DATAVORTEX TECHNOLOGIES

A Data Vortex® Network for Large Scale Graph Analytics

The Data Vortex® IP is a parameterizable network topology that provides a ***self-routing network*** for ***very small messages*** at ***very low latency***. It can support memory semantic load/stores and can be implemented as a network on chip or on board for interconnecting heterogeneous compute elements and memory.

As a system level interconnect, it can ***scale almost linearly to thousands of end points*** in a ***flat***, ***high-bandwidth*** topology. Optimum performance is achieved with packet sizes of 8 bytes to a full cache line.

For large scale graph analytics, Data Vortex Technologies is implementing the network as a dedicated accelerator interconnect to allow accelerators with optimized graph libraries to analyze very large graphs and scale system performance at very high efficiency.

Contact Information:
John LaBry (Chief Commercial Officer)
john.labry@datavortex.com

Specific benefits of a Data Vortex® network for large graph analytics that are not currently achieved by any standard network technology

With a Data Vortex® enabled accelerator network, one can double the performance of the system with each doubling of the processing nodes if analyzing a graph of a specific consistent size.

OR

The Data Vortex® network allows one to double the size of the graph with each doubling of the processing nodes while maintaining the same performance.

ALSO

In large graph analysis, as one increases the number of edge traversals (hops) in the problem to be solved, performance dramatically degrades, and systems can even stall due to the exponential increase in the volume of small messages that must move through traditional networks.

The unique design of the Data Vortex® network for small packet delivery, very low latency, collision free data movement and flat, near linear scalability, prevent that problem from occurring. The performance differential over standard technologies can even increase as the number of hops increases.

Reasons Why the Data Vortex® Network Efficiently Enables Analysis of Very Large Graphs

Data Vortex® Advantages For Graph Analytics.	Inherent arbitration (dynamic self-routing) Uniform low latency (at any scale) Fine grained (very small message packets) Highly scalable and flat Messages never collide or get dropped High bandwidth and high radix
Network Interfaces And Protocols	Accelerator interfaces and protocols allow for easier optimization of message transmission via the Data Vortex® Network Adapter to the Data Vortex® Switch. Messages do not need to be aggregated in order to achieve scaling and performance. Individual small message packets are sent in parallel across all QSFP lanes going into the switch.
Network Congestion Is Nearly Eliminated	Because the Data Vortex® does not have a central arbiter that sets the switch and directs packets, if many nodes simultaneously send packets to the same destination node, that path gets momentary back pressure but will not block any traffic between any other ports on the network. All packets in the switch dynamically self-route around any high traffic connections maintaining a very high throughput.
Scaling the Graph and the Efficiency of the System Performance	FPGA accelerator cards, in single node configurations hosting graph libraries, have proven to significantly outperform general purpose host and server systems. The Data Vortex® provides these accelerators the ability to efficiently connect on a dedicated network scaling to hundreds and eventually thousands of devices. A message passing or shared memory environment can be implemented with this network.
Very Large Graphs	The Data Vortex® network is designed to scale to thousands of nodes. The size of the graph in a Data Vortex® enabled system will be determined by the amount of memory on the accelerator cards. Accelerator card roadmaps show HBM memory increasing to 64GB and perhaps even 128GB. Servers can support 6 to 8 cards. Multiple servers with multiple cards can be implemented in a single network.