MIT LINCOLN LABORATORY

The Bulletin

News, events, and notices transmitted weekly

February 10, 2017

Graph Processor Prototype

Technology Office| Lincoln Laboratory

Financial transactions, social networking, and Internet traffic can all be represented as sparse, highly interconnected graphs. Algorithms can analyze these graphs to detect cyber network intrusions, terrorist social networks, and traffic patterns that indicate a hostile environment. While graph algorithms work well with moderate-size databases, even the fastest modern processors have difficulty providing sufficient throughput for very large databases (billions to trillions of nodes). The difficulty is random memory access and computation dealing with indices, vertices, and edges of the large graphs that soak up a large fraction of the processor resources. Consequently, graph computation can be slowed by several orders of magnitude.

In order to achieve significantly better graph computation performance, an advanced multiprocessor architecture has been developed that is optimized for analysis of large databases. Most graph operations can be implemented as sparse matrix operations. This processor, which is optimized for sparse matrix algebra operations, can be hundreds of times more efficient than conventional processors. To address the memory access and inter-processor communication challenges, the new graph processor architecture does not use cache memory, but rather a specialized network architecture to connect thousands of processing nodes. Each processor node is designed to be capable of a communication rate of over one terabit per second. A specialized systolic sorter module handles most of the overhead computation associated with graph indices, vertices, and edges.

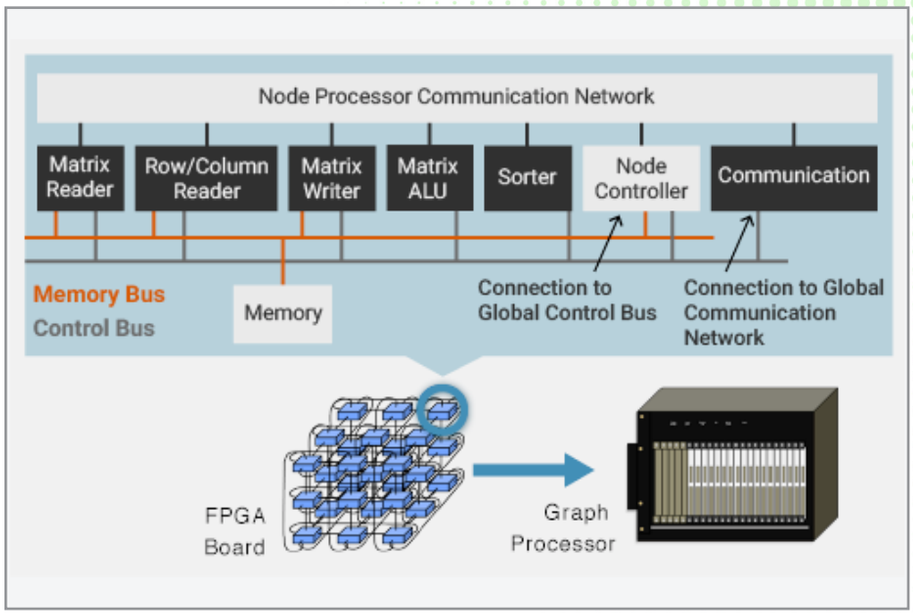

Advanced graph processor prototype development.

An 8-node graph processor prototype using commercial field-programmable gate array (FPGA) boards demonstrated a factor of ten times better performance than commercial processors. The effort is under way to demonstrate even higher performance with 64-node prototype systems using custom FPGA boards, as the performance advantage increases with system size.

For more information, please contact Dr. William Song, Senior Staff, Embedded and Open Systems, Group 102.

MIT LINCOLN LABORATORY SUPERCOMPUTING CENTER