

DRC GraphFind™

Product Brief



DRC's GraphFind rapidly discovers relationships between people, places, events and objects. Simultaneously identifying focal points with weighted strengths of connections. Executed using hardware acceleration hundreds of millions of nodes and edges can be analyzed in seconds.



DRC GraphFind Engine

Background. Large graphs having hundreds of millions of nodes (nodes are entities - people, places, events, objects, etc) and billions of edges (connections between entities) are finding increasing use in diverse application domains in both the public and private sectors. Applications include epidemiology, network routing, logistics operations, social marketing, military, healthcare provider scheduling, power systems management, financial services, fraud detection, finite element analysis, and many others.

Two of the most commonly needed tools are determination of the *shortest path* between node pairs and the *betweenness centrality* of a given node. However, these data intensive operations frequently demand resources beyond all but the largest high-performance clusters. This is particularly true in graph applications that cannot have partitions removed to minimize the search scope.

Solution

DRC GraphFind combines embarrassingly parallel FPGAs with the DRC software/firmware suite resulting in an easy to use, ultra-fast system that can handle graphs having up to 2^{32} nodes and an unlimited number of asymmetrically weighted, directed edges, including hypergraph support.

Shortest Path & Betweenness Centrality

The DRC GraphFind system concurrently solves for shortest path and betweenness centrality maximizing run-time efficiency and system throughput. The shortest path algorithm determines the shortest path between specified node-pairs or between all node-pairs in the input graph partition. The Betweenness Centrality algorithm scores every node in the graph partition for its network centrality—its relative influence on information flow in the graph. Combined, these tools form the backbone for analytics in: social networks, genetics, transportation, artificial intelligence, many-body physics, and other network-centric data-intensive domains.

Single Engine or Appliance

The DRC GraphFind system is configured into an Accelium™ accelerator (a PCIe plug-in board) for integration into an application server or storage system. Also available as an in-line appliance; a plurality of accelerators (depending on the specified chassis) may be installed making DRC GraphFind easily scalable to match the requirements for both economy and throughput.

Storage Integrated

Installing the DRC GraphFind engine in the system that generates or stores graph data ensures that graph analysis is performed as closely as possible to the data source for maximum performance and minimized communications latency.

The Hardware is the Algorithm

The DRC GraphFind algorithms are instanced as FPGA logic blocks—it is hardware, not software. This provides for massive parallelization with up to 96 graph evaluators (depending on architecture) instantiated in each accelerator.

Hierarchical pre-partitioning of the input graph into manageable pGraph (partitioned graph) objects combined with the ability to use both spatial and temporal parallelism allows DRC GraphFind to use adjacency lists rather than adjacency matrices thereby minimizing memory demands without sacrificing throughput and accuracy.

The DRC GraphFind system can either be user interface driven or easily integrated into an application using a simple, robust, and flexible API callable from any language supporting the use of Linux shared object libraries.

Flexible Architecture

Unlike many FPGA-based accelerators, DRC GraphFind provides for build-time flexibility in the FPGA architecture. The default evaluator handles 512 nodes and 4096 edges in each of 32 evaluators. Recent variants have included an 8,000 node, 64,000 edge evaluator. This architecture can be adjusted, at the individual evaluator level, to satisfy specific user requirements—a performance enhancing benefit.

The DRC Difference

Over 200 man-years of experience developing low latency, high capacity solutions gives DRC unique talent in big data applications.

DRC's novel architecture optimizes the data management/analysis balance. The key to ultra-high performance is distributing processing capacity making it available where it's needed rather than centralizing it. DRC moves the processing to the data versus the data to the processor.

Made in America

DRC is a US small business with all staff US citizens, and all engineering and products produced in the USA.



DRC GraphFind Appliance

DRC COMPUTER CORP

3375 Scott Blvd, Suite 206
Santa Clara, CA 95054

PHONE

+1. 408.562.0000

WEB

drccomputer.com