North American Workshop on Silicon Photonics for High Performance Computing

May 17-18 2018, Fort Collins, CO

Silicon Photonic Modulation for High Performance Computing

Alan Mickelson
Electrical, Computer and Energy Engineering
University of Colorado at Boulder
Boulder 80309-0425 CO
alan.mickelson@colorado.edu

Abstract –

High Performance Computers will require orders of magnitude more low-latency inter-chip interconnection density than is presently available. Tera-Flop computing is assumed to require petabyte/s chip-to-chip interconnection networks. Even assuming that silicon photonics modulators soon reach the 100 Gbps/channel milestone, a petabyte/s on-substrate network will require a hundred thousand of such channels to achieve the petabyte/s. These channels can be parsed between into multiple space division channels (waveguides) each containing a multitude of wavelength division multiplexed (color) channels if the product of the number of waveguides and the number of colors equals or exceeds the required 100,000.

Silicon photonic waveguides are compact with less than a half-micron width. Channel to channel optical isolation requires no more than another half micron spacing between waveguides. A thousand channels/mm is then achievable in a silicon photonic (SiP) substrate. The limit on the number of wavelength channels is harder to determine. A laser diode per channel for a hundred thousand channels does not seem rational. Spectral slicing of broadband high luminosity sources is attractive for generating multiple color channels These channels can then be impressed with information in modulator arrays and re-multiplexed for waveguide transmission. De-multiplexing and optoelectronic conversion can take place in detector arrays. The solution is not new but has yet to find commercial application. The spectrally spliced solution is again being seriously considered for metro access and Li-Fi among other applications. Recent results are quite promising.

Crossbar interconnecting multiple (circa 30) accelerator chips that exhibit cut though bandwidths of 120 Tbps and bus rates in excess of 2.4 Tbps is an archetypical model problem for an HPC petabyte/s on-board optical interconnect. Dense arrays of modulators and detectors will be necessary in order to insert and remove the on-chip streams from chip-to-chip in-substrate optical network. Electrode crosstalk will be a serious issue but here attention will be placed on optical considerations, primarily modulation of spectrally sliced broadband sources.

Bio – please provide your bio.

Alan Mickelson is with the Electrical, Computer and Energy Engineering Department of the University of Colorado at Boulder. Alan joined the faculty at UCB in 1984 after a PhD at Caltech and PostDoctoral Fellowships in Armenia and Norway. Alan is actively engaged in research in optical materials, devices and sensors. Over the last decade, Alan has been involved in the modeling, design, fabrication and testing of silicon photonic devices as well as the study of optical interconnection of computing systems. Prof. Mickelson is involved in a company focused on commercialization of silicon photonics. Dr. Mickelson has published 2 textbooks, edited a research monograph and a conference proceeding and written 10 chapters in books. He has published more than 85 papers in the archival literature and more than 95 papers in reviewed conference proceedings.