

Heterogeneous Network Motifs

Ryan A. Rossi
Adobe Research

Nesreen K. Ahmed
Intel Labs

Aldo Carranza
Stanford University

David Arbour
Adobe Research

Anup Rao
Adobe Research

Sungchul Kim
Adobe Research

Eunyee Koh
Adobe Research

ABSTRACT

Many real-world applications give rise to *large heterogeneous networks* where nodes and edges can be of any arbitrary type (e.g., user, web page, location). Special cases of such heterogeneous graphs include homogeneous graphs, bipartite, k-partite, signed, labeled graphs, among many others. In this work, we generalize the notion of network motifs to heterogeneous networks. In particular, small induced typed subgraphs called *typed graphlets* (heterogeneous network motifs) are introduced and shown to be the fundamental building blocks of complex heterogeneous networks. Typed graphlets are a powerful generalization of the notion of graphlet (network motif) to heterogeneous networks as they capture both the induced subgraph of interest and the types associated with the nodes in the induced subgraph. To address this problem, we propose a fast, parallel, and space-efficient framework for counting typed graphlets in large networks. We discover the existence of non-trivial combinatorial relationships between lower-order ($k-1$)-node typed graphlets and leverage them for deriving many of the k -node typed graphlets in $o(1)$ constant time. Thus, we avoid explicit enumeration of those typed graphlets. Notably, the time complexity matches the best untyped graphlet counting algorithm. The experiments demonstrate the effectiveness of the proposed framework in terms of runtime, space-efficiency, parallel speedup, and scalability as it is able to handle large-scale networks.

CCS CONCEPTS

- Computing methodologies → Artificial intelligence; Machine learning; Logical and relational learning;
- Mathematics of computing → Graph algorithms; Approximation algorithms; Combinatorics; Graph theory;
- Information systems → Data mining;
- Theory of computation → Graph algorithms analysis; Streaming, sublinear and near linear time algorithms; Parallel algorithms;

KEYWORDS

Heterogeneous network motifs, typed motifs, heterogeneous graphlets, typed graphlets, colored graphlets, colored motifs, network motifs, heterogeneous networks, labeled graphs, large networks

1 INTRODUCTION

Higher-order connectivity patterns such as small induced subgraphs called graphlets (network motifs)¹ are known to be the fundamental building blocks of simple homogeneous networks [35]

¹The terms graphlet, network motif, and induced subgraph are used interchangeably.

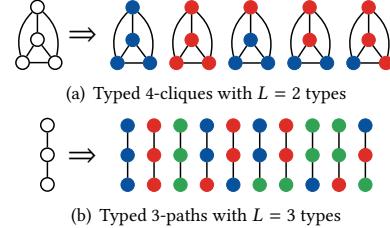


Figure 1: Examples of heterogeneous network motifs

and are essential for modeling and understanding the fundamental components of these networks [3, 4, 10]. Furthermore, graphlets are important for many predictive and descriptive modeling application tasks [2, 21, 31, 34, 35, 40, 50, 55, 62] such as image processing and computer vision [62, 63], network alignment [15, 27, 34, 39], classification [50, 55], visualization and sensemaking [3, 4], dynamic network analysis [23, 26], community detection [10, 38, 41, 53], role discovery [5, 6], anomaly detection [8, 37], and link prediction [44].

However, such (untyped) graphlets are *unable* to capture the rich (typed) connectivity patterns in more complex networks such as those that are heterogeneous (which includes signed, labeled, bipartite, k-partite, and attributed graphs, among others). In heterogeneous graphs, nodes and edges can be of different types and explicitly modeling such types is crucial. Such heterogeneous graphs (networks) arise ubiquitously in the natural world where nodes and edges of multiple types are observed, e.g., between humans [25], neurons [9, 11], routers and autonomous systems (ASes) [45], web pages [59], devices & sensors [16], infrastructure (roads, airports, power stations) [56], economies [49], vehicles (cars, satellites, UAVs) [24], and information in general [46, 54, 60].

In this work, we generalize the notion of network motifs to heterogeneous networks. In particular, we introduce the notion of a small induced typed subgraph called *typed graphlet*.² Typed graphlets are a powerful generalization of the notion of graphlet as they capture both the induced subgraph of interest and the types associated to the nodes in the induced subgraph (Figure 1). These small induced typed subgraphs are the fundamental *building blocks* of *rich heterogeneous networks*. Typed graphlets naturally capture the higher-order typed connectivity patterns in bipartite, k-partite, signed, labeled, k-star, attributed graphs, and more generally heterogeneous networks (Figure 2). As such, typed graphlets are useful for a wide variety of predictive and descriptive modeling applications in these rich complex networks.

²The terms typed, colored, labeled, and heterogeneous graphlet (network motif) are used interchangeably.

Despite their fundamental and practical importance, counting typed graphlets in large graphs remains a challenging and unsolved problem. To address this problem, we propose a fast, parallel, and space-efficient framework for counting typed graphlets in large networks. The time complexity is provably optimal as it matches the best untyped graphlet counting algorithm. Using non-trivial combinatorial relationships between lower-order ($k-1$)-node typed graphlets, we derive equations that allow us to compute many of the k -node typed graphlet counts in $o(1)$ constant time. Thus, we avoid explicit enumeration of many typed graphlets by simply computing the count exactly using the discovered combinatorial relationships. For every edge, we count a few typed graphlets and obtain the exact counts of the remaining typed graphlets in $o(1)$ constant time. Furthermore, we store only the nonzero typed graphlet counts for every edge. To better handle large-scale heterogeneous networks with an arbitrary number of types, we propose an efficient parallel algorithm for typed graphlet counting that scales almost linearly as the number of processing units increase. While we primarily focus on the more challenging problem of local typed graphlet counting, we also demonstrate how to solve the global typed graphlet counting problem. As an aside, the proposed approach is the first such method that counts colored graphlets on the edges (or more generally, between any arbitrary pair of nodes). Nevertheless, the proposed framework overcomes the computational issues with previous work that severely limited its use in practice, e.g., all existing methods are only able to compute typed graphlets on extremely small and unrealistic graphs with hundreds or thousands of nodes/edges.

Theoretically, we show that typed graphlets are more powerful and encode more information than untyped graphlets. In addition, we theoretically demonstrate the worst-case time and space complexity of the proposed framework. It is notably much faster and significantly more space-efficient than the existing state-of-the-art algorithms. In fact, the time complexity of the proposed approach is shown to be equivalent to the best untyped graphlet counting algorithm. Furthermore, unlike previous work, we derive many of the typed graphlets directly in $o(1)$ constant time using counts of lower-order ($k-1$)-node typed graphlets.

Empirically, the proposed approach is shown to be orders of magnitude faster than existing state-of-the-art methods. In particular, we observe between 89 and 10,981 times speedup in runtime performance compared to the best existing method. Notably, on graphs of even moderate size (thousands of nodes/edges), existing approaches fail to finish in a reasonable amount of time (24 hours). In terms of space, the proposed approach uses between 42x and 776x less space than existing methods. We also demonstrate the parallel scaling of the parallel algorithm and observe nearly linear speedups as the number of processing units increases. In addition to real-world graphs from a wide range of domains, we also show results on a number of synthetically generated graphs from a variety of synthetic graph models. Finally, we demonstrate the utility of typed graphlets for exploratory network analysis and understanding using a variety of well-known networks and make a number of new and important observations/findings.

Compared to the homogeneous/untyped graphlet counting problem (which has found many important applications [5, 6, 8, 10, 27, 37, 39, 44, 50, 53, 55]), heterogeneous graphlets are vastly more

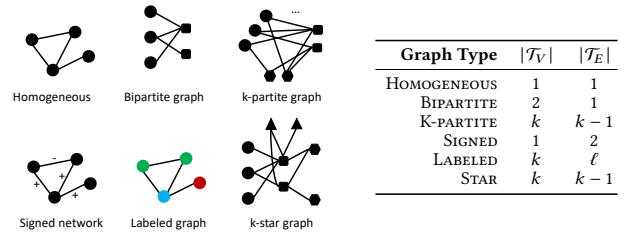


Figure 2: Heterogeneous network motifs are useful for a wide variety of graphs. These are only a few examples of the different graphs that are naturally supported by the proposed framework.

powerful containing a significant amount of additional information. We show this formally using information theory and demonstrate the importance of heterogeneous graphlets empirically using real-world graphs. In real-world graphs, we observe that only a handful of the possible heterogeneous graphlets actually occur. Furthermore, we discover that 99% of the overall counts in a number of real-world networks are from only a few such heterogeneous graphlets. This observation indicates a power-law relationship between the counts of the various heterogeneous graphlets. The rare heterogeneous graphlets (i.e., those that rarely occur) also contain a lot of useful information and these forbidden heterogeneous motifs may indicate anomalies/outliers or simply unique structural behaviors that are fundamentally important but extremely difficult to identify using traditional methods. Moreover, the heterogeneous graphlets found to be important are easily interpretable and provide key insights into the structure and underlying phenomena governing the formation of the complex network that would otherwise be hidden using traditional methods.

This work generalizes the notion of network motif to heterogeneous networks and describes a framework for finding all such *heterogeneous network motifs*. The proposed framework has the following desired properties:

- **Fast:** The approach is fast for large graphs by leveraging *non-trivial combinatorial relationships* to derive many of the typed graphlets in $o(1)$ constant time. Thus, the worst-case time complexity matches that of the best untyped graphlet algorithm. As shown in Table 4, the approach is orders of magnitude faster than the state-of-the-art.
- **Space-Efficient:** The approach is space-efficient by hashing and storing only the typed motif counts that appear on a given edge.
- **Parallel:** The typed graphlet counting approach lends itself to an efficient and scalable parallel implementation. We observe near-linear parallel scaling results in Section 7.3. This indicates the feasibility of the approach for large real-world networks.
- **Scalable for Large Networks:** The proposed approach is scalable for large heterogeneous networks. In particular, the approach scales nearly linearly as the size of the graph increases.
- **Effectiveness:** Typed graphlet counting is shown to be effective for understanding and mining large heterogeneous networks from a variety of different application domains.

Table 1: Summary of notation. Matrices are bold upright roman letters; vectors are bold lowercase letters.

| | |
|-------------------------|---|
| G | graph |
| $V(G)$ | node set of G |
| $E(G)$ | edge set of G |
| H, F | graphlet of G |
| $I_G(H)$ | set of unique instances of H in G |
| N, M | number of nodes $N = V $ and edges $M = E $ in the graph |
| K | size of a network motif (# nodes) |
| L | number of types (<i>i.e.</i> , colors, labels) |
| \mathcal{H} | set of all untyped motifs in G |
| \mathcal{H}_T | set of all typed motifs in G |
| T | # of different typed motifs $T = \mathcal{H}_T $ observed in G with L types |
| T_{\max} | total # of possible typed motifs with L types, hence $T \leq T_{\max}$ |
| T_H | # of different typed motifs for a particular motif $H \in \mathcal{H}$ |
| \mathcal{T}_V | set of node types in G |
| \mathcal{T}_E | set of edge types in G |
| ϕ | type function $\phi : V \rightarrow \mathcal{T}_V$ |
| ξ | type function $\xi : E \rightarrow \mathcal{T}_E$ |
| \mathbf{t} | K -dimensional type vector $\mathbf{t} = [\phi_{w_1} \dots \phi_{w_K}]$ |
| $f_{ij}(H, \mathbf{t})$ | # of instances of motif H that contain nodes i and j with type vector \mathbf{t} |
| \mathbb{F} | an arbitrary typed motif hash function (Section 3.5) |
| c | hash value from \mathbb{F} representing a unique id for an arbitrary typed motif |
| Δ | maximum degree of a node in G |
| Γ_i^t | set of neighbors of node i with type t |
| d_i^t | degree of node i with type t , $d_i^t = \Gamma_i^t $ |
| T_{ij}^t | set of nodes of type t that form typed triangles with i and j |
| S_i^t, S_j^t | set of nodes of type t that form typed 3-node stars centered at i (or j) |
| I | set of non-adjacent nodes with respect to a pair of nodes (i, j) |
| \mathcal{M} | set of typed motif ids that appear in G |
| \mathcal{M}_{ij} | set of typed motif ids for a given pair of nodes (i, j) |
| X_{ij} | nonzero (typed-motif, count) pairs for edge $(i, j) \in E$ |
| Ψ | hash table for checking whether a node is connected to i or j and its “relationship” (<i>e.g.</i> , $\lambda_1, \lambda_2, \lambda_3$) in constant time |

2 HETEROGENEOUS NETWORK MOTIFS

This section introduces a generalization of graphlets (network motifs) called *heterogeneous network motifs* (or simply *typed graphlets*). In particular, Section 2.1 describes the heterogeneous graph model and highlights a number of important properties of it. Next, Section 2.2 formally introduces the notion of *typed graphlet* proposed in this work and discusses its implications. Finally, Section 2.4 discusses a few important special cases of heterogeneous graphs that are naturally supported by the proposed framework in Section 3.

2.1 Heterogeneous Graph Model

In this section, we introduce a general heterogeneous graph model and discuss different problem settings and graph types where typed graphlets are likely to be useful (Figure 2). We use the following heterogeneous graph formulation:

DEFINITION 1 (HETEROGENEOUS NETWORK). A heterogeneous network is defined as $G = (V, E)$ consisting of a set of node objects V and a set of edges E connecting the nodes in V . A heterogeneous network also has a node type mapping function $\phi : V \rightarrow \mathcal{T}_V$ and an edge type mapping function defined as $\xi : E \rightarrow \mathcal{T}_E$ where \mathcal{T}_V and \mathcal{T}_E denote the set of node object types and edge types, respectively. The type of node i is denoted as ϕ_i whereas the type of edge $e = (i, j) \in E$ is denoted as $\xi_{ij} = \xi_e$.

A homogeneous graph is a special case of a heterogeneous graph where $|\mathcal{T}_V| = |\mathcal{T}_E| = 1$. Other special cases of the proposed framework include bipartite graphs, signed networks with $\{+, -\}$ types, and more generally labeled graphs. As such, in general, a heterogeneous network can be represented as a series of matrices and tensors that are coupled (*i.e.*, the tensors and matrices all share at least one mode with each other) [1, 46]. A few other special cases are discussed further in Section 2.4. See Table 1 for a summary of key notation.

2.2 Typed Graphlet Generalization

In this section, we introduce a more general notion of graphlet called *typed graphlet* that naturally extends to both homogeneous and general heterogeneous networks. We use G to represent a graph and H or F to represent graphlets.

2.2.1 Untyped Graphlets. We begin by defining untyped graphlets for graphs with a single type.

DEFINITION 2 (UNTYPED GRAPHLET). An untyped graphlet H is a connected induced subgraph of G .

Given a graphlet in some graph, it may be the case that we can find other topologically identical “appearances” of this structure in that graph. We call these “appearances” *graphlet instances*.

DEFINITION 3 (UNTYPED GRAPHLET INSTANCE). An instance of an untyped graphlet H in graph G is an untyped graphlet F in G that is isomorphic to H .

2.2.2 Typed Graphlets. In heterogeneous graphs, nodes/edges can be of many different types and so explicitly and jointly modeling such types is essential (Figure 2). In this work, we introduce the notion of a *typed graphlet* that explicitly captures both the connectivity pattern of interest and the types. Notice that typed graphlets are a generalization of graphlets to heterogeneous networks.

DEFINITION 4 (TYPED GRAPHLET). A typed graphlet of a graph $G = (V, E, \phi, \xi)$ is a connected induced heterogeneous subgraph $H = (V', E', \phi', \xi')$ of G such that

- (1) (V', E') is a graphlet of (V, E) ,
- (2) $\phi' = \phi|_{V'}$, that is, ϕ' is the restriction of ϕ to V' ,
- (3) $\xi' = \xi|_{E'}$, that is, ξ' is the restriction of ξ to E' .

The terms typed graphlet, colored graphlet, and heterogeneous network motif (graphlet) are used interchangeably.

We can consider the presence of topologically identical “appearances” of a typed graphlet in a graph.

DEFINITION 5 (TYPED GRAPHLET INSTANCE). An instance of a typed graphlet $H = (V', E', \phi', \xi')$ of graph G is a typed graphlet $F = (V'', E'', \phi'', \xi'')$ of G such that

- (1) (V'', E'') is isomorphic to (V', E') ,
- (2) $\mathcal{T}_{V''} = \mathcal{T}_{V'}$ and $\mathcal{T}_{E''} = \mathcal{T}_{E'}$, that is, the multisets of node and edge types are correspondingly equal.

The set of unique typed graphlet instances of H in G is denoted as $I_G(H)$.

Comparing the above definitions of graphlet and typed graphlet, we see at first glance that typed graphlets are nontrivial extensions of their homogeneous counterparts. The “position” of an edge (node)

in a typed graphlet is often topologically important, e.g., an edge at the end of the 4-path (Figure 3(a)) vs. an edge at the center of a 4-path (Figure 3(b)). These topological differences of a typed graphlet are called (automorphism) *typed orbits* since they take into account “symmetries” between edges (nodes) of a graphlet. Typed graphlet orbits are a generalization of (homogeneous) graphlet orbits [39].

2.3 Number of Typed Graphlets

For a single K -node untyped motif (e.g., K -clique), the number of *typed motifs* with L types is:

$$\binom{\binom{L}{K}}{K} = \binom{L+K-1}{K} \quad (1)$$

where L = number of types (colors) and K = size of the network motif (# of nodes). Table 2 shows the number of *typed network motifs* that arise from a single motif $H \in \mathcal{H}$ of size $K \in \{2, \dots, 4\}$ nodes as the number of types varies from $L = 1, 2, \dots, 9$. Notice that Table 2 is for a single K -node motif $H \in \mathcal{H}$ and therefore the total number of typed network motifs for all K -node motifs is a multiple of the amounts given in Table 2. For instance, the total number of typed motif orbits with 4 nodes that arise from 7 types is $10 \cdot 210 = 2100$ since there are 10 connected 4-node homogeneous motif orbits. See Figure 1 for other examples. Unlike homogeneous motifs, it is obviously impossible to show all the heterogeneous motifs counted by the proposed approach since it works for general heterogeneous graphs with any arbitrary number of types L and structure.

Table 2: Number of typed graphlets (for a single untyped graphlet) as the size of the graphlet (K) and types (L) varies.

| Types L | | | | | | | | | |
|-----------|---|---|----|----|----|-----|-----|-----|-----|
| | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 |
| K=2 | 1 | 3 | 6 | 10 | 15 | 21 | 28 | 36 | 45 |
| K=3 | 1 | 4 | 10 | 20 | 35 | 56 | 84 | 120 | 165 |
| K=4 | 1 | 5 | 15 | 35 | 70 | 126 | 210 | 330 | 495 |

2.4 Generalization to Other Graphs

The proposed notion of *heterogeneous network motifs* can be used for applications on bipartite, k-partite, signed, labeled, attributed, and more generally heterogeneous networks. A few examples of such graphs are shown in Figure 2. The proposed framework naturally handles general heterogeneous graphs with arbitrary structure and an arbitrary number of types. It is straightforward to see that homogeneous, bipartite, k-partite, signed, labeled, and star graphs are all special cases of heterogeneous graphs. Therefore, the framework for deriving heterogeneous motifs can easily support such networks. See Figure 2 for a few of the popular special cases of heterogeneous graphs that are naturally supported in the proposed framework. In the case of attributed graphs with $D > 1$ attributes/features, the attributes of a node or edge can be mapped to types using any arbitrary approach such as [6] or [52]. For instance, given an attribute vector $\mathbf{x} \in \mathbb{R}^D$ (which may contain intrinsic attributes or structural features), we derive a type y via $f : \mathbf{x} \rightarrow y$ where f is a function that maps \mathbf{x} to a single type y . Notice that attributed graphs are

a generalization of labeled graphs. Another simple approach is to count typed motifs for each of the D attributes.

3 FRAMEWORK

This section describes the general framework for counting typed graphlets. The typed graphlet framework can be used for counting typed graphlets locally for every edge in G as well as the global typed graphlet counting problem (Definition 13) that focuses on computing the total frequency of a subset or all typed graphlets up to a certain size. This paper mainly focuses on the harder local typed graphlet counting problem defined as:

DEFINITION 6 (LOCAL TYPED GRAPHLET COUNTING). *Given a graph G and an edge $(i, j) \in E$, the local typed graphlet counting problem is to find the set of all typed graphlets that contain nodes i and j and their corresponding frequencies. This work focuses on computing typed graphlet counts for every edge in G .*

This problem is obviously more difficult than the global typed graphlet counting problem. Nevertheless, we discuss the global typed graphlet counting problem in Section 4 and mention a number of improvements and optimization’s that can be leveraged if we are only concerned with the total frequency of the different typed graphlets in G .

Counting *typed k-node motifs* consists of two parts. The first is the actual structure of the typed k -node motif and the second part is the type configuration of the motif. In this work, we propose combinatorial relationships for typed graphlets and use them to derive the counts of many typed graphlet directly in $o(1)$ constant time. Thus, the proposed approach avoids explicit enumeration of the nodes involved in those typed graphlets. Algorithm 1 shows the general approach for counting typed network motifs. Note that we do not make any restriction or assumption on the number of node or edge types. The algorithm naturally handles heterogeneous graphs with arbitrary number of types and structure. See Table 1 for a summary of notation.

3.1 Counting Typed 3-Node Motifs

We begin by introducing the notion of a typed neighborhood and typed degree of a node. These are then used as a basis for deriving all typed 3-node motif counts in worst-case $O(\Delta)$ time (Theorem 1).

DEFINITION 7 (TYPED NEIGHBORHOOD). *Given an arbitrary node i in G , the typed neighborhood Γ_i^t is the set of nodes with type t that are reachable by following edges originating from i within 1-hop distance. More formally,*

$$\Gamma_i^t = \{j \in V \mid (i, j) \in E \wedge \phi_j = t\} \quad (2)$$

Intuitively, a node $j \in \Gamma_i^t$ iff there exists an edge $(i, j) \in E$ between node i and j and the type of node j denoted as ϕ_j is t .

DEFINITION 8 (TYPED DEGREE). *The typed-degree d_i^t of node i with type t is defined as $d_i^t = |\Gamma_i^t|$ where d_i^t is the number of nodes connected to node i with type t .*

Using these notions as a basis, we can define S_i^t , S_j^t , and T_{ij}^t for $t = 1, \dots, L$ (Figure 4). Obtaining these sets is equivalent to computing all 3-node typed motif counts. These sets are all defined with respect to a given edge $(i, j) \in E$ between node i and j with

Algorithm 1 Heterogeneous Network Motifs

Input: a graph $G = (V, E, \Phi, \xi)$

```

1 Initialize  $\Psi$  to all zeros
2 parallel for each  $(i, j) \in E$  in order do
3   Set  $\mathbf{x} = 0$  to a vector of all zeros and  $\mathcal{M}_{ij} = \emptyset$ 
4    $\mathcal{M}_{ij} = \emptyset$            ▷ set of typed motif ids occurring between  $i$  and  $j$ 
5   for  $k \in \Gamma_i$  do if  $k \neq j$  then  $\Psi(k) = \lambda_1$ 
6   for  $k \in \Gamma_j$  do
7     if  $k = i$  then continue
8     if  $\Psi(k) = \lambda_1$  then                                ▷ triangle motif
9        $T_{ij} \leftarrow T_{ij} \cup \{k\}$ ,  $|T_{ij}^{\phi_k}| = |T_{ij}^{\phi_k}| + 1$ , and set  $\Psi(k) = \lambda_3$ 
10       $\langle \mathbf{x}, \mathcal{M}_{ij} \rangle = \text{UPDATE}(\mathbf{x}, \mathcal{M}_{ij}, \mathbb{F}(g_2, \Phi_i, \Phi_j, \Phi_k, 0))$ 
11    else                                     ▷ typed 3-path centered at node  $j$ 
12       $S_j \leftarrow S_j \cup \{k\}$ ,  $|S_j^{\phi_k}| = |S_j^{\phi_k}| + 1$ , and set  $\Psi(k) = \lambda_2$ 
13       $\langle \mathbf{x}, \mathcal{M}_{ij} \rangle = \text{UPDATE}(\mathbf{x}, \mathcal{M}_{ij}, \mathbb{F}(g_1, \Phi_i, \Phi_j, \Phi_k, 0))$ 
14   for  $k \in (\Gamma_i \setminus T_{ij})$  do          ▷ Set of 3-path nodes centered at  $i$ 
15      $S_i \leftarrow S_i \cup \{k\}$  and  $|S_i^{\phi_k}| = |S_i^{\phi_k}| + 1$ 
16      $\langle \mathbf{x}, \mathcal{M}_{ij} \rangle = \text{UPDATE}(\mathbf{x}, \mathcal{M}_{ij}, \mathbb{F}(g_1, \Phi_i, \Phi_j, \Phi_k, 0))$ 
17   Given  $S_i$  and  $S_j$ , derive typed path-based motifs via Algorithm 2
18   Given  $T_{ij}$ , derive typed triangle-based motifs via Algorithm 3
19   for  $t, t' \in \{1, \dots, L\}$  such that  $t \leq t'$  do
20     Derive remaining typed graphlet orbits in constant time via
        Eq. 13-16 and update counts  $\mathbf{x}$  and set of motifs  $\mathcal{M}_{ij}$  if needed
21   for  $c \in \mathcal{M}_{ij}$  do ▷ unique typed graphlets between node  $i$  and  $j$ 
22      $X_{ij} = X_{ij} \cup \{(c, \mathbf{x}_c)\}$  ▷ store nonzero typed graphlet counts
23   Reset  $\Psi$  to all zero
24 end parallel
25 Merge all local typed motifs found by each worker to obtain  $\mathcal{M}$ 
26 return  $X$  and the set of motifs  $\mathcal{M}$  occurring in  $G$ 

```

types ϕ_i and ϕ_j . Since typed graphlets are counted for each edge $(i, j) \in E$, the types ϕ_i and ϕ_j are fixed ahead of time. Thus, there is only one remaining type to select for 3-node typed motifs.

DEFINITION 9 (TYPED TRIANGLE NODES). Given an edge $(i, j) \in E$ between node i and j with types ϕ_i and ϕ_j , let T_{ij}^t denote the set of nodes of type t that complete a typed triangle with node i and j defined as:

$$T_{ij}^t = \Gamma_i^t \cap \Gamma_j^t \quad (3)$$

where $|T_{ij}^t|$ denotes the number of nodes that form triangles with node i and j of type t . Furthermore, every node $k \in T_{ij}^t$ is of type t and thus completes a typed triangle with node i and j consisting of types ϕ_i , ϕ_j , and $\phi_k = t$.

DEFINITION 10 (TYPED 3-STAR NODES CENTERED AT i). Given an edge $(i, j) \in E$ between node i and j with types ϕ_i and ϕ_j . Let S_i^t denote the set of nodes of type t that form 3-node stars (or equivalently 3-node paths) centered at node i (and not including j). More formally,

$$S_i^t = \{k \in (\Gamma_i^t \setminus \{j\}) \mid k \notin \Gamma_j^t\} \quad (4)$$

$$= \Gamma_i^t \setminus (\Gamma_j^t \cup \{j\}) = \Gamma_i^t \setminus T_{ij}^t \quad (5)$$

where $|S_i^t|$ denotes the number of nodes of type t that form 3-stars centered at node i (not including j).

DEFINITION 11 (TYPED 3-STAR NODES CENTERED AT j). Given an edge $(i, j) \in E$ between node i and j with types ϕ_i and ϕ_j , let S_j^t denote the set of nodes of type t that form 3-node stars centered at node j (and not including i) defined formally as:

$$S_j^t = \{k \in (\Gamma_j^t \setminus \{i\}) \mid k \notin \Gamma_i^t\} \quad (6)$$

$$= \Gamma_j^t \setminus (\Gamma_i^t \cup \{i\}) = \Gamma_j^t \setminus T_{ij}^t \quad (7)$$

where $|S_j^t|$ denotes the number of nodes of type t that form 3-stars centered at node j (not including i).

PROPERTY 1.

$$T_{ij} = \bigcup_{t=1}^L T_{ij}^t, \quad S_i = \bigcup_{t=1}^L S_i^t, \quad S_j = \bigcup_{t=1}^L S_j^t \quad (8)$$

This property is shown in Figure 4. These lower-order 3-node typed motif counts are used to derive many higher-order typed motif counts in $o(1)$ constant time (Section 3.3).

DEFINITION 12 (TYPED 3-STARS (FOR AN EDGE)). Given an edge $(i, j) \in E$ between node i and j with types ϕ_i and ϕ_j , the number of typed 3-node stars that contain $(i, j) \in E$ with types ϕ_i, ϕ_j, t is:

$$|S_{ij}^t| = |S_i^t| + |S_j^t| \quad (9)$$

where $|S_{ij}^t|$ denotes the number of typed 3-stars that contain nodes i and j with types ϕ_i, ϕ_j, t .

Moreover, the number of typed triangles centered at $(i, j) \in E$ with types ϕ_i, ϕ_j, t is simply $|T_{ij}^t|$ (Definition 9) whereas the number of typed 3-node stars that contain $(i, j) \in E$ with types ϕ_i, ϕ_j, t is $|S_{ij}^t| = |S_i^t| + |S_j^t|$ (Definition 12). We do not need to actually store the sets S_i^t , S_j^t , and T_{ij}^t for every type $t = 1, \dots, L$. We only need to store the size/cardinality of the sets (as shown in Algorithm 1). For convenience, we denote the size of those sets as $|S_i^t|$, $|S_j^t|$, and $|T_{ij}^t|$ for all $t = 1, \dots, L$, respectively.

Algorithm 2 Heterogeneous Path-based Network Motifs

Input: a graph $G = (V, E, \Phi, \xi)$, an edge (i, j) , sets of nodes S_i and S_j that form 3-paths centered at i and j , respectively, a typed motif count vector \mathbf{x} for (i, j) , and set \mathcal{M}_{ij} of unique typed motifs for i and j .

```

1 for each  $w_k \in S_i$  in order  $w_1, w_2, \dots$  of  $S_i$  do
2   for  $w_r \in \Gamma_{w_k}$  do
3     if  $(w_r = i) \vee (w_r = j)$  continue
4     if  $w_r \notin (\Gamma_i \cup \Gamma_j)$  then          ▷  $\Psi(w_r) = 0$ , then 4-path-edge orbit
5        $\langle \mathbf{x}, \mathcal{M}_{ij} \rangle = \text{UPDATE}(\mathbf{x}, \mathcal{M}_{ij}, \mathbb{F}(g_3, \Phi_i, \Phi_j, \Phi_{w_k}, \Phi_{w_r}))$ 
6     else if  $w_r \in S_i \wedge w_r \leq w_k$  then ▷  $\Psi(w_r) = \lambda_1$ , tailed-tri (tail)
7        $\langle \mathbf{x}, \mathcal{M}_{ij} \rangle = \text{UPDATE}(\mathbf{x}, \mathcal{M}_{ij}, \mathbb{F}(g_7, \Phi_i, \Phi_j, \Phi_{w_k}, \Phi_{w_r}))$ 
8   for each  $w_k \in S_j$  in order  $w_1, w_2, \dots$  of  $S_j$  do
9     for  $w_r \in \Gamma_{w_k}$  do
10      if  $(w_r = i) \vee (w_r = j)$  continue
11      if  $w_r \notin (\Gamma_i \cup \Gamma_j)$  then          ▷  $\Psi(w_r) = 0$ , typed 4-path-edge orbit
12        $\langle \mathbf{x}, \mathcal{M}_{ij} \rangle = \text{UPDATE}(\mathbf{x}, \mathcal{M}_{ij}, \mathbb{F}(g_3, \Phi_i, \Phi_j, \Phi_{w_k}, \Phi_{w_r}))$ 
13     else if  $w_r \in S_j \wedge w_r \leq w_k$  then ▷  $\Psi(w_r) = \lambda_2$ , tailed-tri (tail)
14        $\langle \mathbf{x}, \mathcal{M}_{ij} \rangle = \text{UPDATE}(\mathbf{x}, \mathcal{M}_{ij}, \mathbb{F}(g_7, \Phi_i, \Phi_j, \Phi_{w_k}, \Phi_{w_r}))$ 
15     else if  $w_r \in S_i$  then                  ▷  $\Psi(w_r) = \lambda_1$ , typed 4-cycle
16        $\langle \mathbf{x}, \mathcal{M}_{ij} \rangle = \text{UPDATE}(\mathbf{x}, \mathcal{M}_{ij}, \mathbb{F}(g_6, \Phi_i, \Phi_j, \Phi_{w_k}, \Phi_{w_r}))$ 
17 return set of typed motifs  $\mathcal{M}_{ij}$  occurring between  $i$  and  $j$  and  $\mathbf{x}$ 

```

Algorithm 3 Heterogeneous Triangle-based Network Motifs

Input: a graph $G = (V, E, \Phi, \xi)$, an edge (i, j) , set of nodes T_{ij} that form triangles with i and j , sets of nodes S_i and S_j that form 3-paths centered at i and j , respectively, a typed motif count vector \mathbf{x} for (i, j) , and set \mathcal{M}_{ij} of unique typed motifs for i and j .

```

1 for each  $w_k \in T_{ij}$  in order  $w_1, w_2, \dots$  of  $T_{ij}$  do
2   for  $w_r \in \Gamma_{w_k}$  do
3     if  $w_r \in T_{ij} \wedge w_r \leq w_k$  then  $\triangleright \Psi(w_r) = \lambda_3$ , typed 4-clique
4        $\langle \mathbf{x}, \mathcal{M}_{ij} \rangle = \text{UPDATE}(\mathbf{x}, \mathcal{M}_{ij}, \mathbb{F}(g_{12}, \Phi_i, \Phi_j, \Phi_{w_k}, \Phi_{w_r}))$ 
5     else if  $w_r \in (S_i \cup S_j)$  then  $\triangleright \Psi(w_r) = \lambda_1 \vee \lambda_2$ , chordal-cycle-edge
6       if  $(w_r = i) \vee (w_r = j)$  continue
7        $\langle \mathbf{x}, \mathcal{M}_{ij} \rangle = \text{UPDATE}(\mathbf{x}, \mathcal{M}_{ij}, \mathbb{F}(g_{10}, \Phi_i, \Phi_j, \Phi_{w_k}, \Phi_{w_r}))$ 
8     else if  $w_r \notin (\Gamma_i \cup \Gamma_j)$  then  $\triangleright$  tailed-tri-center
9       if  $(w_r = i) \vee (w_r = j)$  continue
10       $\langle \mathbf{x}, \mathcal{M}_{ij} \rangle = \text{UPDATE}(\mathbf{x}, \mathcal{M}_{ij}, \mathbb{F}(g_8, \Phi_i, \Phi_j, \Phi_{w_k}, \Phi_{w_r}))$ 
11 return set of typed motifs  $\mathcal{M}_{ij}$  occurring between  $i$  and  $j$  and  $\mathbf{x}$ 
```

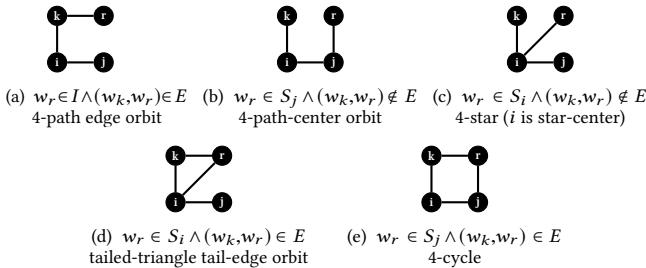


Figure 3: Path-based 4-node orbits derived from lower-order $(k-1)$ -node graphlets in Algorithm 2. For each motif orbit above, we provide the equation for counting the motif orbit. Since $w_k \in S_i$ holds in (a)-(e), it was removed from the equations above for brevity.

3.2 Counting Typed 4-Node Motifs

To derive k -node typed graphlets, the framework leverages the lower-order $(k-1)$ -node typed graphlets. Therefore, 4-node typed graphlets are derived by leveraging the sets $T_{ij} = \Gamma_i \cup \Gamma_j$, $S_j = \Gamma_j \setminus T_{ij}$, and $S_i = \Gamma_i \setminus T_{ij}$ computed from the lower-order 3-node typed graphlets along with the set I of non-adjacent nodes w.r.t. $(i, j) \in E$ defined formally as follows:

$$\begin{aligned} I &= V \setminus (\Gamma_i \cup \Gamma_j \cup \{i, j\}) \\ &= V \setminus (T_{ij} \cup S_i \cup S_j \cup \{i, j\}). \end{aligned} \quad (10)$$

PROPERTY 2.

$$|V| = |I| + |S_i| + |S_j| + |T_{ij}| + |\{i, j\}| \quad (11)$$

The proof is straightforward by Eq. 10 and applying the principle of inclusion-exclusion [4].

3.2.1 General Principle for Counting Typed Graphlets. We now introduce a general typed graphlet formulation. Let $f_{ij}(H, \mathbf{t})$ denote the number of distinct k -node typed graphlet orbits of H with the type vector \mathbf{t} that contain edge $(i, j) \in E$ and have properties

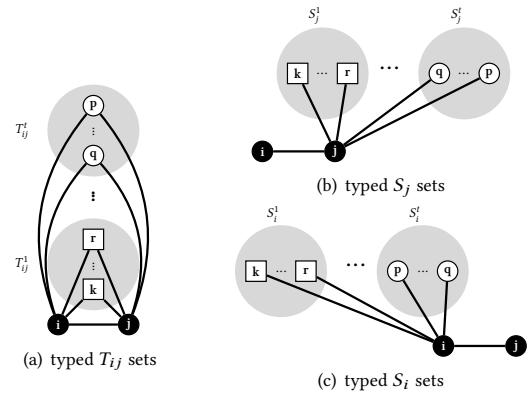


Figure 4: Typed lower-order sets used to derive many higher-order graphlets in constant time. Note node i and j can be of arbitrary types.

$P, Q \in \{S_i, S_j, T_{ij}, I\}$ defined formally as:

$$\begin{aligned} f_{ij}(H, \mathbf{t}) &= \left| \left\{ \{i, j, w_k, w_r\} \mid w_k \in P \wedge w_r \in Q \wedge \right. \right. \\ &\quad \left. \left. \mathbb{I}\{(w_k, w_r) \in E\} \wedge w_r \neq w_k \wedge \right. \right. \\ &\quad \left. \left. \mathbf{t} = [\phi_i \ \phi_j \ \phi_{w_k} \ \phi_{w_r}] \right\} \right| \end{aligned} \quad (12)$$

where $\mathbb{I}\{(w_k, w_r) \in E\} = 1$ if $(w_k, w_r) \in E$ holds and 0 otherwise (i.e., $\mathbb{I}\{(w_k, w_r) \in E\} = 0$ if $(w_k, w_r) \notin E$). For clarity and simplicity, $(w_k, w_r) \in E$ or $(w_k, w_r) \notin E$ is sometimes used (e.g., Table 3) as opposed to $\mathbb{I}\{(w_k, w_r) \in E\} = 1$ or $\mathbb{I}\{(w_k, w_r) \in E\} = 0$.

The equations for deriving every typed graphlet orbit are provided in Table 3. Notice that all typed graphlets with k -nodes are formulated with respect to the node sets $\{S_i, S_j, T_{ij}, I\}$ derived from the typed graphlets with $(k-1)$ -nodes. Hence, the higher-order typed graphlets with order k are derived from the lower-order $(k-1)$ -node typed graphlets.

We classify typed motifs as path-based or triangle-based. Typed path-based motifs are the typed 4-node motifs derived from the sets S_i and S_j of nodes that form 3-node paths centered at node i and j , respectively. Conversely, typed triangle-based motifs are the typed 4-node motifs derived from the set T_{ij} of nodes that form triangles (3-cliques) with node i and j . Naturally, typed path-based motifs are the least dense (motifs with fewest edges) whereas the typed triangle-based motifs are the most dense. Recall $T_{ij} = \Gamma_i \cap \Gamma_j$, $S_j = \Gamma_j \setminus T_{ij}$, $S_i = \Gamma_i \setminus T_{ij}$, and $I = V \setminus (\Gamma_i \cup \Gamma_j \cup \{i, j\}) = V \setminus (T_{ij} \cup S_i \cup S_j \cup \{i, j\})$. Therefore, $|\Gamma_i| \geq |T_{ij}|$ and $|\Gamma_i| \geq |S_i|$. Further, if $|\Gamma_i| = |T_{ij}|$ then $|S_i| = 0$ and conversely if $|\Gamma_i| = |S_i|$ then $|T_{ij}| = 0$.

In this work, we derive equations using new non-trivial combinatorial relationships between lower-order typed $(k-1)$ -graphlets that allow us to derive many of the k -node typed graphlets in $o(1)$ constant time. Notably, we make no assumptions about the number of types, their distribution among the nodes and edges, or any other additional information. On the contrary, the framework is extremely general for arbitrary heterogeneous graphs (see Figure 2 for a number of popular special cases covered by the framework). In addition, we also avoid a lot of computations by symmetry breaking techniques, and other conditions to avoid unnecessary work. Typed graphlets that are computed in constant time include typed 4-path-center orbit, typed 4-node stars, typed chordal-cycle-center

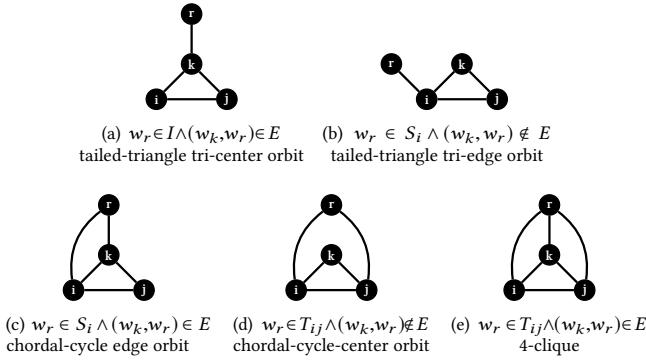


Figure 5: Triangle-based 4-node orbits derived from lower-order ($k-1$)-node graphlets in Algorithm 3. For each motif orbit above, we provide the equation used to count it. Since $w_k \in T_{ij}$ holds for all triangle-based 4-node orbits, it was removed from the equations above for brevity.

orbit, and typed tailed-triangle-tri-edge orbit. Notice that two of the typed tailed-triangle orbits including the tri-center and tail-edge orbit are essentially derived for free while computing the typed graphlets using S_i (typed 4-cycles) and S_j .

3.2.2 Directed Typed Graphlets. The approach is also straightforward to adapt for directed typed motifs. In particular, we simply replace Γ_i^t with $\Gamma_i^{t,+}$ and $\Gamma_i^{t,-}$ for typed out-neighbors and typed in-neighbors, respectively. Thus, we also have $T_{ij}^{t,+}, T_{ij}^{t,-}, S_j^{t,+}, S_j^{t,-}, S_i^{t,+}$, and $S_i^{t,-}$. Now it is just a matter of enumerating all combinations of these sets with the out/in-neighbor sets as well. That is, we essentially have two additional versions of Algorithm 1 and Algorithm 2-3 for each in and out set (w.r.t. to the main for loop). The other trivial modification is to ensure each directed typed motif is assigned a unique id (this is the same modification required for typed orbits). The time and space complexity remains the same since all we did is split the set of neighbors (and the other sets) into two smaller sets by partitioning the nodes in Γ_i^t into $\Gamma_i^{t,+}$ and $\Gamma_i^{t,-}$. Similarly, for T_{ij}^t, S_j^t , and S_i^t .

3.3 Combinatorial Relationships

Using new combinatorial relationships between lower-order typed graphlets, we derive all remaining typed graphlet orbits in $o(1)$ constant time via Eq. 13-16 (See Line 19-20). Since we derive all typed graphlet counts for a given edge $(i, j) \in E$ between node i and j , we already have two types ϕ_i and ϕ_j . Thus, these types are fixed ahead of time. In the case of 4-node typed graphlets, there are two remaining types that need to be selected. Notice that for typed graphlet orbits, we must solve $\frac{L(L-1)}{2} + L$ equations in the worst-case. The counts of all remaining typed graphlets are derived in $o(1)$ constant time using the counts of the lower-order ($k-1$)-node typed graphlets and a few other counts from the k -node typed graphlets. After deriving the exact count of each remaining graphlet with types ϕ_i, ϕ_j, t , and t' for every $t, t' \in \{1, \dots, L\}$ such that $t \leq t'$ (Line 19-20), if such count is nonzero, we compute a motif hash $c = \mathbb{F}(g, \phi_i, \phi_j, t, t')$ for graphlet orbit g , set $\mathcal{M}_{ij} \leftarrow \mathcal{M}_{ij} \cup \{c\}$, and

then set the count of that typed motif in \mathbf{x}_c to the count derived in constant $o(1)$ time.

Typed 4-Path Center Orbit Count (Figure 3(b)). To count the typed 4-path center orbits for a given edge $(i, j) \in E$ with types ϕ_i and ϕ_j , we simply select the remaining two types denoted as t and t' to obtain the 4-dimensional type vector $\mathbf{t} = [\phi_i \ \phi_j \ t \ t']$ and derive the count in $o(1)$ constant time as follows:

$$f_{ij}(g_4, \mathbf{t}) = \begin{cases} (|S_i^t| \cdot |S_j^t|) - f_{ij}(g_6, \mathbf{t}) & \text{if } t = t' \\ (|S_i^t| \cdot |S_j^{t'}|) + (|S_i^{t'}| \cdot |S_j^t|) - f_{ij}(g_6, \mathbf{t}) & \text{otherwise} \end{cases} \quad (13)$$

where $f_{ij}(g_6, \mathbf{t})$ is the typed 4-cycle count for edge $(i, j) \in E$ with type vector \mathbf{t} .

Typed 4-Star Count (Figure 3(c)). To count the typed 4-stars for a given edge $(i, j) \in E$ with types ϕ_i and ϕ_j , we simply select the remaining two types denoted as t and t' to obtain the 4-dimensional type vector $\mathbf{t} = [\phi_i \ \phi_j \ t \ t']$. We derive the typed 4-star counts with the type vector \mathbf{t} for edge $(i, j) \in E$ in constant time as follows:

$$f_{ij}(g_5, \mathbf{t}) = \begin{cases} \binom{|S_i^t|}{2} + \binom{|S_j^t|}{2} - f_{ij}(g_7, \mathbf{t}) & \text{if } t = t' \\ (|S_i^t| \cdot |S_j^{t'}|) + (|S_j^t| \cdot |S_i^{t'}|) - f_{ij}(g_7, \mathbf{t}) & \text{otherwise} \end{cases} \quad (14)$$

where $f_{ij}(g_7, \mathbf{t})$ is the tailed-triangle tail-edge orbit count for edge $(i, j) \in E$ with type vector \mathbf{t} . The only path-based typed graphlet containing a triangle is the tailed-triangle tail-edge orbit shown in Figure 3(d). As an aside, Figure 3(d) is for the tailed-triangle tail-edge orbits centered at node i ; however, the other tailed-triangle tail-edge orbit centered at node j is also computed. Observe that this is the only orbit needed to derive the typed 4-stars in constant time as shown in Eq. 14.

Typed Tailed-Triangle Tri-Edge Orbit Count (Figure 5(b)):

$$f_{ij}(g_9, \mathbf{t}) = \begin{cases} (|T_{ij}^t| \cdot (|S_i^t| + |S_j^t|)) - f_{ij}(g_{10}, \mathbf{t}) & \text{if } t = t' \\ (|T_{ij}^t| \cdot (|S_i^{t'}| + |S_j^{t'}|)) + (|T_{ij}^{t'}| \cdot (|S_i^t| + |S_j^t|)) - f_{ij}(g_{10}, \mathbf{t}) & \text{otherwise} \end{cases} \quad (15)$$

where $f_{ij}(g_{10}, \mathbf{t})$ is the chordal-cycle edge orbit count for edge $(i, j) \in E$ with type vector \mathbf{t} .

Typed Chordal-Cycle Center Orbit Count (Figure 5(d)):

$$f_{ij}(g_{11}, \mathbf{t}) = \begin{cases} \binom{|T_{ij}^t|}{2} - f_{ij}(g_{12}, \mathbf{t}) & \text{if } t = t' \\ (|T_{ij}^t| \cdot |T_{ij}^{t'}|) - f_{ij}(g_{12}, \mathbf{t}) & \text{otherwise} \end{cases} \quad (16)$$

where $f_{ij}(g_{12}, \mathbf{t})$ is the typed 4-clique count for edge $(i, j) \in E$ with type vector \mathbf{t} .

Discussion: Many other combinatorial relationships can be derived in a similar fashion. The equations shown above are the ones required such that the worst-case time complexity matches that of the best known untyped graphlet counting algorithm (See Section 6).

Table 3: Typed Graphlet Orbit Equations. All typed graphlet orbits with 4-nodes are formulated with respect to the node sets $\{S_i, S_j, T_{ij}, I\}$ derived from the typed 3-node graphlets. Recall $T_{ij} = \Gamma_i \cap \Gamma_j$, $S_j = \Gamma_j \setminus T_{ij}$, $S_i = \Gamma_i \setminus T_{ij}$, and $I = V \setminus (\Gamma_i \cup \Gamma_j \cup \{i, j\}) = V \setminus (T_{ij} \cup S_i \cup S_j \cup \{i, j\})$. Therefore, $|\Gamma_i| \geq |T_{ij}|$ and $|\Gamma_i| \geq |S_i|$. Further, if $|\Gamma_i| = |T_{ij}|$ then $|S_i| = 0$ and conversely if $|\Gamma_i| = |S_i|$ then $|T_{ij}| = 0$. In all cases, $w_r \neq w_k$. Note $\rho(H)$ is density.

| MOTIF H | ORBIT | $ E(H) $ | $\rho(H)$ | $f_{ij}(H, t) = \left \left\{ \{i, j, w_k, w_r\} \mid w_k \in P \wedge w_r \in Q \wedge \exists \{(w_k, w_r) \in E\} \wedge w_r \neq w_k \wedge t = [\phi_i \ \phi_j \ \phi_{w_k} \ \phi_{w_r}] \right\} \right $ |
|------------------------|-----------|----------|-----------|---|
| 4-path | edge | 3 | 0.50 | $f_{ij}(g_3, t) = \left \left\{ \{i, j, w_k, w_r\} \mid w_k \in S_i \wedge w_r \in I \wedge (w_k, w_r) \in E \wedge t = [\phi_i \ \phi_j \ \phi_{w_k} \ \phi_{w_r}] \right\} \right $ |
| | center | 3 | 0.50 | $f_{ij}(g_4, t) = \left \left\{ \{i, j, w_k, w_r\} \mid w_k \in S_j \wedge w_r \in S_i \wedge (w_k, w_r) \notin E \wedge t = [\phi_i \ \phi_j \ \phi_{w_k} \ \phi_{w_r}] \right\} \right $ |
| 4-star | | 3 | 0.50 | $f_{ij}(g_5, t) = \left \left\{ \{i, j, w_k, w_r\} \mid w_k \in S_i \wedge w_r \in S_i \wedge (w_k, w_r) \notin E \wedge t = [\phi_i \ \phi_j \ \phi_{w_k} \ \phi_{w_r}] \right\} \right $ |
| 4-cycle | | 4 | 0.67 | $f_{ij}(g_6, t) = \left \left\{ \{i, j, w_k, w_r\} \mid w_k \in S_j \wedge w_r \in S_i \wedge (w_k, w_r) \in E \wedge t = [\phi_i \ \phi_j \ \phi_{w_k} \ \phi_{w_r}] \right\} \right $ |
| tailed-triangle | tail-edge | 4 | 0.67 | $f_{ij}(g_7, t) = \left \left\{ \{i, j, w_k, w_r\} \mid w_k \in S_i \wedge w_r \in S_i \wedge w_r \neq w_k \wedge (w_k, w_r) \in E \wedge t = [\phi_i \ \phi_j \ \phi_{w_k} \ \phi_{w_r}] \right\} \right $ |
| | center | 4 | 0.67 | $f_{ij}(g_8, t) = \left \left\{ \{i, j, w_k, w_r\} \mid w_k \in T_{ij} \wedge w_r \in I \wedge (w_k, w_r) \in E \wedge t = [\phi_i \ \phi_j \ \phi_{w_k} \ \phi_{w_r}] \right\} \right $ |
| | tri-edge | 4 | 0.67 | $f_{ij}(g_9, t) = \left \left\{ \{i, j, w_k, w_r\} \mid w_k \in T_{ij} \wedge w_r \in S_i \wedge (w_k, w_r) \notin E \wedge t = [\phi_i \ \phi_j \ \phi_{w_k} \ \phi_{w_r}] \right\} \right $ |
| chordal-cycle | edge | 5 | 0.83 | $f_{ij}(g_{10}, t) = \left \left\{ \{i, j, w_k, w_r\} \mid w_k \in T_{ij} \wedge w_r \in (S_i \cup S_j) \wedge w_r \neq w_k \wedge (w_k, w_r) \in E \wedge t = [\phi_i \ \phi_j \ \phi_{w_k} \ \phi_{w_r}] \right\} \right $ |
| | center | 5 | 0.83 | $f_{ij}(g_{11}, t) = \left \left\{ \{i, j, w_k, w_r\} \mid w_k \in T_{ij} \wedge w_r \in T_{ij} \wedge w_r \neq w_k \wedge (w_k, w_r) \notin E \wedge t = [\phi_i \ \phi_j \ \phi_{w_k} \ \phi_{w_r}] \right\} \right $ |
| 4-clique | | 6 | 1.00 | $f_{ij}(g_{12}, t) = \left \left\{ \{i, j, w_k, w_r\} \mid w_k \in T_{ij} \wedge w_r \in T_{ij} \wedge w_r \neq w_k \wedge (w_k, w_r) \in E \wedge t = [\phi_i \ \phi_j \ \phi_{w_k} \ \phi_{w_r}] \right\} \right $ |

3.4 From Typed Orbits to Graphlets

Counts of the *typed graphlets* for each edge $(i, j) \in E$ can be derived from the *typed graphlet orbits* using the following equations:

$$f_{ij}(h_3, t) = f_{ij}(g_3, t) + f_{ij}(g_4, t) \quad (17)$$

$$f_{ij}(h_4, t) = f_{ij}(g_5, t) \quad (18)$$

$$f_{ij}(h_5, t) = f_{ij}(g_6, t) \quad (19)$$

$$f_{ij}(h_6, t) = f_{ij}(g_7, t) + f_{ij}(g_8, t) + f_{ij}(g_9, t) \quad (20)$$

$$f_{ij}(h_7, t) = f_{ij}(g_{10}, t) + f_{ij}(g_{11}, t) \quad (21)$$

$$f_{ij}(h_8, t) = f_{ij}(g_{12}, t) \quad (22)$$

where h is the graphlet without considering the orbit (Table 3).

Algorithm 4 Update Typed Graphlets. Add typed graphlet (with id c) to \mathcal{M}_{ij} if $c \notin \mathcal{M}_{ij}$ and increment x_c (frequency of that typed graphlet for a given edge).

```

1 procedure UPDATE(x,  $\mathcal{M}_{ij}$ ,  $c = \mathbb{F}(g, \Phi_i, \Phi_j, \Phi_k, \Phi_r)$ )
2   if  $c \notin \mathcal{M}_{ij}$  then
3      $\mathcal{M}_{ij} \leftarrow \mathcal{M}_{ij} \cup \{c\}$ 
4      $x_c = 0$ 
5      $x_c = x_c + 1$ 
6   return updated set of typed graphlets  $\mathcal{M}_{ij}$  and their counts x

```

3.5 Typed Motif Hash Functions

Given a general heterogeneous graph with L unique types such that $L < 10$, then a simple and efficient typed motif hash function \mathbb{F} is defined as follows:

$$\mathbb{F}(g, t) = g10^4 + t_110^3 + t_210^2 + t_310^1 + t_4 \quad (23)$$

where g encodes the k -node motif orbit (e.g., 4-path center) and t_1, t_2, t_3, t_4 encode the type of the nodes in $H \in \mathcal{H}$ with type vector

$t = [t_1 \ t_2 \ t_3 \ t_4]$. Since the maximum hash value resulting from Eq. 23 is small (and fixed for any arbitrarily large graph G), we can leverage a perfect hash table to allow for fast $o(1)$ constant time lookups to determine if a typed motif was previously found or not as well as updating the typed motif count in $o(1)$ constant time. For k -node motifs where $k < 4$, we simply set the last $4 - k$ types to 0. Note the simple typed motif hash function defined above can be extended trivially to handle graphs with $L \geq 10$ types as follows:

$$\mathbb{F}(g, t) = g10^8 + t_110^6 + t_210^4 + t_310^2 + t_4 \quad (24)$$

In general, any non-cryptographic hash function \mathbb{F} can be used (see [13] for some other possibilities). Thus, the approach is independent of \mathbb{F} and can always leverage the best known hash function. The only requirement of the hash function is that it is invertible \mathbb{F}^{-1} .

Thus far we have not made any assumption on the ordering of types in t . As such, the hash function \mathbb{F} discussed above can be used directly in the framework for counting typed graphlets such that the type structure and position are preserved. However, since we are interested in counting all typed graphlets *w.r.t.* Definition 5, then we map all such orderings of the types in t to the same hash value using a precomputed hash table. This allows us to obtain the unique hash value in $o(1)$ constant time for any ordering of the types in t . In our implementation, we compute $s = t_110^3 + t_210^2 + t_310^1 + t_4$ and then use s as an index into the precomputed hash table to obtain the unique hash value c in $o(1)$ constant time.

3.6 Sparse Typed Motif Format

This section describes a space-efficient representation for typed motifs based on a key observation.

PROPERTY 3. Let T denote the number of unique typed motifs that appear in an arbitrary graph G with L types. Assuming the graph

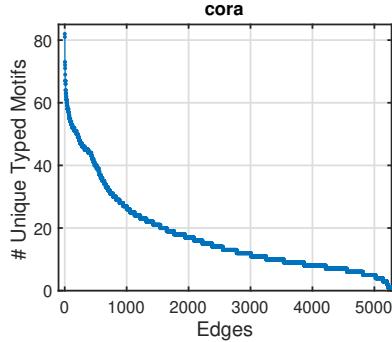


Figure 6: Distribution of unique typed motifs that occur on the edges. This experiment considers all typed motifs of $\{3, 4\}$ -nodes. Among the 1428 possible unique typed motifs that could arise in G , there are only 876 unique typed motifs that actually occur (at least once at an edge in G). Even more striking, the maximum unique typed motifs that occur on any edge in G (cora) is only 82. Overall, the mean number of unique typed motifs over all edges in G is 17, i.e., only about 1.1% of the possible typed motifs. These results indicate the significance of only a few typed motifs as the vast majority of the typed motif counts for any arbitrary edge is zero. Thus, the space required by the approach is nearly-optimal.

G has a skewed degree distribution, then most edges in G appear in only a small fraction of the T actual typed motifs that can occur.

This property is shown empirically in Figure 6. Property 3 implies that using a $M \times T$ matrix to store the typed graphlet counts is far from optimal in terms of the space required due to most of the T typed motif counts being zero for any given edge. Based on this observation, we describe an approach that uses near-optimal space by storing only the typed motifs with nonzero counts for each edge $(i, j) \in E$ in the graph. Typed motif counts are stored in a sparse format since it would be impractical in large graphs to store all typed motifs as there can easily be hundreds of thousands depending on the number of types in the input graph. For each edge, we store only the nonzero typed motif counts along with the unique ids associated with them. The unique ids allow us to map the nonzero counts to the actual typed motifs. We also developed a space-efficient format for storing the resulting typed motif counts to disk. Instead of using the typed motif hash as the unique id, we remap the typed motifs to smaller consecutive ids (starting from 1) to reduce the space requirement even further. Finally, we store a typed motif lookup table that maps a given motif id to its description and is useful for looking up the meaning of the typed motifs discovered.

4 GLOBAL TYPED GRAPHLET COUNTS

While Section 3 focused on counting typed graphlets locally for each edge $(i, j) \in E$ in G , one may also be interested in the total counts of each typed graphlet in G . This is known as the global typed graphlet counting problem. More formally,

DEFINITION 13 (GLOBAL TYPED GRAPHLET COUNTING). Given a graph G with L types, the global typed graphlet counting problem

is to find the set of all typed motifs that occur in G along with their corresponding frequencies.

A general equation for solving the above problem for any arbitrary *typed graphlet* H is given below. Let H denote an arbitrary typed graphlet and \mathbf{x} be an M -dimensional vector of counts of H for every edge $(i, j) \in E$, then the frequency of H in G is:

$$C = \frac{1}{|E(H)|} \mathbf{x}^\top \mathbf{e} \quad (25)$$

where $|E(H)|$ is the number of edges in the typed graphlet H and $\mathbf{e} = [1 \ \dots \ 1]$ is an M -dimensional vector of all 1's. For instance, suppose H is one of the typed 4-cliques in Figure 1, then $|E(H)| = 6$. It is straightforward to see that if we were interested in *only* deriving global counts of the typed graphlets, then we can significantly speedup the proposed algorithms by enforcing constraints to avoid counting a typed graphlet numerous times. Instead, we can add constraints to count such motifs once per edge by enforcing an ordering.

5 PARALLELIZATION

We now describe a parallelization strategy for the proposed typed motif counting approach. While our implementation uses shared memory, the parallelization is described generally such that it can be used with a distributed-memory architecture as well. As such, our discussion is on the general scheme.

The parallel constructs we use are a worker task-queue and a global broadcast channel. Here, we assume that each worker has a copy of the graph and distribute edges to workers to find the typed graphlet counts that node i and j participate. At this point, we view the main while loop as a task generator and farm the current edge out to a worker to find the typed graphlet counts that co-occur between node i and node j . The approach is lock free since each worker uses the same motif hash function to obtain a unique hash value for every typed motif. Thus, each worker can simply maintain the typed motifs identified and their counts for every edge assigned to it. In our own shared memory implementation, we avoid some of the communications by using global arrays and avoiding locked updates to them by using a unique edge id. Counting typed graphlets on the edges as opposed to the nodes also has computational advantages with respect to parallelization and in particular load balancing. Let x_i and x_{ij} denote the node and edge count of an arbitrary graphlet H . Since $|E| \gg |V|$ and $\sum_{i \in V} x_i = \sum_{(i,j) \in E} x_{ij}$, then $\frac{1}{|V|} \sum_{i \in V} x_i < \frac{1}{|E|} \sum_{(i,j) \in E} x_{ij}$. Hence, more work per vertex is required than per edge. Therefore, counting typed graphlets on the edges is guaranteed to have better load balancing than node-centric algorithms.

6 THEORETICAL ANALYSIS

We first show that heterogeneous graphlets contain more information than untyped graphlets. Next, Section 6.1 provides the worst-case time complexity of the proposed approach whereas Section 6.2 gives the space complexity. See Table 1 for a summary of notation.

First, we show the relationship between the count of an untyped motif H in G and the count of all typed motifs in G with induced subgraph H .

PROPOSITION 1. Let \mathbf{x} denote the vector of counts for any untyped network motif $H \in \mathcal{H}$ (e.g., 4-cycle). Further, let \mathbf{X} denote a $M \times T_H$ matrix of typed motif counts for motif H where T_H denotes the number of typed network motifs that arise from L types. Then the following holds:

$$C = \sum_{i=1}^M x_i = \sum_{i=1}^M \sum_{j=1}^{T_H} X_{ij} \quad (26)$$

Consider the counts of an untyped graphlet for a single edge. The above demonstrates that these counts are partitioned among the set of typed graphlets that arise from the untyped graphlet when types are considered.

We now show that typed network motifs contain more information than untyped motifs. Entropy is a measure of the *average information content* of a state [28]. Let $\mathbf{p} \in \mathbb{R}^{T_H}$ denote a typed motif probability distribution ($\mathbf{p}^T \mathbf{e} = 1$), the entropy (average information content) of \mathbf{p} is $\mathbb{H}(\mathbf{p}) = -\sum_i p_i \log p_i$. Hence, $\mathbb{H}(\mathbf{p})$ quantifies the amount of information in the relative frequencies of the typed network motifs (of a given motif $H \in \mathcal{H}$). In the case of untyped motifs, the $C = \sum_{i=1}^M x_i$ untyped motifs are assumed to belong to a single homogeneous motif where all nodes are of the same type. This matches exactly the information we have if types are not considered.

PROPOSITION 2. Assume $\mathbf{p} \in \mathbb{R}^{T_H}$ is an arbitrary typed network motif probability distribution such that $p_i < 1, \forall i$ and \mathbf{q} is the untyped motif distribution where $q_i = 1$ and $q_j = 0, \forall j \neq i$, then $\mathbb{H}(\mathbf{p}) > \mathbb{H}(\mathbf{q})$.

This implies the typed motif distribution \mathbf{p} contains more information than the untyped motif distribution \mathbf{q} . The proof is straightforward.

6.1 Time Complexity

PROPERTY 4.

$$d_i + d_j = 2|T_{ij}| + |S_i| + |S_j| \quad (27)$$

where $d_i = |\Gamma_i|$, $d_j = |\Gamma_j|$, $T_{ij} = \Gamma_i \cap \Gamma_j$, $S_i = \Gamma_i \setminus T_{ij}$, and $S_j = \Gamma_j \setminus T_{ij}$.

PROPERTY 5. The space required to store T_{ij} , S_i , and S_j is less than $d_i + d_j$ iff $|T_{ij}| > 0$.

This is straightforward to see since $|S_i| + |S_j| = |S_i \cup S_j|$ always holds. However, if $|T_{ij}| = 0$, then $|S_i| + |S_j| = d_i + d_j$. Hence, triangles represent the smallest clique, and as shown in [47] can be used to compress the graph. As the density of the graph increases, more triangles are formed, and therefore less space is used. Notice that the worst case is also unlikely to occur because of this fact. For instance, suppose $d_i = \Delta$, $d_j = \Delta$, and $\Delta = n$ (worst case), then $|T_{ij}| = d_i = d_j$, and $|S_i| = 0$, $|S_j| = 0$. Furthermore, if $|S_i| = n$, then $|S_j| = 0$ and $|T_{ij}| = 0$ must hold. Obviously, if node i is connected to all n nodes, then any node $k \in \Gamma_j$ must form a triangle with i ($k \in T_{ij}$).

6.1.1 Typed 3-Node Graphlets.

THEOREM 1. The overall worst-case time complexity for counting all 3-node typed graphlets for a given edge $(i, j) \in E$ is:

$$\mathcal{O}(2|\Gamma_i| + |\Gamma_j|) = \mathcal{O}(\Delta) \quad (28)$$

where $|\Gamma_i|$ and $|\Gamma_j|$ denote the number of nodes connected to node i and j , respectively. Further, Δ is the maximum degree in G .

PROOF. It takes at most $\mathcal{O}(|\Gamma_i|)$ time for Line 5 and Line 14-16 whereas the time complexity for Line 6-13 is $\mathcal{O}(|\Gamma_j|)$. Obviously, since Γ_i is enumerated twice, then we can always ensure i is the node with smallest degree (i.e., $d_i \leq d_j$). The worst-case arises when $|\Gamma_i|$ and $|\Gamma_j|$ is the maximum degree of a node in G denoted as Δ . ■

6.1.2 Typed 4-Node Graphlets. We first provide the time complexity of deriving path-based and triangle-based graphlet orbits in Lemma 6.1-6.2, and then give the total time complexity of all 3 and 4-node typed graphlets in Theorem 2 based on these results. Note that Lemma 6.1-6.2 includes the time required to derive all typed 3-node typed graphlets.

LEMMA 6.1. For a single edge $(i, j) \in E$, the worst-case time complexity for deriving all typed path-based graphlet orbits is:

$$\mathcal{O}\left(\Delta(|S_i| + |S_j|)\right) \quad (29)$$

Note $|S_i|\Delta \geq \sum_{k \in S_i} d_k$ and $|S_j|\Delta \geq \sum_{k \in S_j} d_k$.

LEMMA 6.2. For a single edge $(i, j) \in E$, the worst-case time complexity for deriving all typed triangle-based graphlet orbits is:

$$\mathcal{O}(\Delta|T_{ij}|) \quad (30)$$

Notice $|T_{ij}|\Delta \geq |T_{ij}|\Delta_T \geq \sum_{k \in T_{ij}} d_k$ where Δ is the maximum degree of a node in G and Δ_T is the maximum degree of a node in T_{ij} . Thus, $|T_{ij}|\Delta$ only occurs iff $\forall k \in T_{ij}, d_k = \Delta$ where $\Delta =$ maximum degree of a node in G . In sparse real-world graphs, T_{ij} is likely to be smaller than S_i and S_j as triangles are typically more rare than 3-node paths. Conversely, T_{ij} is also more likely to contain high degree nodes, as nodes with larger degrees are obviously more likely to form triangles than those with small degrees.

From Lemma 6.1-6.2, we have the following:

THEOREM 2. For a single edge $(i, j) \in E$, the worst-case time complexity for deriving all 3 and 4-node typed graphlet orbits is:

$$\mathcal{O}(\Delta(|S_i| + |S_j| + |T_{ij}|)) \quad (31)$$

PROOF. The time complexity of each step is provided below. Hashing all neighbors of node i takes $\mathcal{O}(|\Gamma_i|)$. Recall from Lemma 1 that counting all 3-node typed graphlets takes $\mathcal{O}(2|\Gamma_i| + |\Gamma_j|) = \mathcal{O}(\Delta)$ time for an edge $(i, j) \in E$. This includes the time required to derive the number of typed 3-node stars and typed triangles for all types $t = 1, \dots, L$. This information is needed to derive the remaining typed graphlet orbit counts in constant time. Next, Algorithm 2 is used to derive a few path-based typed graphlet orbit counts taking $\mathcal{O}(\Delta(|S_i| + |S_j|))$ time in the worst-case. Similarly, Algorithm 3 is used to derive a few triangle-based typed graphlet orbit counts taking in the worst-case $\mathcal{O}(\Delta|T_{ij}|)$ time. As an aside, updating the count of a typed graphlet count is $o(1)$ (Algorithm 4).

Now, we derive the remaining typed graphlet orbit counts in constant time (Line 19-20). Since each type pair leads to different typed graphlets, we must iterate over at most $L(L - 1)/2 + L$ type pairs. For each pair of types selected, we derive the typed graphlet orbit counts in $o(1)$ constant time via Eq. 13-16 (See Line 19-20). Furthermore, the term involving L is for the worst-case when there

is at least one node in all L sets (*i.e.*, at least one node of every type L). Nevertheless, since L is a small constant, $L(L - 1)/2 + L$ is negligible.

For a single edge, the worst-case time complexity is $O(\Delta(|S_i| + |S_j| + |T_{ij}|))$. Let \bar{T} and \bar{S} denote the average number of triangle and 3-node stars incident to an edge in G . More formally, $\bar{T} = \frac{1}{M} \sum_{(ij) \in E} |T_{ij}|$ and $\bar{S} = \frac{1}{M} \sum_{(ij) \in E} |S_i| + |S_j|$. The total worst-case time complexity for all M edges is $O(M\Delta(\bar{S} + \bar{T}))$. Note that obviously $\bar{S}M = \sum_{(ij) \in E} |S_i| + |S_j|$ and $\bar{T}M = \sum_{(ij) \in E} |T_{ij}|$. ■

COROLLARY 1. *The worst-case time complexity of counting typed graphlets using Algorithm 1 matches the worst-case time complexity of the best known untyped graphlet counting algorithm.*

PROOF. From Theorem 2 we have that $O(\Delta(|S_i| + |S_j| + |T_{ij}|))$, which is exactly the time complexity of the best known untyped (homogeneous) graphlet counting algorithm [3, 4]. ■

6.2 Space Complexity

Since our approach generalizes to graphs with an arbitrary number of types L , the specific set of typed motifs is unknown. As demonstrated in Table 2, it is impractical to store the counts of all possible k -node typed motifs for any graph of reasonable size as typically done in traditional methods for untyped graphlets [3, 4, 32].

Despite this being obviously impractical due to the amount of space that would be required, the existing state-of-the-art methods such as GC [20] store counts of all possible typed graphlets, and therefore the space complexity of such methods is:

$$O(MT_{\max}) \quad (32)$$

where $M = |E|$ is the number of edges in G and T_{\max} is the number of different possible typed graphlets with L types. Thus, MT_{\max} is the total space to store M vectors of length T_{\max} (*i.e.*, one T_{\max} -dimensional vector per edge). To understand the above space requirements and how it is impractical for any moderately sized graph, suppose we have a graph with 10,000,000 (M) edges and $L = 7$ types. Counting all 3- and 4-node typed graphlet orbits for every edge would require 90.72 GB of space to store the large MT_{\max} matrix (assuming 4 bytes per count/entry). This is obviously impractical for any graph of even moderate size.

In contrast, the total space used by our approach for storing the typed graphlet counts is:

$$O(M\bar{T}) \ll O(MT_{\max}) \quad (33)$$

where $\bar{T} = \frac{1}{M} \sum_{(ij) \in E} |X_{ij}|$ is the average number of typed graphlets with nonzero counts per edge. Note $|X_{ij}|$ is the number of typed graphlets with nonzero counts for edge $(ij) \in E$. Therefore, the total space is only $O(M\bar{T})$. The space of all other data structures used in Algorithm 1 is small in comparison, *e.g.*, Ψ takes at most $O(|V|)$ space, whereas T_{ij} , S_i , and S_j take $O(\Delta)$ space in the worst-case (by Property 1) and can be reused for every edge. In addition, the size of \mathbf{x} is independent of the graph size ($|V| + |E|$) and can also be reused.

7 EXPERIMENTS

The experiments are designed to investigate the runtime performance (Section 7.1), space-efficiency (Section 7.2), parallelization (Section 7.3), and scalability (Section 7.4) of the proposed approach.

From a computational point of view, any practical approach for this problem must have all four desired properties, *i.e.*, fast, space-efficient, parallel with near-linear scaling, and scalable for large networks. To demonstrate the *effectiveness* of the approach, we use a variety of heterogeneous and labeled/attributed graph data from different application domains. All data can be accessed at NetworkRepository [43]. Unless otherwise mentioned, we use a serial version of the proposed typed graphlet approach for comparison with other methods. In addition, our approach counts the frequency of all typed {2, 3, 4}-node graphlets globally for the entire graph and locally for every edge (*i.e.*, local typed graphlet counting [48]). This is in contrast to the existing methods that focus on counting graphlets for every node as opposed to every edge. As an aside, existing methods are serial (as they are difficult to parallelize effectively) and typically only solve the global or local subgraph counting problem. All methods used for comparison are configured to count only 4-node graphlets.

7.1 Runtime Comparison

We first demonstrate how fast the proposed framework is for deriving typed graphlets by comparing the runtime (in seconds) of our approach against ESU (using fanmod) [58], G-Tries [42], and GC [20]. As an aside, the proposed framework is the only approach that derives typed graphlets for every edge (as opposed to counting typed graphlets for every node as done by the existing methods). Other key differences are summarized in Section 8.

For comparison, we use a wide variety of real-world graphs from different domains as well as graphs generated from four different synthetic graph models. In Table 4, we report the time (in seconds) required by each method. The results reported in Table 4 are from a serial version of our approach. The network statistics and properties of the graphs used for evaluation are also shown in Table 4 (columns 2-6). To be able to compare with the existing methods, we included a variety of very small graphs for which the existing methods could solve in a reasonable amount of time. Note ETL indicates that a method did not terminate within 24 hours for that graph. Strikingly, the existing methods are unable to handle medium to large graphs with hundreds of thousands or more nodes and edges as shown in Table 4. Even small graphs can take hours to finish using existing methods (Table 4). For instance, the small citeseer graph with only 3.3k nodes and 4.5k edges takes 46.27 seconds using the best existing method whereas ours finishes in a tiny fraction of a second, notably, 2/100 seconds. This is about 2,100 times faster than the next best method. Similarly, on the small cora graph with 2.7K nodes and 5.3K edges, GC takes 467 seconds whereas G-Tries takes 351 seconds. However, our approach finishes in only 0.03 seconds. This is 10,000 times faster than the next best method. Unlike existing methods, our approach is able to handle large-scale graphs. On flickr, our approach takes about 2 minutes to count the occurrences of all typed graphlets for all 6.8 million edges. Across all graphs, the proposed method achieves significant speedups over the existing state-of-the-art methods as shown in Table 4. These results demonstrate the effectiveness of our approach for counting typed graphlets in large networks.

Table 5 provides the number of unique typed motifs that occur for each induced subgraph. From these results, we make an important

Table 4: Results comparing the proposed approach to the state-of-the-art methods in terms of runtime performance (seconds). Since existing state-of-the-art methods are unable to handle large or even medium-sized graphs as shown below, we had to include a number of very small graphs (e.g., cora, citeseer, webkb) in order to compare with existing methods. Note $\Delta = \max$ node degree; $|\mathcal{T}_V| = \#$ of node types; $|\mathcal{T}_E| = \#$ of edge types.

| | $ V $ | $ E $ | Δ | $ \mathcal{T}_V $ | $ \mathcal{T}_E $ | SECONDS | | | | SPEEDUP (OURS VS.) | | |
|------------------|--------|--------|----------|-------------------|-------------------|-----------|-----------|--------------|----------------|--------------------|----------|--------------|
| | | | | | | GC [20] | ESU [58] | G-Tries [42] | Ours | GC [20] | ESU [58] | G-Tries [42] |
| citeseer | 3.3k | 4.5k | 99 | 6 | 21 | 46.27 | 5937.75 | 144.08 | 0.022 | 2103x | 269897x | 6549x |
| cora | 2.7k | 5.3k | 168 | 7 | 28 | 467.20 | 10051.07 | 351.40 | 0.032 | 14600x | 314095x | 10981x |
| fb-relationship | 7.3k | 44.9k | 106 | 6 | 20 | 1374.60 | 54,837.69 | 3789.17 | 0.701 | 1960x | 78227x | 5405x |
| web-polblogs | 1.2k | 16.7k | 351 | 2 | 1 | 28,986.70 | 26,577.10 | 1,563.04 | 1.055 | 27475x | 25191x | 1481x |
| ca-DBLP | 2.9k | 11.3k | 69 | 3 | 3 | 149.20 | 1,188.11 | 18.90 | 0.100 | 1492x | 11881x | 189x |
| inf-openflights | 2.9k | 15.7k | 242 | 2 | 2 | 9262.20 | 18,839.36 | 458.01 | 0.578 | 16024x | 32594x | 792x |
| soc-wiki-elec | 7.1k | 100.8k | 1.1k | 2 | 2 | ETL | ETL | 26,468.85 | 5.316 | ∞ | ∞ | 45793x |
| webkb | 262 | 459 | 122 | 5 | 14 | 85.82 | 7,158.10 | 187.22 | 0.006 | 14303x | 1193016x | 31203x |
| terrorRel | 881 | 8.6k | 36 | 2 | 3 | 192.6 | 3130.7 | 241.1 | 0.039 | 4938x | 80274x | 6182x |
| pol-retweet | 18.5k | 48.1k | 786 | 2 | 3 | ETL | ETL | ETL | 0.296 | ∞ | ∞ | ∞ |
| web-spam | 9.1k | 465k | 3.9k | 3 | 6 | ETL | ETL | ETL | 210.97 | ∞ | ∞ | ∞ |
| movielens | 28.1k | 170.4k | 3.6k | 3 | 3 | ETL | ETL | ETL | 5.23 | ∞ | ∞ | ∞ |
| citeulike | 907.8k | 1.4M | 11.2k | 3 | 2 | ETL | ETL | ETL | 126.53 | ∞ | ∞ | ∞ |
| yahoo-msg | 100.1k | 739.8k | 9.4k | 2 | 2 | ETL | ETL | ETL | 35.22 | ∞ | ∞ | ∞ |
| dbpedia | 495.9k | 921.7k | 24.8k | 4 | 3 | ETL | ETL | ETL | 56.02 | ∞ | ∞ | ∞ |
| digg | 217.3k | 477.3k | 219 | 2 | 2 | ETL | ETL | ETL | 5.592 | ∞ | ∞ | ∞ |
| bibsonomy | 638.8k | 1.2M | 211 | 3 | 3 | ETL | ETL | ETL | 3.631 | ∞ | ∞ | ∞ |
| epinions | 658.1k | 2.6M | 775 | 2 | 2 | ETL | ETL | ETL | 85.27 | ∞ | ∞ | ∞ |
| flickr | 2.3M | 6.8M | 216 | 2 | 2 | ETL | ETL | ETL | 120.79 | ∞ | ∞ | ∞ |
| orkut | 6M | 37.4M | 166 | 2 | 2 | ETL | ETL | ETL | 1241.01 | ∞ | ∞ | ∞ |
| ER (10K,0.001) | 10k | 50.1k | 26 | 5 | 15 | 183.32 | 5,399.14 | 241.27 | 0.48 | 381x | 11248x | 502x |
| CL (1.8) | 9.2k | 44.2k | 218 | 5 | 15 | 31,668 | 45,399.14 | 5,241.27 | 1.46 | 21690x | 31095x | 3589x |
| KPGM (log 12,14) | 3.3k | 43.2k | 1.3k | 5 | 15 | ETL | ETL | 63,843.86 | 8.94 | ∞ | ∞ | 7141x |
| SW (10K,6,0.3) | 10k | 30k | 12 | 5 | 15 | 21.48 | 5,062.67 | 206.92 | 0.24 | 89x | 21094x | 862x |

* ETL = Exceeded Time Limit (24 hours / 86,400 seconds)

observation. In real-world graphs we observe that certain typed motifs do not occur at all in the graph. We define such typed motifs that do not occur in G as *forbidden typed motifs* as their appearance in the future would indicate something strong. For instance, perhaps an anomaly or malicious activity. Other interesting insights and applications of typed graphlets are discussed and explored further in Section 7.6.

In addition to the large collection of real-world networks used for evaluation, we also generated synthetic graphs from 4 different graph models including: Erdős-Rényi (ER) model [17], Chung-Lu (CL) graph model [14], Kronecker Product Graph Model (KPGM) [30], and Watts-Strogatz Small-World (SW) graph model [57]. Since these graph models do not generate types/labels, we must assign them. Given L node types, we assign types to nodes uniformly at random such that $\frac{N}{L}$ nodes are assigned to every type. In these experiments, we used $L = 5$ types. Results are provided at the bottom of Table 4 in the last 4 rows. Overall, our proposed method is orders of magnitude faster than all other methods. Compared to G-Tries, our method is between 502x and 7141x faster for these graphs. Compared to GC, our method is between 89x and 21690x faster.

7.2 Space Efficiency Comparison

We theoretically showed the space complexity of our approach in Section 6.2. In this section, we empirically investigate the space-efficiency of our approach compared to ESU (using fanmod) [58], G-Tries [42], and GC [20]. For this, we again use a variety of real-world networks. Table 6 reports the space used by each method for a variety of real-world graphs. Strikingly, the proposed approach uses between 42x and 776x less space than existing methods as shown in Table 6. These results indicate that our approach is space-efficient and practical for large networks.

7.3 Parallel Speedup

This section evaluates the parallel scaling of the parallel algorithm (Section 5). As an aside, this work describes the first parallel typed graphlet algorithm as all existing methods are inherently serial. In these experiments, we used a two processor, Intel Xeon E5-2686 v4 system with 256 GB of memory. None of the experiments came close to using all the memory. Parallel speedup is simply $S_p = \frac{T_1}{T_p}$ where T_1 is the execution time of the sequential algorithm, and T_p is the execution time of the parallel algorithm with p processing units (cores). In Figure 7, we observe nearly linear speedup as we increase the number of cores. These results indicate the effectiveness of the parallel algorithm for counting typed graphlets in general heterogeneous graphs.

Table 5: Comparing the number of unique typed motifs that occur for each induced subgraph (e.g., there are 40 different typed triangles that appear in citeseer where each typed triangle count has a different type/color configuration).

| Network data | • | •• | ••• | Y | □ | △ | □□ | □□□ |
|------------------|----|----|-----|-----|----|-----|----|-----|
| citeseer | 56 | 40 | 124 | 119 | 66 | 98 | 56 | 19 |
| cora | 82 | 49 | 202 | 190 | 76 | 157 | 73 | 19 |
| fb-relationship | 50 | 47 | 112 | 109 | 85 | 106 | 89 | 77 |
| web-polblogs | 4 | 4 | 5 | 5 | 5 | 5 | 5 | 5 |
| ca-DBLP | 10 | 10 | 15 | 15 | 15 | 15 | 15 | 15 |
| inf-openflights | 4 | 4 | 5 | 5 | 5 | 5 | 5 | 5 |
| soc-wiki-elec | 4 | 4 | 5 | 5 | 5 | 5 | 5 | 5 |
| webkb | 31 | 21 | 59 | 59 | 23 | 51 | 32 | 8 |
| terrorRel | 4 | 4 | 5 | 0 | 4 | 5 | 5 | 5 |
| pol-retweet | 4 | 4 | 5 | 5 | 5 | 5 | 5 | 4 |
| web-spam | 10 | 10 | 15 | 15 | 15 | 15 | 15 | 15 |
| movielens | 7 | 1 | 6 | 9 | 6 | 3 | 3 | 0 |
| citeulike | 5 | 0 | 3 | 6 | 3 | 0 | 0 | 0 |
| yahoo-msg | 3 | 2 | 3 | 4 | 3 | 3 | 3 | 2 |
| dbpedia | 8 | 0 | 6 | 10 | 5 | 0 | 0 | 0 |
| digg | 4 | 3 | 4 | 5 | 4 | 4 | 4 | 2 |
| bibsonomy | 7 | 1 | 6 | 9 | 6 | 3 | 3 | 0 |
| epinions | 3 | 2 | 3 | 4 | 3 | 3 | 3 | 2 |
| flickr | 3 | 2 | 3 | 4 | 3 | 3 | 3 | 2 |
| orkut | 4 | 3 | 4 | 4 | 3 | 4 | 3 | 2 |
| ER (10K,0.001) | 35 | 30 | 70 | 70 | 69 | 66 | 1 | 0 |
| CL (1.8) | 35 | 35 | 70 | 70 | 70 | 70 | 70 | 68 |
| KPGM (log 12,14) | 35 | 35 | 70 | 70 | 70 | 70 | 70 | 70 |
| SW (10K,6,0.3) | 35 | 35 | 70 | 70 | 70 | 70 | 70 | 69 |

Table 6: Comparing the space used by our approach to the state-of-the-art methods. Note existing methods derive counts for the nodes, whereas the proposed method derives counts for each edge, and since $|E| > |V|$ then methods that compute counts for nodes should require less space.

| | citeseer | cora | movielens | web-spam |
|--------------|--------------|--------------|---------------|----------------|
| GC [20] | 30.1MB | 50.4MB | ETL | ETL |
| ESU [58] | 13.4MB | 46.2MB | ETL | ETL |
| G-Tries [42] | 161.9MB | 448.6MB | ETL | ETL |
| Ours | 316KB | 578KB | 22.5MB | 128.9MB |

* ETL = Exceeded Time Limit (24 hours / 86,400 seconds)

In addition, Table 7 demonstrates the parallel scalability of the approach as the number of types L increases. Notably, the parallel speedup of the approach remains constant indicating that parallel performance is independent of the number of types L assigned to the nodes in the graph.

7.4 Scalability

To evaluate the scalability of the proposed framework as the size of the graph grows (*i.e.*, number of nodes and edges increase), we generate Erdős-Rényi graphs of increasing size (from 100 to 1 million nodes) such that each graph has an average degree of 10. In Figure 8, we observe that our approach scales linearly as the number of nodes and edges grow large. Furthermore, the scalability of the

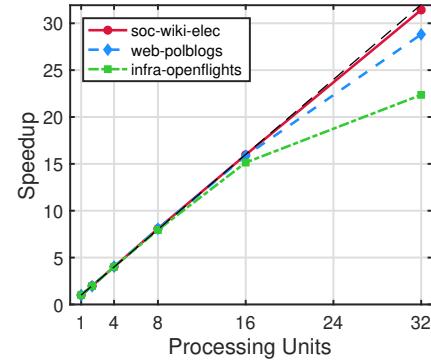


Figure 7: Parallel speedup.

typed graphlet framework is independent of the number of types as shown in Figure 8. As an aside, our approach takes less than 2 minutes to derive all typed $\{2, 3, 4\}$ -node graphlets for a large graph with 1 million nodes and 10 million edges. Note that existing methods are not shown in Figure 8 since they are unable to handle medium to large-sized graphs as shown previous in Table 4.

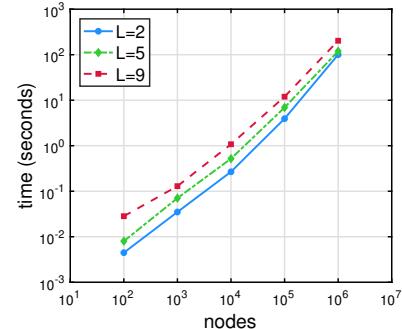


Figure 8: Runtime comparison on Erdős-Rényi graphs with an average degree of 10. Our approach is shown to scale linearly as the size of the graph grows increasingly large.

7.5 Synthetic Graph Experiments

In these experiments, we generate synthetic graphs. For each graph, we vary the number of node types L from 2 to 9, and measure the runtime performance as the number of node types increases as well as the impact in terms of space as L increases. Given $L \in \{2, \dots, 9\}$ node types, we assign types to nodes uniformly at random such that $\frac{N}{L}$ nodes are assigned to every type.

7.5.1 Impact on Performance. We first investigate the runtime performance of our approach as the number of types L increases from 2 to 9. We use both a serial and parallel implementation of our method for comparison. Results are shown in Table 7. Notably, the parallel speedup of the parallel algorithm is constant regardless of L . Therefore, it is not impacted by the increase in L . Furthermore, the runtime of both the serial and parallel algorithm increases slightly as L increases. Notice the additional work depends on the number of unique typed graphlets (the sum of columns 2-9 in Table 7)

and not directly on L itself. The total amount of unique typed graphlets substantially increases as L increases from 2 to 9 as shown in Table 7. This is primarily due to the random assignment of types to nodes. However, in sparse real-world graphs the total unique typed graphlets is typically much smaller as shown in Table 5.

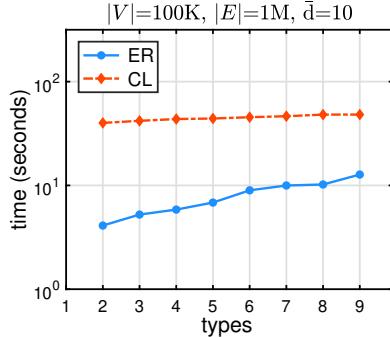


Figure 9: Comparing runtime performance as the number of types increases. In the case of CL graphs with a skewed degree distribution, the runtime is nearly constant as the number of types increases. Note $\bar{d} = \frac{1}{n} \sum d_i$ denotes average node degree.

To further understand how the structure of the graph impacts runtime, we generate an ER and Chung-Lu (CL) graph with 100K nodes, 1M edges, and average degree 10. We vary the number of types L and assign types to nodes uniformly at random as discussed previously. Notice that both the ER and CL graph are generated such that they each have 100K nodes, 1 million edges, with average degree 10. However, both graphs are structurally very different. For instance, the degrees among the nodes in the ER graph are more uniform whereas the degree of the nodes in CL are skewed such that a few nodes have very large degree while the others have relatively small degree. We observe in Figure 9 that for CL graphs with a skewed degree distribution, the runtime of the approach as the number of types increases is essentially constant. This result is important as most real-world graphs also have a skewed degree distribution (social networks, web graphs, information networks, etc.) [18, 19]. However, even in the case where the degrees are more uniform across the nodes, our approach still performs well as shown in Figure 9.

Table 7: Comparing the number of unique typed motifs that occur for each induced subgraph as we vary the number of types L using a KPGM graph with 3.3k nodes, 43.2k edges, average degree = 26, and max degree = 1.3K. Speedup is shown in parenthesis.

| L | time (sec.) | | | | | | | |
|-----|-------------|--------------------|-----|-----|-----|-----|-----|------------------|
| | serial | parallel - 4 cores | | | | | | |
| 2 | 4 | 4 | 5 | 5 | 5 | 5 | 5 | 8.11 (3.89x) |
| 5 | 35 | 35 | 70 | 70 | 70 | 70 | 70 | 8.94 (3.95x) |
| 9 | 165 | 165 | 495 | 495 | 495 | 495 | 495 | 10.37 (3.95x) |

7.5.2 *Impact on Space.* Figure 10 shows the memory (space) required by our approach as the number of types increases from $L \in \{2, \dots, 9\}$. Both ER and CL graphs are shown to have similar space requirements. This is likely due to the random assignment of types to nodes. This assignment represents a type of worst case since every edge is likely to have significantly more distinct typed graphlets compared to sparse real-world graphs. This difference can be seen in Table 5.

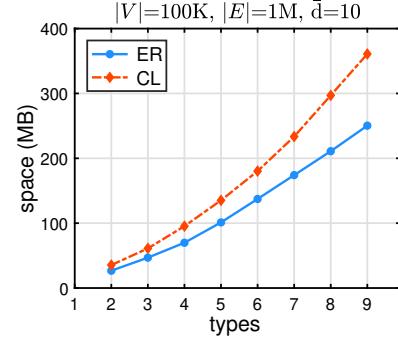


Figure 10: Comparing memory (space) required by our approach as the number of types increases. Note $\bar{d} = \frac{1}{n} \sum d_i$ denotes average node degree.

7.6 Exploratory Analysis

This section demonstrates the use of heterogeneous network motifs for graph mining and exploratory analysis.

7.6.1 *Political retweets.* The political retweets data consists of 18,470 Twitter users. The graph has 61,157 links representing retweets. There are 24,815 triangles in the political retweet network. Triangles in this graph indicate that users retweeted by an individual also retweet each other (*i.e.*, triangle = three users that have all mutually retweeted each other). Triangles may represent users with similar interests. However, triangles alone do not reveal any additional information about the users. To study the (higher-order) structural characteristics of users in this network *w.r.t.* their political orientation, we assign types to nodes based on their political leanings (*i.e.*, left, right). By using the political leanings as the type, we can directly apply the proposed approach to investigate a variety of interesting questions. Interestingly, the 24,815 (untyped) triangles are distributed as follows:

$$\begin{array}{cccc} \text{triangle} & \text{triangle} & \text{triangle} & \text{triangle} \\ \text{---} & \text{---} & \text{---} & \text{---} \\ p = [0.608 & 0.003 & 0.001 & 0.388] \end{array}$$

Notably, we observe that 60.86% and 38.79% of the 24,815 triangles are formed among users with the same political leanings. This implies that three users with the same political leanings are more likely to retweet each other than with users of different political leanings. These results indicate the presence of homophily [33] as users tend to retweet similar others. Furthermore, these homogeneous typed triangles (\triangleleft , \trianglecirc) account for 99.65% of the 24,815 triangles. Intuitively, this implies that the network consists of two tightly-knit communities of users of the same political leanings. The two communities are sparsely connected. Typed triangles obviously

contain significantly more information than untyped triangles. This includes not only information about the local properties but also about the global structure of the network as shown above. Obviously, untyped network motifs are unable to provide such insights as they do not encode the types, attribute values, or class labels associated with a network motif. They only reveal the structural information independent of any important external information associated with the node.

We also investigated typed 4-clique motifs. Strikingly, only 4 of the 5 typed 4-clique motifs that arise from 2 types actually occur in the graph. In particular, the typed 4-clique motif with 2 right users and 2 left users does not even appear in the graph. This typed motif might indicate collusion between individuals from different political parties or some other extremely rare anomalous activity. The other typed 4-cliques that are extremely rare are the typed 4-clique motif with 3 right (left) users and a single left (right) user.

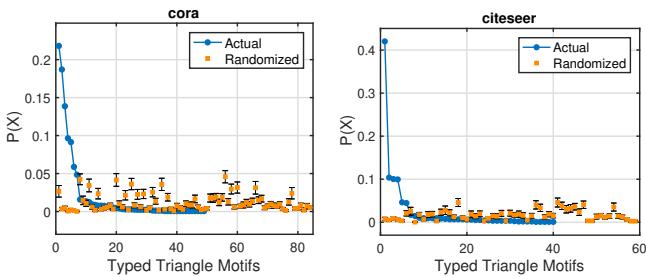


Figure 11: Comparing the actual typed triangle distribution to the randomized typed triangle distribution. We compute 100 random permutations of the node types and run the approach on each permutation then average the resulting counts to obtain the mean randomized typed triangle distribution. There are three key findings. First, we observe a significant difference between the actual and randomized typed triangle distributions. Second, many of the typed triangles that occur when the types are randomized, do not occur in the actual typed triangle distribution. Third, we find the typed triangle distribution to be skewed (approximately power-lawed) as a few typed triangles occur very frequently while the vast majority have very few occurrences.

7.6.2 Cora citation network. The Cora citation network consists of 2708 scientific publications classified into one of seven types (class labels) that indicate the paper topic. The citation network consists of 5429 links. Using the proposed heterogeneous motifs, we find 129 typed 3-node motifs among the 168 possible typed 3-node motifs that could occur. Notably, we observe the most frequent typed triangle motifs are of a single type. Indeed, the first 7 typed triangle motifs with largest frequency in Figure 11 are of a single type. Overall, these 7 typed triangle motifs account for 83.86% of all typed triangle motifs. This finding indicates *strong homophily* among nodes. For instance, the co-authors of a researcher are also typically co-authors, but more strongly, the co-authors are also all typically from the same research area. Homophily has been widely studied in social networks and detecting it remains a challenging problem [29]. Unlike untyped motifs, typed motifs simultaneously

capture the labeling and structural properties that lie at the heart of homophily [33]. Therefore, typed motifs provide a principled foundation for studying homophily in social networks [33]. In Figure 11, we observe a large gap that clearly separates the 7 single-typed triangle motifs from the other typed triangle motifs with heterogeneous types. Furthermore, only 49 out of the 84 possible typed triangle motifs (Table 2) actually occur in the graph. The 25 typed triangle motifs that do not occur are “forbidden heterogeneous motifs” and can provide additional insights into the network and the processes governing the formation with respect to the types (research areas).

Figure 12 investigates a variety of typed 4-node graphlet distributions (from most dense to least dense). Strikingly, only 19 of the 210 possible typed 4-clique graphlets actually occur in G when node types are randomly shuffled. In the case of typed 4-node cycle graphlets, we observe 66 of the actual 210 possible typed 4-node cycle graphlets appear when node types are randomly shuffled.

8 RELATED WORK

While the bulk of work on graphlets (network motifs) have focused on untyped/uncolored graphlets [3, 4, 7, 21, 22, 32, 36, 48, 51, 61], there have been a few recent works for typed graphlets.³ All of these works focus on the problem of counting typed graphlets for *nodes* whereas this paper focuses on the problem of counting typed graphlets for *edges* (or more generally, between a pair of nodes i and j). The earliest such work by Ribeiro *et al.* [42] proposed an approach called G-Tries for finding frequent and statistically significant colored motifs [42, 58]. This problem differs from the one we study in this work that focuses on deriving all such colored graphlets. One recent work by Gu *et al.* [20] used a relaxed definition of colored graphlets for network alignment on very small heterogeneous networks. This work focused mainly on the problem of network alignment for simple heterogeneous networks and not on the approach for deriving typed graphlets. Nevertheless, the method GC used in that work differs from our approach in four significant ways. First, while we leverage combinatorial relationships to derive a number of typed graphlets in $o(1)$ constant time, GC must enumerate all homogeneous graphlets in order to obtain their type/color configuration. Therefore, our approach is significantly faster than GC as that approach requires a lot of extra work to compute the typed graphlets that we can derive in constant time. Second, our approach is significantly more space-efficient and stores only the nonzero counts of the typed graphlets discovered at each edge. Third, we parallelize our approach to handle large sparse networks. As an aside, Gu *et al.* [20] claim the time complexity of their approach is equivalent to counting homogeneous graphlets. This is only true if the homogeneous graphlet algorithm used enumerates all such homogeneous graphlets since this is the only way that step 2 of their approach could be performed, which requires them to obtain the colors of the nodes involved in the k -node graphlet found in step 1. However, there are much faster algorithms for homogeneous graphlets that avoid explicit enumeration of all such graphlets, *e.g.*, [3, 4]. In a similar fashion, our approach avoids enumerating all such graphlets and explicitly obtains the colors

³Typed graphlets are also called colored/heterogeneous graphlets or motifs.

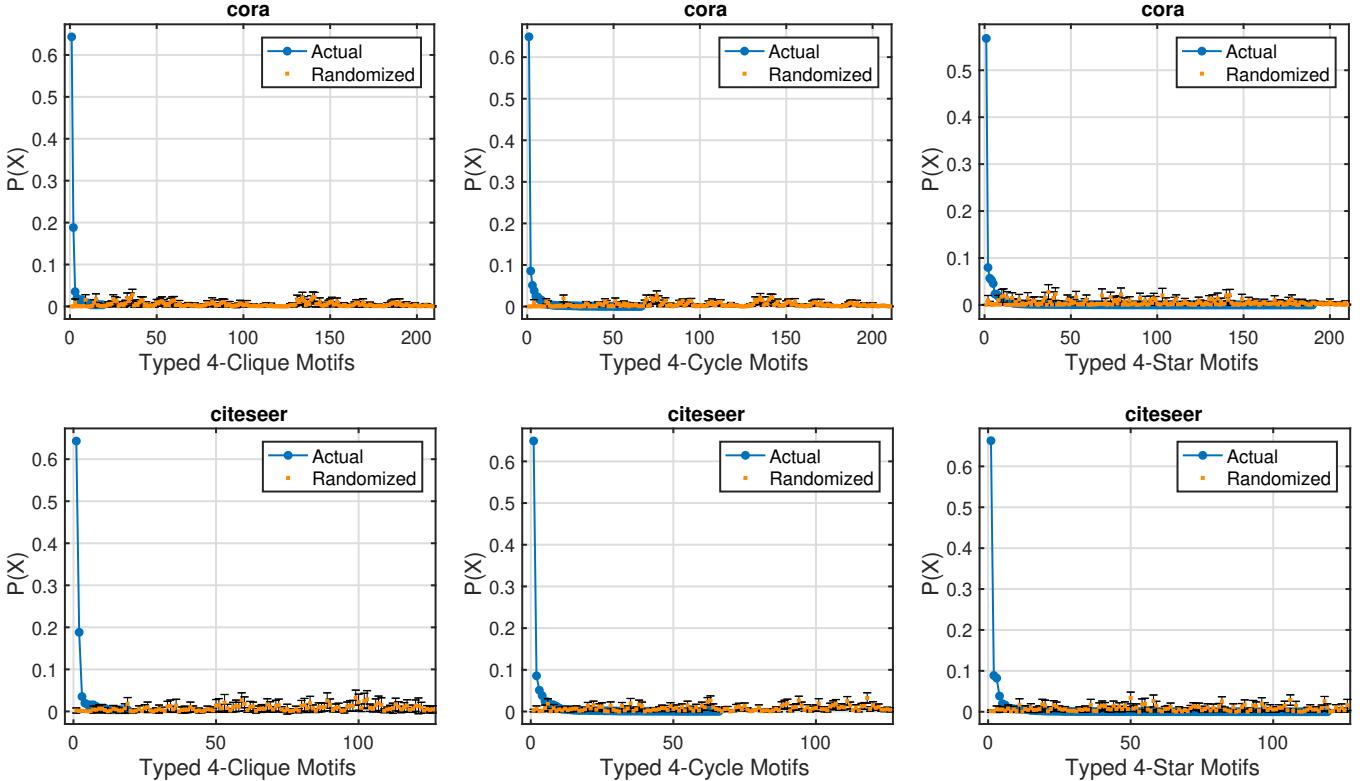


Figure 12: Comparing the actual *typed* 4-clique, 4-cycle and 4-star motif distributions to the randomized typed distributions. We compute 100 random permutations of the node types and run the approach on each permutation then average the resulting counts to obtain the mean randomized typed motif distribution. There are three key findings. First, we observe a significant difference between the actual and randomized typed motif distributions. Second, many of the typed motifs that occur when the types are randomized, do not occur in the actual typed motif distribution. Third, we find the typed motif distribution to be extremely skewed as a few typed motifs occur very frequently while the vast majority have very few occurrences (and many typed network motifs are even forbidden, in the sense that they do not occur at all in the graph).

of the nodes involved in those graphlets (without actually enumerating or knowing the nodes involved in those graphlets), and therefore the computational complexity is actually equivalent to the best known homogeneous graphlet algorithm. Other work by Carranza *et al.* [12] used typed graphlets as a basis for higher-order spectral clustering on heterogeneous networks.

9 CONCLUSION

In this work, we generalized the notion of network motif to heterogeneous networks. We proposed a fast, space-efficient, and parallel framework for counting all k -node typed graphlets. We provide theoretical analysis of combinatorial arguments that capture the relationship between various typed graphlets. Using these combinatorial relationships, the proposed approach is able to derive many typed graphlet counts directly in $o(1)$ constant time by leveraging the counts of a few typed graphlets. Thus, the proposed approach avoids explicit enumeration of any nodes involved in those typed graphlets. For every edge, we count a few typed graphlets and obtain the exact counts of the remaining ones in $o(1)$ constant time. The time complexity of the proposed approaches matches that of

the best untyped graphlet algorithm. Empirically, our approach is shown to outperform the state-of-the-art in terms of runtime, space-efficiency, and scalability as it is able to handle large networks. While existing methods take hours on small graphs with thousands of edges, our typed graphlet counting approach takes only seconds on networks with millions of edges. Finally, unlike other approaches, the proposed approach is able to handle large-scale general heterogeneous networks with an arbitrary number of types and with millions or more edges while lending itself to an efficient and highly scalable (asynchronous & lock-free) parallel implementation.

REFERENCES

- [1] Evrim Acar, Tamara G Kolda, and Daniel M Dunlavy. 2011. All-at-once optimization for coupled matrix and tensor factorizations. *arXiv:1105.3422* (2011).
- [2] Nesreen K. Ahmed, Nick Duffield, Theodore L. Willke, and Ryan A. Rossi. 2017. On Sampling from Massive Graph Streams. In *VLDB*. 1430–1441.
- [3] Nesreen K. Ahmed, Jennifer Neville, Ryan A. Rossi, and Nick Duffield. 2015. Efficient Graphlet Counting for Large Networks. In *ICDM*. 10.
- [4] Nesreen K. Ahmed, Jennifer Neville, Ryan A. Rossi, Nick Duffield, and Theodore L. Willke. 2016. Graphlet Decomposition: Framework, Algorithms, and Applications. *KAIS* (2016), 1–32.

Heterogeneous Network Motifs

- [5] Nesreen K. Ahmed, Ryan A. Rossi, Theodore L. Willke, and Rong Zhou. 2017. A Higher-order Latent Space Network Model. In *Proceedings of the AAAI PAIR (Plan, Activity, and Intent Recognition) Workshop*. 1–7.
- [6] Nesreen K. Ahmed, Ryan A. Rossi, Rong Zhou, John Boaz Lee, Xiangnan Kong, Theodore L. Willke, and Hoda Eldardiry. 2018. Learning Role-based Graph Embeddings. In *StarAI IJCAI*.
- [7] Nesreen K. Ahmed, Theodore L. Willke, and Ryan A. Rossi. 2016. Estimation of Local Subgraph Counts. In *Proceedings of the IEEE International Conference on BigData*. 586–595.
- [8] Leman Akoglu, Hanghang Tong, and Danai Koutra. 2015. Graph based anomaly detection and description: a survey. *DMKD* 29, 3 (2015), 626–688.
- [9] Danielle Smith Bassett and ED Bullmore. 2006. Small-world brain networks. *The neuroscientist* 12, 6 (2006), 512–523.
- [10] Austin R Benson, David F Gleich, and Jure Leskovec. 2016. Higher-order organization of complex networks. *Science* 353, 6295 (2016), 163–166.
- [11] Ed Bullmore and Olaf Sporns. 2009. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience* 10, 3 (2009), 186–198.
- [12] Aldo G. Carranza, Ryan A. Rossi, Anup Rao, and Eunyee Koh. 2018. Higher-order Spectral Clustering for Heterogeneous Graphs. In *arXiv:1810.02959*. 15.
- [13] Lianhua Chi and Xingquan Zhu. 2017. Hashing techniques: A survey and taxonomy. *ACM Computing Surveys (CSUR)* 50, 1 (2017), 11.
- [14] F. Chung and L. Lu. 2002. Connected components in random graphs with given expected degree sequences. *Annals of combinatorics* 6, 2 (2002), 125–145.
- [15] Joseph Crawford and Tijana Milenković. 2015. GREAT: GRaphlet Edge-based network AlignmenT. In *BIBM*. 220–227.
- [16] N. Eagle and A. Pentland. 2006. Reality mining: sensing complex social systems. *Personal and Ubiquitous Computing* 10, 4 (2006), 255–268.
- [17] Paul Erdős and A Rényi. 1960. On the evolution of random graphs. *Publ. Math. Inst. Hungar. Acad. Sci* 5 (1960), 17–61.
- [18] M. Faloutsos, P. Faloutsos, and C. Faloutsos. 1999. On power-law relationships of the internet topology. In *Proceedings of the conference on Applications, technologies, architectures, and protocols for computer communication*. ACM, 262.
- [19] M Girvan and M E J Newman. 2002. Community structure in social and biological networks. *PNAS* 99, 12 (2002), 7821–7826.
- [20] Shaw Gu, John Johnson, Fazole E Faisal, and Tijana Milenković. 2018. From homogeneous to heterogeneous network alignment via colored graphlets. *Scientific reports* 8, 1 (2018), 12524.
- [21] Wayne Hayes, Kai Sun, and Nataša Pržulj. 2013. Graphlet-based measures are suitable for biological network comparison. *Bioinformatics* 29, 4 (2013), 483–491.
- [22] Tomaž Hočevar and Janez Demšar. 2014. A combinatorial approach to graphlet counting. *Bioinformatics* 30, 4 (2014), 559–565.
- [23] Yury Hulovatyy, Huili Chen, and T Milenković. 2015. Exploring the structure and function of temporal networks with dynamic graphlets. *Bioinformatics* 31, 12 (2015), i171–i180.
- [24] Chia-Chen Hung, Hope Chan, and Eric Hsiao-Kuang Wu. 2008. Mobility pattern aware routing for heterogeneous vehicular networks. In *WCNC*. 2200–2205.
- [25] Xiangnan Kong, Jiawei Zhang, and Philip S Yu. 2013. Inferring anchor links across multiple heterogeneous social networks. In *CIKM*. 179–188.
- [26] Lauri Kovanen, Márton Karsai, Kimmo Kaski, János Kertész, and Jari Saramäki. 2011. Temporal motifs in time-dependent networks. *Journal of Statistical Mechanics: Theory and Experiment* 2011, 11 (2011), P11005.
- [27] Mehmet Koyutürk, Yohan Kim, Umut Topkara, Shankar Subramaniam, Wojciech Szpankowski, and Ananth Grama. 2006. Pairwise alignment of protein interaction networks. *JCB* 13, 2 (2006), 182–199.
- [28] Solomon Kullback. 1997. *Information theory and statistics*. Courier Corporation.
- [29] Timothy La Fond and Jennifer Neville. 2010. Randomization Tests for Distinguishing Social Influence and Homophily Effects. In *WWW (WWW '10)*. ACM, New York, NY, USA, 601–610.
- [30] Jure Leskovec, Deepayan Chakrabarti, Jon Kleinberg, Christos Faloutsos, and Zoubin Ghahramani. 2010. Kronecker graphs: An approach to modeling networks. *JMLR* 11, Feb (2010), 985–1042.
- [31] Ryan N Lichtenwalter and Nitesh V Chawla. 2012. Vertex Collocation Profiles: Subgraph Counting for Link Analysis and Prediction. In *WWW*. 1019–1028.
- [32] Dror Marcus and Yuval Shavitt. 2012. RAGE—a rapid graphlet enumerator for large networks. *Computer Networks* 56, 2 (2012), 810–819.
- [33] Miller McPherson, Lynn Smith-Lovin, and James M Cook. 2001. Birds of a Feather: Homophily in Social Networks. *Ann. Rev. of Soc.* 27, 1 (2001).
- [34] Tijana Milenković and Nataša Pržulj. 2008. Uncovering Biological Network Function via Graphlet Degree Signatures. *Cancer Informatics* 6 (2008), 257.
- [35] R Milo, S Shen-Orr, S Itzkovitz, N Kashtan, D Chklovskii, and U Alon. 2002. Network Motifs: Simple Building Blocks of Complex Networks. *Science* 298, 5594 (2002), 824–827.
- [36] Ron Milo, Shai Shen-Orr, Shalev Itzkovitz, Nadav Kashtan, Dmitri Chklovskii, and Uri Alon. 2002. Network motifs: simple building blocks of complex networks. *Science* 298, 5594 (2002), 824–827.
- [37] C.C. Noble and D.J. Cook. 2003. Graph-based anomaly detection. In *SIGKDD*. 631–636.
- [38] G. Palla, I. Derényi, I. Farkas, and T. Vicsek. 2005. Uncovering the overlapping community structure of complex networks in nature and society. *Nature* 435, 7043 (2005), 814–818.
- [39] Nataša Pržulj. 2007. Biological network comparison using graphlet degree distribution. *Bioinfo.* 23, 2 (2007), e177–e183.
- [40] N Pržulj, Derek G Corneil, and Igor Jurisica. 2004. Modeling interactome: scale-free or geometric? *Bioinformatics* 20, 18 (2004), 3508–3515.
- [41] F. Radicchi, C. Castellano, F. Cecconi, V. Loreto, and D. Parisi. 2004. Defining and identifying communities in networks. *PNAS* 101, 9 (2004), 2658–2663.
- [42] Pedro Ribeiro and Fernando Silva. 2014. Discovering colored network motifs. In *Complex Networks V*. Springer, 107–118.
- [43] Ryan A. Rossi and Nesreen K. Ahmed. 2015. The Network Data Repository with Interactive Graph Analytics and Visualization. In *AAAI*. 4292–4293. <http://networkrepository.com>
- [44] Ryan A. Rossi, Nesreen K. Ahmed, Eunyee Koh, Sungchul Kim, Anup Rao, and Yasin Abbasi-Yadkori. 2018. HONE: Higher-Order Network Embeddings. *arXiv:1801.09303* (2018).
- [45] Ryan A. Rossi, Sonia Fahmy, and Nilotpal Talukder. 2013. A Multi-Level Approach for Evaluating Internet Topology Generators. In *Proceedings of IFIP Networking*. 1–9.
- [46] Ryan A. Rossi and Rong Zhou. 2016. Parallel Collective Factorization for Modeling Large Heterogeneous Networks. In *SNAM*. 30.
- [47] Ryan A. Rossi and Rong Zhou. 2018. GraphZIP: A Clique-based Sparse Graph Compression Method. *Journal of Big Data* 5, 1 (2018), 14.
- [48] Ryan A. Rossi, Rong Zhou, and Nesreen K. Ahmed. 2018. Estimation of Graphlet Counts in Massive Networks. In *IEEE Transactions on Neural Networks and Learning Systems (TNNSL)*. 1–14.
- [49] Frank Schweitzer, Giorgio Fagiolo, Didier Sornette, Fernando Vega-Redondo, Alessandro Vespignani, and Douglas R White. 2009. Economic networks: The new challenges. *science* 325, 5939 (2009), 422–425.
- [50] Nino Shervashidze, Tobias Petri, Kurt Mehlhorn, Karsten M Borgwardt, and Svn Vishwanathan. 2009. Efficient graphlet kernels for large graph comparison. In *AISTATS*.
- [51] Nino Shervashidze, Tobias Petri, Kurt Mehlhorn, Karsten M Borgwardt, and Svn Vishwanathan. 2009. Efficient graphlet kernels for large graph comparison. In *International conference on artificial intelligence and statistics*. 488–495.
- [52] Nino Shervashidze, Pascal Schweitzer, Erik Jan van Leeuwen, Kurt Mehlhorn, and Karsten M Borgwardt. 2011. Weisfeiler-lehman graph kernels. *JMLR* 12, Sep (2011), 2539–2561.
- [53] Ryan W Solava, Ryan P Michaels, and T Milenković. 2012. Graphlet-based edge clustering reveals pathogen-interacting proteins. *Bioinformatics* 28, 18 (2012), i480–i486.
- [54] Yizhou Sun, Jiawei Han, Xifeng Yan, Philip S Yu, and Tianyi Wu. 2011. Pathsim: Meta path-based top-k similarity search in heterogeneous information networks. *VLDB* 4, 11 (2011), 992–1003.
- [55] S Vichy N Vishwanathan, Nicol S Schraudolph, Risi Kondor, and Karsten M Borgwardt. 2010. Graph kernels. *JMLR* 11 (2010), 1201–1242.
- [56] Jian-Wei Wang and Li-Li Rong. 2009. Cascade-based attack vulnerability on the US power grid. *Safety Science* 47, 10 (2009), 1332–1336.
- [57] D.J. Watts and S.H. Strogatz. 1998. Collective dynamics of small-world networks. *Nature* 393, 6684 (1998), 440–442.
- [58] Sebastian Wernicke and Florian Rasche. 2006. FANMOD: a tool for fast network motif detection. *Bioinformatics* 22, 9 (2006), 1152–1153.
- [59] Zhijun Yin, Rui Li, Qiaozhu Mei, and Jiawei Han. 2009. Exploring social tagging graph for web object classification. In *SIGKDD*. ACM, 957–966.
- [60] Xiao Yu, Xiang Ren, Yizhou Sun, Quanquan Gu, Bradley Sturt, Urvashi Khandelwal, Brandon Norick, and Jiawei Han. 2014. Personalized entity recommendation: A heterogeneous information network approach. In *WSDM*. 283–292.
- [61] Luming Zhang, Yahong Han, Yi Yang, Mingli Song, Shuicheng Yan, and Qi Tian. 2013. Discovering discriminative graphlets for aerial image categories recognition. *IEEE Transactions on Image Processing* 22, 12 (2013), 5071–5084.
- [62] L. Zhang, R. Hong, Y. Gao, R. Ji, Q. Dai, and X. Li. 2016. Image Categorization by Learning a Propagated Graphlet Path. *TNNLS* 27, 3 (2016), 674–685.
- [63] Luming Zhang, Mingli Song, Zicheng Liu, Xiao Liu, Jiajun Bu, and Chun Chen. 2013. Probabilistic graphlet cut: Exploiting spatial structure cue for weakly supervised image segmentation. In *CVPR*.