

SubGraph2Vec: Highly-Vectorized Tree-like Subgraph Counting

Langshi Chen Jiayu Li Ariful Azad Cenk Sahinalp Madhav Marathe Anil Vullikanti
Indiana University Indiana University Indiana University Indiana University University of Virginia University of Virginia
lc37@indiana.edu jl145@indiana.edu azad@iu.edu cenkahi@indiana.edu marathe@virginia.edu vsakumar@virginia.edu

Andrey Nikolaev Egor Smirnov Ruslan Israfilov Judy Qiu
Intel Corporation Intel Corporation Intel Corporation Indiana University
andrey.nikolaev@intel.com egor.smirnov@intel.com ruslan.israfilov@intel.com xqiu@indiana.edu

Abstract—Subgraph counting aims to count occurrences of a template T in a given network $G(V, E)$. It is a powerful graph analysis tool and has found real-world applications in diverse domains. Scaling subgraph counting problems is known to be memory bounded and computationally challenging with exponential complexity. Although scalable parallel algorithms are known for several graph problems such as Triangle Counting and PageRank, this is not common for counting complex subgraphs. Here we address this challenge and study connected acyclic graphs or trees. We propose a novel vectorized subgraph counting algorithm, named SUBGRAPH2VEC, as well as both shared memory and distributed implementations: 1) reducing algorithmic complexity by minimizing neighbor traversal; 2) achieving a highly-vectorized implementation upon linear algebra kernels to significantly improve performance and hardware utilization. 3) SUBGRAPH2VEC improves the overall performance over the state-of-the-art work by orders of magnitude and up to 660x on a single node. 4) SUBGRAPH2VEC in distributed mode can scale up the template size to 20 and maintain good strong scalability. 5) enabling portability to both CPU and GPU.

Index Terms—Subgraph Counting, Vectorization, Portability

I. INTRODUCTION

Counting tree-like subgraphs from a large network is fundamental in graph problems. It has been used in real-world applications across a range of disciplines, including:

- Social network analysis: Online social network has billion- or trillion-sized network, where a certain group of users may share specific interests [1] [2].
- Bioinformatics: The frequency or distribution of the occurrence of each different testing templates may characterize a protein-protein interaction network [3] [4], where repeated subgraphs are crucial in understanding cell physiology as well as developing new drugs. [5]
- Computing kernel of other algorithms: Sub-tree counting is one of the computing kernels of bounded treewidth subgraph (such as circles, cactus graphs, series-parallel graphs etc.) counting problem[6] and also the kernel of network clustering [7].

Despite subgraph counting plays an important role in discovery of patterns in a graph network, counting the exact number of subgraphs of size k in a n -vertex network takes $O(n^k)$ time [4], which is computationally challenging even

for moderate values of n and k . In fact, determining whether a graph G contains a subgraph to H is a related graph isomorphic problem that is NP-complete [8].

Alon et al. [9] provides an approximate algorithm, color-coding, to estimate the number of subgraphs with statistical guarantees. Although the color-coding algorithm in [3] has a time complexity linear in network size, it is exponential to subgraph size. Therefore, efficient parallel implementations are the only viable way to count subgraphs from large-scale networks. To the best of our knowledge, a multi-threaded implementation named FASCIA [4] is considered to be the state-of-the-art work in this area. Still, it takes FASCIA more than 4 days (105 hours) to count a 17-vertex subgraph from the RMAT-1M network (1M vertices, 200M edges) on a 48-core Intel (R) Skylake processor. While our proposed algorithm named SUBGRAPH2VEC takes only 9.5 minutes to complete the same task on the same hardware.

The primary contributions of this paper are as follows:

- **Algorithmic Design.** We identify and reduce the computation complexity of the sequential color-coding algorithm, which also helps reduce communication overhead in distributed systems.
- **System design and optimization.** We design a data structure as well as a thread execution model to leverage the hardware efficiency of using linear-algebra kernels in terms of vector processing units (VPU) and memory bandwidth.
- **Portability to the distributed system and GPU** We scale out our single node implementation on a distributed system with near-linear strong scalability. In addition, we export the codes to NVIDIA GPU by using NVIDIA cuSPARSE kernels thanks to our modular system design.

The codebase of our work on SUBGRAPH2VEC is made public in our open-sourced repository [10].

II. PRELIMINARIES

A. Motivation

Counting repeated subgraphs (motifs) can be used to measure topological features and further reveal the similarity of

any given two networks. Fig. 1 illustrates such a real-world application in protein-protein interaction network (PPIN), where we use SUBGRAPH2VEC to count tree-like motifs in Fig. 6(a) to estimate their frequencies. We compared the PPI networks of humans[11], yeasts, *C. elegans*[12], and *E. coli*[13], and we have findings on the normalized treelet distributions for the unicellular organisms: *E. coli* and yeasts are very close, while the more complex *C. elegans* (a kind of worm) is significantly different.




Fig. 1: A comparison of treelet distributions of five PPIN networks by SUBGRAPH2VEC

As PPINs usually include many false (positive and negative) and missing interactions[3] in practice, an occurrence of a specific treelet may include additional (or missing) edges in different networks. Counting non-induced subgraph is more suitable to obtain reliable and robust results[3]. In Fig. 6(a), we show 47 unlabeled treelets with similarity. They have the same size of 9 vertices but vary slightly in topology. Note that an induced subgraph of a graph $G(V, E)$ is a subset of the vertices of the graph $G(V, E)$ as well as with any edges connecting pairs of vertices in that subset. There are many more non-induced subgraphs isomorphic to a given topology as they allow missing edges among vertex pairs. Thus, it is challenging to count non-induced subgraphs of a network.

B. Statement of Problem

1) *Subgraph Counting*: Subgraph finding and counting is a widely studied subject. A (non-induced) subgraph of a simple unweighted graph $G(V, E)$ is a graph $H(V_H, E_H)$ satisfying $V_H \subset V$ and $E_H \subset E$. H is an embedding of a template graph T if T is isomorphic to H . The subgraph counting problem is to count the number of all embeddings of a given template T in a network G . We use $emb(T, G)$ to denote the number of all embeddings of template T in network G .

2) *Color coding*: Color coding [9] is an algorithmic technique to discover network motifs. Given a k -node template T , it assigns random colors between 0 and $k-1$ to each vertex of a network graph G , and it counts the number of the occurrences of colorful embedding, which is isomorphic to T while having distinct colors on each vertex. Both theoretical proof [9], [14], [6] and experiments [3], [15] show that, with proper normalization, the count of colorful embeddings is an unbiased estimator of the actual count of embeddings. Alon et al. [9] proved a guarantee of bounding the count by $(1 \pm \epsilon)emb(T, G)$

with a probability of $1 - 2\delta$ after running at most N iterations of the algorithm.




Fig. 2: Illustration of the template partitioning within a colored input $G = (V, E)$

C. Counting Tree-like Subgraph

By applying color coding, [3] provides a *fixed parameter tractable* algorithm to address the subgraph counting problem where T is a tree. It has a time complexity of $O(c^k \text{poly}(n))$, which is exponential to template size k but polynomial to vertex number n . Algorithm 1 describes the standard sequential algorithm with definition and notation shown in Table I, which contains three important steps as follows.

Algorithm 1: Standard Sequential Algorithm

- 1 $N = O(\frac{e^k \log(1/\delta)}{\epsilon^2})$ // required iterations to converge
 - 2 Partition T into sub-templates T_s
 - 3 **for** $j = 1$ **to** N **do**
 - 4 color all $V_i \in G(V, E)$ randomly
 - 5 counting T_s in a dynamic programming procedure
 - 6 $P \leftarrow$ probability that the template is colorful
 - 7 $\alpha \leftarrow$ number of automorphisms of T_0
 - 8 $finalCount[j] \leftarrow \frac{1}{P\alpha} \sum_i \sum_C M_0(i, I_C)$
 - 9 Output the average of all $finalCount$.
-

1) *Random Coloring*: Each vertex $v \in G(V, E)$ is given an integer value of color randomly selected between 0 and $k-1$, where $k \geq |V_T|$ (we consider $k = |V_T|$ for simplicity). $G(V, E)$ is therefore converted to a labeled graph.

2) *Template Partitioning*: For tree-like templates, we can recursively partition T into a chain of sub-templates T_s until the sub-template containing only one vertex. When partitioning a template T , a single vertex ρ is selected as the root while $T_s(\rho)$ refers to the s -th sub-template rooted at ρ . Secondly, one of the edges (ρ, τ) adjacent to root ρ is cut, creating two child sub-templates. The child holding ρ as its root is named *active child* and denoted as $T_{s,a}$. The child rooted at τ of the cutting edge is named *passive child* and denoted as $T_{s,p}$.

3) *Dynamic Programming*: Algorithm 2 describes the dynamic programming procedure to count partitioned template T from the randomly colored $G(V, E)$. For bottom sub-template $|T_s| = 1$, $M_s(i, I_s)$ is 1 only if C_s equals the color randomly assigned to V_i , and otherwise it is 0. For non-bottom cases where $|T_s| > 1$, we obtain $M_s(i, I_s)$ by multiplying the count

Algorithm 2: Dynamic Programming in Standard Sequential Algorithm

```

input :  $G(V, E), T$ 
output:  $M_s$ 
1 forall sub-templates  $T_s$  in reverse order of their partitioning
  do
2   if  $T_s$  consists of a single vertex then
3     forall  $V_i \in V$  do
4        $M_s(i, \text{color of } V_i) \leftarrow 1$ 
5   else
6     //  $T_s$  has an active child  $T_{s,a}$  and a
        passive child  $T_{s,p}$ 
7     forall vertices  $V_i \in V$  do
8       forall color sets  $C_s$  satisfying  $|C_s| = |T_s|$  do
9         forall color sets  $C_{s,a}$  and  $C_{s,p}$  created by
            splitting  $C_s$  satisfying  $|C_{s,a}| = |T_{s,a}|$  and
             $|C_{s,p}| = |T_{s,p}|$  do
           $M_s(i, I_s) \leftarrow \sum_{V_j \in N(V_i)} M_{s,a}(i, I_{s,a}) M_{s,p}(j, I_{s,p})$ 

```

values from its two children, which have been calculated in previous steps of dynamic programming.

TABLE I: Definitions and Notations

Notation	Definition
$G(V, E)$ or G	The input network
A_G	$ V \times V $ sparse adjacency matrix of $G(V, E)$
$N(V_i)$ or $N(i)$	Neighbors of vertex V_i
T, T_s	The input template and the s -th sub-template
$ T_s $	Number of vertices in T_s
$T_{s,a}, T_{s,p}$	Active and passive child of T_s
n	$n = V $ is the number of vertices in G
k	$k = V_T $ is the number of vertices in T
C_s	Color set for T_s
M_s	$ V \times \binom{k}{ T_s }$ dense matrix to store counts for T_s
$M_{s,a}, M_{s,p}$	Dense matrix to store counts for $T_{s,a}, T_{s,p}$
B	$B = A_G M_{s,p}$, the sum of the counts of all neighbors.
I_{C_s} or I_s	Column index of color set C_s

III. RELATED WORK

The graphlet frequency distance was proposed by Przulj et al [16] as a global comparative measure based on the local structural characteristics of different networks. Bordino et al. [7] demonstrates that one can use the relative frequency of subgraphs within networks to distinguish and cluster different networks.

A tree subgraph enumeration algorithm by combining color coding with a stream-based cover decomposition was developed in [17]. To process massive networks, [18] developed a distributed color-coding based tree counting solution upon the MapReduce framework in Hadoop, [19] implemented an MPI-based solution, and [20] [21] pushed the limit of subgraph counting to process billion-edged networks and trees up to 15 vertices. [22] developed a coloring method that achieves a provable confidence value in a small number of iterations.

[9] proved that color-coding could apply to subgraph counting problems, where the template is a tree, a cycle, or any

graph with bounded treewidth. [6] is a color-coding implementation applying to all templates with a treewidth of no more than 2. Beyond counting trees, a sampling and random-walk based technique [23][2] could count graphlets, small induced graph with size up to 4 or 5.

Other subgraph topics include: 1) *subgraph finding*. As in [24], paths and trees with size up to 18 could be detected by using multilinear detection; 2) *Graphlet Frequency Distribution* estimates relative frequency among all subgraphs with the same size [25]; 3) *clustering* networks by using the relative frequency of their subgraphs [26]. *Subgraph Matching* finds and enumerates all isomorphic subgraphs to a given template from input network. [27] contributes an online algorithm to query subgraph templates from billion-node network by using intelligent graph exploration to replace expensive join operations. [28] compares and summarizes subgraph isomorphism algorithms in graph databases. Later on [29] improves the performance of subgraph matching up to three orders of magnitude by postponing the Cartesian products based on the structure of a query to minimize the redundant Cartesian products. [30], [31] provides a pruning method on labeled networks and graphlets to reduce the vertex number by orders of magnitude prior to the actual counting.

IV. ALGORITHMIC DESIGN OF SUBGRAPH2VEC

Unlike standard sequential algorithm in Algorithm 2, we decouple the dynamic programming into two stages, as shown in Algorithm 3: 1) A vertex-neighbour traversal stage, and 2) A counts updating stage. This design brings a two-fold benefit. First, it helps us identify and remove redundant computation and thus reduce the complexity of Algorithm 2, and secondly it enables us to apply the optimization on system design.

In Algorithm 2, it requires $\binom{|T|}{|T_s|} \binom{|T_s|}{|T_{s,a}|}$ times of vertex neighbor traversal for each V_i from line 7 to 9. However, we find that redundancy of traversal exists as shown in Figure 3, where it counts a two-vertex sub-template with a total of




Fig. 3: Identify the redundancy of standard color coding in a two-vertex sub-template T_s .

three colors. The left case and the middle case have the same color set (the same $I_{s,p}$) assigned to their passive child $T_{s,p}$, which causes redundant access to $M_{s,p}(j, I_{s,p})$ when traversing neighbor vertices.

On the contrary, SUBGRAPH2VEC proposes a novel way to accomplish the vertex neighbor traversal described from 1 to 4 of Algorithm 3:

- 1) The vertex neighbor traversal is decoupled from line 9 of Algorithm 2.




Fig. 4: Decouple the vertex neighbor traversal from updating of the count value according to distributive property of addition and multiplication in Equation 1.

Algorithm 3: Dynamic Programming in SUBGRAPH2VEC

input : $G(V, E), M_{s,p}, T_s$
output: M_s : matrix storing traversal results

```

1 for  $V_i \in G(V, E)$  do
2   for color sets  $C_{s,p}$  satisfying  $|C_{s,p}| = |T_{s,p}|$  do
3     for  $V_j \in N(V_i)$  do
4        $B(i, I_{s,p}) \leftarrow B(i, I_{s,p}) + M_{s,p}(j, I_{s,p})$ 
5 for  $V_i \in G(V, E)$  do
6   for color set  $C_s$  satisfying  $|C_s| = |T_s|$  do
7     for color sets  $C_{s,a}, C_{s,p}$  split from  $C_s$  do
8        $M_s(i, I_s) \leftarrow M_s(i, I_s) + M_{s,a}(i, I_{s,a}) B(i, I_{s,p})$ 

```

2) Only $\binom{|T|}{|T_{s,p}|}$ times of traversal is applied on each vertex.

According to distributive property of addition and multiplication, line 9 of Algorithm 2 can be re-written as

$$\sum_{V_j \in N(i)} M_{s,a}(i, I_{s,a}) M_{s,p}(j, I_{s,p}) = M_{s,a}(i, I_{s,a}) \sum_{V_j \in N(i)} M_{s,p}(j, I_{s,p}) \quad (1)$$

, where the first item $M_{s,a}(i, I_{s,a})$ at right-hand only contains count values of V_i while the second item $\sum_{V_j \in N(i)} M_{s,p}(j, I_{s,p})$ only involves traversing neighbors of V_i . This decoupled design enables a caching and re-using of the traversal results (the summation of $M_{s,p}(j, I_{s,p})$ shown in Figure 4), which allows us to reduce the traversal times.

The second stage of SUBGRAPH2VEC is to update the count values by multiplying $M_{s,a}(i, I_{s,a})$ and $B(i, I_{s,p})$, and both of them are local to vertex i , which improves data locality and allows a vectorized computation when compared to the standard sequential Algorithm 2 where the non-consecutive

indices of i and j at line 9 do not meet the requirement of SIMD paradigm.

V. SYSTEM DESIGN OF SUBGRAPH2VEC

We propose a new scheme in SUBGRAPH2VEC to achieve two goals: 1) Vectorize [32] both stages in Section IV. 2) Transform and leverage the vectorized codes into BLAS kernels.

A. Vectorize Vertex-Neighbor Traversal

SUBGRAPH2VEC utilizes an adjacency matrix, notated as A_G , to store input network $G(V, E)$ in Sparse Row Compressed (CSR) format. A_G is a sparse 0-1 matrix satisfying $A_G(i, j) = 1$ if and only if $V_j \in N(V_i)$. Correspondingly, we re-write line 1 to 4 of Algorithm 3 by Algorithm 4, where for each $I_{s,p}$, we schedule loops of V_i to threads while each thread is vectorizing its own work. We observe that j has successive

Algorithm 4: Vectorized Vertex-Neighbor Traversal in SUBGRAPH2VEC

input : $A_G, T_s, M_{s,p}$
output: B

```

1 forall color sets  $C_{s,p}$  satisfying  $|C_{s,p}| = |T_{s,p}|$  do
2   forall  $V_i \in A_G$  do // loop is scheduled to threads
3     forall  $j = A_G.rowIdx[i]$  to  $A_G.rowIndex[i + 1]$  do
4       // thread workload is vectorized
        $B(i, I_{s,p}) \leftarrow B(i, I_{s,p}) + A_G.val[j] M_{s,p}(A_G.colIdx[j], I_{s,p})$ 

```

values from $A_G.rowIdx[i]$ to $A_G.rowIdx[i + 1]$ resulting in coalesced data access to three dense arrays of A_G . Unfortunately, $A_G.colIdx[j]$ does not guarantee successive values due to the sparsity of A_G . However, advanced compilers still provide partial vectorization support to this indexed access pattern. We will introduce our customization in addressing this partial vectorization issue in Section V-C.

B. Vectorize Count Updating

In Algorithm 3, counting templates at line 8 cannot be vectorized because the indices are not successive. To address this issue, we propose a new scheme illustrated in Figure 5.

- 1) Change memory layout from row-major order to column-major order.
- 2) All threads are working on same columns of $M_s, M_{s,a}, M_{s,p}$ concurrently and processing the matrices column by column.
- 3) All rows of a column are evenly distributed to threads.
- 4) Each thread vectorizes the work on its own portion of rows.

Compared to Algorithm 2, the length of vectorization is converted from $\binom{|T|}{|T_{s,p}|}$ to the million-level number of vertices in $G(V, E)$, which is sufficient to utilize the hardware fully and invariant to different sub-templates. Furthermore, the stride one regular memory access is efficient in prefetching data from memory to cache lines.

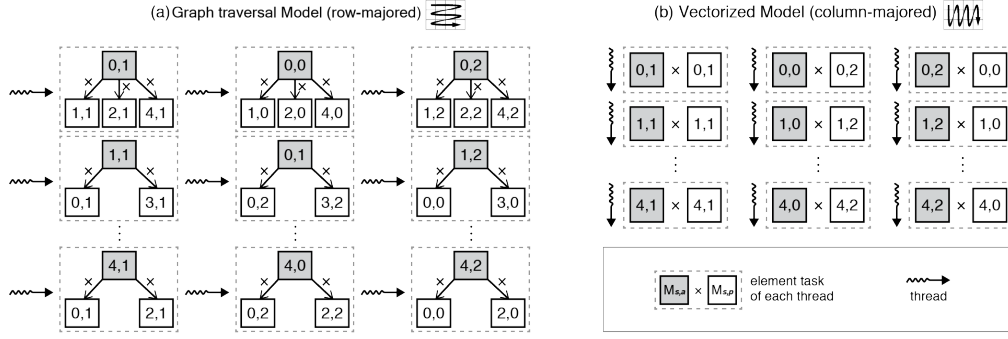


Fig. 5: Comparing the thread execution order, where (a) Graph traversal Model counts data stored in memory with a row-major layout, and (b) SUBGRAPH2VEC (Vectorized Model) counts data stored in memory with a column-major layout.

C. Invocation of Linear Algebra Kernels

Compared to the graph traversal model, SUBGRAPH2VEC is designed to be portable among hardware platforms while keeping high-performance. The vectorized vertex neighbor traversal module in Algorithm 4 mathematically equals to an operation of sparse matrix dense vector multiplication (SpMV), which is an essential sparse linear solver on different hardware platforms. Correspondingly, line 8 of Algorithm 3 equals an element-wised multiplication and addition of dense vectors (eMA). A complete SUBGRAPH2VEC made of SpMV and eMA kernels is described in Algorithm 5, which also applies in-place storage of the SpMV results from the column vector buffer \mathbf{B} back to $\mathbf{M}_{s,p}$ to reduce the memory footprint.

To achieve better kernel performance than by using public libraries, we customize both of SpMV and eMA kernels. For SpMV, we combine a bundle of SpMV operations in Algorithm 5 into a Sparse matrix dense matrix (SpMM) operation shown in Algorithm 6, where the right-hand dense matrix is $\mathbf{M}_{s,p}$. To save peak memory utilization, we also split columns of $\mathbf{M}_{s,p}$ into batches with pre-selected batch size. To improve the load balancing and data locality, we utilize a split compressed sparse column (CSC-Split) format instead of the default compressed sparse row (CSR) format that is widely used by public libraries. CSC-Split format converts the standard CSC format into a fixed number of partitions. Entries of CSC matrix are distributed to a partition when their row IDs fit into a pre-defined range of that partition. Inside a partition, the entries are ordered by their column IDs of CSC format, and therefore entries sharing the same column ID and adjacent row IDs are bundled together to improve the data locality and cache usage. Meanwhile, we store batches of right-hand vectors from $\mathbf{M}_{s,p}$ in a row-major memory layout, and we set up the batch size to the maximal concurrent element number of the hardware SIMD unit. Finally, when a partition is assigned to a thread, the thread processes its entries one by one while vectorizing the computation work on a batch of row entries from $\mathbf{M}_{s,p}$.

To customize the eMA kernel, we utilize Intel (R) AVX intrinsics, where multiplication and addition are implemented by using the fused multiply-add (FMA) instruction, which

Algorithm 5: SUBGRAPH2VEC with Linear Algebra Kernels

input : $\mathbf{A}_G, T, \epsilon, \delta$
output: \mathbf{A} (ϵ, δ)-approximation to $emb(T, G)$

```

1  $N = O(\frac{e^k \log(1/\delta)}{\epsilon^2})$  // required iterations to converge
2 Partition  $T$  into sub-templates  $T_s$ 
3 for  $j = 1$  to  $N$  do
4   forall  $V_i \in G(V, E)$  do
5     Color  $V_i$  by a value randomly drawn from 1 to  $k = |T|$ 
6   for  $s = 0, 1, \dots, S-1$  do
7     forall color sets  $C_{s,p}$  satisfying  $|C_{s,p}| = |T_{s,p}|$  do
8        $\mathbf{B} \leftarrow \mathbf{A}_G \mathbf{M}_{s,p}(:, I_{s,p})$  // SpMV kernel
9        $\mathbf{M}_{s,p}(:, I_{s,p}) \leftarrow \mathbf{B}$  // Sum of neighbor counts
10      forall color sets  $C_s$  satisfying  $|C_s| = |T_s|$  do
11         $\mathbf{M}_s(:, I_s) \leftarrow 0$ 
12        forall color sets  $C_{s,a}$  and  $C_{s,p}$ , created by splitting  $C_s$  satisfying  $|C_{s,a}| = |T_{s,a}|$  and  $|C_{s,p}| = |T_{s,p}|$  do
13           $\mathbf{M}_s(:, I_s) \leftarrow \mathbf{M}_s(:, I_s) + \mathbf{M}_{s,a}(:, I_{s,a}) \odot \mathbf{M}_{s,p}(:, I_{s,p})$  // eMA kernel
14       $P \leftarrow$  probability that the template is colorful
15       $\alpha \leftarrow$  number of automorphisms of  $T_0$ 
16       $finalCount[j] \leftarrow \frac{1}{P\alpha} \sum_i \sum_C \mathbf{M}_0(i, I_C)$ 
17 Output the average of all  $finalCount$ .
```

cuts the computation instructions by half. In addition, there are already substantial research work in developing high-performance linear algebra kernels. The invocation of linear algebra kernels in SUBGRAPH2VEC benefits from: 1) using formats and kernel implementations tailored for different input datasets; 2) increasingly improved kernel performance on various hardware platforms.

VI. EXPERIMENTS AND RESULTS

A. Datasets and Templates

The datasets in our experiments are listed in Table II, where *E.coli* [13], *Worm* [12] and *Yeast* [12] are from Biology;


Algorithm 6: CSCSplit SpMM for sub-template T_s in SUBGRAPH2VEC

```


input :  $A_G, M_{s,p}, sIdx, BSize, SplitPars$ 
output: Out
1 forall  $Par \in SplitPars$  do // partition per
   thread
2   forall  $e \in Par$  do // workload per thread
3     for  $j = sIdx, \dots, sIdx + BSize$  do
4        $Out(e.rowId, j) \leftarrow Out(e.rowId, j) +$ 
          $A_G(e.rowId, e.colId)M_{s,p}(e.colId, j)$ 
         // rowId, colId are row and
         column indices in CSC-Split
         compressed sparse format

```

Graph500 Scale=20, 21, 22 are collected from [33]; *Miami*, *Orkut*, and *NYC* are from [34] [35] [36]; RMAT are widely used synthetic datasets generated by the RMAT model [37], where we increase parameter K to generate datasets with increasingly skewed degree distribution. Figure 6(a) shows all the templates with 9 nodes. The template in Figure 6(b) is from [4], and the templates with more than 12 nodes are randomly selected. The script used to generate the templates can be found in our open-sourced repository [10].



(a) For protein-protein interaction networks comparison



(b) For Social networks, Graph500 and Synthetic data comparison

Fig. 6: Templates used in the experiments

B. Hardware and Software

In the experiments, we use: 1) a single node of a dual-socket Intel(R) Xeon(R) CPU E5-2670 v3 (architecture Haswell), 2) a single node of a dual-socket Intel(R) Xeon(R) Platinum 8160 CPU (architecture Skylake-SP) processors, and 3) a single node of Tesla V100 SXM2 paired with an Intel(R) Xeon(R) CPU E5-2630 v4. More details of the testbed hardware as well as the computation environment are released in the Artifact Description file.

We use the following implementations.

- **FASCIA** implements the graph traversal model of color-coding algorithm with multi-threading on a single CPU [4], which serves as a performance baseline.

- **SUBGRAPH2VEC** implements SUBGRAPH2VEC on CPU by using our in-house CSC-Split format with a SpMM kernel and eMA kernel (threaded by OpenMP). It is the default implementation of SUBGRAPH2VEC and supports distributed systems.
- **SUBGRAPH2VEC-MKL** implements SUBGRAPH2VEC on a single CPU by using CSR based SpMV kernel from Intel MKL and eMA kernel (threaded by OpenMP). It also supports distributed systems.
- **SUBGRAPH2VEC-cuSPARSE** SUBGRAPH2VEC on GPU by using CSR based SpMV kernel from NVIDIA cuSPARSE and eMA kernel (threaded by CUDA). Supports distributed systems by using the CSR format API from distributed mode of SUBGRAPH2VEC-MKL.

Binaries on CPU are compiled by the Intel(R) C++ compiler for Intel(R) 64 target platform from Intel(R) Parallel Studio XE 2019, with compilation flags of “-O3”, “-xCORE-AVX2”, “-xCORE-AVX512”, and the Intel(R) OpenMP. Binaries on GPU are compiled by CUDA release 9.1 (V9.1.85). The distributed binaries are compiled by Intel MPI 2019. We use, by default, a thread number equal to the physical core number of CPU, i.e., 48 threads on a Skylake node and 24 threads on a Haswell node. The threads are bind to cores with a spread affinity. For GPU, we use a thread block with a size of 1024 for the eMA kernel. For kernel invoked by Intel MKL and NVIDIA cuSPARSE, we use the default setup. We mainly use the SUBGRAPH2VEC to evaluate our work except for Section VI-G, where SUBGRAPH2VEC with public library kernels are evaluated against FASCIA.

C. Overall Performance Improvement

We first examine the performance improvement of SUBGRAPH2VEC over the state-of-the-art FASCIA on a Skylake node. The best performance we can obtain is by using a customized matrix format and SpMM kernel. Note that we scale the template size up to the memory limitation on our Skylake testbed for each dataset in Table III. The reduction of execution time is significant, particularly for template sizes larger than 14. For instance, FASCIA spends four days to process a million-vertex dataset RMAT-1M with template u17 while SUBGRAPH2VEC only spends 9.5 minutes. For relatively smaller templates such as u12, SUBGRAPH2VEC still achieves 10x to 100x of improvement on datasets Miami, Orkut, and RMAT-1M.

In Table III, we observe that the improvement is approximately proportional to the average degree of datasets. For instance, SUBGRAPH2VEC achieves 10x and 20x improvements on datasets Miami (average degree of 49) and Orkut (average degree of 76), respectively. It implies that our optimization works better on dense graph network when compared to FASCIA.

The three Graph500 datasets in Table II have comparable average degrees but growing vertex number and edge number. For the same template, SUBGRAPH2VEC obtains similar improvements over FASCIA across the three datasets




Fig. 7: The hardware utilization on one Skylake node for template u12.

TABLE II: Datasets used in the experiments ($K=10^3$, $M=10^6$)

Data	Vertices	Edges	Avg Deg	Max Deg	Abbreviation	Source
Human STRING10	10,971	214,298	19.53	2009	Human	Biology [11]
EcoliGO-BP	1,474	6,896	9.36	72	Ecoli	Biology [13]
WI-2004	1,239	1,736	2.8	74	Worm1	Biology [12]
WI-2007	1,498	1,817	2.43	86	Worm2	Biology [12]
Combined-APMS	1,622	9,070	11.18	127	Yeast1	Biology [12]
LC-multiple	1,536	2,925	3.81	40	Yeast2	Biology [12]
Graph500 Scale=20	600K	31M	48	67K	GS20	Graph500 [33]
Graph500 Scale=21	1M	63M	51	107K	GS21	Graph500 [33]
Graph500 Scale=22	2M	128M	53	170K	GS22	Graph500 [33]
Miami	2.1M	200M	49	10K	MI	Social network [34]
Orkut	3M	230M	76	33K	OR	Social network [35]
NYC	18M	960M	54	429	NY	Social network [36]
RMAT-1M	1M	200M	201	47K	RT1M	Synthetic data [37]
RMAT(K=3)	4M	200M	52	26K	RTK3	Synthetic data [37]
RMAT(K=5)	4M	200M	73	144K	RTK5	Synthetic data [37]
RMAT(K=8)	4M	200M	127	252K	RTK8	Synthetic data [37]

TABLE III: Execution time (s) of SUBGRAPH2VEC (S) versus FASCIA (F) with increasing template sizes from U12 to U17. Tests run on a Skylake node.

Dataset	Impl	u12	u13	u14	u15-1	u15-2	u16	u17
Miami	F	163	400	944	2663	2435		
Miami	S	18	38	55	160	150		
Orkut	F	642	2006	4347	1.5e4	1.2e4		
Orkut	S	30	67	80	238	230		
RMAT 1M	F	1535	5378	1.2e4	3.4e4	3.2e4	1.1e5	3.8e5
RMAT 1M	S	16	32	34	97	97	224	573
Graph500 20	F	132	452	923	3379	2679		
Graph500 20	S	7	14	21	63	56		
Graph500 21	F	289	1044	2036	7535	5914		
Graph500 21	S	12	26	36	105	102		
Graph500 22	F	764	2814	5477	1.9e4	1.6e4		
Graph500 22	S	26	53	74	220	194		
RMAT K=3	F	1191	4890	9711	3.0e4	3.2e4		
RMAT K=3	S	39	110	170	377	262		
RMAT K=5	F	2860	9906	2.0e4	9.0e4	5.4e4		
RMAT K=5	S	29	60	82	233	240		
RMAT K=8	F	5620	2.0e4	3.3e4	9.4e4	8.5e4		
RMAT K=8	S	25	51	67	217	234		

implying that SUBGRAPH2VEC has a scalable performance improvement with respect to the dataset size.

Finally, we compare RMAT datasets with increasingly skewed degree distribution, which causes a thread-level workload imbalance. The results show that SUBGRAPH2VEC has comparable execution time regardless of the degree distribution. On the contrary, FASCIA spends significantly (2x to 3x) more time on datasets with skewed degree distribution.

D. Benefit of System Design

To evaluate the benefit of our system design, we remove the optimization in Section V while only applying multi-threading in both stages of Algorithm 3 as a baseline in Figure 8. We observe that on Miami dataset, SUBGRAPH2VEC obtains 8x

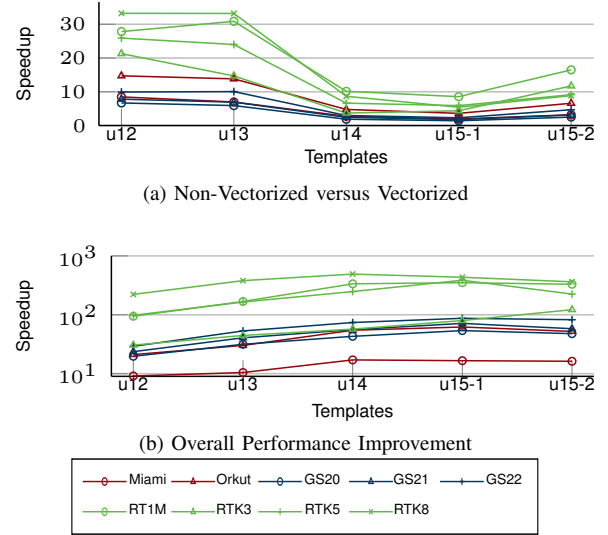


Fig. 8: The performance improvement brought by vectorization and the overall performance improvement. The Scaling tests run on a Skylake node.

speedup by average and up to 30x for some templates. The reason behind that is a significant improvement in hardware utilization, which is evaluated from: 1) The CPU and VPU utilization, 2) An overall efficiency by using the roofline model.

1) **CPU and VPU Utilization:** Figure 7(a) first compares the CPU utilization defined as the average number of concurrently running physical cores. For Miami, FASCIA achieves 60% of CPU utilization. However, the CPU utilization drops down below 50% on Orkut and NYC. Conversely, SpMM kernel keeps a high CPU utilization from 65% to 78% for

all datasets. The eMA kernel has a growing CPU utilization from Miami (46%) to NYC (88%). We have two explanations: 1) the SpMM kernel splits and regroups the nonzero entries by their row IDs, which mitigates the imbalance of nonzero entries among rows; 2) the eMA kernel has its computation workload for each column of $M_{s,a}$, $M_{s,p}$ evenly dispatched among threads.

Secondly, we examine the code vectorization in Figure 7. VPU in a Skylake node is a group of 512-bit registers. The scalar instruction also utilizes the VPU, but it cannot fully exploit its 512-bit length. Figure 7 refers to the portion of instructions vectorized with a full vector capacity. For all of the three datasets, FASCIA only has 6.7% to 12.5% VPU utilization, implying that the codes are not vectorized. While for SpMM and eMA kernels of SUBGRAPH2VEC, the VPU utilization is 100%. A further metric of packed float point instruction ratio (Packed FP) justifies the implication that FASCIA has zero vectorized instructions, but SUBGRAPH2VEC has all of its float point operations vectorized.

2) **Roofline Model:** The roofline model in Figure 9 reflects hardware efficiency. The horizontal axis is the operational intensity (FLOP/byte), and the vertical axis refers to the measured throughput (FLOP/second). The solid roofline is the maximal performance the hardware can deliver under a certain operational intensity. Because of the low operational intensity, the performance of FASCIA and SUBGRAPH2VEC are bounded by the memory bandwidth, and we consider it as a memory-bound roofline. For a relatively small dataset like Miami, both of FASCIA and SUBGRAPH2VEC are close to the memory-bound roofline because the data can be fit into the 33 MB L3 cache. For dataset Orkut, whose data size is beyond the capacity of L3 cache, SUBGRAPH2VEC is much closer to the memory-bound roofline than that of FASCIA because of its regular and vectorized memory access pattern.




Fig. 9: Apply roofline model to FASCIA and SUBGRAPH2VEC. Dataset Miami, Orkut for template u15-1. Tests are done on a Skylake node.

3) **Memory Bandwidth and Cache Usage:** In Table IV, we compare SpMM and eMA of SUBGRAPH2VEC to FASCIA. It shows that the eMA kernel has the highest bandwidth value around 110 GB/s for the three datasets, which is due to the highly vectorized codes and regular memory access patterns. The data is prefetched into cache lines, which mitigates the cache miss rate as low as 0.1%. The SpMM kernel also enjoys

TABLE IV: Memory and Cache Usage of FASCIA, SpMM, and eMA of SUBGRAPH2VEC on a Skylake Node

Orkut	Bandwidth	L1 Miss Rate	L2 Miss Rate	L3 Miss Rate
FASCIA	8 GB/s	9.6%	5.3%	46%
SpMM	59.5 GB/s	6.7%	42.8%	45%
eMA	116 GB/s	0.32%	22.2%	9.0%
NYC	Bandwidth	L1 Miss Rate	L2 Miss Rate	L3 Miss Rate
FASCIA	7 GB/s	2.4%	8.1%	87%
SpMM	96 GB/s	7.7%	76%	74%
eMA	122 GB/s	0.1%	99%	14.8%

a decent bandwidth usage of around 70 to 80 GB/s by average when compared to FASCIA.

E. Parallelization on a Single Node

We perform a strong scaling test using up to 48 threads on Skylake node in Figure 10. We choose RMAT generated datasets with increasing skewness parameters of $K = 3, 5, 8$.




Fig. 10: Strong scaling test of RMAT datasets with increasing cores on a Skylake node.

As the performance is bounded by memory, which has 6 memory channels per socket, we have a total of 12 memory channels on a Skylake node that bounds the thread scaling. Eventually, SUBGRAPH2VEC obtain a 7.5x speedup at 48 threads when $K = 3$. When increasing the skewness of datasets to $K = 5, 8$, the thread scalability of SUBGRAPH2VEC drop down because the skewed data distribution brings workload imbalance.

F. From Single Node to Distributed System

A distributed SUBGRAPH2VEC extends the memory capacity of a single node that enables to run larger templates. As an example, dataset RT1M in Table III can only run templates with size up to u17 on a single node. However, we can scale up the template size to u20 with a cluster of 16 nodes as shown in Figure 12.

Our distributed SUBGRAPH2VEC has a good strong scalability, which even achieves super linear speedup in Figure 13 from 1 node to 8 nodes. According to the analysis in Section VI-D2, SUBGRAPH2VEC is memory bounded, and increasing the number of nodes scales out not only computation resources but also memory bandwidth and cache resources. Having less data on each node can increase the percentage of data held by the last level of the CPU cache.




Fig. 11: Execution Time of SUBGRAPH2VEC on three platforms. On Haswell and Skylake nodes, we use CSR based SpMV kernel from Intel MKL; On Volta GPU V100, we use CSR based SpMV kernel from NVIDIA cuSPARSE




Fig. 12: Scaling up the templates up to u20 by distributed SUBGRAPH2VEC on 16 Haswell nodes.




Fig. 13: Strong scaling test on distributed SUBGRAPH2VEC. Four datasets with template u14 running on 16 Haswell nodes.

G. Portability to Other Platforms

Hardware platforms such as NVIDIA GPU already have highly-optimized public libraries of linear algebra kernels. For SpMV operation, we have the *mkl_sparse_s_mv* kernel from Intel MKL library on CPU and the *cusparseScsrnv* kernel from NVIDIA cuSPARSE library on GPU. For eMA kernel, we can use a combination of *vsMul* and *vsAdd* kernels from Intel MKL or hand-implement such kernels whenever the kernel is absent because of its simplicity. Hence, we have ported SUBGRAPH2VEC to GPU by keeping the CPU codes other than SpMV and eMA on the host side while invoking cuSPARSE and CUDA kernels for the two linear algebra operations.

In Figure 11, we port the performance of SUBGRAPH2VEC to three platforms by using CSR-SpMV libraries kernels. When the template size is small, SUBGRAPH2VEC-cuSPARSE has comparable or even better performance than SUBGRAPH2VEC-MKL. However, the performance of SUBGRAPH2VEC-cuSPARSE drops down when the template size grows up. As NVIDIA-V100 only has 16GB of device memory, it is probable that the large memory footprint of M_s brought by large template size cannot fit into the device memory, and the bi-directional data transfer between the host and device memory compromises the performance of SUBGRAPH2VEC-cuSPARSE. Nevertheless, both of Intel MKL and NVIDIA cuSPARSE are not open-sourced, and we cannot

conclude on their performance gap. Also, the three hardware platforms have different theoretical peak performances and memory bandwidths. This result is only meant to demonstrate the portability of our SUBGRAPH2VEC across hardware platforms.

H. Error Discussion

We implement the standard color coding algorithm that Alon et al. [3] prove to run at most N iterations to control approximation quality as in Algorithm 1. In practice, the subgraph counting with color-coding requires only 100 iterations for a 7 node template on *H.pylori* with an error of less than 1% in FASCIA [15]. SUBGRAPH2VEC with its pruning and vectorization optimization only differs from the FASCIA due to the restructuring of the computation from Algorithm 2 to Algorithm 3. It should give identical results with exact arithmetic in Equation 1.




Fig. 14: Relative error on dataset Graph500 Scale=20.

However, when dealing with large graphs, the counted value will exceed the range of integer variables. As a consequence, both FASCIA and our SUBGRAPH2VEC use 32-bit floating-point numbers to avoid overflow. Hence, slightly different results are observed between FASCIA and SUBGRAPH2VEC due to the rounding error consequent from floating-point arithmetic operations. Figure 14 reports such relative errors between SUBGRAPH2VEC and FASCIA in the range of 10^{-6} across all the tests on a Graph500 GS20 dataset with increasing template sizes, which is negligible.

VII. CONCLUSION

In this paper, we fully vectorize a sophisticated algorithm of subgraph analysis, and the novelty is a co-design approach with pattern identification of linear algebra kernels that leverage hardware vectorization of Intel CPU and NVIDIA GPU architectures. The overall performance achieves significant improvements over the state-of-the-art work by orders of magnitude by average and up to 660x (RMAT1M with u17) within a shared-memory multi-threaded system. The

distributed memory system runs large tree subgraphs with sizes up to 20.

This work demonstrates that fundamental algorithms, such as subgraph mining, could be efficiently explored. More significantly, by implementing it with hardware vectorization, we point a direction where scaling performance of complex graph applications with random access to the vast memory region and dynamic programming workflow is possible and can be done with sparse linear algebra kernels in contrast to conventional graph traversal. An interesting future work will be exploring machine learning as the irregularity of memory access remains a roadblock to improve hardware utilization. It will lead to new research and improvements on portability to other emerging hardware accelerators and heterogeneous architectures.

VIII. ACKNOWLEDGEMENT

We gratefully acknowledge the support from NSF CIF21 DIBBS 1443054: Middleware and High Performance Analytics Libraries for Scalable Data Science, Science and NSF EEC 1720625: Network for Computational Nanotechnology (NCN) – Engineered nanoBio node, NSF OAC 1835631 CINES: A Scalable Cyberinfrastructure for Sustained Innovation in Network Engineering and Science, and Intel Parallel Computing Center (IPCC) grants. We would like to express our special appreciation to the FutureSystems team.

REFERENCES

- [1] J. Ugander, L. Backstrom, and J. Kleinberg, “Subgraph frequencies: Mapping the empirical and extremal geography of large graph collections,” in *Proceedings of the 22nd international conference on World Wide Web*, pp. 1307–1318, ACM, 2013.
- [2] X. Chen and J. Lui, “Mining graphlet counts in online social networks,” *ACM Transactions on Knowledge Discovery from Data (TKDD)*, vol. 12, no. 4, p. 41, 2018.
- [3] N. Alon, P. Dao, I. Hajirasouliha, F. Hormozdiari, and S. C. Sahinalp, “Biomolecular network motif counting and discovery by color coding,” *Bioinformatics*, vol. 24, no. 13, pp. i241–i249, 2008.
- [4] G. M. Slota and K. Madduri, “Fast approximate subgraph counting and enumeration,” in *Parallel Processing (ICPP), 2013 42nd International Conference on*, pp. 210–219, IEEE, 2013.
- [5] F. Battiston, V. Nicosia, M. Chavez, and V. Latora, “Multilayer motif analysis of brain networks,” *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 27, no. 4, p. 047404, 2017.
- [6] V. T. Chakaravarthy, M. Kapralov, P. Murali, F. Petrini, X. Que, Y. Sabharwal, and B. Schieber, “Subgraph counting: Color coding beyond trees,” in *Parallel and Distributed Processing Symposium, 2016 IEEE International*, pp. 2–11, IEEE, 2016.
- [7] I. Bordino, D. Donato, A. Gionis, and S. Leonardi, “Mining large networks with subgraph counting,” in *2008 Eighth IEEE International Conference on Data Mining*, pp. 737–742, IEEE, 2008.
- [8] S. A. Cook, “The complexity of theorem-proving procedures,” in *Proceedings of the third annual ACM symposium on Theory of computing*, pp. 151–158, ACM, 1971.
- [9] N. Alon, R. Yuster, and U. Zwick, “Color-coding,” *Journal of the ACM (JACM)*, vol. 42, no. 4, pp. 844–856, 1995.
- [10] “Subgraph2vec.” <https://github.com/DSC-SPIDAL/harp/tree/subgraph2vec>.
- [11] E. Hodzic, R. Shrestha, K. Zhu, K. Cheng, C. C. Collins, and S. Cenik Sahinalp, “Combinatorial detection of conserved alteration patterns for identifying cancer subnetworks,” *GigaScience*, vol. 8, no. 4, p. giz024, 2019.
- [12] “Interactome database.” <http://interactome.dfci.harvard.edu/>.
- [13] “Ecolinet.” <https://www.inetbio.org/ecolinet/downloadnetwork.php>.
- [14] N. Alon, P. Dao, I. Hajirasouliha, F. Hormozdiari, and S. C. Sahinalp, “Biomolecular network motif counting and discovery by color coding,” *Bioinformatics*, vol. 24, no. 13, pp. i241–i249, 2008.
- [15] G. M. Slota and K. Madduri, “Parallel color-coding,” *Parallel Computing*, vol. 47, pp. 51–69, 2015. bibtex: slota_parallel_2015.
- [16] N. Pržulj, D. G. Corneil, and I. Jurisica, “Modeling interactome: scale-free or geometric?,” *Bioinformatics*, vol. 20, no. 18, pp. 3508–3515, 2004.
- [17] Z. Zhao, M. Khan, V. A. Kumar, and M. V. Marathe, “Subgraph enumeration in large social contact networks using parallel color coding and streaming,” in *Parallel Processing (ICPP), 2010 39th International Conference on*, pp. 594–603, IEEE, 2010.
- [18] Z. Zhao, G. Wang, A. R. Butt, M. Khan, V. S. A. Kumar, and M. V. Marathe, “Sahad: Subgraph analysis in massive networks using hadoop,” in *IEEE International Parallel & Distributed Processing Symposium*, 2012.
- [19] G. M. Slota and K. Madduri, “Complex network analysis using parallel approximate motif counting,” in *Parallel and Distributed Processing Symposium, 2014 IEEE 28th International*, pp. 405–414, IEEE, 2014.
- [20] Z. Zhao, L. Chen, M. Avram, M. Li, G. Wang, A. Butt, M. Khan, M. Marathe, J. Qiu, and A. Vullikanti, “Finding and counting tree-like subgraphs using mapreduce,” *IEEE Transactions on Multi-Scale Computing Systems*, vol. 4, no. 3, pp. 217–230, 2018.
- [21] L. Chen, B. Peng, S. Ossen, A. Vullikanti, M. Marathe, L. Jiang, and J. Qiu, “High-performance massive subgraph counting using pipelined adaptive-group communication,” *Big Data and HPC: Ecosystem and Convergence*, vol. 33, p. 173, 2018.
- [22] G. Guelsoy, B. Gandhi, and T. Kahveci, “Topac: alignment of gene regulatory networks using topology-aware coloring,” *Journal of bioinformatics and computational biology*, vol. 10, no. 01, p. 1240001, 2012.
- [23] N. K. Ahmed, J. Neville, R. A. Rossi, and N. Duffield, “Efficient graphlet counting for large networks,” in *Data Mining (ICDM), 2015 IEEE International Conference on*, pp. 1–10, IEEE, 2015.
- [24] S. Ekanayake, J. Cadena, U. Wickramasinghe, and A. Vullikanti, “Midas: Multilinear detection at scale,” in *2018 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pp. 2–11, IEEE, 2018.
- [25] N. Pržulj, “Biological network comparison using graphlet degree distribution,” *Bioinformatics*, vol. 23, no. 2, pp. e177–e183, 2007.
- [26] M. Rahman, M. A. Bhuiyan, and M. Al Hasan, “An efficient graphlet counting method for large graph analysis,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 26, no. 10, pp. 2466–2478, 2014.
- [27] Z. Sun, H. Wang, H. Wang, B. Shao, and J. Li, “Efficient subgraph matching on billion node graphs,” *Proceedings of the VLDB Endowment*, vol. 5, no. 9, pp. 788–799, 2012.
- [28] J. Lee, W.-S. Han, R. Kasperovics, and J.-H. Lee, “An in-depth comparison of subgraph isomorphism algorithms in graph databases,” in *Proceedings of the VLDB Endowment*, vol. 6, pp. 133–144, VLDB Endowment, 2012.
- [29] F. Bi, L. Chang, X. Lin, L. Qin, and W. Zhang, “Efficient subgraph matching by postponing cartesian products,” in *Proceedings of the 2016 International Conference on Management of Data*, pp. 1199–1214, ACM, 2016.
- [30] T. Reza, M. Ripeanu, N. Tripoul, G. Sanders, and R. Pearce, “Prunejuice: pruning trillion-edge graphs to a precise pattern-matching solution,” in *Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis*, p. 21, IEEE Press, 2018.
- [31] T. Reza, C. Klymko, M. Ripeanu, G. Sanders, and R. Pearce, “Towards practical and robust labeled pattern matching in trillion-edge graphs,” in *2017 IEEE International Conference on Cluster Computing (CLUSTER)*, pp. 1–12, IEEE, 2017.
- [32] “Improve performance with vectorization.” <https://software.intel.com/en-us/articles/improve-performance-with-vectorization>.
- [33] GraphChallenge, “Graph challenge mit,” 2019.
- [34] C. L. Barrett, R. J. Beckman, M. Khan, V. Anil Kumar, M. V. Marathe, P. E. Stretz, T. Dutta, and B. Lewis, “Generation and analysis of large synthetic social contact networks,” in *Winter Simulation Conference*, pp. 1003–1014, Winter Simulation Conference, 2009.
- [35] J. Leskovec and A. Krevl, “SNAP Datasets: Stanford large network dataset collection.” <http://snap.stanford.edu/data>, June 2014.
- [36] J. Yang and J. Leskovec, “Defining and evaluating network communities based on ground-truth,” *Knowledge & Information Systems*, vol. 42, no. 1, pp. 181–213, 2012.
- [37] D. Chakrabarti, Y. Zhan, and C. Faloutsos, “R-mat: A recursive model for graph mining,” in *Proceedings of the 2004 SIAM International Conference on Data Mining*, pp. 442–446, SIAM, 2004.