ciena
Experience. Outcomes.

WHITE PAPER

High-performance Coherent in CMOS: Foundational Technology for WDM Optical Networking

The next generation of semiconductor chip design for coherent optical transmission

Throughout 2018, the optical networking industry has ramped up deployments based on third-generation coherent technology (defined as coherent DSPs supporting 56 to 69Gbaud[1]). As new applications and market segments for coherent technology are emerging, the next two generations to hit the market will happen in a similar time scale rather than sequentially. One generation will be based on much higher baud for high-performance applications where maximum spectral efficiency is a key requirement, and the other on footprint-optimized solutions capped at 400 Gb/s to service short-reach edge applications such as 400ZR. For high-performance applications, the next generation of chips will push the capabilities of coherent transmission forward yet again, delivering greater capacity and reach over fiber optic connections. For footprint-optimized, short-reach applications, coherent chip design is also pushing boundaries by reducing the power and area needed to support up to 400 Gb/s transmission in a single device. By curtailing the reach, these solutions will support the implementation of coherent transmission technology in compact, pluggable form factors. Given the demanding requirements for both new generations, it is worth considering the foundational technology that underpins them and the design challenges faced in realizing these high-performance semiconductor devices.

A decade ago, the first coherent transceiver CMOS chips were deployed for 40 Gb/s systems, paving the way for 100 Gb/s coherent WDM technology deployments in 2010. Successive generations doubled the optical line rate transmitted per wavelength from a single coherent chip each time. Advancements in chip-level system design have helped drive the industry's ability to deploy higher fiber capacities and, in turn, lowered the cost/bit for data transport. As well, successive coherent chip generations from some vendors allowed network operators to build greater flexibility in provisioning services from their fiber plants. These devices delivered not just higher top-end capacity of the optical line rate, but also a broad range of selectable line rates from a single device. The options for scaling and improving cost efficiency of networks changed radically over the last year with the release of Ciena systems based on programmable third-generation coherent chips (WaveLogic Ai), providing the ability to tune capacity to address all applications with a single design.

1 "Third Generation Coherent Technology", Andrew Schmitt, Cignal AI, July 25, 2018

Highlights

- Optical networks continue to evolve with greater performance and flexibility based on a foundation of CMOS chip designs as the latest designs are heading to market in 2019
- The next generation of coherent chip designs will segment into two different paths, one targeting high performance and flexibility, and the other targeting low power and area
- New semiconductor design techniques, both to advance performance in spectral efficiency and circuit optimization for power efficiency, pave the way for these upcoming new generations of coherent DSP chips
- The path forward for future generations of coherent chip designs with respect to CMOS process technology must be carefully considered

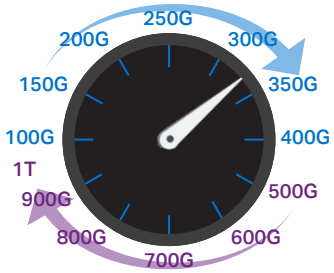

Figure 1. Programmable capacity

Coherent optical evolution
Learn more

Designing for spectral efficiency

The next generation of coherent chips will continue to rely on the use of advanced Digital Signal Processing (DSP), including the use of more complex modulation formats (more bits/symbol), allowing greater data transmission capacity (more bits/s) for a given signal processing rate (baud or symbol/sec) through to the optical fiber. Minimizing the increase in baud from generation to generation has allowed the use of lower-bandwidth photonics and sub-components while gaining greater data throughput. This helps to reduce chip design constraints and reduces system costs. For example, coherent CMOS devices have quadrupled in capacity, going from 100 Gb/s to 400 Gb/s, but the baud (signal processing rate) has only had to double, resulting in better spectral efficiency on the fiber, particularly at shorter reaches.

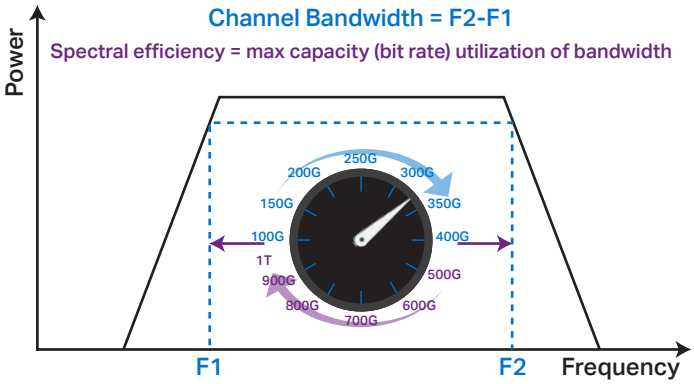

Figure 2. Spectral efficiency

As data rates per wavelength continue to increase, DSP techniques mitigate the impact of non-linearities in the photonic components connecting coherent chips to the fiber. In this case, the frequency response of the photonic components is usually well-characterized and non-linearities are mostly deterministic so the methods used to compensate are relatively straightforward. More recently, DSP techniques have been employed to mitigate dynamic non-linear responses due to transmission fiber medium such as cross-phase modulation and self-phase modulation. These mitigations are more complex in nature and require more rigorous approaches. Compensating for non-linearities is vital to extract fractions of dB and maximize system performance.

Other elements of the coherent chip will affect the spectral efficiency of the optical system. The high-performance, high-speed Analog-to-Digital Converters (ADCs) and Digital-to-Analog Converters (DACs) drive signals to the optical line side of the DSP converter to and from the digital logic at the heart of the chip. They interface with the analog electrical output of the photonic components, which, in turn, modulate and demodulate the optical signal on the fiber and convert from photonic to electronic signals. Enhancing the analog bandwidth of the ADCs and DACs mitigates noise and distortion inherent in all chip-level circuits and contributes to higher spectral efficiency of the optical system. In addition to maximizing transmission performance, chip-level digital and analog design considerations play a role in optimizing for power dissipation.

Designing for power efficiency

CMOS chips for high performance computing—such as machine learning or data analytics, or those for mobile platforms—have intermittent off-state functions. In these types of designs, techniques such as power gating or power islands can be used, effectively powering off parts of the chip while not in use or in standby mode. Typically, coherent optical devices are 100 percent in-service or 'always-on' devices. Since virtually all functions of these devices are continually operating, power gating may not inherently seem applicable. However, new coherent chip designs with multiple functions can also benefit from power gating. Other power-saving techniques such as clock gating and efficient clock routing can be effectively applied. Given the extremely high data throughput of coherent DSPs, this translates to high-frequency synchronization signals. Carefully managing the design and distribution of the clock circuits will help to optimize power efficiency.

And what does the future hold, beyond the imminent generations?

State-of the-art semiconductor design is '7nm' CMOS, which moved to volume production in 2018 from some CMOS foundries. Design implementation of chips is more challenging due to the finer geometries of the base transistor structures in 7nm and the constraints they impose on the physical layout of the chip. But the gains will be significant for coherent chips in 7nm. The near-term generations developed in this process node will benefit from higher transistor density and lower power dissipation compared to the previous node. Designs in 7nm will be supporting greater levels of functionality and higher data throughput in smaller areas. Power dissipation will also be lower compared to the same functionality in the previous process node.

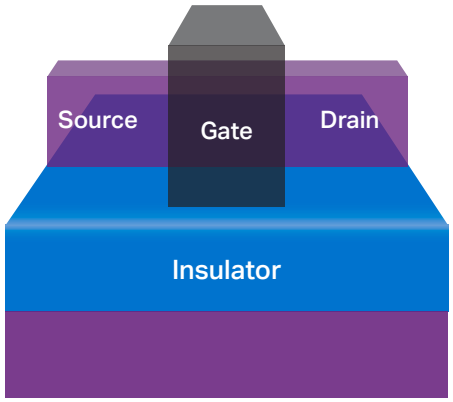

Figure 3. FinFET transistor

As the wider semiconductor industry moves forward to introduce the next process node, currently termed '5 nm,' the motivation for doing designs in this node, for any application, will once again be for greater density and functionality, or the same functionality in a much smaller chip area, with lower power. Gains in lower power dissipation with each successive CMOS node are reducing to an incremental amount while the cost increase for moving to the next node generally rises dramatically. For coherent optical DSP devices, chip developers must weigh the return of lower power against the even greater challenges for design robustness and significantly higher costs to develop in 5nm.

The Ciena Community
Get answers to your questions