

# GraphMineSuite: Enabling High-Performance and Programmable Graph Mining Algorithms with Set Algebra

Maciej Besta<sup>1\*</sup>, Zur Vonarburg-Shmaria<sup>1</sup>, Yannick Schaffner<sup>1</sup>, Leonardo Schwarz<sup>1</sup>,  
Grzegorz Kwasniewski<sup>1</sup>, Lukas Gianinazzi<sup>1</sup>, Jakub Beranek<sup>2</sup>, Kacper Janda<sup>3</sup>, Tobias Holenstein<sup>1</sup>,  
Sebastian Leisinger<sup>1</sup>, Peter Tatkowski<sup>1</sup>, Esref Ozdemir<sup>1</sup>, Adrian Balla<sup>1</sup>, Marcin Copik<sup>1</sup>,  
Philipp Lindenberger<sup>1</sup>, Marek Konieczny<sup>3</sup>, Onur Mutlu<sup>1</sup>, Torsten Hoefler<sup>1\*</sup>

<sup>1</sup>ETH Zurich, Zurich, Switzerland; <sup>2</sup>VSB, Ostrava, Czech Republic; <sup>3</sup>AGH-UST, Krakow, Poland; \*Corresponding authors

# What is graph mining?

- Graph mining is the process of finding and extracting useful information from graphs, i.e. sssp, triangle counting, k-cliques, maximal cliques, etc.
- Many real world applications: social sciences, bioinformatics, chemistry, medicine, cybersecurity, and many others
- Issue #1: graphs can be very large and require a lot of compute power
- Solution #1: Parallelism!
- Issue #2: Too many choices!
  - Hard to keep up and find relevant baseline graph mining algorithms to improve upon, a plethora of relevant datasets, numerous design choices
- Solution #2: **GraphMineSuite (GMS)** - a benchmarking suite for high-performance graph mining algorithms.

# GraphMine Suite



## Benchmark specification

### 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, ...)

Details: Section 4 **S**

## 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

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

### Traditional

Details: Sections 5 & 7 **M**

- 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

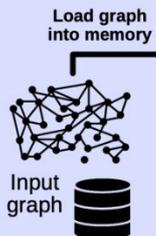
### Aspects

Details: Section 6 **C**

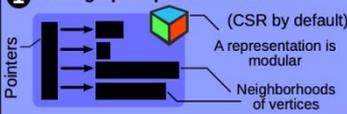
- Performance (work, depth),
- Storage, → Tradeoffs.

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

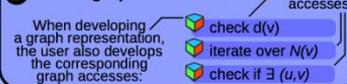
: a dark background and a cube indicate that a particular part of the design can be substituted by the developer with their own implementation



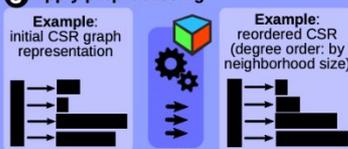
### 1 Build graph representation



### 2 Define graph accesses



### 3 Apply preprocessing



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 */
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()
```

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

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



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 intrinsics), and **design** ideas (e.g., using novel form of load balancing).

# Benchmark specification: Graph Problems

	Graph problem	Corresponding algorithms	E.?	P.?	Why included, what represents? (selected remarks)
Graph Pattern Matching	• Maximal Clique Listing [48]	Bron-Kerbosch [24] + optimizations (e.g., pivoting) [29, 51, 117]	👍 5+	🚫	Widely used, NP-complete, example of backtracking
	• $k$ -Clique Listing [41]	Edge-Parallel and Vertex-Parallel general algorithms [41], different variants of Triangle Counting [104, 107]	👍 5+	🚫	P (high-degree polynomial), example of backtracking
	• Dense Subgraph Discovery [5]	Listing $k$ -clique-stars [63] and $k$ -cores [54] (exact & approximate) VF2 [40], TurboISO [58], Glasgow [89], VF3 [26, 28], VF3-Light [27] BFS and DFS exploration strategies, different isomorphism kernels	👍 5+	🚫	Different relaxations of clique mining Induced vs. non-induced, and backtracking vs. indexing schemes Useful when one is interested in many different motifs
	• Subgraph isomorphism [48]		👍	🚫	
• Frequent Subgraph Mining [5]	👍		🚫		
Graph Learning	• Vertex similarity [75]	Jaccard, Overlap, Adamic Adar, Resource Allocation, Common Neighbors, Preferential Attachment, Total Neighbors [101]	👍 5+	🚫	A building block of many more complex schemes, different methods have different performance properties
	• Link Prediction [114]	Variants based on vertex similarity (see above) [7, 80, 83, 114], a scheme for assessing link prediction accuracy [121]	👍 5+	🚫	A very common problem in social network analysis
	• Clustering [103]	Jarvis-Patrick clustering [65] based on different vertex similarity measures (see above) [7, 80, 83, 114]	👍 5+	🚫	A very common problem in general data mining; the selected scheme is an example of overlapping and single-level clustering
	• Community detection	Label Propagation and Louvain Method [108]	👍	🚫	Examples of convergence-based on non-overlapping clustering
Vertex Ordering	• Degree reordering	A straightforward integer parallel sort Computing triangle counts per vertex Exact and approximate [54] [70]	👍	👍	A simple scheme that was shown to bring speedups
• Triangle count ranking	👍 5+		👍	Ranking vertices based on their clustering coefficient	
• Degeneracy reordering	👍 5+		👍	Often used to accelerate Bron-Kerbosch and others	

**Table 3: Graph problems/algorithms considered in GMS. “E.?” (Extensibility)” indicates how extensible given implementations are in the GMS benchmarking platform: “👍” indicates full extensibility, including the possibility to provide new building blocks based on set algebra (1 – 5, 5+). “👍”: an algorithm that does not straightforwardly (or extensively) use set algebra. “P.?” (Preprocessing) indicates if a given algorithm can be seamlessly used as a preprocessing routine; in the current GMS version, this feature is reserved for vertex reordering.**

# Set Algebra

- Many graph algorithms are/can be formulated with set algebra
- GMS allows users to implement their own sets, set operations, set elements, and set algebra based graph representations.
- Allows users to break complex graph mining algorithms into simple building blocks, and work on these building blocks independently.

```
1 class Set {
2 public:
3 //In methods below, we denote "*this" pointer with A
4 //(1) Set algebra methods:
5 Set diff(const Set &B) const; //Return a new set  $C = A \setminus B$ 
6 Set diff(SetElement b) const; //Return a new set  $C = A \setminus \{b\}$ 
7 void diff_inplace(const Set &B); //Update  $A = A \setminus B$ 
8 void diff_inplace(SetElement b); //Update  $A = A \setminus \{b\}$ 
9 Set intersect(const Set &B) const; //Return a new set  $C = A \cap B$ 
10 size_t intersect_count(const Set &B) const; //Return  $|A \cap B|$ 
11 void intersect_inplace(const Set &B); //Update  $A = A \cap B$ 
12 Set union(const Set &B) const; //Return a new set  $C = A \cup B$ 
13 Set union(SetElement b) const; //Return a new set  $C = A \cup \{b\}$ 
14 Set union_count(const Set &B) const; //Return  $|A \cup B|$ 
15 void union_inplace(const Set &B); //Update  $A = A \cup B$ 
16 void union_inplace(SetElement b); //Update  $A = A \cup \{b\}$ 
17 bool contains(SetElement b) const; //Return  $b \in A$  ? true:false
18 void add(SetElement b); //Update  $A = A \cup \{b\}$ 
19 void remove(SetElement b); //Update  $A = A \setminus \{b\}$ 
20 size_t cardinality() const; //Return set's cardinality
21 //(2) Constructors (selected):
22 Set(const SetElement *start, size_t count); //From an array
23 Set(); Set(Set &&); //Default and Move constructors
24 Set(SetElement); //Constructor of a single-element set
25 static Set Range(int bound); //Create set  $\{0, 1, \dots, bound - 1\}$ 
26 //(3) Other methods:
27 begin() const; //Return iterators to set's start
28 end() const; //Return iterators to set's end
29 Set clone() const; //Return a copy of the set
30 void toArray(int32_t *array) const; //Convert set to array
31 operator==; operator!=; //Set equality/inequality comparison
32
33 private:
34 using SetElement = GMS::NodeId; //(4) Define a set element
35 }
```

Algorithm 1: The set algebra interface provided by GMS.

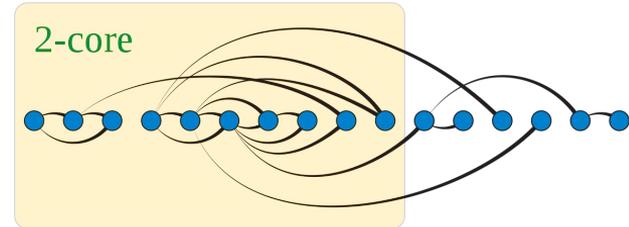
# GMS Set Implementations

GMS offers three default set implementations:

- **RoaringSet**
  - Implemented with a “roaring bitmap” that allows for mild compression rates but inexpensive decompression
- **SortedSet**
- **HashSet**

# Use Case # 1: Degeneracy Order & k-Cores

- The **degeneracy** of a graph  $G$  is the smallest  $d$  such that every subgraph in  $G$  has a vertex of degree at most  $d$ .
  - A measure of graph sparsity
- A **degeneracy ordering (DGR)** is an ordering of vertices of  $G$  such that each vertex has  $d$  or fewer neighbors that come later in this ordering
  - DGR can be obtained by repeatedly removing a vertex of minimum degree in a graph.
- A **k-core** of  $G$  is a maximal connected subgraph of  $G$  whose all vertices have degree at least  $k$ .
  - A  $k$ -core can be obtained by iterating over vertices in the DGR order, and removing vertices with out-degree less than  $k$



# Use Case # 1: Degeneracy Order & k-Cores cont.

- Issue: Not easily parallelizable,  $O(n)$  iterations!
- Solution: GMS offers a  $(2+\epsilon)$ -approximate degeneracy order (ADG),  $O(\log n)$  iterations for any  $\epsilon > 0$ !

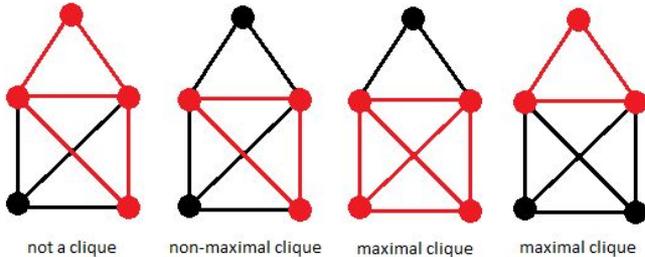
---

```
1 //Input: A graph  $G$  1. Output: Approx. degeneracy order (ADG)  $\eta$ .
2  $i = 1$  // Iteration counter
3  $U = V$  //  $U$  is the induced subgraph used in each iteration  $i$ 
4 while  $U \neq \emptyset$  do:
5    $\widehat{\delta}_U = \left( \sum_{v \in U} |N_U(v)| \right) / |U|$  //Get the average degree in  $U$  2
6   //  $R$  contains vertices assigned priority in this iteration:
7    $R = \{v \in U : |N_U(v)| \leq (1 + \epsilon)\widehat{\delta}_U\}$  2
8   for  $v \in R$  in parallel 2 5 do:  $\eta(v) = i$  //assign the ADG order
9    $U = U \setminus R$  5+ //Remove assigned vertices
10   $i = i+1$ 
```

---

# Use Case # 2: Maximal Clique Listing

- A maximal clique of a graph  $G$  is a fully connected subgraph of  $G$  that cannot be further extended by including one more adjacent vertex.



```
1 /* Input: A graph  $G$  1. Output: all maximal cliques. */
2
3 //Preprocessing: reorder vertices with DGR or ADG.
4  $(v_1, v_2, \dots, v_n) = \text{preprocess}(V, /* \text{selected vertex order} */) 3$ 
5
6 //Main part: conduct the actual clique enumeration.
7 for  $v_i \in (v_1, v_2, \dots, v_n)$  do: //Iterate over  $V$  in a specified order
8 //For each vertex  $v_i$ , find maximal cliques containing  $v_i$ .
9 //First, remove unnecessary vertices from  $P$  (candidates
10 //to be included in a clique) and  $X$  (vertices definitely
11 //not being in a clique) by intersecting  $N(v_i)$  with vertices
12 //that follow and precede  $v_i$  in the applied order.
13  $P = N(v_i) \cap \{v_{i+1}, \dots, v_n\} 5+$ ;  $X = N(v_i) \cap \{v_1, \dots, v_{i-1}\} 5+$ ;  $R = \{v_i\}$ 
14
15 //Run the Bron-Kerbosch routine recursively for  $P$  and  $X$ .
16 BK-Pivot( $P, \{v_i\}, X$ )
17
18 BK-Pivot( $P, R, X$ ) //Definition of the recursive BK scheme
19 if  $P \cup X == \emptyset 5+$ : Output  $R$  as a maximal clique
20  $u = \text{pivot}(P \cup X) 5+$  //Choose a "pivot" vertex  $u \in P \cup X$ 
21 for  $v \in P \setminus N(u) 5+$ : // Use the pivot to prune search space
22 //New candidates for the recursive search
23  $P_{\text{new}} = P \cap N(v) 5+$ ;  $X_{\text{new}} = X \cap N(v) 5+$ ;  $R_{\text{new}} = R \cup \{v\} 5+$ 
24 //Search recursively for a maximal clique that contains  $v$ 
25 BK-Pivot( $P_{\text{new}}, R_{\text{new}}, X_{\text{new}}$ )
26 //After the recursive call, update  $P$  and  $X$  to reflect
27 //the fact that  $v$  was already considered
28  $P = P \setminus \{v\} 5+$ ;  $X = X \cup \{v\} 5+$ 
```

# Use Case # 3: k-Clique Listing

- A **k-Clique** of a graph  $G$  is a fully connected subgraph of  $G$  of  $k$  vertices.

---

```
1 /*Input: A graph  $G$  1,  $k \in \mathbb{N}$  Output: Count of  $k$ -cliques  $ck \in \mathbb{N}$ . */
2
3 //Preprocessing: reorder vertices with DGR or ADG.
4 //Here, we also record the actual ordering and denote it as  $\eta$ 
5 ( $v_1, v_2, \dots, v_n; \eta$ ) = preprocess( $V$ , /* selected vertex order */) 3
6
7 //Construct a directed version of  $G$  using  $\eta$ . This is an
8 //additional optimization to reduce the search space:
9  $G$  = dir( $G$ ) 3 //An edge goes from  $v$  to  $u$  iff  $\eta(v) < \eta(u)$ 
10  $ck = 0$  //We start with zero counted cliques.
11 for  $u \in V$  in parallel do: 2 //Count  $u$ 's neighboring  $k$ -cliques
12    $C_2 = N^+(u)$ ;  $ck += \text{count}(2, G, C_2)$ 
13
14 function count( $i, G, C_i$ ):
15   if ( $i == k$ ): return  $|C_k|$  5+ //Count  $k$ -cliques
16   else:
17      $ci = 0$ 
18     for  $v \in C_i$  5+ do: //search within neighborhood of  $v$ 
19        $C_{i+1} = N^+(v) \cap C_i$  5+ //  $C_i$  counts  $i$ -cliques.
20        $ci += \text{count}(i+1, G, C_{i+1})$ 
21   return  $ci$ 
```

---

TAKEAWAY: GMS offers lots of modularity in implementing graph mining algorithms, specifically set algebra based modularity!

# Theoretical Analysis

	$k$ -Clique Listing <i>Node Parallel</i> [41]	$k$ -Clique Listing <i>Edge Parallel</i> [41]	★ $k$ -Clique Listing with ADG (§ 6)	ADG (Section 6)	Max. Cliques Eppstein et al. [51]	Max. Cliques Das et al. [42]	★ Max. Cliques with ADG (§ 7.3)	Subgr. Isomorphism <i>Node Parallel</i> [26, 40]	Link Prediction <sup>†</sup> , JP Clustering
<b>Work</b>	$O\left(mk \left(\frac{d}{2}\right)^{k-2}\right)$	$O\left(mk \left(\frac{d}{2}\right)^{k-2}\right)$	$O\left(mk \left(d + \frac{\epsilon}{2}\right)^{k-2}\right)$	$O(m)$	$O\left(dm3^{d/3}\right)$	$O\left(3^{n/3}\right)$	$O\left(dm3^{(2+\epsilon)d/3}\right)$	$O\left(n\Delta^{k-1}\right)$	$O(m\Delta)$
<b>Depth</b>	$O\left(n + k \left(\frac{d}{2}\right)^{k-1}\right)$	$O\left(n + k \left(\frac{d}{2}\right)^{k-2} + d^2\right)$	$O\left(k \left(d + \frac{\epsilon}{2}\right)^{k-2} + \log^2 n + d^2\right)$	$O\left(\log^2 n\right)$	$O\left(dm3^{d/3}\right)$	$O\left(d \log n\right)$	$O\left(\log^2 n + d \log n\right)$	$O\left(\Delta^{k-1}\right)$	$O(\Delta)$
<b>Space</b>	$O(nd^2 + K)$	$O\left(md^2 + K\right)$	$O\left(md^2 + K\right)$	$O(m)$	$O(m + nd + K)$	$O(m + pd\Delta + K)$	$O(m + pd\Delta + K)$	$O(m + nk + K)$	$O(m\Delta)$

**Table 4: Work, depth, and space for some graph mining algorithms in GMS.**  $d$  is the graph degeneracy,  $K$  is the output size,  $\Delta$  is the maximum degree,  $p$  is the number of processors,  $k$  is the number of vertices in the graph that we are mining for,  $n$  is the number of vertices in the graph that we are mining, and  $m$  is the number of edges in that graph. <sup>†</sup> Link prediction and the JP clustering complexities are valid for the Jaccard, Overlap, Adamic Adar, Resource Allocation, and Common Neighbors vertex similarity measures. ★Algorithms derived in this work.

- Obtained better bounds for maximal cliques
- Obtained similar bounds for  $k$ -clique, but scales better depending on graph

# Evaluation: Datasets, Methodology, Architectures

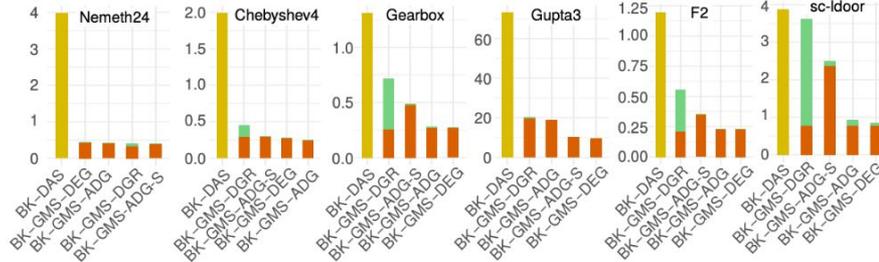
- GMS uses a wide selection of public datasets for flexibility
- GMS compares GMS variants to the most optimized state-of-the-art algorithms.
- GMS runs algorithms on the maximum number of cores available on machine.

Table 5: Some considered real-world graphs. Graph class/origin: [so]: social network, [wb]: web graph, [st]: structural network, [sc]: scientific computing, [re]: recommendation network, [bi]: biological network, [co]: communication network, [ec]: economics network, [ro]: road graph. Structural features:  $m/n$ : graph sparsity,  $\widehat{d}_i$ : maximum in-degree,  $\widehat{d}_o$ : maximum out-degree,  $T$ : number of triangles,  $T/n$ : average triangle count per vertex,  $T$ -skew: a skew of triangle counts per vertex (i.e., the difference between the smallest and the largest number of triangles per vertex). Here,  $\widehat{T}$  is the maximum number of triangles per vertex in a given graph. Dataset: (W), (S), (K), (D), (C), and (N) refer to the publicly available datasets, explained in § 8.1. For more details, see § 4.2.

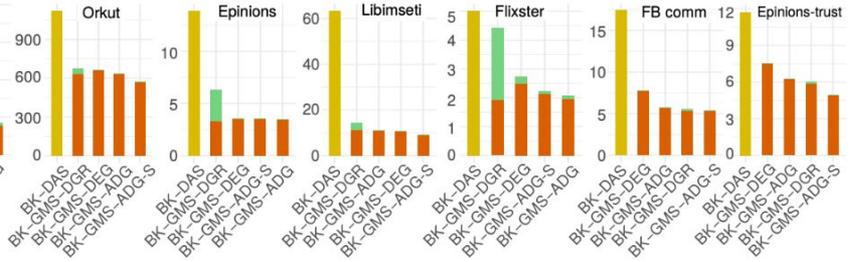
Graph †	$n$	$m$	$\frac{m}{n}$	$\widehat{d}_i$	$\widehat{d}_o$	$T$	$\frac{T}{n}$	Why selected/special?
[so] (K) Orkut	3M	117M	38.1	33.3k	33.3k	628M	204.3	Common, relatively large
[so] (K) Flickr	2.3M	22.8M	9.9	21k	26.3k	838M	363.7	Large $T$ but low $m/n$ .
[so] (K) Libimseti	221k	17.2M	78	33.3k	25k	69M	312.8	Large $m/n$
[so] (K) Youtube	3.2M	9.3M	2.9	91.7k	91.7k	12.2M	3.8	Very low $m/n$ and $T$
[so] (K) Flixster	2.5M	7.91M	3.1	1.4k	1.4k	7.89M	3.1	Very low $m/n$ and $T$
[so] (K) Livemocha	104k	2.19M	21.1	2.98k	2.98k	3.36M	32.3	Similar to Flickr, but a lot fewer 4-cliques (4.36M)
[so] (N) Ep-trust	132k	841k	6	3.6k	3.6k	27.9M	212	Huge $T$ -skew ( $\widehat{T} = 108k$ )
[so] (N) FB comm.	35.1k	1.5M	41.5	8.2k	8.2k	36.4M	1k	Large $T$ -skew ( $\widehat{T} = 159k$ )
[wb] (K) DBpedia	12.1M	288M	23.7	963k	963k	11.68B	961.8	Rather low $m/n$ but high $T$
[wb] (K) Wikipedia	18.2M	127M	6.9	632k	632k	328M	18.0	Common, very sparse
[wb] (K) Baidu	2.14M	17M	7.9	97.9k	2.5k	25.2M	11.8	Very sparse
[wb] (N) WikiEdit	94.3k	5.7M	60.4	107k	107k	835M	8.9k	Large $T$ -skew ( $\widehat{T} = 15.7M$ )
[st] (N) Chebyshev4	68.1k	5.3M	77.8	68.1k	68.1k	445M	6.5k	Very large $T$ and $T/n$ and $T$ -skew ( $\widehat{T} = 5.8M$ )
[st] (N) Gearbox	154k	4.5M	29.2	98	98	141M	915	Low $\widehat{d}$ but large $T$ ; low $T$ -skew ( $\widehat{T} = 1.7k$ )
[st] (N) Nemeth25	10k	751k	75.1	192	192	87M	9k	Huge $T$ but low $\widehat{T} = 12k$
[st] (N) F2	71.5k	2.6M	36.5	344	344	110M	1.5k	Medium $T$ -skew ( $\widehat{T} = 9.6k$ )
[sc] (N) Gupta3	16.8k	4.7M	280	14.7k	14.7k	696M	41.5k	Huge $T$ -skew ( $\widehat{T} = 1.5M$ )
[sc] (N) ldoor	952k	20.8M	21.5	76	76	567M	595	Very low $T$ -skew ( $\widehat{T} = 1.1k$ )
[re] (N) MovieRec	70.2k	10M	142.4	35.3k	35.3k	983M	14k	Huge $T$ and $\widehat{T} = 4.9M$
[re] (N) RecDate	169k	17.4M	102.5	33.4k	33.4k	286M	1.7k	Enormous $T$ -skew ( $\widehat{T} = 1.6M$ )
[bi] (N) sc-ht (gene)	2.1k	63k	30	472	472	4.2M	2k	Large $T$ -skew ( $\widehat{T} = 27.7k$ )
[bi] (N) AntColony6	164	10.3k	62.8	157	157	1.1M	6.6k	Very low $T$ -skew ( $\widehat{T} = 9.7k$ )
[bi] (N) AntColony5	152	9.1k	59.8	150	150	897k	5.9k	Very low $T$ -skew ( $\widehat{T} = 8.8k$ )
[co] (N) Jester2	50.7k	1.7M	33.5	50.8k	50.8k	127M	2.5k	Enormous $T$ -skew ( $\widehat{T} = 2.3M$ )
[co] (K) Flickr (photo relations)	106k	2.31M	21.9	5.4k	5.4k	108M	1019	Similar to Livemocha, but many more 4-cliques (9.58B)
[ec] (N) mbeacxc	492	49.5k	100.5	679	679	9M	18.2k	Large $T$ , low $\widehat{T} = 77.7k$
[ec] (N) orani678	2.5k	89.9k	35.5	1.7k	1.7k	8.7M	3.4k	Large $T$ , low $\widehat{T} = 80.8k$
[ro] (D) USA roads	23.9M	28.8M	1.2	9	9	1.3M	0.1	Extremely low $m/n$ and $T$

# Results: Maximal Clique

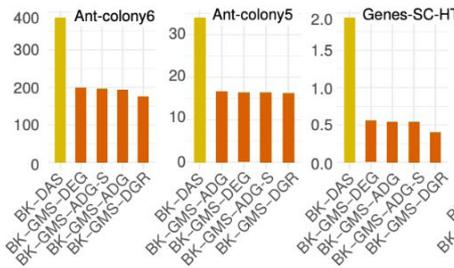
**Structural, scientific, various networks**



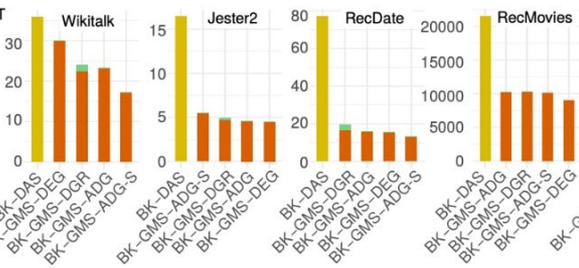
**Social networks**



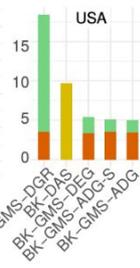
**Biological networks**



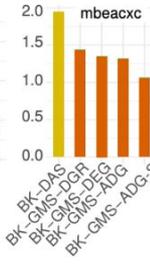
**Communication, recommendation networks**



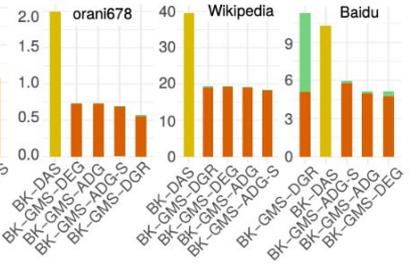
**Road graph**



**Economics networks**



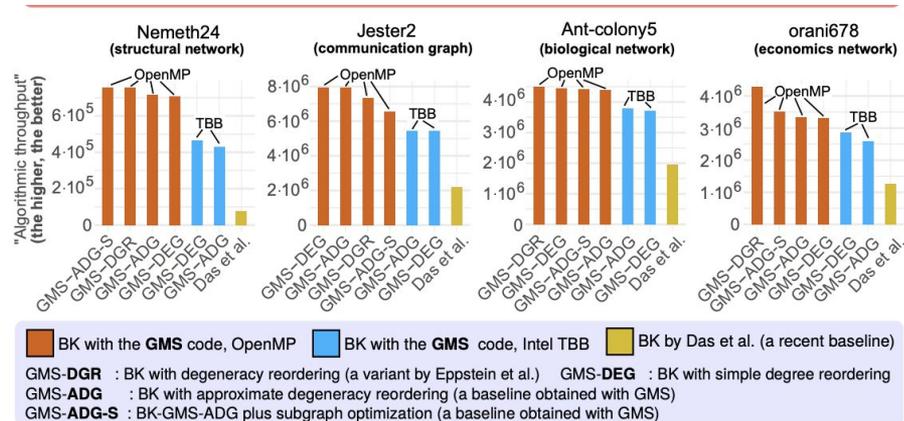
**Web graphs**



■ Bron-Kerbosch (GMS code)    
 ■ Fraction needed for reordering    
 ■ Bron-Kerbosch by Das et al. (BK-DAS, main competitor)    
 BK-GMS-DGR : BK with degeneracy reordering    
 BK-GMS-ADG : BK with approximate degeneracy reordering (proposed in this work)    
 BK-GMS-DEG : BK with simple degree reordering    
 BK-GMS-ADG-S : BK-GMS-ADG plus subgraph optimization (proposed in this work)

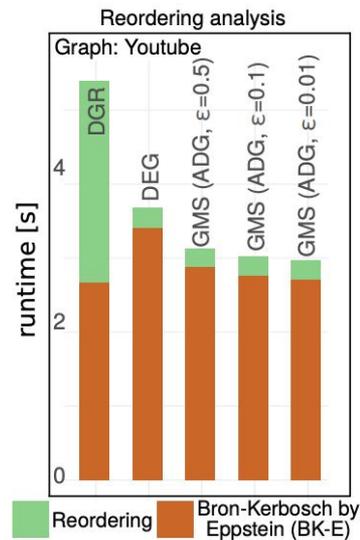
# Results: Maximal Clique cont.

- Achieved a significant improvement in maximal clique using ADG or DGR.
- GMS variants often faster than main competitor by >50%, sometimes even >9x.
  - Consistent over graphs of different structural characteristics
- Another view (**algorithmic efficiency**): number of maximal cliques found per second



# More Results

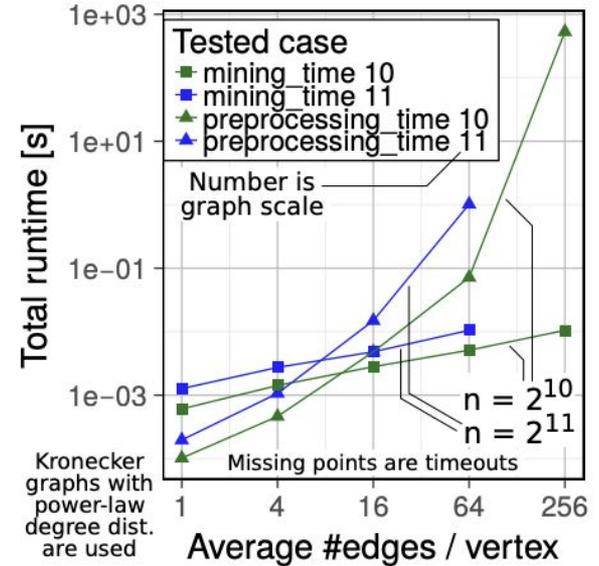
- Up to 10% speedup on k-clique algorithm with different parameters.
- ADG outperforms DGR as a preprocessor for Bron-Kerbosch.



**Figure 4: Speedups of ADG for different  $\epsilon$  over DEG/DGR, details in § 8.4. System: Ault.**

# Additional Analyses

- Subtleties of higher-order structures
  - Graphs similar in number of vertices, number of edges, sparsity, degree distribution etc., can have very different higher-order structures, such as number of 4-cliques. Choose datasets wisely!
- Using synthetic graphs can affect whether vertex reordering or mining dominates



(a) Analysis of synthetic graphs.

Thank you! Any questions?