

DecoMine: A Compilation-Based Graph Pattern Mining System with Pattern Decomposition

Jingji Chen
chen3385@purdue.edu
Purdue University
West Lafayette, IN, USA

Xuehai Qian
qian214@purdue.edu
Purdue University
West Lafayette, IN, USA

ABSTRACT

Graph pattern mining (GPM) is an important application that identifies structures from graphs. Despite the recent progress, the performance gap between the state-of-the-art GPM systems and an efficient algorithm—pattern decomposition—is still at least an order of magnitude. This paper clears the fundamental obstacles of adopting pattern decomposition to a GPM system.

First, the performance of pattern decomposition algorithms depends on how to decompose the whole pattern into subpatterns. The original method performs complexity analysis of algorithms for different choices, and selects the one with the lowest complexity upper bound. Clearly, this approach is not feasible for average or even expert users. To solve this problem, we develop a GPM compiler with conventional and GPM-specific optimizations to generate algorithms for different decomposition choices, which are evaluated based on an accurate cost model. The executable of the GPM task is obtained from the algorithm with the best performance. Second, we propose a novel partial-embedding API that is sufficient to construct advanced GPM applications while preserving pattern decomposition algorithm advantages. Compared to state-of-the-art systems, our new GPM system, DecoMine, developed based on the ideas, reduces the execution time of GPM on large graphs and patterns from days to a few hours with low programming effort.

CCS CONCEPTS

• **Computing methodologies** → **Parallel computing methodologies**; • **Mathematics of computing** → *Graph theory*; • **Software and its engineering** → *Compilers*.

KEYWORDS

graph pattern mining, pattern decomposition, graph analytics compiler

ACM Reference Format:

Jingji Chen and Xuehai Qian. 2023. DecoMine: A Compilation-Based Graph Pattern Mining System with Pattern Decomposition. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 1 (ASPLOS '23)*, March 25–29, 2023, Vancouver, BC, Canada. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3567955.3567956>



This work is licensed under a Creative Commons Attribution 4.0 International License.

ASPLOS '23, March 25–29, 2023, Vancouver, BC, Canada

© 2023 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9915-9/23/03.

<https://doi.org/10.1145/3567955.3567956>

1 INTRODUCTION

Graph pattern mining (GPM) [55] is an important workload [18, 46, 51, 56] with the goal of finding all subgraphs, known as *embeddings*, in the input graph that “match”, i.e., *isomorphic* to user-specified patterns. GPM has important applications in functional modules discovery [48], biochemical structures mining [40] and anomaly detection [6] and many others. The key challenge of GPM is to enumerate a large number of subgraphs. In WikiVote, a small graph with 7k vertices, the number of embeddings matching the 5-chain pattern can reach 71 billion based on our experiment. For frequent subgraph mining (FSM), mining all frequent size-3 patterns with certain criteria on a median-size LiveJournal graph by Peregrine [26], one of the recent systems, can take more than 12 hours.

Due to the importance of GPM for various applications, several domain-specific GPM systems have been proposed in recent years [14, 15, 17, 26, 42, 52, 55, 58]. The intuitive APIs, i.e., requesting the count of a given pattern or defining operations for each identified embedding of a pattern returned from the system in user-defined functions (UDFs), allow non-expert users to construct GPM applications while achieving high performance. We observe that the performance gap between the state-of-the-art systems and the efficient algorithm is still large. For example, using one core, the fastest execution time (617.2s) of mining all size-5 patterns from the WikiVote graph using GraphPi [52] is 48.6× slower than the manually-optimized pattern decomposition algorithm [44] (12.7s).

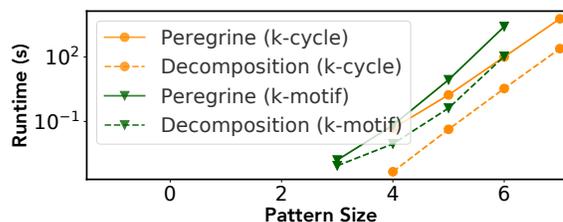


Figure 1: Pattern Size vs. Runtime

The key idea of *pattern decomposition* is to decompose a target pattern into smaller subpatterns, and then compute the count of each. The count of the target pattern can be calculated using the subpattern counts with low additional cost. The method is very fast because, empirically, the embedding enumeration cost increases rapidly with pattern size. Figure 1 shows the execution time of Peregrine [26] when counting *k*-motif (all size-*k* patterns) and *k*-cycle (a cycle-structure pattern with *k* vertices/edges) on the EmailEuCore graph [35, 64] with increasing pattern size (*k*). In comparison, we show the execution time of an implementation based on the pattern

decomposition algorithm. There are two fundamental obstacles to adopt the algorithm to an existing GPM system.

Obstacle 1: APIs incompatible for pattern decomposition.

A GPM system should not only perform pattern counting but also allow application-specific operations to be expressed in the user-defined functions (UDFs) for each pattern embedding identified by the system. Such API is problematic for pattern decomposition: designed for pattern counting, the algorithm counts without materializing any whole pattern embedding. However, the current API forces the system to materialize and returns each whole pattern embedding to UDFs. As a result, the algorithmic and performance advantages are largely wiped out. Thus, the *fundamental problem* is to seek an API that can intrinsically work with pattern decomposition and preserve most of its advantages, while still being sufficient to express applications and easy to use.

Obstacle 2: Automating pattern decomposition.

Pattern decomposition algorithm decomposes the target pattern by selecting a *vertex cutting set* V_C —removing the vertices in the set will break the pattern graph p into K connected components. The K subpatterns, i.e., p_1, \dots, p_K , are generated by merging the vertex cutting set with each of the K components. The algorithm identifies the subgraphs matching V_C , denoted as e_C . For each e_C , it counts the number of subgraphs containing e_C that can match each subpattern p_i , denoted as M_i . With certain exceptions, $M_1 \times \dots \times M_K$ is the number of embeddings matching p containing e_C . This procedure is illustrated in Figure 6. We see that different choices of V_C lead to different sets of subpatterns and different algorithms. The original paper [44] performs *complexity analysis of manually-optimized algorithms* for different V_C choices on specific patterns, and selects the one with the lowest complexity upper bound. This approach is not feasible for average or even expert users. In a GPM system, the V_C selection process should be *automated*.

Specifically, the system should search good implementations, instead of just applying the original pattern decomposition algorithm, which can perform even worse. A bad selection can be more than one order of magnitude slower than a good selection. For example, for pattern p_5 in Figure 11 (a), the worst selection is more than 60× slower than the best one on the EmailEuCore graph. The key challenge is predicting the performance of the implementations with an accurate cost model.

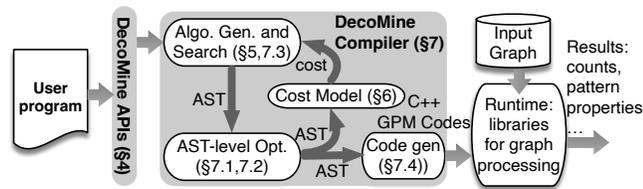


Figure 2: DecoMine System Overview

Our solution: DecoMine—a new compilation-based GPM system with a novel API that automates algorithm selection using the accurate cost model. We build the DecoMine compiler that can *automatically* generate algorithms for *arbitrary* patterns with different V_C and pattern vertex matching orders, i.e., the order that the vertices in the pattern graph are enumerated. Specifically,

DecoMine compiler generates the intermediate representation, i.e., abstract syntax tree (AST), of algorithm candidates, which are evaluated using the cost model. For each AST, the compiler performs not only conventional optimizations such as Loop Invariant Code Motion (LICM) and Common Subexpression Elimination (CSE) but also a novel *pattern-aware loop rewriting optimization* to eliminate redundant computation. The system overview is shown in Figure 2. We make three novel contributions that *bridge pattern decomposition algorithm and system*.

- **Partial-embedding API: the natural API for pattern decomposition.** Instead of forcing the system to return whole embedding of pattern p , the partial-embedding API only requires the system to return “partial” embeddings for p_1, \dots, p_K —*any* valid pattern decomposition. It avoids the overhead of materializing whole pattern embedding while still allowing UDFs to be specified on partial embeddings. We find that this API can correctly express advanced GPM applications such as Frequent Subgraph Mining (FSM) as long as the system guarantees two *properties* when passing partial embeddings. Partial-embedding API is an *elegant* abstraction because it is inherently compatible with pattern decomposition while *not requiring users to determine how pattern is decomposed*—the responsibility of the DecoMine compiler.

- **Generalized pattern decomposition algorithm.** The current pattern decomposition algorithms only support *pattern counting for patterns with no more than five vertices* [44], but the DecoMine compiler needs to generate algorithms for *arbitrary* patterns. To this end, we propose a generalized decomposition algorithm that returns partial embeddings during execution.

- **Accurate cost models: capturing real-world graph characteristics.** The pattern enumeration procedure can be expressed as nested loops. Based on the cost model, we need to estimate the number of iterations that a loop will execute, which depends on the input graph. For example, for a graph with n vertices, the loop to process the 1-hop neighbors can be estimated to execute np iterations, where p is the probability that a vertex is connected to another vertex. It is the method adopted by AutoMine [42], i.e., a recent compilation-based GPM system with *no support for pattern decomposition* and used as one of the baselines. We propose two new cost models that significantly improve the accuracy.

The *locality-aware* cost model is a simple extension to AutoMine’s method that assigns a higher connection probability if two vertices are within k -hop, where k can be specified as a threshold. It coarsely considers the real graph properties by adjusting the probability. The second *approximate-mining based* cost model is based on a *key observation*: we can *estimate the number of loop iterations at a loop level by the approximate count of the corresponding pattern reaching that level*. Refer to Figure 9, the 4-vertex pattern can be enumerated with matching order ① (A,B,C,D) or ② (D,B,C,A) among others possibilities. The number of iterations for the loop to match vertex C can be estimated by the *count* of smaller pattern (A,B,C) for ① and (D,B,C) for ②. We develop a *fast* method to obtain the *approximate* and *relative* count of all patterns in input graph up to a certain size with random sampling and approximate pattern mining [25].

- **Alternative solutions.** A few systems [14, 15], provide low-level but more flexible APIs that allow users to construct sophisticated algorithms *manually*. The low-level APIs of Sandslash [14]

can express pattern decomposition algorithms. However, the manual implementation based on the APIs makes it *almost the same as constructing the native algorithms for specific patterns*.

Evaluation highlights. The implementation of DecoMine has about 10,000 lines of codes. We perform comprehensive performance evaluation comparing DecoMine with recent GPM systems—Arabesque [55], RStream [58], our own implementation of AutoMine [42] (AutoMineInHouse) which is not open sourced¹, Peregrine [26], Pangolin [15], Fractal [17], and GraphPi [52]. Our results show that DecoMine can be *up to 827× and 575× faster than AutoMineInHouse and Peregrine*, respectively. For *large graphs*, on Friendster and RMAT-100M, two large graphs with more than one billion edges, DecoMine reduces the execution time of 4-motif counting *from tens of hours to less than two hours compared to Peregrine*. For *large patterns*, AutoMineInHouse fails to mine all 6-motif patterns in *one week* on a small graph with 7k vertices; while DecoMine can finish the same task with roughly *one hour* on a median-size graph with 3.8M vertices. Most importantly, DecoMine effectively closes the gap between the GPM system and efficient native algorithms. With multiple threads, DecoMine achieves better performance.

Sources of performance improvements. The primary reason for the drastic performance improvements is due to the better algorithm, but the proposed novel techniques *bring algorithmic advantages to a general GPM system*. DecoMine fully automates the algorithm generation, optimization, and selection from the search space, ensuring the ease-of-use and high performance at the same time. We believe that DecoMine is a significant advance over the state-of-the-art of GPM systems.

2 GRAPH PATTERN MINING BACKGROUND

2.1 Definitions

- **Graph.** $G = (V, E)$ contains a vertex set V and edge set E . E is a subset of $V \times V$ and $(u, v) \in E$ if u and v are connected by an edge.
- **Labeled graph.** Graph vertices may have labels captured by a mapping $f_L : V \rightarrow L$ where L is the set of vertex labels.
- **Edge induced subgraph.** It contains a subset of the edges and all corresponding vertices. Formally, $G' = (V', E')$ is a edge-induced subgraph iff. $E' \subseteq E$ and $V' = \{v \in V | (v, u) \in E' \text{ for some } u\}$.
- **Vertex induced subgraph.** It contains a subset of the vertices and all edges induced by them. Formally, $G' = (V', E')$ is a vertex-induced subgraph iff. $V' \subseteq V$ and $E' = \{(u, v) \in E | u, v \in V'\}$.
- **Isomorphism.** Two graphs $G_0 = (V_0, E_0)$ and $G_1 = (V_1, E_1)$ are *isomorphic* iff. there exists a one-to-one mapping $f : V_0 \rightarrow V_1$ such that $(u, v) \in E_0 \iff (f(u), f(v)) \in E_1$.
- **Graph pattern mining (GPM).** It takes an undirected graph as input, enumerate all its subgraphs that are isomorphic to a given pattern graph, and process them to gather information. When the pattern graph has vertex labels, only the isomorphic subgraphs with matching labels are considered.
- **Embedding.** In GPM, it refers to the subgraphs that are isomorphic to the pattern graph.

- **Edge (vertex) induced embedding.** The embedding is an edge (vertex) induced subgraph.

Figure 3 shows an example of embedding of the pattern graph in input graph.

2.2 Pattern Enumeration Method and Optimization

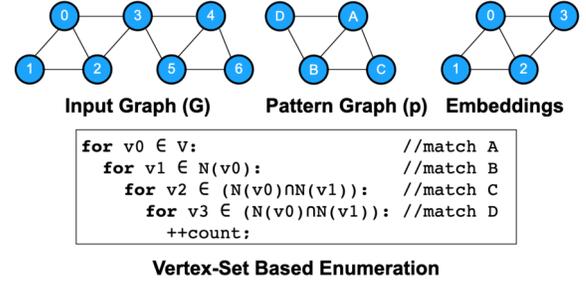


Figure 3: GPM Example and Implementation

Recent GPM systems [26, 41, 42, 52] use a pattern-aware vertex-set-based method to construct the embeddings based on the pattern graph. It significantly outperforms pattern-oblivious method used in early systems [17, 55, 58] that enumerate all subgraphs and perform the expensive isomorphic check. Thus, DecoMine adopts the vertex-set-based method, which uses nested loops to extend subgraphs incrementally from a single vertex until reaching the embedding of the pattern graph. Figure 3 shows one possible implementation of nested loops to enumerate embeddings for the pattern, where $N(v)$ refers to the vertex set containing all neighbors of v . It also illustrate the concept of *pattern vertex matching order*. A different matching order, i.e., (D,B,C,A), will lead to different nested loops that perform the equivalent enumeration. A *cost model* can be utilized to search for pattern vertex matching orders with high performance.

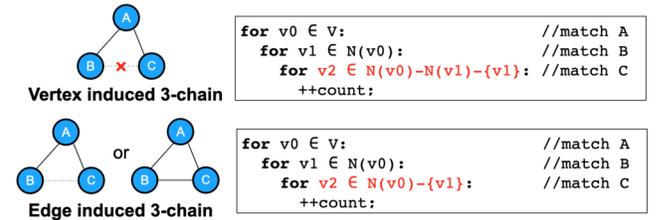


Figure 4: Vertex/Edge Induced Embedding Enumeration

Vertex vs. edge induced embedding enumeration. Both kinds of embedding can be enumerated with a vertex-set-based method. Figure 4 shows the enumeration of 3-chain pattern (A,B,C). In a vertex induced 3-chain, B and C cannot be connected, otherwise (A,B,C) forms a triangle, not a 3-chain. It can be realized by excluding “ $-N(v1)$ ” when matching C: $v2$ cannot be the neighbor of both A and B. In an edge induced 3-chain, As long as the edge (A,B) and (A,C) are included, whether B and C are connected does not

¹AutoMineInHouse achieves comparable performance with the original implementation reported in [42], detailed numbers are reported in Table 2.

matter: both cases lead to valid edge induced 3-chain embeddings. It can be realized by removing the constraint “ $-N(v1)$ ” for $v2$.

DecoMine supports both kinds of pattern enumeration. Unless explicitly mentioned, embedding refers to edge induced embedding, which is assumed by pattern decomposition algorithm. For vertex induced embedding, the pattern count (cnt) can be obtained from the counts of edge induced embeddings. In Figure 4, we have the following relation: $cnt(vertex\ induced\ 3-chain) = cnt(edge\ induced\ 3-chain) \times 3 \times cnt(edge\ induced\ triangle)$. The count of triangle is multiplied by 3 because one triangle can lead to three 3-chains by removing different edges. Our cost model can evaluate the performance of enumerating vertex induced embedding between the two options: (1) direct method similar to Figure 4; or (2) indirect method by enumerating multiple edge induced embeddings with pattern decomposition. It is possible that option (2) is estimated to be slower, for which the compiler falls back to option (1) without using pattern decomposition.

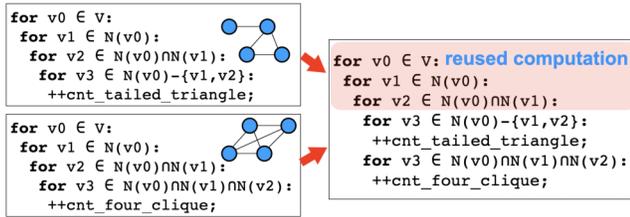


Figure 5: Computation Reuse

Optimization 1: symmetry breaking. This technique ensures that the same embedding is only enumerated once by enforcing restrictions on vertices IDs of the embeddings [22, 26, 41]. In Figure 3, if a subgraph $(v0, v1, v2, v3)$ matches pattern vertex (A, B, C, D) , then $(v1, v0, v2, v3)$, $(v0, v1, v3, v2)$, $(v1, v0, v3, v2)$ can also match these pattern vertices—the same embedding will be enumerated for 4 times. The redundancy can be eliminated with the restrictions on vertex IDs, e.g., $v0 < v1$ and $v2 < v3$, so that only one of the matchings is preserved. The preserved matching is called automorphism canonical. For a pattern, the set of restrictions can be generated by finding equivalent vertices according to pattern automorphisms [22], which is used in Peregrine [26] and GraphZero [41]. There may exist multiple possible sets of restrictions, GraphPi [52] utilizes a cost model to select the breaking restrictions that lead to the best performance.

Optimization 2: computation reuse. It is profitable to schedule the common computations shared by the same or different patterns together [42]. In Figure 5, the first three loops of the enumeration process of edge-induced 4-cliques and tailed-triangles (left) are the same. The compiler can merge these loops so that the common computation is only performed once (right). Computation reuse can benefit applications that need to enumerate multiple patterns such as Frequent Subgraph Mining. With pattern decomposition, since the algorithm needs to enumerate subpattern embeddings for even a single pattern, this optimization may lead to more benefits.

2.3 Prior GPM Systems

General-purpose GPM systems. Arabesque [55] enumerates all subgraphs and relies on expensive isomorphism checks to identify pattern embeddings. RStream [58] is an out-of-core GPM system that combines streaming processing and relational joins to optimize for disk I/O accesses. Peregrine [26] utilizes pattern properties (e.g., symmetry) to improve system efficiency. Pangolin [15] and Sandslash [14] provide flexible and low-level interfaces so that users can implement more advanced algorithms manually. Fractal [17], G-miner [13] and G-thinker [61] are three recent GPM systems focusing on distributed execution. Automine [42], GraphZero [41] and GraphPi [52] are compilation-based systems that generate high-performance GPM programs based on high-level pattern specifications.

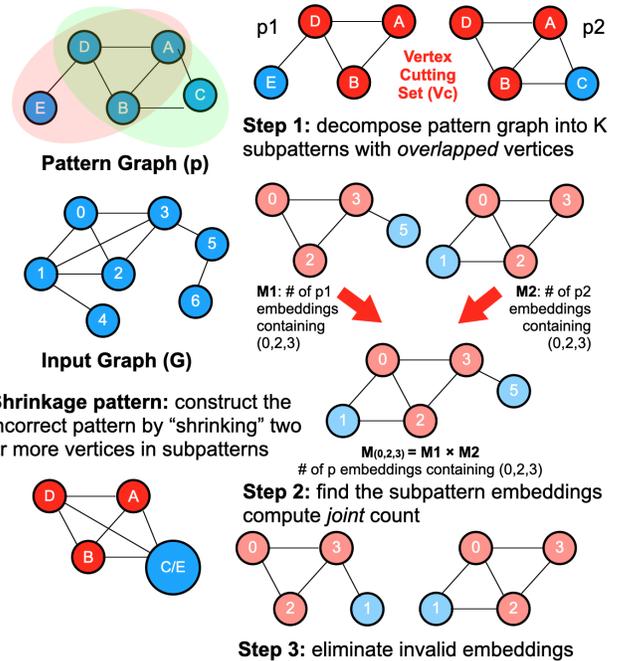


Figure 6: Pattern Decomposition Algorithm Overview

Specialized GPM implementations. There are a large number of *hand-optimized* implementations targeting various GPM applications on different platforms. OPT [31] is a disk-based triangle solver that discovers all triangle-shape subgraphs. [53] is a cache-friendly triangle solver designed for single-node multi-core environments. PDTL [21] and DistTC [24] are fast distributed triangle discovering implementations. Frequent subgraph mining (FSM) is another important class of GPM workloads aimed to find frequent patterns of interest. gSpan [62] is an FSM solver based on DFS. GRAMI [19] avoids redundant subgraph enumeration in FSM to improve performance. ScaleMine [1] is a distributed FSM solver with a novel load-balancing approach. DistGraph [54] is a distributed FSM system that leverages graph partitioning to support massive datasets. Other specialized GPM implementations also include those for clique-shape subgraph finding [16, 38, 60],

approximate pattern counting [9, 27, 43, 59], and subgraph counting/listing [2, 3, 7, 8, 23, 30, 32, 39, 44, 49, 50]. These specialized implementations achieve high-performance thanks to various algorithmic optimizations but incur considerable programming effort.

3 DECOMINE SYSTEM

3.1 Pattern Decomposition Method

This section introduces the pattern decomposition algorithm [44, 46] at a conceptual level. The highlights of the algorithm with a concrete pattern graph (p) and an input graph (G) are shown in Figure 6. The algorithm contains three key steps.

- **Pattern decomposition.** The original pattern is decomposed into subpatterns by choosing a *vertex cutting set* (V_C), i.e., a subset of pattern graph vertices, of which the removal breaks p into several connected components. The connected components can be merged with V_C to produce subpatterns. Since V_C does not exist in a clique pattern, this pattern cannot benefit from the pattern decomposition. However, clique counting is typically fast and not the performance bottleneck. The example has two subpatterns $p_1: (A,B,D,E)$ and $p_2: (A,B,C,D)$ with the vertex cutting set (A,B,D) .
- **Vertex set enumeration and joint counting.** This step enumerates the embedding matching V_C and then calculates the number of subpattern embeddings from each V_C embedding. In the example, M_1 and M_2 are the number of embeddings matching p_1 and p_2 , respectively. They both contain (v_0, v_2, v_3) matching V_C . (v_0, v_2, v_3, v_5) and (v_0, v_1, v_2, v_3) are two examples of embeddings matching p_1 and p_2 . Then the M_1 and M_2 subpattern embeddings are conceptually *joined*, leading to $M_1 \times M_2$ full pattern (p) embeddings. The embedding of p are not materialized.
- **Invalid embedding elimination.** It identifies and eliminate *invalid* embeddings of p . It happens when the join of valid subpattern embeddings does not produce a valid embedding of the whole pattern. The example shows two embeddings of p_1 and p_2 with four vertices v_0, v_1, v_2 , and v_3 , but the joint embedding is not a valid embedding of p . The reason is that the selection of C and E happen to be the same. Pinar *et al.* [44] proposes to explicitly construct the “invalid” patterns, known as *shrinkage pattern*, by shrinking at least two vertices in p in different subpatterns. We can enumerate the shrinkage patterns, get the count, and exclude them from the count calculated in step 2. The final count of p is further adjusted by dividing pattern-dependent multiplicity as previous works [42].

3.2 Design Principles

To adopt the algorithm to a GPM system, we take a compilation-based approach to automatically generate algorithms for arbitrary patterns with different V_C and pattern vertex matching orders. The various implementation candidates are evaluated based on a cost model so that the one with the best performance can be chosen. This approach addresses two key factors affecting performance.

- **Incompatible with symmetry breaking.** Two symmetry embeddings of a subpattern may lead to different embeddings of the whole pattern, if only one subpattern embedding is counted with symmetry breaking, the whole pattern embedding will be

under-counted. The compilation-based approach allows a novel pattern-aware loop rewriting optimization (Section 7.2) that partially recovers the performance lost due to the lack of symmetry breaking in subpatterns.

- **No guaranteed performance improvement.** Some subpatterns after the decomposition may be very frequent, which can actually *increase* enumeration cost. An accurate cost models (Section 6) can be used to choose the implementations that speedup execution. The accuracy is crucial for our approach since the cost model is required to prevent performance degradation.

4 PARTIAL-EMBEDDING API

4.1 Motivation

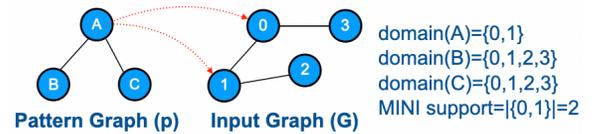


Figure 7: Building FSM Application in a GPM System

A GPM system typically provides high level API such as requesting the system to compute the embedding count of a give pattern; or user-defined function (UDF) so that users can define the operations to be executed on each embedding identified by the system. We use Frequent Subgraph Mining (FSM) [11] as an example to illustrate the usage of UDF. In FSM, *domain* [26] of a pattern vertex is the set of input graph vertices that maps to it; *MINI support* [11] is the size of the smallest domain across all pattern vertices. FSM discovers *frequent patterns*, whose supports are no less than a user-specified threshold.

Figure 7 shows the domains and MINI support of the given pattern graph on an input graph. To construct an FSM application, when an embedding of p is identified by the system, the UDF updates domain and MINI support of pattern vertices. The problem of using this API with pattern decomposition is, the algorithm, which works based on subpattern embedding counts, is forced to materialize the embedding of the whole pattern from a set of subpattern embeddings. This procedure will completely diminish the algorithm advantages. For FSM, it is actually not necessary because only the *mapping* of pattern vertex to input vertex needs to be obtained.

4.2 API Specification

Figure 8(a) shows the APIs of DecoMine. The key property of the APIs is that they are all *compatible* with pattern decomposition algorithms in a sense that they do not lead to unnecessary performance overhead, while still being sufficient to build advanced applications. The first API `get_pattern_count()` simply returns the count of a given pattern. Since pattern decomposition algorithm is designed for pattern counting, the API does not introduce additional overhead.

The other two APIs are centered around the concept of *partial embedding*—an embedding of a subpattern. The UDF function `process_partial_embedding` allows the users to define operations on each partial embedding pe passed from the system. The

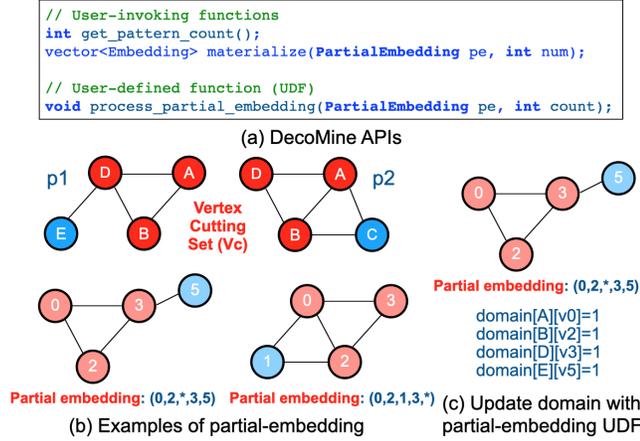


Figure 8: Partial-Embedding APIs

number of whole pattern embeddings that can be expanded from pe is also passed in $count$. The function `materialize` allows users to materialize a certain number (num) of whole pattern embeddings from a partial embedding. Typically `materialize` function is called inside `process_partial_embedding` using the partial embedding passed from the system. Figure 8(b) shows two examples of partial embedding corresponding for subpatterns $p1$ and $p2$. Each partial embedding is represented as a tuple of vertices, each matches a vertex in pattern graph in alphabetical order. The $*$ means the sub-pattern does not have certain vertex (or vertices). For FSM, the domain can be updated in `process_partial_embedding` for the subpattern vertices, see Figure 8(c). But we cannot update the domain of the missing vertex (C in the example). How do we ensure the correctness of the application?

DecoMine ensures two properties:

- **Completeness:** if a partial-embedding matching a subpattern is passed from the system, then all other partial-embeddings of matching the same subpattern will be also passed;
- **Coverage:** the set of subpatterns must fully cover all vertices of the pattern graph.

For FSM, the completeness property ensures the correct calculation of domains, while the coverage property that we can obtain the domain of all vertices in the pattern graph. Note that the coverage property is ensured with pattern decomposition algorithm due to the way the subpatterns are constructed: they share the vertices in V_C and cover all pattern graph vertices.

We claim that the partial-embedding API is an *elegant* solution that works intrinsically with pattern decomposition algorithm while providing intuitive user interface. The users only need to construct the applications based on the concept of subpattern and partial-embedding, but *are not exposed to algorithmic details such as how to decompose a pattern or how to leverage shrinkage patterns*, which are the responsibility of DecoMine compiler.

4.3 Applicability

With the API, DecoMine is able to benefit applications beyond pattern counting as long as they can be implemented with partial

embedding materialization. For FSM, by only materializing the fragments of embeddings, whose number can be much smaller than whole embeddings, the mapping between pattern and input graph vertices are correctly discovered. Another example of using partial materialization is a graph query like “listing all types (labels) of vertices that are the centers of size-10 star-shape subgraphs (i.e., with no less than 10 neighbors)”. Users can discover all center vertices from partial embeddings of a star-shape pattern to record the labels. For applications that fundamentally require complete materialization (e.g., writing all embeddings to a disk), DecoMine has to fall back to a non-decomposition-based method, resulting in performance similar to existing systems.

5 GENERALIZED PATTERN DECOMPOSITION

In this section, we propose a generalized decomposition algorithm designed for the partial-embedding API. The algorithm shares the matching process of vertices in V_C among all subpatterns (the embedding matching V_C is denoted as e_C), and perform additional on-the-fly enumeration to obtain how many embeddings of each subpattern can be extended from e_C . This means that e_1 and e_2 (embeddings of p_1 and p_2) are constructed in two steps: first generate e_C , and then e'_1 and e'_2 , which are obtained by the extension from the common e_C .

Algorithm 1 DecoMine Algorithm Template

```

1: Input: Pattern  $p$ , Cutting Set  $V_C$ , Matching orders of all subpatterns
2: Decompose  $p$  using  $V_C$  to generate  $K$  subpatterns, and the shrinkage patterns
3:  $pattern\_cnt \leftarrow 0$ 
4: for all  $e_C = (v_0, v_1, \dots, v_{|V_C|-1})$  matching the cutting set  $V_C$  do
5:    $num\_shrinkages \leftarrow$  new HashTable()
6:   clear( $num\_shrinkages$ )
7:    $M \leftarrow 1$ 
8:   for all  $i \leftarrow 1 \dots K$  do
9:      $M_i \leftarrow$  the number of  $pe$  extending  $e_C$  and matching the  $i$ -th subpattern
10:     $M \leftarrow M \times M_i$ 
11:    $pattern\_cnt \leftarrow pattern\_cnt + M$ 
12:   for all  $e$  extending  $e_C$  and matching one of the shrinkage patterns do
13:      $pattern\_cnt \leftarrow pattern\_cnt - 1$ 
14:     for all  $i \leftarrow 1 \dots K$  do
15:        $pe_i \leftarrow$  extract_subpattern_embedding( $e, i$ )
16:        $num\_shrinkages[pe_i] \leftarrow num\_shrinkages[pe_i] + 1$ 
17:   for all  $i \leftarrow 1 \dots K$  do
18:     for all  $pe$  extending  $e_C$  and matching the  $i$ -th subpattern do
19:        $count \leftarrow M/M_i - num\_shrinkages[pe]$ 
20:       if  $count > 0$  then
21:         process_partial_embedding( $pe, count$ )

```

Algorithm 1 shows the algorithm template from which the compiler can generate concrete pattern decomposition algorithms for a combination of a vertex cutting set V_C and a matching order. A *matching order* is defined as $(o_{vc}, o_1, \dots, o_k, o_{s1}, \dots, o_{sn})$, where o_{vc} is the matching order of V_C used in Line 4; each o_i determines the order that extends from a match of V_C to the i -th subpatterns (k in total) used in Line 9; and each o_{sj} determines the order that extends to the j -th shrinkage pattern (n in total) used in Line 12. The template generates algorithms when application uses `process_partial_embedding`. For pattern counting, the pattern count is returned without executing Line 14-21. For each e_C , three computation steps are performed in an iteration (Line 6-21).

The first step (Line 7-10) enumerates the embeddings that can be extended from e_C to match the subpatterns using o_C , and gets a count M_i for each subpattern. Multiplying them together (Line 10) obtains the number of embeddings M after joining all partial-embeddings. By accumulating all M across different e_C (line 11),

and deducting the redundant and invalid embeddings (line 12), we can get the correct counts of embeddings of the whole pattern. Due to space limit, we omit the codes to eliminate multiplicity.

The second step (Line 12-16) constructs the hash table for fast invalid embedding count query. It enumerates the invalid embeddings e that can be extended from e_C based on each shrinkage pattern i with matching order o_{si} . For each invalid embedding e , `extract_subpattern_embedding` returns the partial-embedding pe_i of subpattern $_i$ contained in e . The same pe_i may be contained in multiple invalid embeddings, so each pe_i should be discounted multiple times from the number of embeddings of the whole pattern that can be extended from the partial-embeddings of subpattern $_i$. The discount number for each pe_i is recorded in the hash table `num_shrinkages` (line 16).

The third step (line 17-21) deducts the discount number recorded in the second step for each subpattern. M/M_i is the number of embeddings that can be extended from a partial-embedding of subpattern $_i$ before the deduction.

To reduce the overhead, for each hash table entry, we add a 64-bit integer field `entry_valid`, and maintain another hash-table-wise integer `global_valid`. They are both initialized to 0. An entry is valid only if `entry_valid` is equal to `global_valid`. Once the hash table is cleared, we only increase `global_valid` by 1 without modifying any hash table entries. If overflow happens, we reinitialize all these `valid` integers to 0. Thus, the complexity of clearing (Line 6) is reduced to $O(1)$. This optimization mainly benefits large cutting sets, which result in a large number of clearings.

6 DECOMINE COST MODELS

6.1 Problem, Existing Solution and Improvement

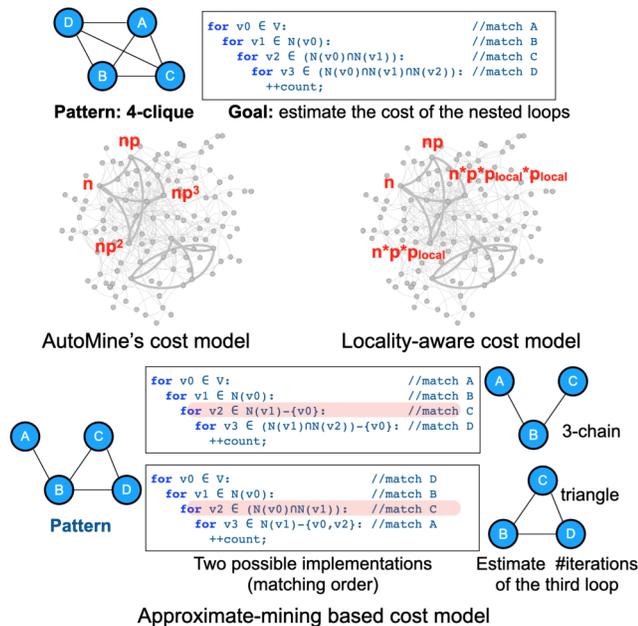


Figure 9: Cost Model Insights

Given a pattern, the pattern enumeration is performed by nested loops. The problem is to estimate the execution time of the nested loops, which can be done by estimating the number of iterations executed in each loop. At the top of Figure 9, we show the nested loops for the 4-clique pattern.

Automine is the first compilation-based GPM system facing the same problem. Its cost model assumes a random input graph: the algorithm runs on a random graph with n vertices, where every vertex pair is directly connected by a fixed probability p . For the 4-clique pattern, the number of iterations of the 1-st, 2-nd, 3-rd, and 4-th loop are n , np , np^2 , and np^3 , as shown in Figure 9. Unfortunately, the accuracy of the simple cost model is poor. Consider the livejournal graph with 4.8M vertices, the connection probability is 1.8×10^{-6} —the average degree divided by the number of vertices. The estimated cost is $n \times np \times np^2 \times np^3 \approx 1.9 \times 10^{-8}$. However, there are in fact 9.9B 4-cliques in the graph and the actual cost is $(9.9B) \times 4! = 238B$. There is a huge gap between the actual cost and the estimation.

A simple improvement is to consider the locality property of real-world graphs. If two vertices are k -hop neighbors with k less than some threshold α , we increase the probability that they are connected to p_{local} that is much larger than p . We empirically choose $\alpha = 8$, and a p_{local} that is close to real-world graphs (e.g., 0.27 for the LiveJournal). The system allows users to set these parameters. In Figure 9, when estimating the size of $N(v_0) \cap N(v_1)$, since vertices in $N(v_1)$ are 2-hop neighbors of v_0 , $N(v_0) \cap N(v_1)$ is estimated as $|N(v_1)|p_{local} = np_{local}$.

6.2 Approximate-Mining Based Cost Model

To further improve the accuracy of cost prediction, we propose to estimate the number of loop iterations at a level by the approximate count of the corresponding pattern reaching that level. To illustrate the insights, consider two possible implementations matching the same pattern in Figure 9. To estimate the number of iterations for the 3-rd loop, the first implementation follows the matching order (A,B,C) while the second implementation follows the matching order (D,B,C). Thus, the first three loop levels in the two implementations can be estimated by the count of 3-chain and triangle, respectively.

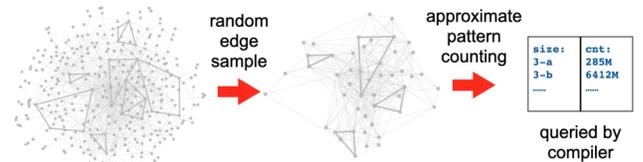


Figure 10: Approximate-Mining Cost Model Methodology

Based on this idea, we propose to (1) randomly sample a fixed number of edges (e.g., 32M) from input graph; and then (2) perform the neighbor sampling [25] to get the approximate counts of patterns up to a certain size in the sampled graph. These counts are stored in a table that is later queried by the compiler to get the cost estimation. The procedure is shown in Figure 10. The counts in the table are approximate and relative—sufficient for the purpose of

cost estimation. Unlike vertex sampling, which may drop critical structures like hub nodes, edge samplings can preserve the hub nodes with high probability. Hub nodes are adjacent to a large number of edges, as long as one of them is sampled, the hub vertex will be preserved. In practice, collecting approximate counts for patterns up to 5 vertices is mostly enough. If the counts are missing due to the large pattern, DecoMine can quickly run the profiling on demand and cache the results.

With a small number of sampled edges, the approximate mining algorithms can accurately obtain the count of frequent patterns, while underestimating that of the infrequent ones [9, 10] since they are less likely to get sampled. For a cost model, accurately estimating the frequent patterns is much more important because they correspond to loops with more iterations, which dominate the execution time.

6.3 Comparison

We randomly generate 100 implementations by choosing different cutting sets/matching orders of three applications: 5-motif and two large patterns (p4 and p5) shown in Figure 11 (a). Figure 11 (b) shows the relationship between the actual runtime on the EmailEuCore graph [35, 64] and the estimated cost by the cost model. We use linear correlation coefficient R (the larger the better) to describe how well the cost model predicts the runtime. It shows that the locality-aware model is much better than AutoMine’s model while approximate-mining based model provide the highest accuracy.

The end-to-end performance comparison of different cost models shown in Figure 11 (c), where LA and AM stand for locality-aware and approximate mining cost model, respectively. We see that the accuracy improvement translates to drastic end-to-end performance improvement—the implementations selected by the locality-aware model and the approximate-mining based model is up to 46.38× and 62.35× (on average 7.12× and 10.92×) faster than those selected by AutoMine’s cost model.

The profiling step of approximate-mining based model is very fast, for CiteSeer (4.5K), MiCo (1.1M), Patents (16.5M), LiveJournal (42.9M), and Friendster (1.8B) (number of edges in parentheses), the profiling only takes 1.96s, 3.50s, 6.64s, 7.14s, and 7.10s. The benefits of the more accurate cost estimation always overshadow the small profiling cost.

7 DECOMINE COMPILER

The overall compiler workflow is shown in Figure 12. Based on the template in Algorithm 1, the front-end generates the abstract syntax tree (AST) of the algorithm for each combination of the vertex cutting set and the vertex matching order; the middle-end performs algorithm-oblivious optimizations on the AST, whose performance is estimated by a cost model. The back-end generates the C++ GPM codes based on the AST with the lowest cost. The compilation overhead is very low. For 6-motif, a complicated applications with 112 patterns, compilation takes less than 50ms while running the application on the WikiVote graph (7K vertices) takes 42 minutes.

7.1 AST Intermediate Representation

We use the Abstract Syntax Tree (AST) as the Intermediate Representation (IR) to capture the vertex-set-based matching process

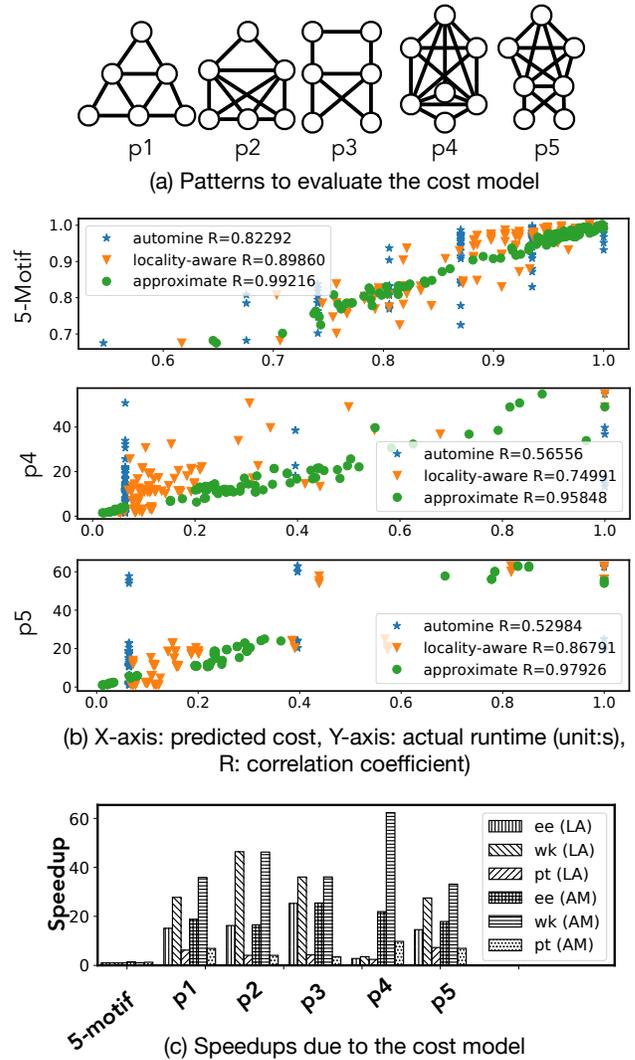


Figure 11: Cost Model Comparison

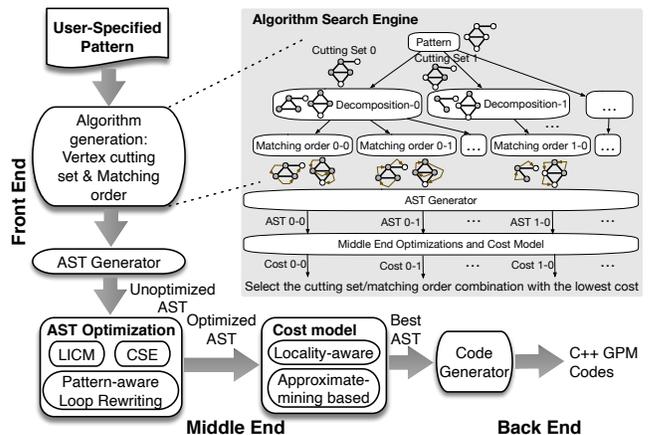


Figure 12: DecoMine Compiler and Algorithm Search

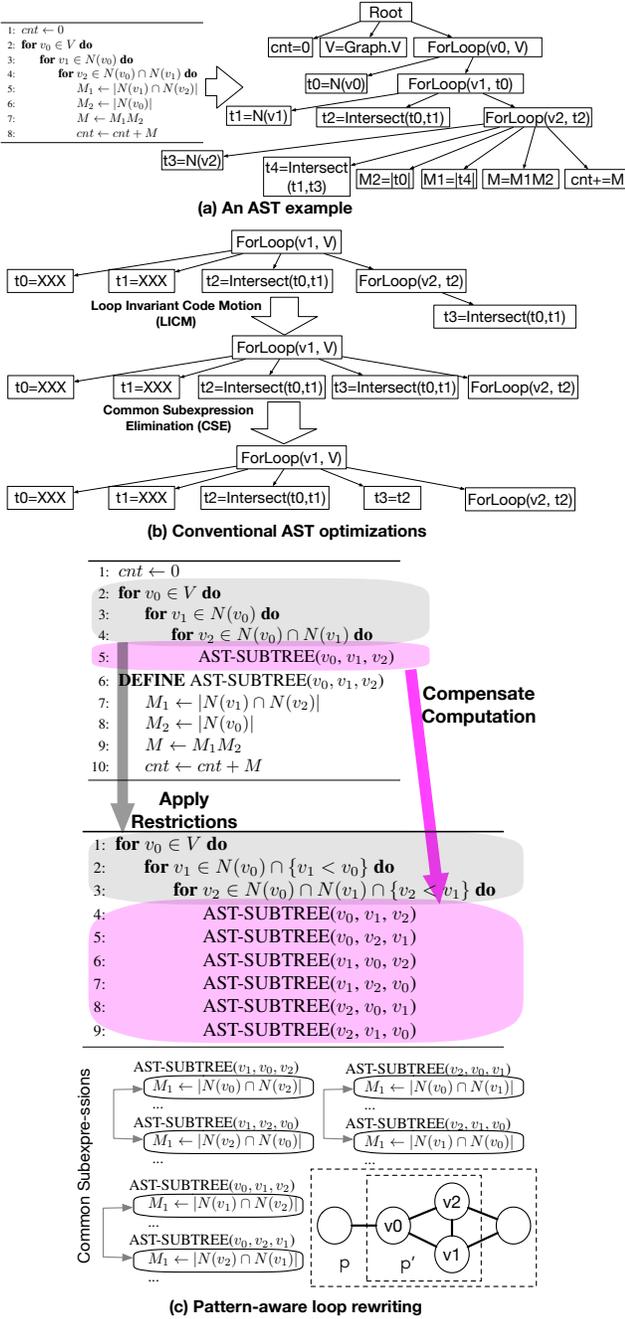


Figure 13: AST-Level Optimizations

of subpatterns. The AST of a vertex-set-based pattern matching process contains five types of nodes. 1) Loop nodes, which iterate a vertex-ID variable over a vertex set. 2) Vertex-set operation nodes, which perform an operation to generate a new vertex-set variable. Currently, supported operations include set intersection, subtraction, copy assignment and trimming (eliminating the elements smaller/larger than a lower/upper-bound). We also support a special vertex-set loading operation that takes a vertex-ID variable

as input and loads the corresponding neighbor vertex set. 3) Arithmetic operation nodes, which perform an operation (e.g., addition, multiplication) on scalar variable(s) and produce an output. 4) Hash table operation nodes, which perform query/update operations to a hash table variable. 5) A virtual root node. Figure 13 shows the AST for a nested loop (a) and two conventional optimizations (b). Global variables are defined within the root node. We require that the update to global variables should be associative and commutative, which is critical to the correctness of our parallelization strategy. It is satisfied in GPM applications since the ordering of embedding enumeration does not matter.

7.2 Pattern-Aware Loop Rewriting

For a pattern graph p , if a subgraph p' is symmetric, we can remove the redundant enumeration within p' , while preserving the same enumeration for p as no symmetry breaking. The key idea is to only enumerate the subgraphs that match p' once by enforcing restrictions and *compensate* the skipped computations for p . Figure 13 (c) shows an example, where p contains p' , a triangle, which is symmetric, we can first enumerate the subgraphs matching the triangle once but compensate the computation to extend matched triangles to embeddings of p for *six* times—the multiplicity of triangle pattern.

Since pattern enumeration is expressed as nested loops, this optimization is essentially *pattern-aware loop rewriting* (PLR). Figure 13 (c) shows the transformation based on the example, the extension of triangles matched in the first three loops are expressed as an AST subtree. With restrictions, $v_1 < v_0$ and $v_2 < v_1$, the AST subtree is immediately scheduled five additional times for equivalent triangles eliminated by the restrictions. Compared to codes without symmetry breaking, PLR improves performance for two reasons: 1) eliminate redundant enumeration for p' ; and more interestingly, 2) increase the opportunity of CSE when multiple iterations are scheduled together in the compensation. We can see 2) from Figure 13 (c), when the AST-SUBTREE are grouped together, the paired intersections only need to perform once.

When the whole pattern p is symmetric, PLR still performs strictly more computation than applying standard symmetry breaking on p . We attempt to apply PLR to the first k loops of the vertex cutting set enumeration that are shared by all subpatterns. If PLR is possible, the cost of the optimized AST is estimated by the cost model to determine: 1) whether PLR is beneficial; and 2) the best k value. The reason that PLR may lead to worse performance is that, the benefit of eliminating redundant enumeration in p' is overshadowed by code expansion for compute compensation. Note that there might be symmetry unexploited outside the subpattern p' . We leave exploiting such symmetry as our future work.

7.3 Algorithm Search

DecoMine's algorithm search is performed on optimized ASTs in a search space determined by two algorithm-level decisions: (1) how to decompose a pattern; and (2) the matching order defined in Section 5. Generating all decomposition candidates of a pattern with n vertices and m edges consists of two steps. The first step generates all cutting set candidates by brute force. All 2^n subsets of the n vertices are enumerated, and each of them is checked to

see whether it can break the pattern. This step takes $O(2^n(n+m))$ runtime. The second step constructs the subpatterns and shrinkage patterns for each cutting set.

7.4 Code Generation

The compiler back-end generates C++ implementations. We choose to parallelize the outer-most loop, whose iterations are statically divided and assigned to computation threads before execution. To achieve load balance, the compiler uses fine-grained work stealing to allow idle threads to steal iterations that are statically assigned but have not been executed by others. We use privatization to ensure thread-safe updates to global variables. Each thread maintains a private copy for each global variable to accumulate the updates during the execution. The copy is updated to the corresponding global variable using mutex lock. The associative and commutative updates to global variables (Section 7.1) ensures the correctness of the implementation. When multiple patterns are enumerated concurrently, the compiler performs computation reuse optimization across different patterns.

7.5 Support to Constraints on Labels

With the generalized pattern decomposition method, DecoMine is able to support label constraints that can be represented by one or multiple sub-constraints, each of which applies to a part of an embedding. If a constraint $F(e)$ (e is an embedding) can be expressed as $F_1(e_1) \wedge F_2(e_2) \wedge \dots \wedge F_k(e_k)$, in which e_i ($1 \leq i \leq k$) is a fragment of e , DecoMine can carefully choose a cutting set so that each sub-constraint can be resolved by only partially materializing e . For example, a constraint (on the pattern in Figure 6) like “A, B and C must have different labels and B, D, E must have the same label” can be supported. If $F(e)$ cannot be converted into this representation, or such a cutting set doesn’t exist, the system has to fall back to non-decomposition methods since resolving the constraint depends on all embedding vertices.

8 EXPERIMENTS

8.1 Evaluation Methodology

System configuration. We conducted experiments on an 8-node cluster. Each node has two 8-core Intel Xeon E5-2630 V3 CPUs and 64GB DDR4 RAM. The experiments of Arabesque and Fractal use the whole cluster. All other systems except for Pangolin-GPU are tested on one node with an additional 2TB NVMe SSD. Pangolin-GPU uses an NVIDIA V100-32GB GPU. Unless specified, the reported runtime is the average of three runs excluding graph loading time. By default, DecoMine uses the approximate mining based cost model.

Graph mining applications. *Motif Counting (MC)* [45] aims at counting all connected vertex-induced patterns with a particular size. *Frequent Subgraph Mining (FSM)* [1, 11, 19, 28] discovers all frequent labeled patterns in an input graph. *Pseudo Clique Mining (PC)* [57] counts vertex-induced pseudo clique patterns of a given size. A pattern is a pseudo clique if the number of its edges is no less than $n(n-1)/2 - k$, in which n is the number of vertices in the pattern, and k is a pre-defined parameter. All pseudo cliques with n vertices can be obtained by deleting at most k edges from an n -clique. In our experiments, we choose $k = 1$.

Table 1: Graph Datasets [36]

Graph	Abbr.	V	E	L
CiteSeer [5, 20, 47]	cs	3.3K	4.5K	6
EmailEuCore [35, 64]	ee	1.0K	16.1K	42
WikiVote [33]	wk	7.1K	100.8K	N/A
MiCo [19]	mc	96.6K	1.1M	29
Patents [34]	pt	3.8M	16.5M	N/A
LiveJournal [4, 37]	lj	4.8M	42.9M	N/A
Friendster [63]	fr	65.6M	1.8B	N/A
RMAT-100M[12]	rmat	100M	1.6B	N/A

Graph datasets. Table 1 shows the graph datasets. The largest dataset is Friendster with roughly 1.8 billion edges. The RMAT-100M graph is synthesized by the RMAT generator[12, 29] using default parameters. We preprocessed all datasets to delete duplicated edges and self-loops.

In-house Automine implementation. We implemented our own

Table 2: In-House Automine vs. Automine Runtime in [41]

App	Graph	Our Impl.	Original Impl.
3-MC	wk	27.3ms	34.5ms
	mc	161ms	230ms
	pt	0.9s	1.9s
	lj	9.0s	13.4s
4-MC	wk	7.0s	11.5s
	mc	31.7s	45.2s
	pt	24.3s	82.1s
5-MC	lj	457m	367m
	wk	4345s	5300s
	mc	2.91h	5.56h
	pt	54m	117m

Automine (AutoMineInHouse) with all optimizations in [42]. Table 2 compares the performance of AutoMine reported in [41] and AutoMineInHouse using a similar machine. AutoMineInHouse is faster in almost all cases except for the 4-motif counting on the lj graph, which may be due to the minor hardware discrepancy.

8.2 Overall Performance

Table 3 compares DecoMine to AutoMineInHouse, RStream [58], and Arabesque [55]. The runtimes exclude graph loading and profiling time as they can be amortized with multiple applications. In both k -MC and k -PC, k refers to the pattern size. In FSM, the support thresholds are 300 and 3000. Since FSM requires the labeled input graph, we only evaluate it on CiteSeer, EmailEuCore, and MiCo. Similar to previous works [15, 26], we only discover frequent patterns with less than four edges. We do not run the Pseudo-Clique Counting experiments for RStream and Arabesque due to the lack of reference implementations. DecoMine outperforms other systems significantly, achieving speedups of up to 827 \times , 882, 667 \times , and 72, 143 \times over Automine, RStream, and Arabesque, respectively. An important reason why existing systems like RStream and Arabesque are so slow is that they have to materialize all embeddings, which is usually not required by applications. For example, MC/PC only

Table 3: Comparing with Automine/RStream/Arabesque. T: Timeout (12 h) C: Crashed (out of memory/disk space)

App.	G	DecoMine	AutomineInHouse	RStream	Arabesque (8-node)
3-MC	cs	0.14ms	0.16ms (1.1x)	142ms (1,014x)	10.1s (72,143x)
	ee	0.87ms	7.3ms (8.4x)	21.0s (24,138x)	10.2s (11,724x)
	wk	7ms	27.3ms (3.9x)	17.9m (153,428x)	12.1s (1,729x)
	pt	332ms	931ms (2.8x)	104.1m (18,813x)	96.4s (290x)
	mc	48ms	161ms (3.4x)	144.8m (181,000x)	21.1s (440x)
	lj	2.7s	9.0s (3.3x)	T	24.3m (540x)
4-MC	cs	0.17ms	4.8ms (28x)	3.7s (21,765x)	9.9s (58,235x)
	ee	9ms	920ms (102x)	132.4m (882,667x)	19.1s (2,122x)
	wk	60ms	7.0s (117x)	T	402.2s (6,703x)
	pt	1.5s	24.3s (16x)	T	68.3m (2,732x)
	mc	1.3s	31.7s (24x)	T	42.8m (1,975x)
	lj	33.1s	456.5m (827x)	T	C
5-MC	cs	2.1ms	332ms (158x)	146.4s (69,714x)	11.4s (5,429x)
	ee	416ms	104.8s (252x)	T	19.4m (2,798x)
	wk	8.1s	72.4m (536x)	T	C
	pt	36.9s	53.9m (88x)	T	C
	mc	111.6s	174.6m (94x)	T	C
	lj	167.7m	T	T	C
6-MC	cs	270ms	35.9s (133x)	108.7m (24,156x)	48.7s (180x)
	ee	106.6s	259.0m (146x)	T	C
	wk	42.2m	T	T	C
	pt	63.0m	T	T	C
7-PC	cs	0.3ms	0.5ms (1.7x)		
	ee	719ms	67.1s (93x)		
	wk	735ms	90.8s (24x)		
	pt	499ms	15.7s (31x)		
8-PC	cs	0.3ms	0.5ms (1.7x)		
	ee	1.3s	433.1s (322x)		
	wk	1.2s	463.0s (387x)		
	pt	582ms	86.2s (148x)		
FSM-300	cs	2.6ms	7.7ms (3.0x)	522ms (201x)	10.3s (3,962x)
	ee	0.3ms	0.3ms (1.0x)	3.6s (12,000x)	9.6s (32,000x)
	mc	210.8s	242.7s (1.2x)	149.1m (42x)	C
FSM-3K	cs	0.3ms	0.3ms (1.0x)	77.9ms (260x)	9.6s (32,000x)
	ee	0.3ms	0.3ms (1.0x)	210ms (700x)	9.8s (32,667x)
	mc	513ms	30.0s (58x)	141.9m (16,596x)	157.9s (308x)

need the counts of patterns while FSM only needs the pattern vertex domains. In contrast, the flexibility of DecoMine allows users to match the need of applications more precisely and thus avoid unnecessary materialization for higher performance.

For MC and PC, on the same graph, for larger pattern sizes, DecoMine can achieve higher speedups over Automine and RStream. Compared to Arabesque, the speedup is higher when the pattern size is small due to Arabesque’s startup overhead. When the pattern size is larger, this performance issue is mitigated since the overhead is amortized and thus the speedup over Arabesque decreases.

For FSM, the performance of DecoMine and Automine on the cs (support=3K) and ee graphs is similar. Both datasets are tiny graphs. Hence, almost all labeled patterns are filtered away even with a relatively low support threshold, i.e., 3K. As a result, only trivial startup computation overhead is left, leading to a similar performance of Automine and DecoMine.

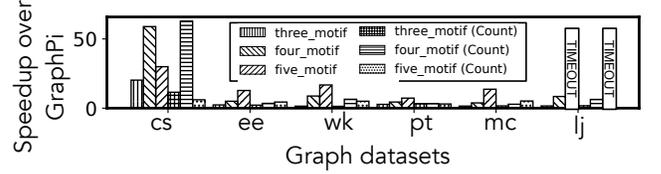
We also compare DecoMine with Peregrine [26], Pangolin [15] and Fractal [17]. We run the four systems for MC and FSM. We evaluate MC on cs/pt/mc and FSM on mc/lj (lj with randomly synthesized labels). As shown in Table 4, DecoMine is consistently faster, and achieves up to 575 \times , 328 \times and 42, 143 \times speedup over Peregrine, Pangolin-CPU and Fractal, respectively. Pangolin ran out of memory and crashed for many benchmarks due to its memory-demanding BFS exploration. Pangolin-GPU’s performance is competitive with DecoMine. Nevertheless, it runs on a significantly more expensive device (NVIDIA V100-32GB). It is around 8K USD,

6.15 \times costly compared with the CPUs used by DecoMine (two Intel Xeon E5-2630 v3, 1.3K USD). More importantly, Pangolin-GPU fail to run on large graphs/patterns due to the GPU memory limitation (marked as "C").

Table 4: DecoMine vs. Peregrine and Pangolin

App.	G	DecoMine	Peregrine	Pangolin(CPU/GPU)	Fractal (8-node)
3-MC	cs	0.14ms	5.8ms	5.0ms / 0.1ms	5.9s (42,143x)
	pt	332ms	1.4s	1.4s / 0.2s	79.7s
	mc	48ms	60ms	280ms / 14.1ms	12.9s
4-MC	cs	0.17ms	21.2ms	15.3ms / 0.7ms	6.0s
	pt	1.5s	11.2s	329.5s / 8.0s	141.6s
	mc	1.3s	5.3s	242.7s / 3.7s	58.4s
5-MC	cs	2.1ms	41.7ms	688.3ms (328x) / 1.3ms	6.1s
	pt	36.9s	513.6s	C / C	4517.0s
	mc	111.6s	5,635.1s	C / C	1240.0s
6-MC	cs	270ms	0.8s	14.9s / C	4.6s
	pt	63.0m	T	C / C	T
FSM-300	cs	210.8s	C	C / C	280.2s
	FSM-1K	3.1s	1,782.2s (575x)	C / C	169.1s
	FSM-3K	513ms	189.3s	C / C	109.4s
FSM-1.0M	cs	27.8s	T	C / C	T
	FSM-1.5M	27.8s	T	C / C	3.4h
	FSM-2.0M	27.8s	T	C / C	3.3h

We compare DecoMine with GraphPi [52] in Figure 14 for motif mining applications. We do not use FSM since GraphPi does not supported it. The version marked with “(count)” indicates that the pattern counting mathematical optimization of GraphPi is enabled,

**Figure 14: Comparing with GraphPi**

which quickly evaluates the number of iterations of some certain loops. Our system consistently outperforms GraphPi, and achieves up to 62.8 \times speedup. GraphPi’s mathematical optimization significantly improves its performance but is only suitable for pattern counting.

Comparing with the efficient native algorithm. We compare DecoMine with ESCAPE [44] in Table 5, the fastest expert-tailored single-thread decomposition-based 4/5-motif counting implemen-

Table 5: DecoMine vs. Native Algorithm

App.	Graph	DecoMine	GraphPi	ESCAPE
#Cores		16	1	1
4-MC	ee	9ms	95ms	397ms
	wk	60ms	879ms	5.8s
	pt	1.5s	19.9s	62.4s
5-MC	ee	416ms	5.4s	26.5s
	wk	8.1s	121.0s	617.2s
	pt	36.9s	557.5s	1719.5s

tation [46] with pattern-specific optimizations. With one thread, DecoMine is on average 4.0 \times slower than ESCAPE, while GraphPi,

even with their mathematical optimization that only applies to pattern counting, is on average $17.3\times$ slower (up to $48.6\times$). With multicores, DecoMine runs faster than ESCAPE. Note that ESCAPE still outperforms DecoMine in single-thread performance because it adopts an algorithmic optimization that converts each size-5 pattern to a DAG to reduce the exploration space.

8.3 Pattern-Aware Loop Rewriting

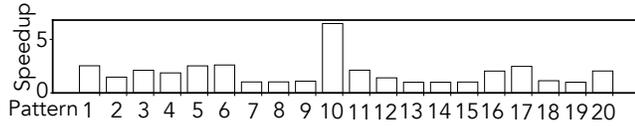


Figure 15: Speedups due to PLR

We generate the counting application for each size-5 pattern except for the 5-clique (since 5-clique cannot benefit from pattern decomposition) (in total 20 patterns) with and without PLR enabled, and run them on the Patents graph. Figure 15 shows the speedup brought about by PLR. The optimization improves the performance by up to $6.5\times$, and benefits more than a half of size-5 patterns.

8.4 Scalability

Table 6: Large Graphs

G.	V	E	System	Runtime
fr	65.6M	1.8B	DecoMine	1.4h
			Peregrine	29.1h
			GraphPi	15.4h
rmat	100M	1.6B	DecoMine	1.7h
			Peregrine	39.7h
			GraphPi	10.2h

We evaluate DecoMine’s capability to scale to large graphs by comparing it with Peregrine and GraphPi, for 4-motif mining on fr and rmat, both with more than one billion edges. For GraphPi, the pattern counting mathematical optimization is enabled to achieve all its performance benefits. We report the runtimes in Table 6. For large workloads like motif mining on large graphs, DecoMine can significantly reduce execution time from tens of hours in recent systems to less than two hours.

Table 7: Large Patterns (d: days)

G.	App.	DecoMine	Peregrine	GraphPi
ee	6-cycle	3.4s	102.7s	64.8s
	7-cycle	249.4s	6131.9s	3,674.7s
	8-cycle	5.7h	5.6d	2.8d
pt	6-cycle	370.2s	6913.9s	1960.0s
	7-cycle	4.4h	2.9d	23.3h
wk	6-cycle	136.2s	5754.9s	3,248.6s
	7-cycle	4.8h	>1 week	4.0d

We use cycle mining on ee, pt and wk to evaluate the scalability of DecoMine to large patterns. k -cycle mining counts the number

of size- k cycles of the input graph. We keep increasing the pattern size (i.e., k) until our system cannot finished it within 24 hours. The runtime results are reported in Table 7. For comparison purposes, we also report the execution times of Peregrine and GraphPi for the same workloads. The results show that DecoMine can easily mine large patterns, e.g., 8-cycle on the ee graph, within a few hours while the other two systems usually takes *days* to complete. For 7-cycle on wk, DecoMine reduces the execution time *from more than one weeks on Peregrine to less than five hours*.

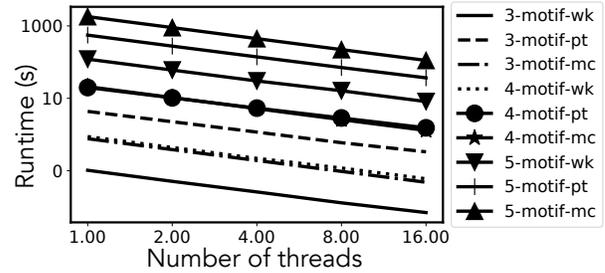


Figure 16: Scalability with MT

The scalability of DecoMine with multi-thread is shown in Figure 16. DecoMine achieves almost linear scalability. The single-thread and 16-thread runtimes of DecoMine for 5-Motif on Patents are 557.5 seconds and 36.9 seconds, respectively—a speedup of $15.11\times$.

8.5 FSM with Various Support Thresholds

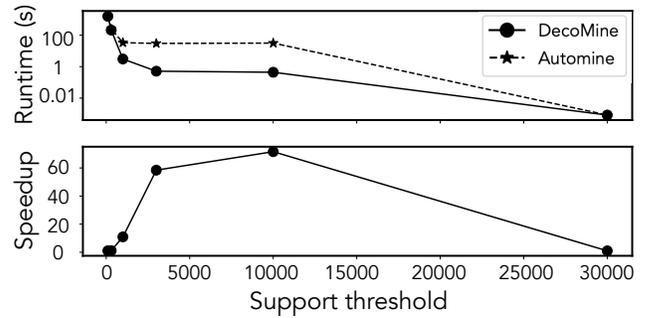


Figure 17: FSM Sensitivity

We run DecoMine and AutomineInHouse for FSM with various support thresholds ranging from 100 to 30K on the MiCo graph to analyze the performance sensitivity with respect to thresholds. The runtimes and DecoMine’s speedups over AutomineInHouse are reported in Figure 17. DecoMine is consistently faster than AutoMineInHouse in all settings. We also see that the speedups are small with extremely large and small thresholds, while the a peak speedup of roughly $70\times$ is reached when the threshold is 10K. For extremely large thresholds like 30K, almost all patterns are filtered away by the exceedingly high threshold, and hence only trivial computation cost is left, leading to the similar performance of DecoMine and AutomineInHouse. On the other side, when the threshold is very small (e.g., 100), there are a huge number of light-weight labeled patterns to be processed and hence FSM is mainly

bottlenecked by the per-pattern overhead (e.g., clearing data structures, launching computation threads), which cannot be accelerated by pattern decomposition.

8.6 Workloads with Label Constraints

To demonstrate DecoMine’s ability to handle workloads with complicated label constraints, we evaluate a graph query “counting the subgraphs matching pattern p in Figure 6, in which vertices matching A, B, C must have different labels and vertices matching B, D, E must have the same label” on DecoMine and Peregrine. Thanks to DecoMine’s ability to resolve constraints on partially-materialized embeddings, the system is able to achieve significantly higher performance than Peregrine: for (cs,ee,mc,lj) graphs, the runtimes of DecoMine are (0.35ms,43ms,11.9s,288.4s), while the runtimes of Peregrine are (2.2ms,975ms,2030.9s,>12h).

8.7 Compilation Cost

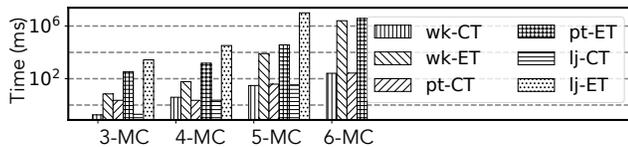


Figure 18: Compilation Time (CT) vs. Execution Time (ET)

The compilation overhead of DecoMine is negligible compared to the execution time. We compare the compilation and execution time for 3/4/5/6-MC on wk, pt, and lj in Figure 18. The trends on other workloads are similar and we omit them due to space limit. The compilation times are *orders of magnitude smaller* than execution times. Even for 6-MC, the most complicated workloads evaluated in this paper (with 112 size-6 patterns), the compilation takes less than 300ms while the execution on a median-size graph (e.g., wk) takes tens of minutes.

8.8 Detailed Cost Model Analysis

To analyze the contribution of the improved cost models in a pattern decomposition based system, in Figure 19, we compare the runtimes of Automine with a perfect cost model (AM-OPT) (obtained by trying all matching orders on evaluated graphs) and DecoMine with Automine’s cost model (DM-Auto), the locality-aware model (DM-LA) and the approximate mining based model (DM-AM) for patterns p_1 , p_2 and p_3 in Figure 11 (a). There are two key observations. First, even with an ideally perfect cost model, Automine (AM-OPT) is significantly slower than DecoMine with an improved cost model (DM-LA/AM). It confirms that algorithmic advantage is the fundamental reason for DecoMine’s high performance. Second, with an inaccurate cost model (DM-Auto), DecoMine can be indeed slower than a system without decomposition (AM-OPT) because a bad cutting set is selected. It confirms the importance of an accurate cost model in DecoMine.

9 CONCLUSION

This paper proposes a new compilation-based GPM system that adopts efficient pattern decomposition algorithms. The compiler

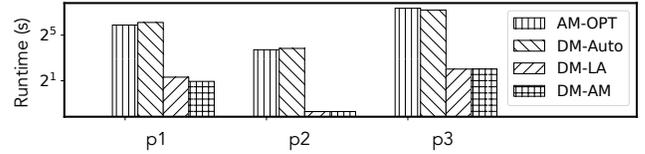


Figure 19: Comparing DecoMine with Automine with an Optimal Cost Model (wk graph)

generates algorithms with conventional and GPM-specific optimizations for different decomposition choices, which are evaluated based on an accurate cost model. The executable of the GPM task is obtained from the algorithm with the best performance. We propose a novel partial-embedding API that is sufficient to construct advanced GPM applications while preserving pattern decomposition algorithm advantages. The key insight is that the system does not need to materialize the embedding of the whole pattern while maintaining correctness with two properties. Compared to state-of-the-art systems, DecoMine reduces the execution time of GPM on large graphs and patterns from days to a few hours with low programming effort.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their insightful comments and suggestions. This work is supported by National Science Foundation (Grant No. CCF-2127543, CCF-1750656, and CCF-1717754).

REFERENCES

- [1] Ehab Abdelhamid, Ibrahim Abdelaziz, Panos Kalnis, Zuhair Khayyat, and Fuad Jamour. 2016. Scalemine: Scalable parallel frequent subgraph mining in a single large graph. In *SC'16: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 716–727.
- [2] Nesreen K Ahmed, Jennifer Neville, Ryan A Rossi, and Nick Duffield. 2015. Efficient graphlet counting for large networks. In *2015 IEEE International Conference on Data Mining*. IEEE, 1–10.
- [3] Khaled Ammar, Frank McSherry, Semih Salihoglu, and Manas Joglekar. 2018. Distributed evaluation of subgraph queries using worstcase optimal lowmemory dataflows. *arXiv preprint arXiv:1802.03760* (2018).
- [4] Lars Backstrom, Dan Huttenlocher, Jon Kleinberg, and Xiangyang Lan. 2006. Group formation in large social networks: membership, growth, and evolution. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*. 44–54.
- [5] David A Bader, Henning Meyerhenke, Peter Sanders, and Dorothea Wagner. 2012. Graph Partitioning and Graph Clustering. In *10th DIMACS Implementation Challenge Workshop*.
- [6] Luca Becchetti, Paolo Boldi, Carlos Castillo, and Aristides Gionis. 2008. Efficient semi-streaming algorithms for local triangle counting in massive graphs. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*. 16–24.
- [7] Bibek Bhattacharai, Hang Liu, and H Howie Huang. 2019. Ceci: Compact embedding cluster index for scalable subgraph matching. In *Proceedings of the 2019 International Conference on Management of Data*. 1447–1462.
- [8] Fei Bi, Lijun Chang, Xuemin Lin, Lu Qin, and Wenjie Zhang. 2016. Efficient subgraph matching by postponing cartesian products. In *Proceedings of the 2016 International Conference on Management of Data*. 1199–1214.
- [9] Marco Bressan, Flavio Chierichetti, Ravi Kumar, Stefano Leucci, and Alessandro Panconesi. 2018. Motif counting beyond five nodes. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 12, 4 (2018), 1–25.
- [10] Marco Bressan, Stefano Leucci, and Alessandro Panconesi. 2019. Motivo: fast motif counting via succinct color coding and adaptive sampling. *Proceedings of the VLDB Endowment* 12, 11 (2019), 1651–1663.
- [11] Björn Bringmann and Siegfried Nijssen. 2008. What is frequent in a single graph?. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer, 858–863.
- [12] Deepayan Chakrabarti, Yiping Zhan, and Christos Faloutsos. 2004. R-MAT: A recursive model for graph mining. In *Proceedings of the 2004 SIAM International*

- Conference on Data Mining*. SIAM, 442–446.
- [13] Hongzhi Chen, Miao Liu, Yunjian Zhao, Xiao Yan, Da Yan, and James Cheng. 2018. G-Miner: an efficient task-oriented graph mining system. In *Proceedings of the Thirteenth EuroSys Conference*. 1–12.
 - [14] Xuhao Chen, Roshan Dathathri, Gurbinder Gill, Loc Hoang, and Keshav Pingali. 2021. SandSlash: a two-level framework for efficient graph pattern mining. In *Proceedings of the ACM International Conference on Supercomputing*. 378–391.
 - [15] Xuhao Chen, Roshan Dathathri, Gurbinder Gill, and Keshav Pingali. 2020. Pangolin: an efficient and flexible graph mining system on CPU and GPU. *Proceedings of the VLDB Endowment* 13, 10 (2020), 1190–1205.
 - [16] Maximilien Danisch, Oana Balalau, and Mauro Sozio. 2018. Listing k-cliques in sparse real-world graphs. In *Proceedings of the 2018 World Wide Web Conference*. 589–598.
 - [17] Vinicius Dias, Carlos HC Teixeira, Dorgival Guedes, Wagner Meira, and Srinivasan Parthasarathy. 2019. Fractal: A general-purpose graph pattern mining system. In *Proceedings of the 2019 International Conference on Management of Data*. 1357–1374.
 - [18] Alexandra Duma and Alexandru Topirceanu. 2014. A network motif based approach for classifying online social networks. In *2014 IEEE 9th IEEE International Symposium on Applied Computational Intelligence and Informatics (SACI)*. IEEE, 311–315.
 - [19] Mohammed Elseidy, Ehab Abdelhamid, Spiros Skiadopoulos, and Panos Kalnis. 2014. Grami: Frequent subgraph and pattern mining in a single large graph. *Proceedings of the VLDB Endowment* 7, 7 (2014), 517–528.
 - [20] Robert Geisberger, Peter Sanders, and Dominik Schultes. 2008. Better Approximation of Betweenness Centrality. In *ALENEX*. SIAM, 90–100.
 - [21] Ilias Giachasiel, George Panagopoulos, and Eiko Yoneki. 2015. PDL: Parallel and distributed triangle listing for massive graphs. In *2015 44th International Conference on Parallel Processing*. IEEE, 370–379.
 - [22] Joshua A Grochow and Manolis Kellis. 2007. Network motif discovery using subgraph enumeration and symmetry-breaking. In *Annual International Conference on Research in Computational Molecular Biology*. Springer, 92–106.
 - [23] Wook-Shin Han, Jinsoo Lee, and Jeong-Hoon Lee. 2013. Turboiso: towards ultrafast and robust subgraph isomorphism search in large graph databases. In *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data*. 337–348.
 - [24] Loc Hoang, Vishwesh Jatala, Xuhao Chen, Udit Agarwal, Roshan Dathathri, Gurbinder Gill, and Keshav Pingali. 2019. DistTC: High performance distributed triangle counting. In *2019 IEEE High Performance Extreme Computing Conference (HPEC)*. IEEE, 1–7.
 - [25] Anand Padmanabha Iyer, Zaoxing Liu, Xin Jin, Shivaram Venkataraman, Vladimir Braverman, and Ion Stoica. 2018. {ASAP}: Fast, Approximate Graph Pattern Mining at Scale. In *13th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 18)*. 745–761.
 - [26] Kasra Jamshidi, Rakesh Mahadasa, and Keval Vora. 2020. Peregrine: a pattern-aware graph mining system. In *Proceedings of the Fifteenth European Conference on Computer Systems*. 1–16.
 - [27] Madhav Jha, C Seshadhri, and Ali Pinar. 2015. Path sampling: A fast and provable method for estimating 4-vertex subgraph counts. In *Proceedings of the 24th international conference on world wide web*. 495–505.
 - [28] Chuntao Jiang, Frans Coenen, and Michele Zito. 2013. A survey of frequent subgraph mining algorithms. *The Knowledge Engineering Review* 28, 1 (2013), 75–105.
 - [29] Farzad Khorasani, Rajiv Gupta, and Laxmi N. Bhuyan. 2015. Scalable SIMD-Efficient Graph Processing on GPUs. In *Proceedings of the 24th International Conference on Parallel Architectures and Compilation Techniques (PACT '15)*. 39–50.
 - [30] Hyeonji Kim, Juneyoung Lee, Sourav S Bhowmick, Wook-Shin Han, JeongHoon Lee, Seongyun Ko, and Moath HA Jarrah. 2016. Dualsim: Parallel subgraph enumeration in a massive graph on a single machine. In *Proceedings of the 2016 International Conference on Management of Data*. 1231–1245.
 - [31] Jinha Kim, Wook-Shin Han, Sangyeon Lee, Kyungyeol Park, and Hwanjo Yu. 2014. OPT: A new framework for overlapped and parallel triangulation in large-scale graphs. In *Proceedings of the 2014 ACM SIGMOD international conference on Management of data*. 637–648.
 - [32] Longbin Lai, Lu Qin, Xuemin Lin, and Lijun Chang. 2015. Scalable subgraph enumeration in mapreduce. *Proceedings of the VLDB Endowment* 8, 10 (2015), 974–985.
 - [33] Jure Leskovec, Daniel Huttenlocher, and Jon Kleinberg. 2010. Signed networks in social media. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 1361–1370.
 - [34] Jure Leskovec, Jon Kleinberg, and Christos Faloutsos. 2005. Graphs over time: densification laws, shrinking diameters and possible explanations. In *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining*. 177–187.
 - [35] Jure Leskovec, Jon Kleinberg, and Christos Faloutsos. 2007. Graph evolution: Densification and shrinking diameters. *ACM transactions on Knowledge Discovery from Data (TKDD)* 1, 1 (2007), 2–es.
 - [36] Jure Leskovec and Andrej Krevl. 2014. SNAP Datasets: Stanford Large Network Dataset Collection. <http://snap.stanford.edu/data>.
 - [37] Jure Leskovec, Kevin J Lang, Anirban Dasgupta, and Michael W Mahoney. 2009. Community structure in large networks: Natural cluster sizes and the absence of large well-defined clusters. *Internet Mathematics* 6, 1 (2009), 29–123.
 - [38] Can Lu, Jeffrey Xu Yu, Hao Wei, and Yikai Zhang. 2017. Finding the maximum clique in massive graphs. *Proceedings of the VLDB Endowment* 10, 11 (2017), 1538–1549.
 - [39] Shuai Ma, Yang Cao, Jimpeng Huai, and Tianyu Wo. 2012. Distributed graph pattern matching. In *Proceedings of the 21st international conference on World Wide Web*. 949–958.
 - [40] Avi Ma'ayan. 2009. Insights into the organization of biochemical regulatory networks using graph theory analyses. *Journal of Biological Chemistry* 284, 9 (2009), 5451–5455.
 - [41] Daniel Mawhirter, Sam Reinehr, Connor Holmes, Tongping Liu, and Bo Wu. 2021. GraphZero: A High-Performance Subgraph Matching System. *ACM SIGOPS Operating Systems Review* 55, 1 (2021), 21–37.
 - [42] Daniel Mawhirter and Bo Wu. 2019. AutoMine: harmonizing high-level abstraction and high performance for graph mining. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles*. 509–523.
 - [43] Rasmus Pagh and Charalampos E Tsourakakis. 2012. Colorful triangle counting and a mapreduce implementation. *Inform. Process. Lett.* 112, 7 (2012), 277–281.
 - [44] Ali Pinar, C Seshadhri, and Vaidyanathan Vishal. 2017. Escape: Efficiently counting all 5-vertex subgraphs. In *Proceedings of the 26th International Conference on World Wide Web*. 1431–1440.
 - [45] Nataša Pržulj. 2007. Biological network comparison using graphlet degree distribution. *Bioinformatics* 23, 2 (2007), e177–e183.
 - [46] Pedro Ribeiro, Pedro Paredes, Miguel EP Silva, David Aparicio, and Fernando Silva. 2019. A Survey on Subgraph Counting: Concepts, Algorithms and Applications to Network Motifs and Graphlets. *arXiv preprint arXiv:1910.13011* (2019).
 - [47] Ryan A. Rossi and Nesreen K. Ahmed. 2015. The Network Data Repository with Interactive Graph Analytics and Visualization. In *AAAI*. <http://networkrepository.com>
 - [48] Matthew C Schmidt, Andrea M Rocha, Kanchana Padmanabhan, Zhengzhang Chen, Kathleen Scott, James R Mihelcic, and Nagiza F Samatova. 2011. Efficient α , β -motif finder for identification of phenotype-related functional modules. *BMC bioinformatics* 12, 1 (2011), 440.
 - [49] Marco Serafini, Gianmarco De Francisci Morales, and Georgos Siganos. 2017. Qfrag: Distributed graph search via subgraph isomorphism. In *proceedings of the 2017 symposium on cloud computing*. 214–228.
 - [50] Yingxia Shao, Bin Cui, Lei Chen, Lin Ma, Junjie Yao, and Ning Xu. 2014. Parallel subgraph listing in a large-scale graph. In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data*. 625–636.
 - [51] Daron R Shaw. 1999. The methods behind the madness: Presidential electoral college strategies, 1988-1996. *The Journal of Politics* 61, 4 (1999), 893–913.
 - [52] Tianhui Shi, Mingshu Zhai, Yi Xu, and Jidong Zhai. 2020. GraphPi: high performance graph pattern matching through effective redundancy elimination. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. 1–14.
 - [53] Julian Shun and Kanat Tangwongsan. 2015. Multicore triangle computations without tuning. In *2015 IEEE 31st International Conference on Data Engineering*. IEEE, 149–160.
 - [54] Nilotpal Talukder and Mohammed J Zaki. 2016. A distributed approach for graph mining in massive networks. *Data Mining and Knowledge Discovery* 30, 5 (2016), 1024–1052.
 - [55] Carlos HC Teixeira, Alexandre J Fonseca, Marco Serafini, Georgos Siganos, Mohammed J Zaki, and Ashraf Aboulnaga. 2015. Arabesque: a system for distributed graph mining. In *Proceedings of the 25th Symposium on Operating Systems Principles*. 425–440.
 - [56] Shahadat Uddin, Liaquat Hossain, et al. 2013. Dyad and triad census analysis of crisis communication network. *Social Networking* 2, 01 (2013), 32.
 - [57] Takeaki Uno. 2010. An efficient algorithm for solving pseudo clique enumeration problem. *Algorithmica* 56, 1 (2010), 3–16.
 - [58] Kai Wang, Zhiqiang Zuo, John Thorpe, Tien Quang Nguyen, and Guoqing Harry Xu. 2018. Rstream: Marrying relational algebra with streaming for efficient graph mining on a single machine. In *13th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 18)*. 763–782.
 - [59] Pinghui Wang, Junzhou Zhao, Xiangliang Zhang, Zhenguo Li, Jiefeng Cheng, John CS Lui, Don Towsley, Jing Tao, and Xiaohong Guan. 2017. MOSS-5: A fast method of approximating counts of 5-node graphlets in large graphs. *IEEE Transactions on Knowledge and Data Engineering* 30, 1 (2017), 73–86.
 - [60] Jingen Xiang, Cong Guo, and Ashraf Aboulnaga. 2013. Scalable maximum clique computation using mapreduce. In *2013 IEEE 29th International Conference on Data Engineering (ICDE)*. IEEE, 74–85.
 - [61] Da Yan, Guimu Guo, Md Mashur Rahman Chowdhury, M Tamer Özsu, Wei-Shinn Ku, and John CS Lui. 2020. G-thinker: A distributed framework for mining subgraphs in a big graph. In *2020 IEEE 36th International Conference on Data Engineering (ICDE)*. IEEE, 1369–1380.

- [62] Xifeng Yan and Jiawei Han. 2002. *gspan*: Graph-based substructure pattern mining. In *2002 IEEE International Conference on Data Mining, 2002. Proceedings.* IEEE, 721–724.
- [63] Jaewon Yang and Jure Leskovec. 2015. Defining and evaluating network communities based on ground-truth. *Knowledge and Information Systems* 42, 1 (2015), 181–213.
- [64] Hao Yin, Austin R Benson, Jure Leskovec, and David F Gleich. 2017. Local higher-order graph clustering. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining.* 555–564.

Received 2022-03-31; accepted 2022-06-16