first half of the collision. However, if the collision takes place over a distance where the optical power significantly changes (e.g. due to fiber loss), the cancellation is incomplete. This degradation is particularly severe if part of the collision occurs in an optical amplifier where the power abruptly changes by 20 to 30 dB. This leads to unacceptable pulse arrival-time shifts. In 1996 it was shown that dispersion management, used in non-return-to-zero (NRZ) based systems experiments since 1993 [33], was also beneficial to the reduction of timing jitter in soliton systems [71]. However, the resulting “dispersion-managed solitons” were no longer true stationary solitons (whose shape stays constant during nonlinear pulse propagation) but were only periodically stationary (i.e., their shape was restored after every period of the dispersion map). In later “soliton” experiments, the pulses were not even periodically stationary but tended to mimic chirped return-to-zero (CRZ) linear pulses [72]. Indeed, it is the contention of some [73] that CRZ and dispersion managed soliton formats are effectively the same.

The “death knell” for solitons came in the early 2000s with the introduction of commercial 40-Gb/s differential phase-shift keyed (DPSK) systems [74] and even higher-spectral-efficiency differential quadrature phase-shift keyed (DQPSK) field trials at 100 Gb/s [75], discussed in Sec. 2.2.8. Such bit rates and spectral efficiencies were inaccessible to soliton transmission. In addition, solitons, being intrinsically a return-to-zero (RZ) format with pulse widths several times smaller than a bit period, would never have survived the quest for high spectral efficiencies, where digital (Nyquist) pulse shaping takes the time-frequency support of optical signals to their fundamental limits, cf. Sec. 2.3.6. One of the last large-scale soliton (or was it CRZ?) system trials was reported in 2003 [76]. Comprehensive information on solitons can be found in [30,68,77] and references therein.

Even though solitons did not make a lasting commercial impact, the beautiful physics concepts developed around them have survived to this day and have produced other important results to fiber-optic communications. Probably the most impactful example is the time-domain theory of nonlinear interference noise (NLIN), whose “pulse collision” picture draws heavily from soliton theory [78], cf. Sec. 2.3.2. Another example, whose practical importance is not yet evident, are current studies on the nonlinear Fourier transform (NLFT), which decomposes signals into solitonic bases in order to facilitate digital nonlinearity compensation [79], cf. Sec. 3.2.

2.2.7 Harnessing WDM capacity through optical networking

An important aspect coupled with the rise of WDM systems was harnessing the available transmission capacity through *optical networking* and *bandwidth management*. In fact, the capacity provided by WDM systems in the mid-1990s significantly exceeded the capacity of electronic switches and routers, which allowed for the aggregation of multiple electronic signals into a single wavelength for efficient transport, circuit-based through synchronous optical networking (SONET) and the synchronous digital hierarchy (SDH), or packet-based through packet-over-SONET. In 1997, the highest-capacity Cisco IP router had an aggregate capacity of 10 Gb/s with router blades supporting 2.5 Gb/s, while a WDM system could carry up to 40 Gb/s of aggregate traffic [1]. Management of this bandwidth directly in the optical layer became an attractive proposition as eliminating unnecessary high-speed electronics in the path of an optical signal reduced complexity and costs.

The functions of aggregating, disaggregating, and switching data paths in the optical domain, referred to as *optical add-drop multiplexing* (OADM) were initially little more than a set of optical demultiplexing and remultiplexing filters with a fiber patch panel in between, which 20 years ago supported mostly multiplexing of wavelengths on a ring topology. However, the need for reconfiguration of optical paths was already a topic of research interest, with work on various optical switches and even on early wavelength selective switches (WSSs) demonstrated by the late 1990s [80]. An important US government funded project in this context was MONET, demonstrating in a field trial the operation of a transparent, reconfigurable optical network [56,57].

Progressively more complex and more automated *reconfigurable optical add-drop multiplexing* (ROADM) nodes were deployed in the early 2000s, which also started to use irregular frequency grids, such as the asymmetric interleaver architecture of Lucent Technologies' LambdaXtreme that allowed maximizing system capacity by alternating 10-Gb/s and 40-Gb/s direct-detection wavelengths [81]. The widespread commercial introduction of full optical switching flexibility, however, coincided with the introduction of coherent systems, as reviewed in Sec. 2.3.7.

2.2.8 From device physics to communications engineering

The year 1996 saw a major milestone in high-capacity WDM research. The 1-Tb/s barrier was conquered not by just one group, but by three groups independently [82–84]. Surmounting this perhaps psychological barrier convinced the research community that continued exponential progress was indeed achievable. In 1997, commercial WDM systems supported 16 wavelengths at 2.5 Gb/s for an aggregate capacity of 40 Gb/s, hence were a factor of 25 - 50 behind concurrent research records, corresponding to a time lag between research and commercial products of ~6 years (cf. Fig. 2). In general, the time lag between research and commercialization can be seen from the figure to have consistently been between about 4 and 8 years.

By 1999 commercial WDM systems could already carry 80 wavelengths at 2.5 Gb/s, or 40 wavelengths at 10 Gb/s, for an overall WDM system capacity of 400 Gb/s. Interestingly, this aggregate WDM capacity can be carried on a single wavelength by the most advanced transponders commercially available since 2016 [8]. Similar to this 17-year period between turning what used to be an aggregate system capacity into a per-channel interface rate, it took 19 years of research until the 1-Tb/s WDM capacity records of 1996 could be first carried on a single laser wavelength in a research setting [10].

By 2001, research records had broken the 10-Tb/s barrier [85,86]. However, these experiments started to signal a perceptible decrease in the growth rate of single-fiber throughput, a decrease that would continue over the following decade. In fact, the growth rate of WDM capacities would ultimately slow from 3 dB/year to less than 1 dB/year. It took almost 10 years from the first 10-Tb/s record experiments to the equivalent capacity in commercial products, and 6 years to increase the 10-Tb/s research records to merely 25 Tb/s [87]. The reasons for this slow-down were rather fundamental:

- (i) system bandwidths had reached the maximum possible extent of practical single-band optical amplifiers bands (such as the full C-band of EDFAs, supporting a relative bandwidth of a few percent, cf. Sec. 3.3), and
- (ii) transform-limited optical signals could be generated with bandwidths that exceeded the frequency granularity of optical components. (For example, emerging 40-Gb/s (at the time spectrally unshaped) NRZ formats occupied an 80-GHz null-to-null optical bandwidth, which exceeded the then available optical filter bandwidths of between 25 and 50 GHz.)

This made increasing WDM capacity no longer a matter of device engineering (better frequency-stabilizing lasers, opto-electronically generating higher-speed signals, improving the frequency selectivity of optical filters, etc.) to achieve a denser channel spacing, but rather asked for more sophisticated, bandwidth-saving modulation techniques to achieve higher SEs. With the transition of optical systems design from a *physics- and device-dominated discipline* to a *communication engineering discipline*, such aspects as advanced modulation formats, forward error correction (FEC) coding, digital signal processing (DSP), and eventually digital coherent detection became the next frontiers of WDM capacity expansion. The goal from then on was to increase SEs beyond the inherent SE of at most 1 b/s/Hz provided by simple binary formats, most notably NRZ-OOK.

The most successful species in the zoo of “advanced modulation formats” investigated for high-speed optically amplified fiber-optic communications between 2000 and 2008 [88] were differentially encoded phase modulation formats, both binary *differential phase shift keying* (DPSK) and quaternary *differential quadrature phase shift keying* (DQPSK) [89]. They both provided better receiver sensitivities than NRZ-OOK, which was important to mitigate the reach reduction that 40-Gb/s systems experienced relative to 10-Gb/s systems in the early 2000s. In the case of DQPSK, a higher SE could also be attained compared to NRZ-OOK. This was accomplished by differentially encoding the data and implementing optical delay interferometers with balanced photodiodes at the receiver to convert phase differences to intensity modulation [88,89]. The relative simplicity of DPSK and DQPSK receivers were one of the reasons why these modulation formats were the method of choice for “hero” research experiments until around 2007. Both DPSK and DQPSK were commercialized during that time, e.g., as part of Lucent Technologies’ LambdaXtreme system that employed 40-Gb/s DPSK wavelengths in a 64-wavelength system in 2003 [74]. Some of the most important DQPSK research experiments were performed in 2007: The WDM capacity record of 25 Tb/s was established using DQPSK [87], which notably should be the last capacity record achieved in a research setting using direct bit error counting (cf. Sec. 2.3.3); perhaps even more importantly, a series of laboratory WDM experiments at a net bit-rate-per-channel of 100 Gb/s led to a field experiment on a deployed LambdaXtreme system as part of a live-traffic-bearing Verizon installation on a 500-km link between Tampa and Miami, Florida [75]. The transmission of real-time video traffic on a 100-Gb/s wavelength demonstrated that deployed networks were upgradable to 100 Gb/s, which together with Ethernet starting to push towards its next logical rate of 100GbE provided the impetus for the early introduction of the first commercial single-carrier 100-Gb/s line cards by Alcatel-Lucent, shipped to Verizon in late 2009 [90]. These line cards were already built on digital coherent technology, arguably the most disruptive technology introduced in optical networking over the past two decades, and were developed in parallel to the high-speed and high-interface-rate direct detection systems that drove 100-Gb/s, starting the third major era in WDM system developments.

2.3 The era of amplified coherent systems

2.3.1 CMOS catches up

Coherent receivers were extensively studied in the 1980s [91,92], as their higher receiver sensitivity allowed for longer regenerator spacings in the span-by-span regenerated systems of the time, but due to the difficulties associated with analog frequency and phase locking were not commercialized for fiber-optic systems then. The success of the EDFA in the early 1990s made span-by-span regenerated systems obsolete, which also shelved most work on coherent receivers, except for a set of notable 1991/92 papers by Derr [93,94] that introduced the concept of digital frequency and phase locking in coherent optical receivers and, despite the fact that they fell victim to their time and went largely un-noticed for about 15 years, form the link between analog coherent receivers of the pre-EDFA era and fully digital coherent receivers of today. In fact, the resurgence of coherent detection was both *need driven* (to

unlock systems with higher SEs and with less problems from chromatic and polarization-mode dispersion that were plaguing 40-Gb/s systems in the early 2000s) and *opportunity driven*, as the speed of CMOS processing and CMOS-integrated analog-to-digital converters (ADCs) and digital-to-analog converters (DACs) had caught up with 10-GBaud symbol rates. The new CMOS capabilities first enabled maximum-likelihood sequence estimation (MLSE) [95], commercially introduced at 10 Gb/s in 2004 [96] as well as digital electronic dispersion pre-compensation in 2005 [97,98]. Digital coherent receivers combine the advantages of analog homodyne detection (minimum electrical receiver bandwidth requirements) with the simplicity of analog heterodyne detection (no need for analog optical phase locking). Using a free-running local oscillator (LO) laser, they beat the signal roughly to baseband and convert the full optical signal field, i.e., its real and imaginary parts, also called its in-phase (I) and quadrature (Q) components, in both polarizations, to the digital electronics domain, giving them the name “intradyn receiver”. Digital access to the full optical field enables *quadrature modulation* (I/Q) and *polarization-division multiplexing* (PDM) to increase spectral efficiency by a factor of $2 \times 2 = 4$, allowing a 40-Gb/s receiver to be based on 10-Gb/s componentry, compatible with the capabilities of CMOS electronics of the time. Access to the full optical field in digital form also opens up the possibility of adaptive digital compensation of CD, PMD, optical filtering effects, and even distortions from fiber nonlinearities, a possibility that had been exploited in optical coherence tomography (OCT) research already in an Optics Express paper published in 2001 [99]; similarly, integrated digital coherent receivers developed for telecommunications applications were used to demonstrate advanced OCT capabilities in 2016 [100]. These synergetic examples impressively illustrate the cross-fertilizing potential of inter-disciplinary optics research.

The very significant benefits of digital coherent receivers come at the cost of non-negligible extra receiver complexity, including the need for an LO laser at the receiver, and the need for powerful (and power hungry) DSP, commercially implemented within an application-specific integrated circuit (ASIC). However, as all coherent transponders are essentially built on the same basic opto-electronic front-end architecture, counter to the rather application-specific architectures needed for other advanced formats such as direct-detection DPSK and DQPSK, higher component volumes were anticipated for coherent solutions early on which led to significant investments in coherent receiver components.

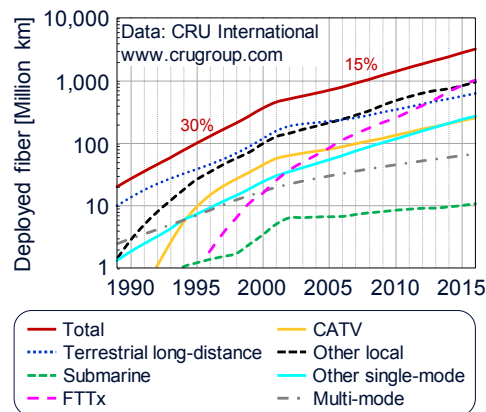


Fig. 5. Deployment of new fiber since the early 1990 across various sectors [132].

Several research groups studied digital coherent detection starting in 2004, with the goal to build 40-Gb/s transponders using polarization-multiplexed QPSK at ~ 10 GBaud [101–107]. The first intradyne transponder was commercially implemented in 2008 by Nortel at 40 Gb/s [98,108], closely followed by a parallel development by Alcatel-Lucent at 28 GBaud to enable the first commercial 100-Gb/s single-wavelength transponder, delivered to Verizon in

late 2009 [90]. In conjunction with the standardization of 100-Gb/s Ethernet (100GbE) in 2010, the adoption of 100-Gb/s technology was very rapid, from less than 10 commercial deployments in 2010 to about 600 worldwide deployments in 2014, mostly in long-haul networks. Broad deployment of 100 Gb/s in metro networks lagged behind long-haul deployments, a normal occurrence; today, metro networks largely use digital coherent transponders as well, and even shorter-reach (~100-km) datacenter interconnect (DCI) systems are likely to eventually employ digital coherent technology, even if simpler solutions based on direct detection and self-coherent techniques may temporarily still play a role in this space [109–112]. An overview of digital coherent transponder technology is given, e.g., in [113–118].

2.3.2 Impact of coherent transmission on fiber types and link designs

Digital coherent transponders are able to compensate trans-oceanic amounts of CD with little to no penalty, depending on the size and architecture of the digital filters used within their DSP chain. In addition, it turned out that high-speed coherent optical transmission, in contrast to electronically precompensated direct detection, which was also studied for DM-free systems [97,119], favors dispersion-uncompensated links with high local fiber CD [120–123]. This led to the elimination of DM from newly deployed coherent optical links and resulted in modern fibers having an even larger CD than SSMF (around 20 ps/km.nm), with ultra-low losses of just over 0.14 dB/km and large effective areas to reduce nonlinear distortions [124–126].

The introduction of digital coherent detection also had a major impact on *network planning*: Direct-detection systems using DM needed sophisticated modeling using computationally intensive split-step Fourier transform simulations to predict system performance with reasonable accuracy. (Accurate system performance predictions are very important, as each dB of prediction error from system planning tools needs to be accounted for by a dB of extra system margin, which is operationally equivalent to a dB in actual receiver performance!) Non-DM coherent links, however, lend themselves to fairly simple analytic and semi-analytic descriptions that accurately replace nonlinear distortions by an equivalent additive Gaussian noise source and facilitate accurate system modeling [127–130], with the resulting Gaussian Noise (GN) and Enhanced Gaussian Noise (EGN) models based on frequency domain analyses, and the freely available on-line tool known as the Nonlinear Interference Noise (NLIN) Wizard [131] based on a time-domain approach. Future automated systems using software-defined networking (SDN) are expected to greatly benefit from these simple models for on-the-fly network management and optimization, cf. Sec. 3.6.

The move from MMF (Sec. 2.1) to SSMF, DSF, and NZDF (Sec. 2.2.1), and then back to high-CD fiber for coherent systems nicely reflects the interplay between fiber types and transponder design that has taken place ever since the first lightwave systems were installed. Following new transponder capabilities, new fiber has continuously been deployed around the world, but if possible with backwards compatibility in mind, as the labor costs associated with laying new fiber in a terrestrial context often dominate all other costs associated with system deployments, and operators must therefore re-use already deployed fiber for new generations of systems. As of today, the overall fiber deployment has exceeded 4 billion kilometers (cf. Fig. 5), which could be wrapped around the globe approximately 100,000 times! – The amount of globally deployed optical fiber has been growing at a rate of ~15%/year since ~2000, driven to a significant extent by short-reach applications such as fiber-to-the-home (FTTH) and datacenter applications. Note that a single big data center can house as much as 50,000 km of fiber [133].

2.3.3 Off-line system performance assessments

From an experimental point of view, advanced communications engineering concepts such as FEC and DSP ended the possibilities of directly counting bit errors in the lab to determine

system performance: While prior to the use of FEC, BERs were routinely measured down to levels of 10^{-10} or 10^{-11} to identify lurking BER floors, starting with hard-decision FEC, BERs were measured only down to the “FEC threshold,” i.e. the BER at the input of an FEC decoder that presumably allows for reliable error correction. Depending on the underlying FEC, pre-FEC BERs are between roughly 10^{-2} and 10^{-5} . While the use of pre-FEC BERs as a system performance metric works reasonably well for hard-decision FEC as long as errors occur statistically independently (in practice ensured by large-scale bit interleaving), this metric fails for more modern systems using soft-decision FEC, for which pre-FEC BERs are not only inaccurate [134] but lose their conceptual meaning and must be replaced by information theoretic metrics such as the normalized generalized mutual information (NGMI) as an accurate pre-FEC metric to predict successful FEC decoding [135–137]. In addition, the extensive DSP required by digital coherent receivers makes it impossible to perform high-speed experiments in real time. Instead, a block of typically a few million digital samples of the coherently received optical field is digitized and stored using a real-time oscilloscope. (Such instruments exist with sampling rates of up to 256 GSamples/s today.) The captured data is then processed *off-line* on a computer, using more or less practically ASIC-implementable DSP algorithms. As such, the resulting performance indicators only capture a relatively short segment of the actual optical signal, and great care must be taken to avoid experimental pitfalls. In particular, transient effects from, e.g., lighting strikes [138] are hard if not impossible to investigate using off-line processing. True *real-time* DSP with full-fledged FEC is reserved either for actual product testing (for which ASICs have been developed) [9,108] or for field-programmable gate array (FPGA)-based demonstrators that either work at an order-of-magnitude lower speed than coherent ASICs [105] or use complex setups involving many FPGAs to get to line speed [139].

2.3.4 High-speed coherent interfaces

Owing to its inherent four-fold parallelization in quadrature and polarization, coherent transponders let symbol rates drop by a factor of 4 relative to their direct-detection binary counterparts. While this was essential to enable commercial coherent systems around 2008, based on CMOS technology between 10 Gb/s and 28 Gb/s per lane, it also allowed experimental interface rates based on higher-speed discrete electronic components to quickly scale beyond 100 Gb/s. With an additional factor of 2 in bit rate through 4-level electrical modulation per quadrature, 56-GBaud 16-ary quadrature amplitude modulation (16-QAM) at a net 400-Gb/s was demonstrated in 2009 [140]; by 2013 coherent research experiments were operating at 107 GBaud [141], the symbol rate that had been demonstrated for direct detection systems back in 2005 [142]; and by 2015 the single-carrier Tb/s barrier was broken using 64-QAM [10]. After 19 years of research, a single laser carrier could carry the information that required an entire WDM research setup in 1996! Although this progress in interface research is impressive, it only corresponds to a growth of 14%/year if we disregard the one-time doubling in interface rates due to PDM (cf. Fig. 2), yet required the use of higher-order QAM to encode more bits per symbol [143]. This inherently comes at the cost of a reduced transmission reach, which limits the applicability of such high-speed interfaces in long-haul networks. It is important to note in this context that the *reach* of a transponder is primarily given by its *modulation format*, irrespective of its symbol rate, i.e., a 100-Gb/s QPSK transponder will to first order have the same reach as a 1-Tb/s QPSK transponder. Performance differences with symbol rates lead to second-order effects, which can either be *fundamental* (such as fiber nonlinearities slightly favoring lower symbol rates [144,145]) or *technological* (such as high-speed implementation penalties degrading system performance at higher symbol rates [118]).

Commercial optical interface rates have continued to scale at ~20%/year (cf. Fig. 1). However, in order to do so, they had to resort to higher-order QAM, which inherently limits their reach. The 400-Gb/s single-wavelength commercial optical interfaces available today

use 64-QAM, with a reach in the ~100-km range [9]. In fact, having to resort to higher-order modulation in interface scaling is a consequence of commercial *symbol rates* only scaling at 10%/year, cf. Fig. 1. The fairly slow scaling of commercial symbol rates is due to digital coherent detection requiring the *integration* of ADCs and DACs with the CMOS DSP, owing to the enormous bandwidth needs between the converters and the DSP. For example, a digital coherent receiver operating at 50 GBaud with 2-fold over-sampling of its 4 ADCs at a nominal 8-bit resolution sports a $2 \times 4 \times 8 \times 50$ GBaud = 3.2 Tbit/s bus between the ADCs and the DSP, far too much to think of self-standing components at reasonable power consumption.

2.3.5 Optical superchannels

While it is increasingly difficult to build higher-rate optical interfaces with long-haul reach to even support a 20% annual speed increase per laser carrier, IP routers, based on CMOS packet processing (hence following Moore's Law), continue to scale at ~40% per year [1]. While it was possible in the early 1990s to aggregate multiple IP router ports onto a common high-speed WDM wavelength, the situation is reversed today: More than one wavelength is generally needed to transport the traffic originating from a single router port, leading to the generally undesirable situation of "*inverse multiplexing*." One of the more important developments of the decade in this context was the introduction in 2009 of the concept of "optical superchannels," which are optical interfaces consisting of multiple optical carriers that are co-generated, co-propagated, and co-detected. While initially studied based on orthogonal frequency division multiplexing (OFDM) [146–148] and optically shaped subcarriers [149,150], the availability of fast DACs eventually allowed for the digitally generated Nyquist-shaped superchannel subcarriers deployed in networks today [151,152]. One of the benefits of optical superchannels in the context of optically routed networks is that flexible-grid optical switches (cf. Sec. 2.3.7) allow modulated carriers to be closely spaced, leading to an improved "intrachannel" SE, as individual subcarriers are not individually filtered and can be co-processed at transmitter or receiver to further mitigate impairments across the superchannel, cf. Fig. 6. Superchannels are now widely used in research experiments and have been implemented in commercial systems to reach bit rates above the respective capabilities of the coherent DSP ASICs. Today's 400-Gb/s transponders for long-haul transport of the just recently standardized 400GbE router interfaces are based, e.g., on 2x200-Gb/s 8-QAM or 4x100-Gb/s QPSK superchannel technologies. The *need for parallelism* to support optical network scaling already points to the main theme of current and future optical communications research, cf. Sec. 3.

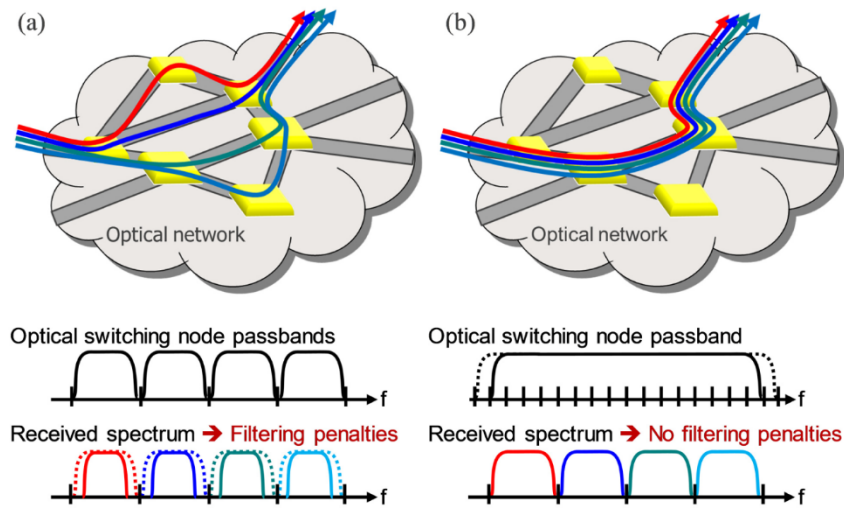


Fig. 6. (a) Routing of an inverse-multiplexed signal through an optical network yields filtering penalties on its subcarriers; (b) optical superchannels avoid these penalties through joint routing.

2.3.6 System capacity reaching its limits

While higher-order QAM has been instrumental to the scaling of single-wavelength interface rates, it has been even more critical to the scaling of WDM system capacity. As with transmission reach, it is primarily the *modulation format* that dictates the achievable SE, and hence the WDM capacity for a given system bandwidth. In general, m -QAM can carry up to $2 \log_2(m)$ information bits/symbol, where the factor of 2 accounts for PDM. The actually achievable SE is always less than this number. On the *digital side*, the SE is reduced by FEC overhead, by pilot and framing bits and symbols, and by deviating from a uniform probability of sending one of m symbols per pulse. Using the latter method, known as probabilistic constellation shaping (PCS), higher-energy symbols are sent less frequently than lower-energy ones, thus gaining up to 1.5 dB in SNR performance and allowing for quasi-continuous adaptation of SE and reach to the respective channel conditions [8,153]. On the *analog side*, the SE is reduced relative to $2 \log_2(m)$ by the spectral extent of the modulated signal with a finite spectral roll-off, as well as by guard bands between WDM channels that are necessary to avoid WDM crosstalk and filtering penalties in ROADMs networks. The spectral extent of the modulated signal is given by the temporal shape of an isolated symbol pulse, which is typically controlled digitally using DACs at the transmitter to obtain square-root raised-cosine pulses, loosely referred to as Nyquist pulse shaping [118,151,152].

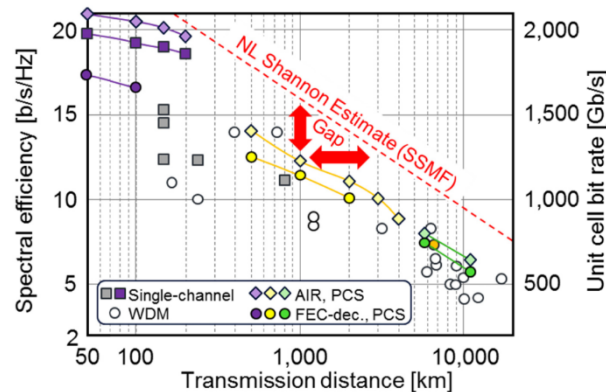


Fig. 7. Experimentally achieved record SEs (in two polarizations) versus transmission distance (markers), shown with the Shannon limit estimate on SSMF (dashed).

Figure 7 summarizes experimental research records in terms of their SE and their obtained transmission reach, from <100-km short-reach systems up to >10,000-km trans-Pacific systems. The trade-off between SE and reach is seen to follow a logarithmic relationship, with each doubling in transmission reach reducing the (polarization-multiplexed) SE by 2 bits/s/Hz, as examined in detail in [1] and references cited therein. The best achieved experimental records across all transmission ranges make use of PCS with constellations as large as 4096-QAM [17,154,155] for short-reach applications, going down to 256-QAM for terrestrial [156], and 64-QAM for submarine distances, the latter performed in 2017 on a deployed trans-Atlantic cable operated by Facebook [8]. It is worth pointing out two experimental aspects that put such record results in perspective:

- Some experiments assume the availability of an ideal FEC and quantify system performance based on the achievable information rate (AIR), which over-estimates the SE by ~ 1 b/s/Hz compared to experiments based on practically implementable FEC (yellow and green diamonds versus circles in Fig. 7, as well as purple diamonds versus squares).
- Some experiments use only a single modulated wavelength, from whose spectral extent a potentially achievable SE is being estimated. This neglects all linear and nonlinear WDM effects (such as required guard bands to avoid linear crosstalk or WDM fiber nonlinearities) and can lead to significantly over-estimated SEs, as impressively evidenced by the short-reach results (purple in Fig. 7), where a single-channel (squares) and a WDM experiment (circles) on the exact same experimental platform yielded more than 2 b/s/Hz of SE difference [17,155].

The scaling of SE over the years is shown in Fig. 8, revealing a $\sim 20\%$ /year growth rate that is rapidly saturating, with current records at 8.65 b/s/Hz *per polarization* for research experiments [17] (red) and 8 b/s/Hz *per carrier* (PDM) for products [9].

The continued slow growth of fiber capacity since ~ 2000 (cf. Fig. 1) was cause for concern, and speculation abounded that perhaps transmission systems were approaching some fundamental limits, defying the notion of “infinite bandwidth” that up until then was generally assumed an optical fiber would provide for all practical purposes. Already in the late 1990s a small working group at Bell Labs, consisting of experts in lightwave systems, information theory, communication theory, nonlinear optics, partial differential equations, and coding, tried to understand the capacity limits of a nonlinear transmission medium such as a silica fiber. The problem proved to be analytically intractable at the time but some progress was made by making a number of empirical assumptions about the mathematical manifestation of the Kerr nonlinearity. In 2001 a short paper [158] produced what, by now,

are very familiar figures colloquially known as the “nonlinear Shannon limit” in which the capacity initially follows the linear Shannon curve but at higher powers rolls over due to nonlinear effects, Fig. 9(a). However, due to the uncertainties of the nonlinear assumptions as well as due to several simplifications and partly unrealistic system approximations made in [158] (as well as in several follow-on papers, cf. the critical review and discussion provided in [122,159]), there was quite some skepticism about the quantitative accuracy of the results. It was not until the end of the decade that a more rigorous and comprehensive solution of the nonlinear Shannon capacity problem was proposed [122,159]. The numerical simulation-based results reported in these analyses, which have since been quantitatively verified using the (semi)analytical techniques mentioned in Sec. 2.3.2 [160], were startling in that it became apparent that the best experimental capacity results at the time were already within a factor of 3-5 of the fundamental limit estimate (cf. Fig. 9(b) [122]). Along the same lines, Fig. 7 reveals the small gap of today’s experimental records to the nonlinear Shannon limit estimate on SSMF, both in terms of SE and in terms of transmission reach, across a wide range of transmission distances.

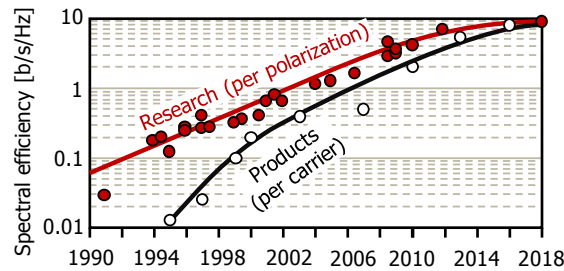


Fig. 8. Scaling of spectral efficiencies in research and products (Figure after [157]).

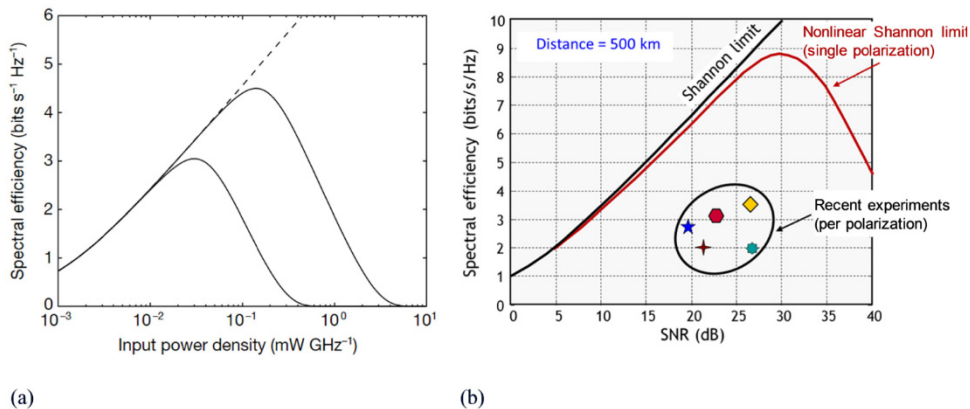


Fig. 9. (a) First attempt to derive a Shannon limit for the nonlinear optical fiber channel [158]; (b) first quantitatively accurate estimate of the fiber-optic Shannon limit in the context of modern coherent systems [122,159].

In addition to these fundamental Shannon considerations, practical issues about the optical damage threshold of deployed fiber due to the fiber fuse effect provided concern within the community [161]. The implications of closely approaching the limit of fiber capacity, in view of the continuously exponentially increasing demands for transmission capacity, were discussed in a plenary talk at the European Conference on Optical Communications in Vienna in 2009 [25], which laid out the transition into the 4th era in optical systems research dealing with spatial multiplexing, cf. Sec. 3.

2.3.7 Flexible mesh networking

The need to provide a flexible networking layer to manage the many WDM channels was recognized in the mid-1990s within the MONET project [56,57] and the ACTS-PELICAN project [162] which performed early field trials of optical networking functions. While some reconfigurable networking was incorporated into lightwave systems in the 2000s, the widespread commercial introduction of full optical switching flexibility coincided with the introduction of coherent systems in the 2010s, and networks since then have embraced mesh networking with higher-degree switching, flexible add/drop options, and the ability to provide remote provisioning, protection and restoration functions. While the expansion of networking functions may be seen to be purely co-incidental with the introduction of coherent transponders, coherent systems do provide some key benefits to flexible networking: they eliminate the need for dispersion compensation at add/drop sites; they allow the transmission format and rate to be adapted to link distance and available channel bandwidth; and they inherently provide steep receive filtering through digital processing at the receiver [163].

Today's ROADMs provide multi-degree mesh connectivity [163], flexible wavelength allocation [165,171], and typically incorporate dynamic equalization of gain and signal power to improve system OSNR in order to maximize reach. While initial WSSs that make up modern ROADM nodes only supported the fixed channel allocation plans of the standard 50-GHz ITU grid used in the 1990s [80,164,167,168], optical networks have evolved by subdividing the grid [165,166,171]. Currently, the grid is defined on a 6.25-GHz granularity with a minimum 12.5-GHz slot width to match flexible data rates and modulation formats for a fine-grained rate/reach optimization on any transmission link, but in principle the technology allows for channel widths to be adjusted with sub-GHz granularity. A flexible grid can allow for, e.g., 30% of additional capacity on smaller (metro) networks [170]. As discussed in Sec. 2.3.5 (Fig. 6), a flexible grid also enables optical superchannels to improve SE by allowing closer-spaced subcarriers.




Fig. 10. Client and line interfaces.

Note in this context that flexible optical networking has two very different cross-layer implications, whose resolution within the industry is still a topic of discussion (cf. Fig. 10):

- (a) If one takes a *router-centric view* and considers the client interface rate (i.e., the bit rate at which an IP router connects to the optical network, white boxes in Fig. 10) as a fixed parameter, a flexible optical transponder must change both its symbol rate *and* its modulation format to adapt its reach. While this strategy always fully utilizes router ports at their maximum speed, it places significantly more stringent requirements on the *optical transponder* (as changing symbol rates within a coherent ASIC is more involved than only changing modulation formats) as well as on the *optical line system* (as changing symbol rates implies changing signal bandwidths, which can lead to stranded capacity, especially in mesh networks, and the need for complicated defragmentation [169]).
- (b) If one takes a *line-system centric view*, one can keep the symbol rate constant and only change the modulation format to adapt transmission reach. This not only simplifies the

coherent transponder design but also leaves the optical signal bandwidth fixed, which eases flexible-bandwidth channel allocation hardware and software on the line system side. However, this approach inherently changes the client interface rate, which results in back-pressure for the IP routers connecting to the WDM system.

Today's ROADMs support mesh architectures with node degrees greater than eight, and can provide colorless, directionless and contentionless (CDC) add/drop, which permits a transponder to transmit on *any wavelength* ("colorless") in *any nodal direction* ("directionless") and *without blocking* ("contentionless") [163]. State-of-the-art ROADMs today support nodal degrees up to eight, which can be used to connect eight fiber pairs leaving the node in physically different directions. Nodes that topologically require degree-eight connectivity are rare; but such a ROADM can also be used for *four* topological directions with *two* fiber pairs per direction, in the spirit of a fiber overlay (SDM) solution, cf. Sec. 3. Using a CDC ROADM, any arbitrary set of optical frequencies from any direction can be dropped to a shared set of drop-side modules without drop-side blocking.

Regardless of their flexibility, ROADM architectures must support *robust connections* in the case of *component failure*. This requires that traffic in one topological direction of the ROADM node has to be independent of the traffic in the other topological directions [163,165]. In rings, this was referred to as *east-west separability*. As a consequence, optical elements handling traffic coming from and going towards a given topological direction (e.g., east) must be capable of continuing to function if any component supporting any other direction (e.g., west) fails, is being replaced, or is undergoing maintenance. The practical implication is that a single switching device, cannot simultaneously provide both the add and the drop function for a given through-path in the network; rather, the switching function must be split between switching elements attached to each fiber direction. However, a single module can be used per direction (e.g., west). With reference to Fig. 11, instead of implementing a node using a single switching element (a), one needs to split the switching function into $1 \times N$ WSSs associated with each nodal direction (b). Transmit and receive WSSs for a given nodal direction can be combined in the same module or linecard, as they serve a single (bidirectional) direction [163].




Fig. 11. (a) Conceptual view of multi degree ROADM supporting colorless directionless contentionless (CDC) functionality. Colored lines represent wavelengths being added, dropped, and passing through the node. Wavelengths can be added from any to any direction from any transponder on any color unless that color is already used for a through path; (b) Typical implementation of the functional node of (a), whose functional modules are explained in Fig. 12.

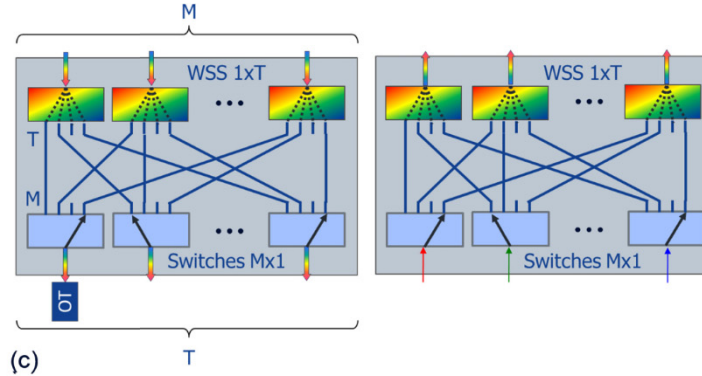
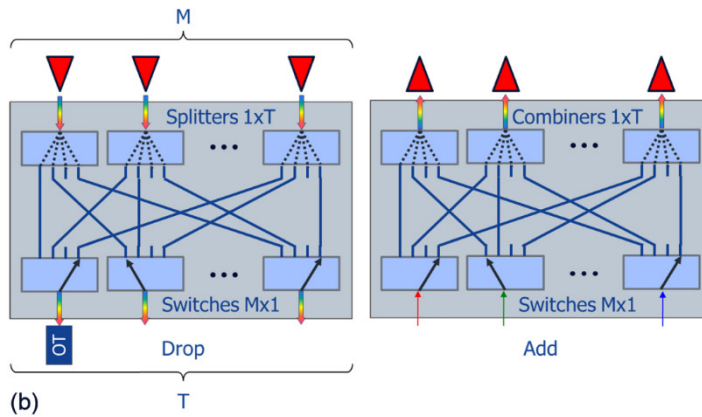



Fig. 12. Functional elements of the ROADM node of Fig. 11; (a) WSSs; (b) DxT multicast-switch-based add/drop architecture; (c) DxT WSS-based add/drop architecture.

While ROADM nodes are generally described as contentionless and colorless, it is important to remember that once a signal is placed on a wavelength, it will remain on that wavelength until it reaches a receiver. Hence, if a certain wavelength is already in use on *any* section of an intended path across a mesh network, the path cannot be used for that wavelength without having to resort to *wavelength conversion*. While in 2001 all-optical wavelength conversion was still seen as holding significant promise [174] and has been widely researched [175,176], it has not resulted in any practical network elements or any real deployments. If required, wavelength conversion must therefore be implemented using optical transponders that receive on one wavelength and transmit on another. As transponders are costly, wavelength conversion becomes an undesirable operational step that routing and wavelength assignment (RWA) algorithms try to avoid. If needed, a pool of local

transponders can be used on stand-by to provide wavelength conversion as well as protection/restoration functionalities. As little as 15 – 20% of overhead capacity can provide sufficient restoration and protection within a large network [163]. The general avoidance of wavelength conversion in transparent optical networks conceptually results in optical switches instantiating a set of *parallel switching planes*, one for each wavelength, which effectively reduces the connectivity and increases blocking probabilities (see discussion in Sec. 3.5).

Internally, a CDC-F node architecture [163] is supported by high-port-count (e.g., 20) *wavelength-selective switches* (WSSs), *multicast switches* (MCSs), and *arrays of amplifiers* in the add and drop paths, as shown in Fig. 11(b). A $1 \times (D + P)$ WSS can send any set of wavelengths at its input to any of its $D + P$ output ports. Of each WSS's output ports, $(D-1)$ ports are routed to the mesh outputs of the degree- D node (typically $D = 8$), 1 port is reserved for future expansion, and P ports provide add/drop functionality (typically $P = 12$). For a total of $P \times T$ add/drop transponders to be supported at the node, each WSS add/drop port goes to either a $D \times T$ MCS or WSS pair (typically $D \times T = 8 \times 16$), allowing for $P \times T = 12 \times 16 = 192$ CDC drop ports, equivalent to 25% add/drop traffic assuming 96-wavelength systems to be interconnected. Using the extension port, the add/drop ratio can be increased. Current WSSs consist of a free-space system that disperses the optical spectrum across a liquid-crystal-on-silicon (LCoS) phased array panel [171], which steers different frequencies to different output ports while re-multiplexing them to different fibers. In the example of Fig. 12(a), the WDM multiplex enters on the yellow fiber on the left and is dispersed at the diffraction grating, with each dispersed signal hitting a different segment of the spatial steering element. The orange-colored signal is then steered to that spot on the diffraction grating that diffracts that wavelength back into its destination output fiber. The LCoS has many more pixels than the number of channels, which allows for a flexible wavelength allocation [165,171]. An MCS is an $D \times T$ broadcast-and-select switch that is not wavelength selective; functionally, it switches any of its D input ports to any of its T output ports. It consists of a D -wide array of $1 \times T$ power splitters, the outputs of which are shuffled and connected to the inputs of an T -wide array of $D \times 1$ selection switches, cf. Figure 12(b). While MCSs are working with coherent receivers that select the respective wavelength channel by tuning the local oscillator to the desired wavelength and apply steep filtering in the digital domain [177], they may be replaced by the $D \times T$ architecture of WSSs shown in Fig. 12(c), which improves power efficiency by eliminating the amplifier arrays shown in Fig. 12(b) [172,173].

An important aspect of the use of ROADMs in networks pertains to their management. This is typically done through generalized multi-protocol label switching (GMPLS), which has evolved with the capabilities of the switching hardware to provide advanced network services [178], including provisioning, restoration, automatic path calculation, and dynamic re-routing and re-optimization, using distributed processing for path computation, signaling, and routing, and making the network resilient to multiple failures while maintaining a single GMPLS/optical transparency domain.

3. Future research directions

3.1 Scaling disparities and the capacity crunch

Network traffic growth has been examined from many different angles, and its continuing exponential nature has been amply verified (see, e.g., [1,157,179–181] and references cited therein). Typical numbers for long-term traffic growth are in the 60%/year (2 dB/year) range, although actual numbers vary widely between applications, operators, and geographies [157]. Complementing these and other existing traffic growth studies, Fig. 13 takes a different viewpoint, looking at the *actually deployed optical transponder hardware*, typically reported by analysts as global annual deployments across the entire industry. As the figure shows the cumulative sum of the reported annual capacity deployments, the first few (light-colored) data

points in each analyst's record are not representative of actually installed capacities, since they do not include the installed base in the starting year of the respective record. Beyond those initial points, the figure shows a $\sim 45\%$ annual increase of the aggregate deployed optical transponder capacity over the past decade, consistent across 4 major analyst firms, both in the metro/long-haul part of the network and across all optical ports, including client and short-reach interfaces. The former reveals a total deployed metro/long-haul capacity of ~ 100 Petabit/s, the latter of ~ 2 Exabit/s. Note that Cisco's Virtual Network Index (VNI) only reports ~ 200 Terabit/s of traffic for 2017 (~ 100 ExaBytes/month), with a growth at a mere $\sim 24\%$ per year [181]. The discrepancy arises because Cisco's VNI only accounts for end-to-end IP traffic, while the deployed WDM capacity captures *all* traffic types, accounts for over-provisioning of operational WDM systems to accommodate peak-to-average traffic variations and diurnal fluctuations, and comprises the fact that an end-to-end transported information bit typically touches many WDM transponder ports on its way from source to destination. While the globally averaged WDM capacity growth of $\sim 45\%$ per year seems at a first glance incompatible with the much lower numbers of $\sim 20\text{-}30\%$ reported by major telecom operators [182,183], the fact that large webscale operators see traffic growth well in excess of 45% balances the equation.

The continuing exponential global traffic growth of $\sim 45\%$ in contrast to interface rates and fiber capacities only scaling at $\sim 20\%$ per year (cf. Fig. 1) reveals an increasingly critical disparity, leading towards the "capacity crunch" anticipated about a decade ago [25,179]. This disparity in growth rates is fundamentally rooted in the inherent scaling differences between digital integrated electronics following Moore's Law (driving the development of devices used to generate, process, and store information) and analog high-speed optoelectronic technologies (driving the development of devices used to transport information), as discussed in detail in [1].

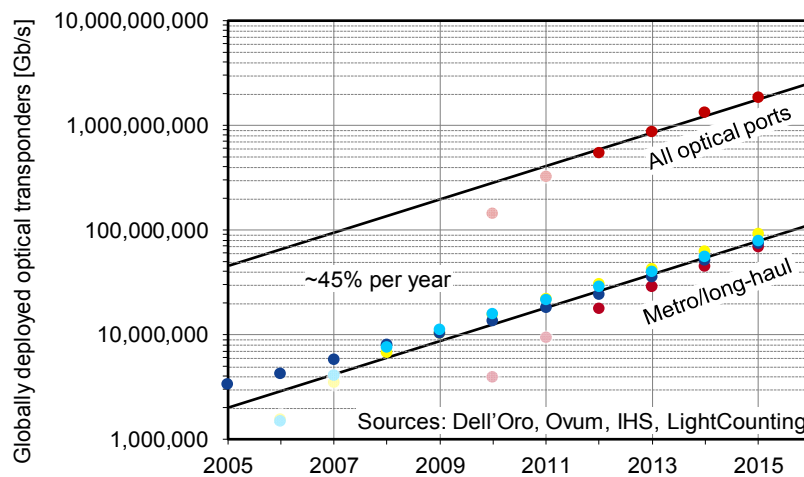


Fig. 13. Global network traffic growth estimated from deployed optical transponders.




Fig. 14. (a) Aggregate data center footprint in millions of square feet for four major Cloud providers as a function of time, showing a 37% annual growth rate; (b) growth rate of Broadcom Ethernet switch capacities and line rates as a function of time, showing 41% and a 22% annual growth rates [192], along with the 71% annual traffic growth rate reported by Google [184].

As a second, alternative viewpoint for traffic growth, and complementing existing analyses, we look at the *size of the Cloud infrastructure* and its evolution. The Cloud is based heavily on centralized architectures using what is often described as a small number of web-scale datacenters, that host data stores, services, and applications. These massive data centers, each costing hundreds of millions of dollars, offer performance and economic benefits through localized aggregation of resources and efficient management of power and heat. As applications shift to the Cloud, its data processing, storage, and transport requirements are growing rapidly. Growth in traffic of 70% per year has been reported by Google [184] (cf. Fig. 14), and YouTube uploads show similar growth rates [1]. Cloud providers have also commented that internal data flows are growing faster than external user data flows.

In order to assess the “size of the Cloud,” we systematically analyzed the evolution of the physical floor plan (area) of datacenters, assuming that the growth in datacenter area corresponds to the growth rate in the number of servers and in overall power dissipation. This provides a quantitative representation of the increase in processing power (and hence in data traffic) beyond that provided by per-server technology scaling. While there are many Cloud companies, we selected Google (Alphabet), Facebook, Amazon (AWS) and Microsoft (Azure) for this study, on the basis that they have existed at sufficient scale and over sufficient time to give reasonable data sets. We restricted this analysis to data centers in the United States, as these represent some of the earliest Cloud deployments, and as it was practicable to determine the locations for these data centers. Since many Cloud services are being shifted from US servers to internationally deployed servers, it is likely that global growth rates are actually higher than the ones given in this study. Additionally, our study does not include the use of shared facility datacenters, which are common in Northern California and in Virginia. For each of the four Cloud companies, we first determined the locations of their US datacenters, based on publicly available information, including the companies’ websites [185–188], blog posts, news articles which often document impact on local communities, trade publications and other public information such as a Greenpeace report [189] on carbon footprint, which itself uses information such as building permits. Once the locations of the data centers were determined, the size of the data center buildings was obtained from measurements on Google Maps satellite pictures, including information on multiple-floor buildings [190]. Historical satellite images available through Google Earth provided information on data center evolution. The aggregate data center area for the sum of the four Cloud providers is shown in Fig. 14(a) as a function of time, revealing growth rates

in the area of data centers of 37% per year. Together with server (30% per year) and switch technology (40%/year) scaling, we arrive at an overall traffic scaling of 70%, consistent with reported numbers.

AWS has indicated [191] that their new datacenters are typically 80,000 servers with 32 MW of power, which for a typical 120,000-sqft data center means approximately one server per 1.5 sqft and 400 W. Conservatively, we can assume that the data centers we considered are at or below 2 sqft per server meaning that we can estimate that these four Cloud providers have on the order of 10 million servers in the US, with 4 GW of total power consumption, and that both numbers are growing at 37% per year.

Note from Fig. 14(b) that the capacity of packet switches has been scaling at ~40% per year [192], while the per-lane interface rate to these switches follows the same 20% scaling as that of optical interfaces shown in Fig. 1. Even though 100GbE and 400GbE standards exist, these all use parallel lanes (inverse multiplexing), owing to limitations in electrical and optoelectronic interface speeds, a problem we already encountered in the context of long-haul coherent interface scaling and superchannels, Secs. 2.3.4 and 2.3.5. Over time, this increasing degree of parallelism is inducing a significant disparity that is driving CMOS switch chips to an I/O limited regime.

3.2 Approaching Shannon per dimension

As discussed in Sec. 2.3.6, WDM capacities are quickly approaching their fundamental Shannon limits, and ways to resolve the looming capacity crunch can only rely on the five physical properties (loosely referred to as the five physical “dimensions”) available for modulation and multiplexing of electromagnetic waves shown in Fig. 15: *time*, *frequency*, *quadrature*, *polarization*, and *space*. With these five dimensions at hand, the overall fiber channel capacity, i.e., the maximum capacity that can be reliably communicated over an equivalent additive white Gaussian noise channel, can be written as

$$C = M \times B \times 2 \times \log_2(1 + \text{SNR}), \quad (2)$$

where the logarithmic term captures the maximum possible SE of a single-polarization complex optical signal (in both in-phase and quadrature components). The preceding factor of 2 accounts for polarization multiplexing, and the final capacity is obtained by multiplying the dual-polarization SE with the system bandwidth B and the number of parallel spatial paths M .

As is evident from Fig. 7, and as reflected in the long-term scaling slow-down clearly visible in Figs. 1 and 2, the gap to Shannon in terms of SE has become small, and efforts to increase the SNR through lower-noise optical amplification [193], lower-loss and lower-nonlinearity fiber [194], and digital nonlinearity compensation [79,195,196] yield only logarithmic (and hence quickly diminishing) returns [1]. Fundamentally, technologies that aim at improving the SNR will in general only yield small capacity gains, on the order of a few ten percent. Big (linear) capacity gains can only be achieved using the “pre-log” factors B and M in Eq. (2).

Access systems like digital subscriber line (DSL) modems that operate over twisted pairs of copper wires can increase capacity by reducing transmission distances (i.e., by bringing high-speed fiber closer and closer to the copper access points). Such a shortening in transmission reach is not an option for core networks, as these are expected to interconnect locations on a fixed-size globe. One option to scale core networks could be to use shorter regeneration spans, though, as visualized in Fig. 16 [197]. In order to reach the capacity-distance combination located beyond the Shannon limit of a single fiber (in this case 20 b/s/Hz over 1,500 km), one can use transponders operating at 20 b/s/Hz (using, e.g., PCS 4096-QAM), with a transponder spacing of ~20 km, i.e., using ~75 transponders to close the link. One could also start with a system that achieves the highest possible SE at the target

distance of 1,500 km (10 b/s/Hz in this example), and use just two such systems in parallel. The difference in required transponders clearly favors a parallel approach.

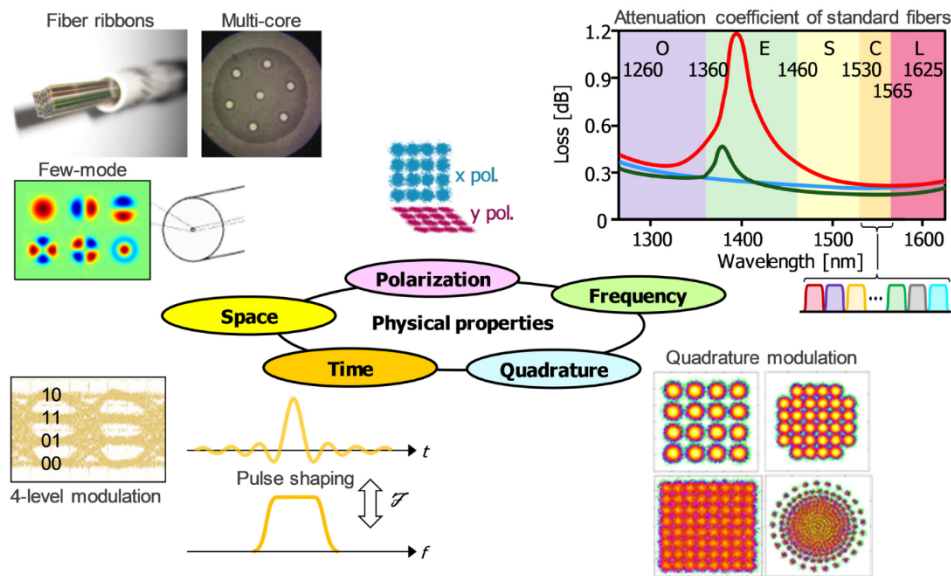


Fig. 15. Physical dimensions available for modulation and multiplexing in optical communications. (Figure after [198].)

As polarization and time/frequency are already fully utilized in the results of Fig. 7, the only pre-log multiplexing options according to Eq. (2) and Fig. 15 are wider frequency bands (B) and more spatial parallelism (M). We will therefore discuss these two scalability options in the light of currently ongoing research.

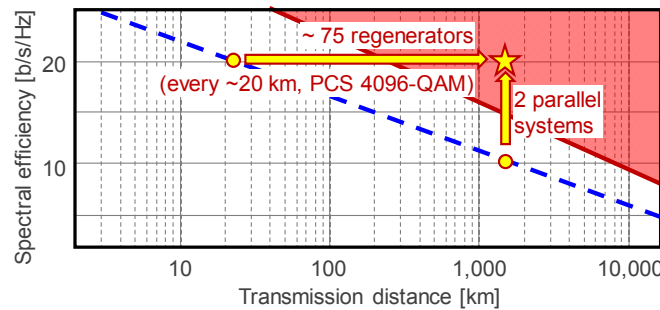


Fig. 16. The benefit of parallelism over opto-electronic regeneration. (Figure after [197].)

3.3 Scaling through ultra-wide-band systems

Ideally, widening the system bandwidth linearly increases system capacity as a pre-log factor B , cf. Equation (2), provided that such a scaling is supported by the underlying system components. Figure 17 shows typical loss coefficients across the low-loss window of commercial optical fiber with (red) and without (blue) the characteristic hydroxyl group OH absorption peak around 1380 nm. A factor of ~ 12 in bandwidth could potentially be gained if operating deployed fiber from the O-band all the way to the L-band (~ 1260 nm – 1625 nm, i.e., 53.5 THz), as opposed to using the C-band only (~ 1530 nm – 1565 nm, i.e., 4.4 THz), as is done in the vast majority of today's commercially deployed systems. However, it is very unlikely that this factor of 12 actually translates into a similar capacity gain due to several

fundamental and practical problems associated with ultra-broadband systems that make a factor of up to ~ 5 in capacity more realistic, as discussed in detail in [1].

Ultra-wideband systems comprise two separate, equally important design aspects: (i) optical fiber with low loss across a wide-band window, and (ii) optical subsystems such as optical amplifiers, lasers, and filters able to operate, seamlessly if possible, across the entire system bandwidth. These two aspects have different bearing on *green-field* situations (where new fiber is installed as part of a new system deployment) and in *brown-field* situations (where new systems are installed on an already existing fiber infrastructure). The former applies to many submarine and DCI situations, while the latter is typical for terrestrial long-haul and metro networks.

In brown-field situations, operators that do not have access to unused fiber strands have a strong incentive to exploit their available fiber resources across as wide a system bandwidth as possible, as both leasing and deploying new fiber is an expensive proposition. New fiber deployments are dominated not by the costs of the fiber but by installation labor, and prices vary widely, depending on the deployment scenario, with $\sim \$20,000$ per km assuming available duct space being a realistic assumption; deploying a 1,000-km fiber ($\sim \$20M$) therefore costs more than the WDM system operating over it. If new cables are being deployed, which is indeed constantly the case, cf. Fig. 5, care must be taken that the new fiber is compatible with already existing fiber, as signals must in general be able to traverse a mixture of old and new fiber types. This leads to the requirement of a “smooth upgrade path” whenever different fiber types are considered to become part of an already existing network [198]. The deployment of radically new fiber then becomes problematic from an overall network operations point of view. An example is photonic crystal hollow-core fiber, which could in principle be designed to operate across a ~ 37 -THz bandwidth in the $2\text{-}\mu\text{m}$ wavelength range, outside the standard $1.55\text{-}\mu\text{m}$ telecom window (cf. dashed green potentially achievable loss prediction shown in Fig. 17 [199,200]). Note that the wavelength scaling of the x-axis of Fig. 17 can be deceiving; the dashed green fiber opens up less bandwidth than the wavelength region from the O-band to the L-band in standard telecom fiber. Potentially wider-band and even lower-loss nested antiresonant nodeless hollow-core fibers that would also include the traditional telecom bands have also been studied, with a potentially achievable loss profile shown by the orange curve in Fig. 17 [200]. Note, though, that even if achievable in practice, a lower fiber loss does not resolve the capacity scalability problem owing to the logarithmic dependence of the capacity on the SNR, cf. Equation (2): As noted in [1], in order to merely *double* the capacity of a system with a SE of 4 b/s/Hz per polarization, the dB-loss coefficient of a fiber would have to be reduced by a factor of 64, much below the most optimistic predictions made for hollow-core fiber [200].




Fig. 17. Commercially achievable low-loss window of conventional single-mode fiber in the wavelength range from 1260 to 1625 nm (O-band to L-band) using legacy fiber with (red) and without the hydroxyl absorption peak (blue). The dashed green and orange curves are redrawn from [200] and represent model predictions for low-loss photonic crystal hollow-core fiber (green) and nested antiresonant nodeless hollow-core fiber (orange). Note that the wavelength scaling (x-axis) does not represent the frequency bandwidth (double arrows).

Regarding the second aspect of ultra-wide-band systems, the availability of optical subsystems across the entire system bandwidth, it is important to consider the *relative*

bandwidth B_{rel} of the targeted solution, defined as the absolute system bandwidth B divided by the system's center frequency f_c ,

$$B_{rel} = \frac{B}{f_c}. \quad (3)$$

In most fields of engineering, including microwave and optical, the complexity of components and subsystems grows with their relative bandwidth. The C-band (corresponding to the gain bandwidth of EDFAs) has a relative bandwidth of 2.3%; a recent 100-nm wide demonstration of semiconductor optical amplifiers for coherent transmission [15] spanned 13 THz of bandwidth (as indicated by the blue shaded region in Fig. 17), corresponding to a relative bandwidth of 6.6%; the frequency region from O-band to L-band occupies a relative bandwidth of 25%, as does the 2- μm region indicated in Fig. 17; octave-spanning components have a relative bandwidth of 67%, which is roughly the relative bandwidth of the nested antiresonant nodeless hollow-core fiber indicated by the dashed orange curve in Fig. 17 [200]. Building amplifiers, tunable lasers, or tunable filters that individually span such large relative bandwidths is complicated, which forces a *banded approach* for the subsystems needed to construct an ultra-broadband system, as indicated in the inset to Fig. 17. Involving *different component technologies* across their sub-bands, banded systems typically have a higher cost-per-bit than single-band (e.g., C-band) systems, which can only be justified if such systems postpone the deployment of new fiber. It becomes clear from these considerations that an expansion in wavelength has only a limited number of practically attractive use cases, and that “wavelength parallelism” is not true parallelism in the sense that a truly parallel system should deploy the *exact same* system components in parallel, which is key to bringing down cost through both *volume* and *integration*. Therefore, scaling in the frequency domain cannot solve the long-term capacity scaling problem, which requires cost-effective capacity scaling factors of 100 or even 1000.

3.4 Scaling through parallel spatial paths

As evident from Fig. 15 and from the preceding discussions on the limitations of bandwidth scaling, *parallelism in space* is the only option to significantly scale system capacities by appreciable factors in the long run. Space-division multiplexing (SDM) denotes the use of parallel spatial paths, to complement wavelength scalability (WDM) using the WDM x SDM matrix shown in Fig. 18 [1].

Table 2. Possible system evolutions over the next 10 and 20 years.

	2017	2027	2037
Symbol rate [GBaud]	50	120	300
Bit rate [Gb/s]	200 – 400	600 – 1,600	2,000 – 6,000
System bandwidth [THz]	5	5 – 12	5 – 20
Capacity per spatial path [Tb/s]	20 – 40	25 – 160	32 – 400
Unit cells per spatial path	100	40 – 100	16 – 66
Target system capacity [Pb/s]	0.02 – 0.04	1 – 2	50 – 100
Required number of spatial paths	1	6 – 80	125 – 3125

Each row of the WDM x SDM matrix represents wavelength multiplexing within one spatial path, and each column represents multiple parallel spatial paths at the same carrier frequency. Each unit square (“*unit cell*”) represents an optical signal modulated onto a single optical carrier using a single optical modulator and detected using a single optical receiver. Assuming that opto-electronic modulation and detection hardware continues to scale along its long-term trajectory, Fig. 1 predicts commercial (CMOS ASIC integrated) symbol rates of 120 GBaud by 2027 and 300 GBaud by 2037, which implies a bandwidth per unit cell on the order of 120 and 300 GHz using digital electronic pulse shaping. According to Fig. 7, and depending on the target transmission distance, this corresponds to a bit rate per unit cell of

around 1 Tb/s in 2027 and up to 5 Tb/s in 2037. Depending on whether systems will use simple C-band technologies or expand to the entire S + C + L-band region, the requirement for Petabit/s systems in 2027 and 100-Petabit/s systems in 2037 dictates tens (2027) to hundreds and even thousands (2037) of parallel spatial paths. Table 2 summarizes this potential system evolution.

Whether a logical interface is constructed out of unit cells on a common optical path but at different wavelengths (*spectral superchannel*), at a common wavelength but across several parallel paths (*spatial superchannel*), or at a mixture of both (*hybrid superchannel*), cf. Fig. 18, depends first and foremost on the availability of spatial paths throughout the network at the time of transponder deployment. *Terrestrial long-haul and metro networks*, which rely on the re-use of already deployed fiber, possibly with very different numbers of fibers available on different links of the network, are likely to prefer spectral superchannels to build out the network by gradually adding more spatial paths once the available spectrum on the existing paths is filled. On the other hand, *submarine cables and DCI systems* are typically confronted with a green-field situation, hence can deploy the required number of parallel fibers at the time of system installation, usually at little extra cost in the context of the overall project [64]. In fact, DCI systems are already employing a massive number of parallel spatial paths today (in some cases with more than 10,000 parallel fibers linking regional datacenter buildings), and both 100GbE and 400GbE standards include parallel single-mode (PSM) interfaces [201]. In the long-haul space, some operators are close to deploying fully-loaded WDM systems on parallel fibers, and ROADMs are offered with more degrees than required to accommodate physically diverse paths, thereby opening the possibility to support multiple parallel fibers per physical nodal direction (cf. Sec. 2.3.7). For submarine systems, which as mentioned in Sec. 2.2.5 are supply-power constrained, it has been shown [62–64] that massive SDM with ~50 parallel fibers per direction leads to higher-capacity, lower-cost system architectures, even if state-of-the-art technology without any further SDM integration is being used.

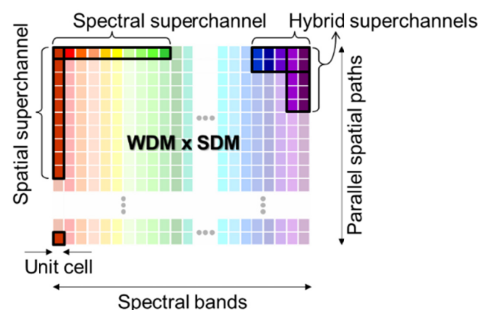


Fig. 18. A matrix of unit cells in frequency and space defines a WDM x SDM system, whose logical channels are spectral, spatial, or hybrid superchannels [1].




Fig. 19. Array integration across all system elements [198].

Across technologies, parallelism has always benefited from the *integration of simple unit cells*, e.g., in integrated circuits and multi-core processors. This lets integration aspects be a central theme associated with SDM systems research, in addition to purely architectural aspects leveraging the WDM x SDM matrix. In an optical communications system, integration can be done for all its subsystems, including transponder arrays, amplifier arrays, optical switch arrays, and the integration of parallel transmission paths in more compact

fibers, such as multi-core fiber (MCF) [202,203] and few-mode fiber (FMF) [204–206], cf. Figure 19. While SDM research has been centered around the use of new fiber (FMF, MCF) over the past years, it is important to understand that SDM systems do *not* mandate the use of such new fiber, a frequent misconception among carriers who understandably dislike the idea of abandoning their vast installed fiber base. In fact, we speak of SDM whenever the WDM x SDM matrix is leveraged through a holistic WDM x SDM system design, involving architectural and/or integration aspects to reduce the cost and energy consumption per bit compared to simply deploying parallel stand-alone WDM systems.

While both spectral and spatial superchannels benefit from array integration, the two approaches have some fundamental differences: *Spectral superchannels* offer the benefit of closer subcarrier spacings in optically routed networks as well as the possibility to digitally compensate for nonlinear crosstalk among its subcarriers, albeit at only small system performance gains [196]. Generating spectral superchannels with a significant number of subcarriers requires managing multiple source wavelengths, either by integrating multiple lasers or by demultiplexing an externally supplied comb (with sufficient power per comb line) and subsequently remultiplexing the modulated signals onto a common fiber within the spectral superchannel transponder, which becomes particularly problematic when the superchannel's spectral tunability is to be maintained, as it either requires a WSS per transponder or a lossy passive combiner, cf. Figures 20(a) and 20(b). In addition, the wavelength variability of the integrated components and the need for gain-flattened amplification within a superchannel transponder can pose practical problems. From a networking and network evolution point of view, though, spectral superchannels have distinctly preferable features, as discussed in Sec. 3.5. On the other hand, *spatial superchannels* by definition operate on a single laser carrier, which makes integrated designs easier, cf. Figures 20(c) and 20(d). For example, only a single laser is needed per superchannel transponder, all integrated components operate at the same wavelength, and gain-flattening of amplifiers within the superchannel transponder is not required; a single source wavelength can be passively split and used across all modulators and across all receivers as LO within the superchannel, with lower laser power requirements compared to a spectral superchannel [157]. Spatial superchannels can also more readily leverage DSP algorithms among their unit cells [207]. Importantly, linear crosstalk arising from the dense integration of spatial paths can be digitally compensated in a spatial superchannel architecture through multiple-input-multiple-output (MIMO) DSP [208–210].




Fig. 20. Architectures for spectral (a, b) and spatial (c, d) superchannel transponders with individual lasers (a, c) and a common optical power supply (b, d).

Significant research has been performed on SDM transmission over the past decade, as shown by the multitude of SDM-specific fibers that have been produced and the explosion of papers published in this field (cf. Fig. 21), many important of which have been reported in Optics Express, including an early Special Issue in 2011 with key original contributions [211]. Since then, SDM research has encompassed widely varying transmission media, including uncoupled-core MCF, coupled-core MCF, MMF, and FMF, examining both linear and nonlinear propagation effects [203,212–217]. In the context of these studies, coupled MCF was shown to exhibit slightly better nonlinear propagation performance compared to single-core fiber with identical core properties [218,219]. A major benefit of the uncoupled-core-fiber approach is that it does not rely on DSP to recover the data streams from the individual cores and so could be implemented with conventional transceivers. However, the complexity and cost of manufacturing such fiber may not be competitive with conventional or reduced-cladding bundles of single-mode fiber [220]. All other approaches require custom MIMO-DSP ASICs to unravel the mixing that occurs among cores or modes. Thus the number of spatial paths through such fibers will be limited both by the complexity and the interface rate of such ASICs. While the former seems manageable [221], the latter requires massive array integration and close coupling between the opto-electronic arrays and the CMOS DSP. A compromise is given by hybrid solutions, with few-moded, nominally uncoupled fiber cores [222,223]. Thus far, almost all research SDM experiments used off-line processing with one notable exception [224]. The latest tour-de-force SDM system experiment transmitted 10 Petabits/s over a 6-mode, 19-core fiber using both C + L bands. The experiment comprised 84,246 SDM/WDM channels [16].

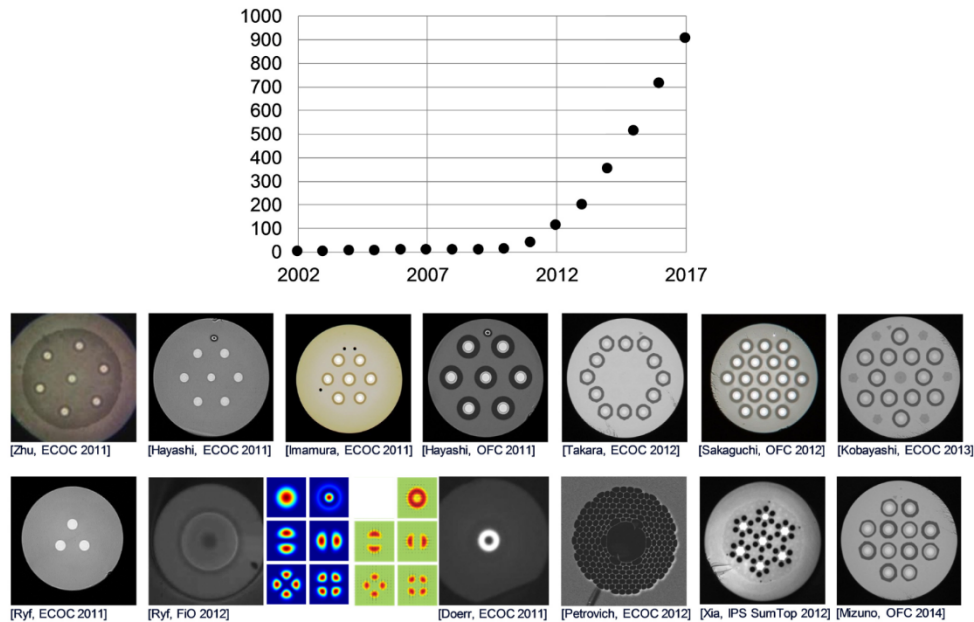


Fig. 21. Cumulative number of papers published on SDM in top journals and conferences, and various examples of SDM-specific fiber cross-sections and mode profiles. (Figure after [27].)

Regardless of which of the above SDM techniques or superchannel architecture will eventually be chosen for commercialization, *array integration* of optical components will be essential. For example, in order to construct a 10-Tb/s superchannel, ten 1-Tb/s unit cells must be integrated. If 1-Tb/s unit cells cannot be achieved in highly integrated array form, one has to transition to a correspondingly larger number of lower-rate unit cells, e.g., 100 x 100 Gb/s instead of 10 x 1 Tb/s, as the product of array size and per-component rate must equal the desired aggregate superchannel rate. As shown in Fig. 22, and as further discussed in [1], there are three pieces to the integration puzzle: (i) Opto-electronic array integration, (ii) optics-electronics integration, i.e., close integration (hybrid or monolithic) of the opto-electronic array and the CMOS DSP ASIC, and (iii) holistic DSP integration, i.e., co-designing the DSP to compensate for performance shortcomings due to the high integration density. An example of the latter is the possibility to digitally compensate for crosstalk in densely integrated modulator arrays [210]. The result will be a *fiber-in/fiber-out (FIFO) coherent processing engine* that interfaces to both client side and line side optically, as depicted in Fig. 22 [1]. Very similar FIFO architectures are expected to enter switch chips as well, owing to the disparity between switch scaling and interface scaling discussed along with Fig. 14.

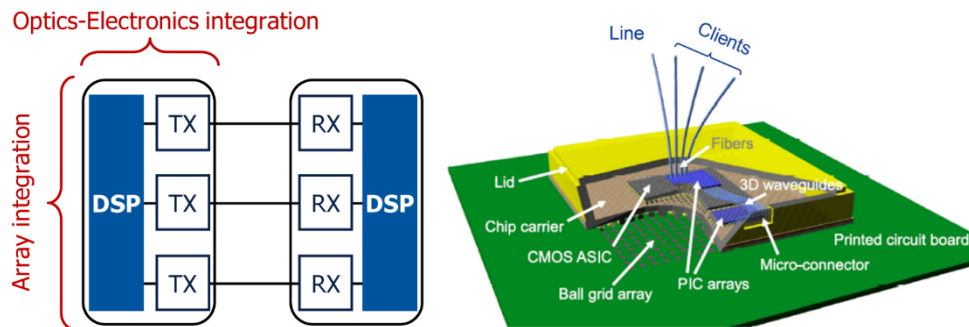


Fig. 22. Array integration, optics-electronics co-integration, and holistic DSP-opto-electronic co-design will be needed for any superchannel transponder to realize fiber-in/fiber-out (FIFO) engines [1].

3.5 Networking aspects of spectral and spatial superchannels

In addition to the technology and array integration aspects discussed above, the choice of spectral versus spatial superchannels has a significant impact on optical networking.

From a routing point of view, it is important to note that the two dimensions making up the WDM x SDM matrix are *conceptually not equivalent*: A connection cannot independently choose any available wavelength on a link but must use the *same* wavelength end-to-end (unless wavelength conversion is performed, which is a costly undertaking, cf. Sec. 2.3.7). In contrast, *any* available spatial path may be chosen independently on each link, at least in the absence of spatial crosstalk. (If crosstalk couples $M_{XT} \leq M$ spatial paths together, routing is restricted to always occur in crosstalk groups of M_{XT} paths. As a simple example, consider a PDM system on a single-mode fiber, where $M_{XT} = M = 2$, implying that both polarizations must be routed together as a group within a network.) Fig. 23 visualizes this situation: If a fraction $\gamma \leq 1$ of the aggregate link capacity in the desired egress direction of a ROADM node is utilized by other traffic, the probability that the particular wavelength range used by a spectral superchannel is already occupied on any one of the M parallel egress links is γ . Hence, the probability that this wavelength range is already occupied on all M parallel egress paths is γ^M , which is the *blocking probability* for the spectral superchannel. In contrast, a spatial superchannel occupying $M_{SC} \leq M$ spatial paths at a given wavelength will encounter blocking at a probability $\gamma^{M/M_{SC}} \geq \gamma^M$, as the number of available egress options is now restricted to the integer part of M / M_{SC} . A lower blocking probability not only means more capacity but also simplifies provisioning algorithms for the spectral superchannel.

In terms of its internal architecture, a ROADM requires the node to first perform a *wavelength separation* of the ingress signal before switching. As architecturally explained along with Fig. 12(a), wavelengths are first transformed to W individual spatial channels inside a WSS by means of a diffractive element and are then *spatially* switched before being diffracted back into a common egress fiber. Importantly, the number of wavelength-equivalent spatial paths W is dictated by the spectral resolution (“steepness”) of the ROADM’s optical filtering function, commonly defined using a 0.5-dB to 6-dB roll-off; the value of 6 dB is chosen because the filter characteristics of adjacent channels typically cross at their 6-dB points. With today’s typical values of 40-GHz 0.5-dB passbands on a 50-GHz grid, implying a 5-GHz spectral resolution, a ROADM supporting a 5-THz system bandwidth (C-band) thus needs $W \sim 1000$ wavelength-equivalent spatial paths. This number depends only on the ratio of system bandwidth to spectral resolution and *not* on the width of the (super) channels to be switched. Extending this architecture to the WDM x SDM case, a WSS would

have to support $W \times M$ wavelength-equivalent spatial paths, which quickly grows to a large number: Using current ROADM architectures and a degree-8 node as an example, one fiber per degree today requires eight 1x20 WSSs to provide full CDC functionality. For the future scenarios of Tab. 2, scaling to 22 x 8 spatial paths (by 2027) requires 176 1x440 WSSs, and scaling to 625 x 8 spatial paths (by 2037) requires 5000 1x125000 WSSs! Approaches to mask the complexity of such large numbers of spatial channels from the switching function, such as joint switching, have been proposed [225], but significant challenges remain in terms of degraded spectral resolution and increased crosstalk into undesired ROADM ports, which impact system performance and total system capacity. Additional spatial paths inside the switch also increase its physical size, even for multi-mode switching where in many cases it is necessary to transform the modal structure of the input signals to single-mode beams prior to switching [225].

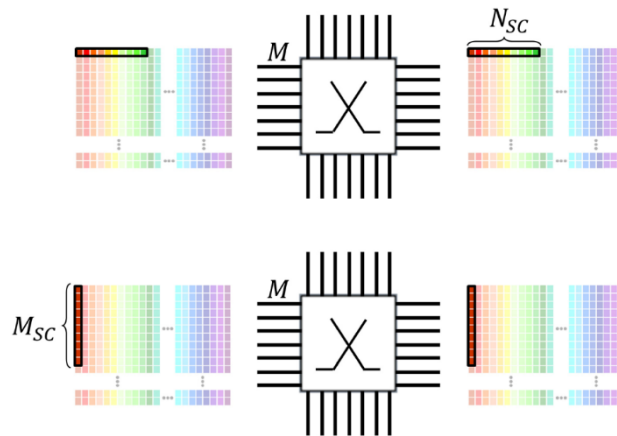


Fig. 23. Comparison of the blocking probability of spectral (top) versus spatial (bottom) super channels. The presence of multiple equivalent paths in the spectral super channel system results in much lower blocking probabilities.

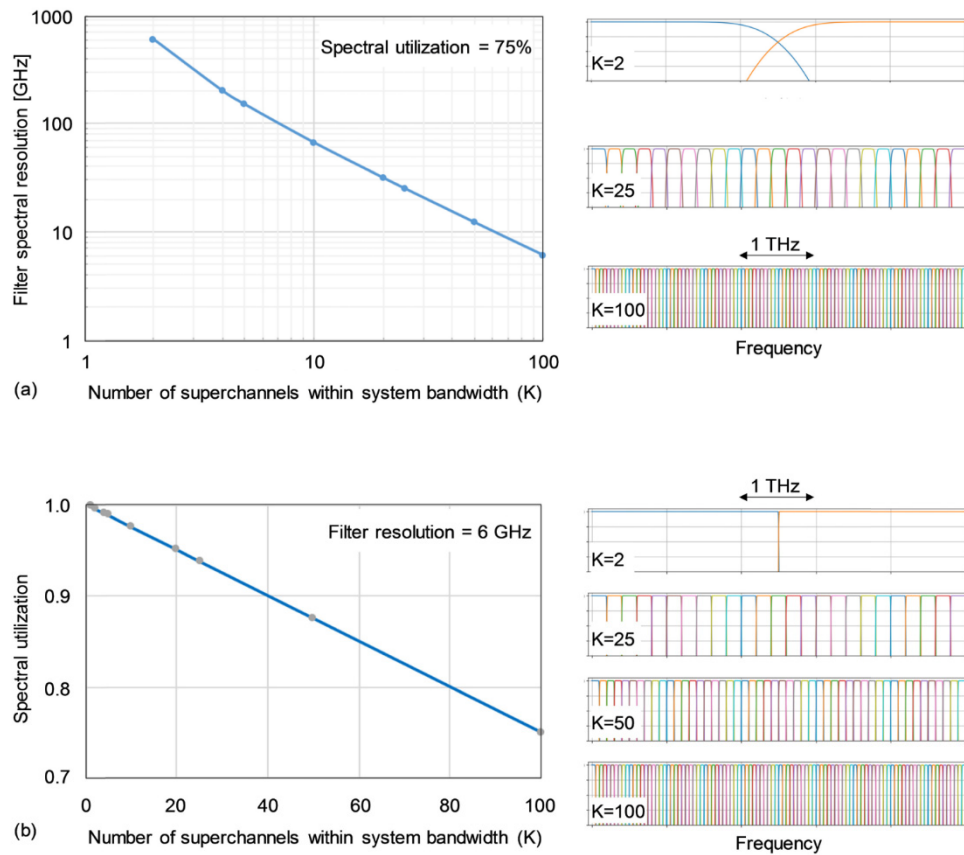


Fig. 24. (a) Required WSS spectral resolution for a constant spectral utilization of 75% as a function of the number of spectral superchannels within the system bandwidth. Example spectra are given for $K = 2, 25$ and 100 superchannels. (b) Spectral utilization as a function of the number of spectral superchannels with int system bandwidth for a constant WSS spectral resolution of 6 GHz. Example spectra are given for $K = 2, 25, 50$ and 100.

While *spatial superchannels* require high-resolution switching with the spectral resolution of a single unit slice, *spectral superchannels* can reduce the spectral resolution of the underlying system components through wider absolute guard bands at the same relative guard bandwidth. (In order to detect individual subcarriers *within* a superchannel, digital coherent detection can be used [177], but true subcarrier add/drop requires either steep optical filtering [226] or interferometric opto-electronic processing [227,228].) The implications of reducing the number of superchannels K across a given system bandwidth are visualized in Fig. 24 with spectral resolution and spectral utilization as parameters, the latter defined as the available spectrum within a 0.5-dB filter transmission bandwidth divided by the total system bandwidth. Figure 24(a) shows that the required spectral resolution at a constant spectral utilization scales inversely with $(K-1)$. As a consequence, W is reduced and the switch architecture is simplified. Alternatively, as illustrated in Fig. 24(b), if the spectral resolution of the ROADMs filtering components is kept constant, higher spectral utilization is obtained for a reduced number of superchannels.

In the limit where a spectral superchannel occupies the *entire system bandwidth* ($K = 1$), no spectral switching is required ($W = 1$) and the node-internal number of spatial paths reduces to just $(D + P)^2 M$. Once spectral superchannel systems evolve to this limit of *pure fiber switching*, several additional advantages can be expected on a system level: (i) switches do not need to be able to split wavelengths at all (except for subcarrier add/drop functions

performed at the edge of the network), leading to lower-loss and broader-band operation (e.g., photonic cross-connects (PXC) can cover 1270 nm – 1650 nm of bandwidth today); (ii) the SE of superchannels can then be increased to its maximum potential, as filtering guard bands are no longer required; (iii) fiber cuts affect entire amplifiers as opposed to fractions of an amplifier's bandwidth, making amplifier transients due to fiber cuts no longer an issue; (iv) the characteristics of a path are largely fixed once that path is configured, making the physical-layer performance independent of optical path re-configuration and consequently allowing for operation at a reduced (almost zero) margin as well as easier system automation.

To gain an understanding of the scale and devices required for a future ROADM, we consider the case of spectral superchannels in the year 2037 according to Tab. 2, assuming 625 fibers per nodal direction, and interfaces filling the entire system bandwidth (i.e., pure fiber switching). The switching architecture would then be entirely based on PXC, potentially with flexible wavelength multiplexing on the add/drop side, should optical subcarrier grooming be required, cf. Fig. 25. The switching architecture for the enormous space switch lends itself to a strictly non-blocking Clos network [229], where the first and third stages are switches associated with one of the ROADM directions, and the mid-stage contains multiple parallel cross-connects. (While there are other possible switch architectures, the Clos architecture has the benefit of keeping the number of spatial paths (i.e., fibers) between stages at twice the number of input fibers.) A degree-8 node in the year 2037 with 25% add/drop functionality will therefore need 6250 fibers. Implementing this directly in the Clos of Fig. 25 with $M = 625$, $D = 8$ and $P = 2$, leads to 625x1250 and 10x10 cross-connects for edge and center stage, respectively. While these are practical devices, it is worth noting that the Clos network allows for an internal resizing; for example, we can maintain 625 fibers per direction but increase the logical degree ($D + P$) by a factor of 5 (i.e., we partition the number of fibers per degree into 5 groups), so the effective M becomes 125, and the effective D and P become 40 and 10, respectively. This leads to 125x250 and 50x50 cross-connects for edge and center stages, which are moderate-sized cross-connects that are available even today.




Fig. 25. Possible future spatial switching node showing amplifier and dynamic gain equalization (DGE) on a per-spatial-link basis and a spatial photonic cross-connect (PXC) composed of a Clos network with the first and third stages associated with nodal degrees and central stages providing degree connectivity. Add/drop functionality uses extra spatial degrees and may retain wavelength switching to provide flexible subcarrier multiplexing.

3.6 Networking automation and control

In addition to the hardware aspects discussed above, we expect to see increasing levels of *network automation*, ultimately allowing “plug-and-play” capabilities devoid of any human intervention. In the physical layer, this will result in “zero touch” by humans, but also through artificial intelligence and machine learning will lead to “zero thought” network deployments and operations: The network, whose components will be automatically added by robots as needed, will autonomously provide the bandwidth-managed connections required for any service, without any human intervention or planning. (Importantly, none of this will result in increased network capacity but will help to reduce the cost of operating the network.)

In the context of full “zero-touch” network automation, it is important to realize that an autonomous network contains three essential functional ingredients: *(i)* sensors, *(ii)* actuators, and *(iii)* control agents. All three must act together to enable the desired functionalities. In the context of modern coherent networks, sensors can leverage the embedded functionality of coherent optical transponders, whose adaptive algorithms automatically learn the physical parameters of the network they operate on [230,231], or stand-alone sensing elements may be deployed [232]. From an optical physical-layer perspective, actuators are flexible line cards (adapting their rate-reach trade-off to accommodate varying transmission parameters) [8] as well as optical switches supporting such dynamism. Finally, in order to establish a “network brain”, universal abstractions with open interfaces are needed to allow SDN that brings together network elements across the networking stack and across diverse functionalities [233].

In parallel with network automation, we currently find a multitude of efforts to provide a *disaggregation* of network functions, which allows resource pools to be flexibly allocated to dynamically changing network requirements, even though disaggregation is often associated with a move to generic lower cost hardware (“*white-boxes*”) with an abstracted SDN control plane. Latest developments are summarized, e.g., in a recent Special Issue of the *Journal of Lightwave Technology* [234]. We note in this context that problems associated with disaggregation are far from being resolved. For example, maintaining proper network functionality and service reliability becomes more difficult when disaggregating analog systems (such as WDM systems and optical networks) as opposed to digital systems (such as IP routers and networks). These difficulties become even more pronounced when considering networking beyond trusted domains, a function that is performed in IP networks by the border gateway protocol (BGP), and whose extension to the physical layer remains an open problem. Hand-off of analog signals such as coherent optical channels to another SDN domain could degrade the performance of the channels transported within that domain and would be beyond the direct control of that domain. Here again, a move towards full-band spectral superchannels would be beneficial as purely spatial channels remain physically confined and can be more easily managed. Nonetheless we expect many inter-domain boundaries to continue to be implemented via optical-electrical-optical regeneration gateways as demarcation points. There will also likely be a need at some domain boundaries for sophisticated processing elements in the data plane working at line speed, which can classify and filter ingress traffic, a functionality currently embedded in IP routers and BGPs.

While we have primarily considered traffic growth in terms of capacity, changes in the *nature* of traffic are also likely to happen. As expectations towards networks transition from a best-effort thinking to a mission-critical high-reliability paradigm, we expect a growing importance of network robustness and of guarantees of data plane properties such as latency and jitter. The importance of such metrics may lead to a re-introduction of circuit-switching models with much stricter latency guarantees than packet-based networks.

Resolving these and similar aspects of network flexibility and autonomy in holistic, cross-layer way on the path towards a “zero-thought” network will be a topic of future research.

4. Conclusion

On the occasion of the 20th anniversary of Optics Express, we reviewed the evolution of fiber-optic communication systems over the previous 20 years and extrapolated their potential evolution over the coming 20 years. Well aware that our long-reaching views will be associated with significant uncertainties, we showed that network scalability over the next 20 years is in principle supported, both in terms of transport and optical switching, by the massive exploitation and integration of parallel spatial paths in the form of WDM x SDM matrix of unit cells as the only scalable solution. We analyzed both spatial and spectral superchannels and showed that spatial superchannels provide significant array integration benefits, including the digital compensation of spatial crosstalk. Spatial superchannels lend themselves more to point-to-point systems with dedicated fiber installations, such as DCI and submarine systems. Spectral superchannels, on the other hand, seem more beneficial for terrestrial mesh networking, which needs to exploit the diverse range of already installed fiber infrastructure as much as possible. In addition, the evolution of ROADMs in terrestrial mesh networks benefits from spectral superchannels, which will seamlessly merge into full-fiber spectral channels in a purely space-switched core, with the additional benefit of easier management and automation than today’s wavelength multiplexed channels. Both superchannel types will be enabled by fiber-in-fiber-out coherent optical engines that will holistically co-integrate opto-electronic modulation and detection arrays for client and line traffic with CMOS DSP.

Acknowledgments

We would like to acknowledge valuable data from and insightful discussions with C. Antonelli, S. Chandrasekhar, R.-J. Essiambre, D. Fishman, N. Fontaine, A. Mecozzi, A. Pilipetskii, R. Ryf, T. Sasaki, W. Thompson, R. Tkach, and S. Zsigmond.