High-Performance Graph Mining

Goal: construct a high--performance algorithm solving a selected graph mining problem

Different symbols indicate which elements of GMS are responsible for a given part of the construction process of a graph mining algorithm

Part 1: Design

Key questions:

- **S** What are **relevant** mining **algorithms** and **datasets**?
- **C** How to **assess the scalability** of a new algorithmic idea?

Part 2: Implementation & tuning

Key questions:

- **I** **P** How to **quickly benchmark** new parallel graph mining algorithms, preprocessing schemes, data layouts, various optimizations?
- **I** **P** How to **effectively use** different **parallel architectures**?

Part 3: Analysis

Key questions:

- **S** What are **state-of-the-art comparison baselines**?
- **P** **C** How to **analyze the performance**, storage requirements, and other aspects of a new algorithm?

Part 4: Evaluation

Key questions:

- **M** What are **insightful performance metrics for graph mining**?
- **I** How to **effectively evaluate algorithms**?

Challenges & questions

Solutions & answers

GraphMine Suite

Benchmark specification

Details: Section 4 **S**

Graph problems & algorithms

- ➜ **Pattern matching** (e.g., clique listing)
- ➜ **Learning** (e.g., link prediction, clustering)
- ➜ **Optimization** (e.g., coloring, minimum cuts)
- ➜ **Reordering** (e.g., degeneracy reordering)

Datasets

- ➜ **Sparse** & **dense**, ➜ many & few **cliques**,
- ➜ High & low **skew** of degree distribution,
- ➜ Many & few **dense** (non-clique) **subgraphs**,
- ➜ different **origins** (purchases, roads, ...)

Reference implementations

Details: Section 5 **I**

Implementations

- ➜ Algorithms,
- ➜ Optimizations,
- ➜ Preprocessing routines,
- ➜ Load balancing,
- ➜ Graph representations,
- ➜ Data layouts,
- ➜ Graph compression,
- ➜ Parallelizations

Features

- ➜ Parallel, ➜ Modular,
- ➜ Scalable, ➜ Fast, ➜ ...

Implemented in ➜ **Benchmarking platform** ➜ **Used by**

Benchmarking platform

Details: Sections 3 & 5 **P**

Features

- ➜ Simple to use,
- ➜ Extensible,
- ➜ Modular,
- ➜ Public.

Key idea for high modularity: use set algebra. Sets and set operations become "modules" that can be implemented in different ways, and still they can be seamlessly combined.

Performance metrics

Details: Sections 5 & 7 **M**

Traditional

- ➜ Run-time, ➜ Scalability,
- ➜ L3 misses (machine efficiency).

Key idea in a novel metric: count the number of graph patterns mined per second (algorithmic efficiency).

Concurrency analysis

Details: Section 6 **C**

Aspects

- ➜ Performance (work, depth),
- ➜ Storage, ➜ Tradeoffs.

Platform pipeline stages (toolchain execution) with details on **extensibility** and **modularity**

: a dark background and a cube indicate that a particular part of the design can be substituted by the developer with their own implementation

Load graph into memory

Input graph

1 Build graph representation

Pointers

By default, we use **CSR**

A representation is modular: the user can provide a new representation

Neighborhoods of vertices

2 Define graph accesses

When developing a graph representation, the user also develops the corresponding graph accesses:

Graph accesses

- check d(v)
- iterate over *N(v)*
- check if ∃ *(u,v)*

3 Apply preprocessing

Example: initial CSR graph representation

Example: reordered CSR (degree order: by neighborhood size)

The user can plug in different preprocessing schemes. We provide a **ready library of reordering schemes**, such as degeneracy or degree reordering (example above).

4 Run graph algorithm

The user can plug in a graph algorithm. We offer >40 reference implementations.

5 Define algorithm building blocks

/* **Example**: Triangle Counting. "**tc**" is the count of triangles */

The user can plug in variants of fine algorithm blocks such as scheduling policies. GMS facilitates it with appropriate modular implementations

```
tc = 0; init_sets( )
#pragma omp parallel for schedule (...)
for v in V:
  for w in N(v):
    tc += |N(v) ∩ N(w)|
tc /= 3; cleanup( )
```

Most simplicity is enabled by using fine building blocks based on **set algebra** **5+**

Gather data

Visualize

How does GMS facilitate extensibility at a given stage?

1. Modular design of classes & files associated with graph representations
2. Well-defined interface (based on set algebra) of routines for graph accesses
3. Enabling running different preprocessing routines with a single function call
4. Modular design of classes & files associated with graph algorithms
5. Clear structure of code facilitating manipulation with fine parts such as scheduling policy of single loops
5+. **Set algebra** based modularity for various parts of algorithms

The user can experiment with **algorithmic** ideas (e.g., new algorithms or data structures), **architectural** ideas (e.g., using SIMD or instrinsics), and **design** ideas (e.g., using novel form of load balancing).