roshmere

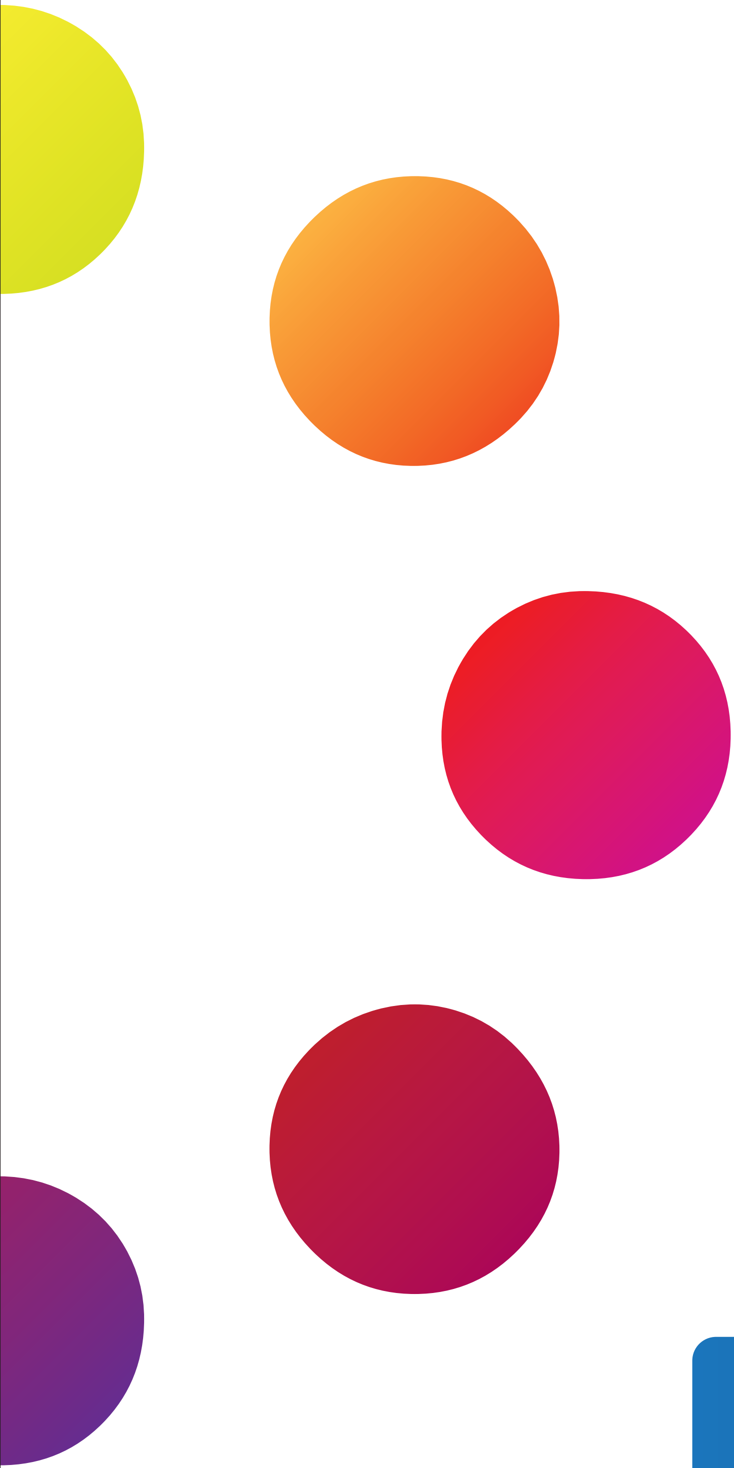

WHITE PAPER

Spearheading the Evolution of Lightwave Transmission Systems

Although the lightwave links envisioned as early as the 80s had ushered in coherent communication, the advent of optical amplifiers in the early 90s postponed the development of coherent transmission by more than a decade. Without an adequately developed data-driven commerce, physical bandwidth of the newly developed fibers and amplifiers, at the time, appeared nearly infinite, justifying simple, incoherent optical transceivers in place of complex coherent devices. Implausible in today's world, an OC-48 (2.5Gbps) channel, when combined with the conventional EDFA band, provided nearly inexhaustible 100Gbps aggregate throughput over a single fiber in the early 90s. By the mid-90s, with the introduction of OC-192 (10Gbps), Nortel Networks deployed band-scalable links that approached Tbps capacity, fully igniting the internet economy. Consequently, by the time of the first internet peak in 2001, the glut of deployed (dark) fiber provided little incentive to consider benefits of coherent modulation. Surprisingly, the very same (Nortel) team forged ahead to pioneer OC-768 digital signal processing (DSP) electronics that finally paved the way to modern coherent communication era. Unfortunately, while the Nortel engineering team was ahead of its time, its business leadership was not ready for the challenges of the early 2000s market environment.

Similar to the earlier, disruptive technology leaps, the introduction of the coherent lightwave modem has initiated a new cycle of network revolution. Transcending the original motivation (electronic dispersion compensation), the coherent transceiver is now seen as the foundation of a scalable network capacity across all connectivity domains. In conventional (terrestrial) links, coherent technology increases both the capacity and spectral efficiency, while simultaneously reducing the transmission cost; in submarine systems, coherent modulation decreases the repeater node count. However, a combination of the DSP cost and dissipation still poses a significant barrier to the deployment in metro and access domains, despite the obvious technical benefits that coherent technology enjoys over its incoherent counterpart.

The deployment barrier is most visible in networks driven by data centers (DC), where, in addition to the common constraints, the cost and scalability play the dominant role. Unlike the conventional lightwave link, whose cost can be amortized over a decade or more, data-center interconnects (DCI) must follow the server upgrade timeline conforming to a much shorter cycle. In contrast to the earlier DC infrastructure that focused on very large, remote facilities, rapid end-user connectivity growth now calls for a plethora of small and medium-size DCs built and operated within user proximity. Thus, the value of the new transmission node is measured by its ability to process and pass data within short-reach DC network. A server-driven upgrade of the intra-DC processing capacity also means that the associated connectivity must be scaled frequently, typically on a 3-to-5-year scale - nearly three times faster than the conventional telecommunication link. More importantly, the rapid-cycle capacity scaling of the DCI link must also be performed at a fraction of the traditional link cost: a decade-long amortization period is simply not supported within the new network architecture.

Not surprisingly, the rapid capacity scaling challenge has already been addressed previously - in a different industry. Indeed, the wireless industry was forced to rapidly scale its capacity within a strictly limited, highly regulated physical bandwidth. Therefore, as the end-user demanded increased connectivity, a rapid cycle of innovations resulted in modems operating on a complex, spectrally efficient channel. Present demand from the lightwave DC connectivity directly mirrors this challenge: finite fiber bandwidth, while orders of magnitude larger than the wireless window, also dictates rapid and agile scaling in channel spectral efficiency. Recognizing this need, fiber links have evolved from the original quaternary phase shift keying (QPSK) format from the last decade to the 64-QAM modulation expected in near-future deployments.

However, while the increase in lightwave modulation complexity retraces the path already covered by the radio-frequency communications industry, a significant set of challenges particular to lightwave systems makes its route significantly more perilous. In particular, the most fundamental challenge is posed by the physical difference between the lightwave and wireless channel carriers. While the relative RF/lightwave carrier uncertainty, defined as the ratio of the frequency uncertainty ("linewidth") and its mean, is comparable (~10-9), the respective absolute uncertainty is much larger. The wireless (RF) carrier is generated by a GHz-scale oscillator, while the lightwave channel rides on the emission of a laser centered at 193 THz – nearly five orders of magnitude difference. As a result, the wireless channel stability is measured in sub-Hz terms, while optical carrier frequency lies within ~100kHz of the oscillator (i.e. laser) linewidth. The latter, unfortunately, implies that the coherent lightwave receiver must process (demodulate) the information while contending with ~105 higher frequency uncertainty than its wireless counterpart. The challenge is easily visualized by comparing the laser linewidth requirement and the modulation complexity, as illustrated in Fig. 1. In practical terms, if one is not willing to pay for excessive computation overhead (dissipation) to process the carrier uncertainty, any capacity scaling in coherent links requires a significant increase in the carrier frequency stability.

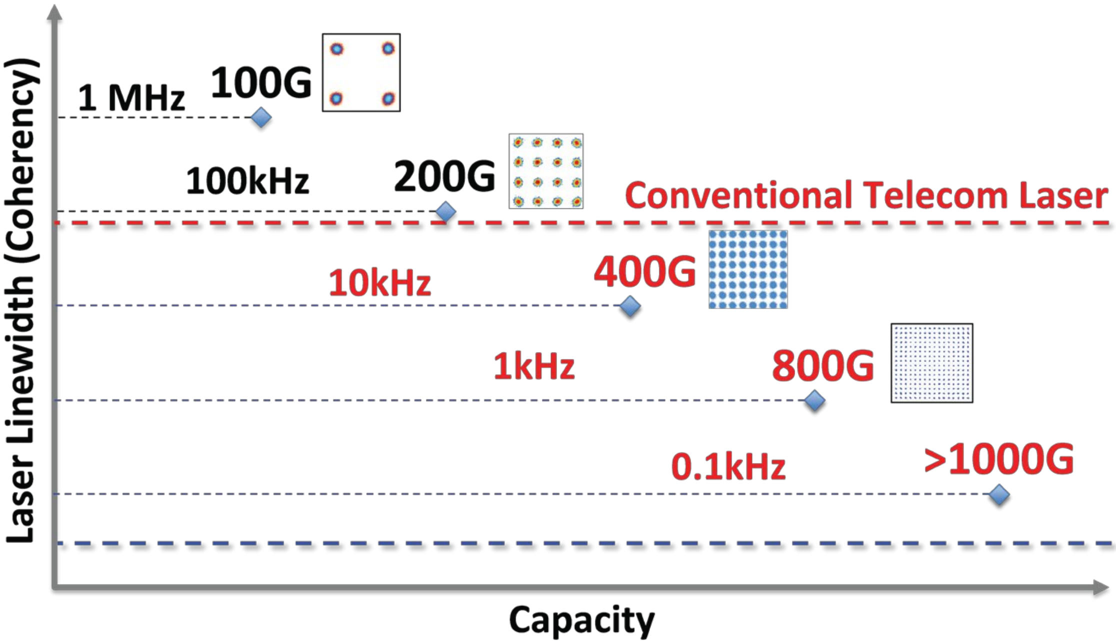

Fig. 1: *Conventional telecommunication laser linewidth (~100 kHz) provides natural path to channel constellation scaling up to 64QAM. While it is possible to envision a DSP solution to offset carrier uncertainty at this or higher format complexities, neither the overhead, nor the associated dissipation are desirable outcomes. In contrast, a 100 Hz-linewidth emitter provides a direct path to a single-carrier, Tbps-class channel that can be supported by low-dissipation DSP evolution.*

However, the increase in carrier stability is a necessary, but not sufficient condition that must be satisfied to follow the capacity path already demonstrated by the wireless industry. Even if the new laser technology were available, to fulfill the full potential of the offered spectral efficiency, the advanced oscillator would still need to be matched by ever-improving optoelectronics and digital signal processing (DSP) solutions. Furthermore, not only does the higher-complexity modulation format need to be mastered during each upgrade cycle, it is also imperative that it be economically viable. Consequently, the last requirement necessarily eliminates the traditional assemblies of discrete or mixed (discrete/monolithic) optical (E/O and O/E) modules with new DSP chipsets that used to be amortized over 10-15 year periods. Instead, in order to be able to follow the data-demand driven capacity scaling, illustrated in Fig. 1, the electro-optic and electronic development cycles must conform to the DC server upgrade cycle. While the computer industry has already demonstrated that high-complexity CMOS can follow the 3-year cycle, this has never been the case in the lightwave domain. Once developed, the traditional optical assembly has been too expensive to be replaced rapidly, generating both the economic and the engineering bottlenecks. Recognizing this basic limitation, Roshmere's strategy is straightforward: design and fabricate coherent transceivers that strictly conform to the CMOS development cycle and interface them with disruptive laser emitters. In simple terms, a new-generation coherent link must be assembled with optics that is treated no differently than the traditional CMOS chipsets – its performance must follow the rapidly ascending trend, whereas the cost must be sufficiently low enough to warrant its disposal and replacement during each upgrade cycle, as illustrated in Fig. 2.

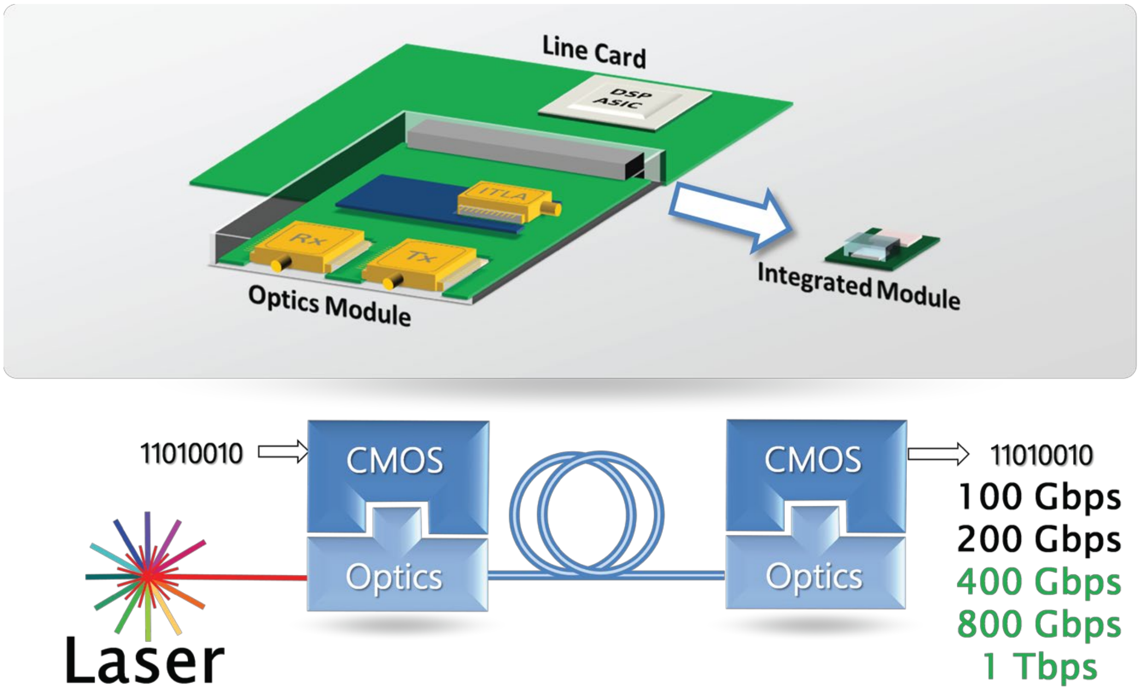

Fig. 2: Rapidly scaled coherent link: *in contrast to traditional (hybrid) transceiver assembly (top) that relies on dedicated optical modules and a DSP chipset, our solution rests on integrated optics and CMOS on a single silicon substrate. A 100 Hz-class emitter lights the silicon assembly, guaranteeing multiple-cycle capacity upgrades that can be initiated with 64QAM and be carried to a single carrier high order QAM channel according to the capacity demand.*