Storing and Searching Massive Scale-free Graphs



Timothy D. R. Hartley, Umit V. Catalyurek

Depts. of Biomedical Informatics, Electrical and Computer Engineering The Ohio State University

Motivation

MSSG is a high-performance parallel, out-of-core graph database, capable of storing and guerying graphs with billions of edges.

Scale-free graphs

- Metabolic pathways
- Gene regulatory networks
- Taxonomies of chemical compounds
- Social networks
- Unique properties
- Scale-free nature of interconnections
- Hubs highly connected vertices



Scale-Free Graph Edge Degree Distribution

Challenges

Aggregate Disk I/O

- In order to aggregate disk I/O, the graph data must be spread across more than one physical disk.
- Partitioning/declustering of input graph data is extremely important, since massive scale-free graphs can have very highly-connected vertices at their core.
- Cache Locality
- Temporal locality
- Only a small portion of the input graph will fit into main memory.
- Spatial locality Amortize the I/O time.

System Architecture



Ingestion Service:

- Partitioning / declustering occurs during ingestion
 - The input graph's vertices or edges are assigned to nodes in a streaming fashion. with a configurable window size.

GraphDB Service:

 Simple interface to hide DB-specific complexity

Query Service:

- Orchestrates actions needed to answer user aueries
- Multiple implementation opportunities BFS
- Best-first search
- Pattern search
- Neighborhood guality guantification



Scale-Free Graph Subset

Andv Yoo Center for Applied Scientific Computing Lawrence Livermore National Laboratory

GrDB: Graph Database

GrDB is a purpose-built, external-memory graph data structure which allows MSSG to leverage cache locality into improved average graph access speed.

- GrDB
- Designed to efficiently store streaming scale-free graphs
- Provides on-demand access to the stored graph via graph gueries
- Achieves a practical compromise between high-performance dynamiclength and easy to use, fixed-length records

GrDB Example





GrDB is able to store scale-free graphs in a streaming manner with a minimum of disk overhead.

Performance: Graph Analysis

- 64 dual 2.4GHz AMD Opteron 250 Cluster
- 8 GB RAM per node
- 500 GB local disks in RAID 0 per node
- Infiniband
- Breadth-first search
- Important graph analysis kernel
- Community analysis
- Betweenness centrality
- Bidirectional BFS Fringe sizes untenable in unidirectional





GrDB is comparable to the highperformance, but impractical compressed adjacency-list format.



MSSG can finish 90% of searches in a 120 billion edge graph in under 60 seconds.

Department of Biomedical Informatics