Ultra-Low Latency Optical Networks for Next Generation Supercomputers

Benjamin A. Small, Odile Liboiron-Ladouceur, Assaf Shacham, and Keren Bergman
Department of Electrical Engineering, Columbia University in the City of New York

Carl Gray, Cory Hawkins, David C. Keezer, Kevin P. Martin, and D. Scott Wills
School of Electrical and Computer Engineering, Georgia Institute of Technology

Gary D. Hughes
Laboratory for Physical Sciences, University of Maryland

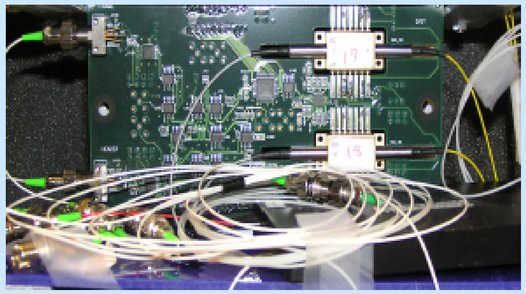

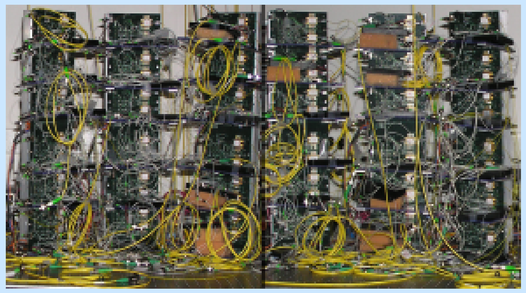

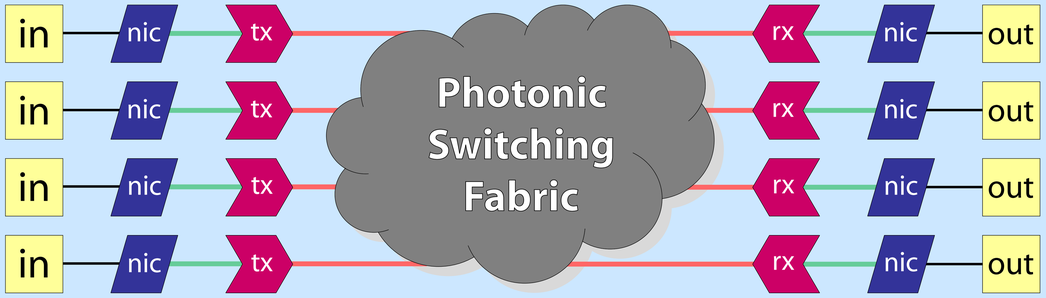

We have designed and implemented the components necessary to construct a complete optical packet switching (OPS) interconnection network. Interfaces have been developed for converting digital packets that emulate those from high-performance computing (HPC) processor and memory elements to optical-domain multiple-wavelength packets. Furthermore, a complete OPS interconnection network has been demonstrated as a prototype of the Data Vortex architecture to illustrate the feasibility and scalability of such OPS networks for high-bandwidth ultra-low latency applications. The photonic switching nodes, of which there are 36 in the 12×12 Data Vortex implementation, are based on semiconductor optical amplifier (SOA) switching elements. These optical devices are used as current-controlled switches which can transmit more than 1 Tbps of optically encoded data due to their relatively uniform wavelength response.

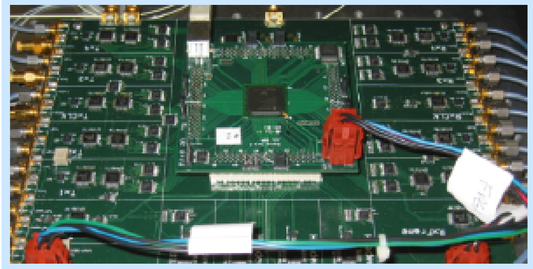

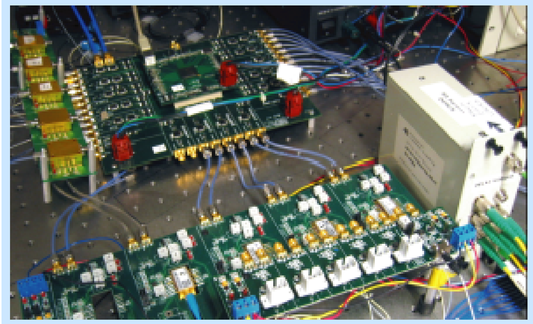

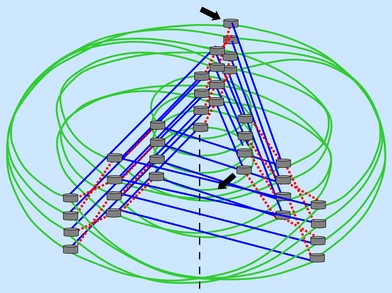

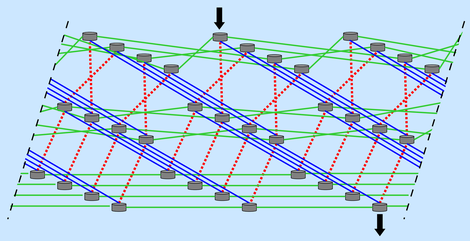

The Data Vortex topology relies on a banyan-like addressing structure while incorporating a unique deflection routing scheme as an internalized buffering mechanism, making it ideal for realization with photonic devices. The implemented 12×12 system is of size $(C,H,A) = (3,4,3)$, and the simulations below illustrate performance compared to a photonic butterfly architecture.

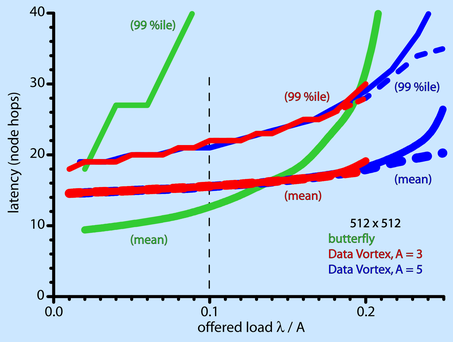

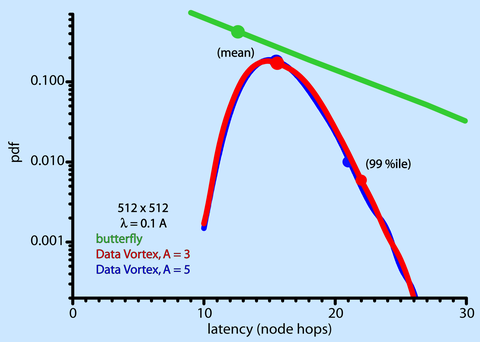

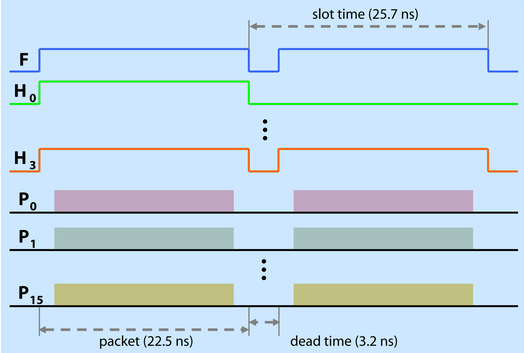

The Data Vortex architecture maximizes network throughput by utilizing a unique high-bandwidth multiple-wavelength packet format. Some wavelengths are designated for header information, and others are left for the high-speed packet payload. One wavelength is reserved for a payload clock.

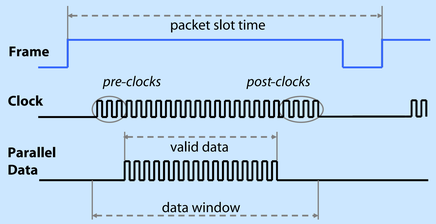

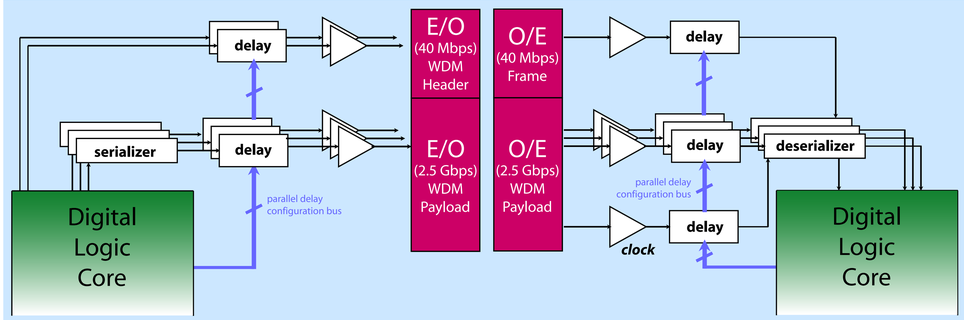

The Digital Logic Core generates packet headers and payloads which are then directed to optical transmitters, resulting in multiple-wavelength packets. These packets are injected into the OPS network, and upon ejection the wavelength-parallel payload is decoded by a series of receivers for processing. Tunable delay elements compensate for possible path length nonuniformities.