

Motifs in Temporal Networks

Ashwin Paranjape*
Stanford University
ashwinpp@stanford.edu

Austin R. Benson*
Stanford University
arbenson@stanford.edu

Jure Leskovec
Stanford University
jure@cs.stanford.edu

ABSTRACT

Networks are a fundamental tool for modeling complex systems in a variety of domains including social and communication networks as well as biology and neuroscience. Small subgraph patterns in networks, called network motifs, are crucial to understanding the structure and function of these systems. However, the role of network motifs in temporal networks, which contain many timestamped links between the nodes, is not yet well understood.

Here we develop a notion of a temporal network motif as an elementary unit of temporal networks and provide a general methodology for counting such motifs. We define temporal network motifs as induced subgraphs on sequences of temporal edges, design fast algorithms for counting temporal motifs, and prove their runtime complexity. Our fast algorithms achieve up to 56.5x speedup compared to a baseline method. Furthermore, we use our algorithms to count temporal motifs in a variety of networks. Results show that networks from different domains have significantly different motif counts, whereas networks from the same domain tend to have similar motif counts. We also find that different motifs occur at different time scales, which provides further insights into structure and function of temporal networks.

1. INTRODUCTION

Networks provide an abstraction for studying complex systems in a broad set of disciplines, ranging from social and communication networks to molecular biology and neuroscience [20]. Typically, these systems are modeled as static graphs that describe relationships between objects (nodes) and links between the objects (edges). However, many systems are not static as the links between objects dynamically change over time [8]. Such *temporal networks* can be represented by a series of timestamped edges, or *temporal edges*. For example, a network of email or instant message communication can be represented as a sequence of timestamped directed edges, one for every message that is sent from one person to another. Similar representations can be used to model computer networks, phone calls, financial transactions, and biological signaling networks.

*These authors contributed equally to this work.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WSDM '17, February 6–10, 2017, Cambridge, United Kingdom.

© 2017 ACM. ISBN 978-1-4503-4675-7/17/02...\$15.00

DOI: <http://dx.doi.org/10.1145/3018661.3018731>

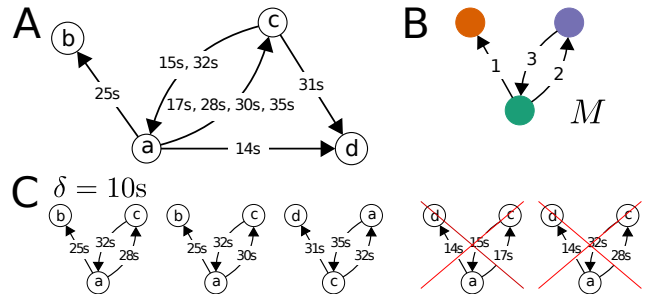


Figure 1: Temporal graphs and δ -temporal motifs. A: A temporal graph with nine temporal edges. Each edge has a timestamp (listed here in seconds). B: Example 3-node, 3-edge δ -temporal motif M . The edge labels correspond to the ordering of the edges. C: Instances of the δ -temporal motif M in the graph for $\delta = 10$ seconds. The crossed-out patterns are not instances of M because either the edge sequence is out of order or the edges do not all occur within the time window δ .

While such temporal networks are ubiquitous, there are few tools for modeling and characterizing the underlying structure of such dynamical systems. Existing methods either model the networks as strictly growing where a pair of nodes connect once and stay connected forever [2, 10, 17] or aggregate temporal information into a sequence of snapshots [1, 6, 23]. These techniques fail to fully capture the richness of the temporal information in the data.

Characterizing temporal networks also brings a number of interesting challenges that distinguish it from the analysis of static networks. For example, while the number of nodes and pairs of connected nodes can be of manageable size, the number of temporal edges may be very large and thus efficient algorithms are needed when analyzing such data. Another interesting challenge is that patterns in temporal networks can occur at different time scales. For example, in telephone call networks, reciprocation (that is, a person returning a call) can occur on very short time intervals, while more intricate patterns (e.g., person A calling person B , who then calls C) may occur at larger time scales. Lastly, there are many possible temporal patterns as the order as well as the sequence of edges play an important role.

Present work: Temporal network motifs. Here, we provide a general methodology for analyzing temporal networks. We define temporal networks as a set of nodes and a collection of directed temporal edges, where each edge has a timestamp. For example, Fig. 1A illustrates a small temporal network with nine temporal edges between five ordered pairs of nodes.

Our analytical approach is based on generalizing the notion of network motifs to temporal networks. In static networks, network motifs or graphlets are defined as small induced subgraphs occurring in a bigger network structure [4, 19, 29]. We extend static mo-

tifs to temporal networks and define δ -temporal motifs, where all the edges in a given motif M have to occur inside the time period of δ time units. These δ -temporal motifs simultaneously account for ordering of edges and a temporal window in which edges can occur. For example, Fig. 1B shows a motif on three nodes and three edges, where the edge label denotes the order in which the edges appear. While we focus on directed edges with a single timestamp in this work, our methodology seamlessly generalizes to common variations on this model. For example, our methods can incorporate timestamps with durations (common in telephone call networks), colored edges that identify different types of connections, and temporal networks with undirected edges.

We then consider the problem of counting how many times does each δ -temporal motif occur in a given temporal network. We develop a general algorithm for counting temporal network motifs defined by any number of nodes and edges that avoids enumeration over subsets of temporal edges and whose complexity depends on the structure of the static graph induced by the temporal motif. For motifs defined by a constant number of temporal edges between 2 nodes, this general algorithm is optimal up to constant factors—it runs in $O(m)$ time, where m is the number of temporal edges.

Furthermore, we design fast variations of the algorithm that allow for counting certain classes of δ -temporal motifs including star and triangle patterns. These algorithms are based on a common framework for managing summary counts in specified time windows. For star motifs with 3 nodes and 3 temporal edges, we again achieve a running time linear in the input, i.e., $O(m)$ time. Given a temporal graph with τ induced triangles in its induced static graph, our fast algorithm counts temporal triangle motifs with 3 temporal edges in $O(\tau^{1/2}m)$ worst-case time. In contrast, any algorithm that processes triangles individually takes $O(\tau m)$ worst-case time. In practice, our fast temporal triangle counting algorithm is up to 56 times faster than a competitive baseline and runs in just a couple of hours on a network with over two billion temporal edges.

Our algorithmic framework enables us to study the structure of several complex systems. For example, we explore the differences in human communication patterns by analyzing motif frequencies in text message, Facebook wall post, email and private online message network datasets. Temporal network motif counts reveal that text messaging and Facebook wall posting are dominated by “blocking” communication, where a user only engages with one other user at a time, whereas email is mostly characterized by “non-blocking” communication as individuals send out several emails in a row. Furthermore, private online messaging contains a mixture of blocking and non-blocking behavior.

Temporal network motifs can also be used to measure the frequency of patterns at different time scales. For example, the difference in δ -temporal motif counts for $\delta = 60$ minutes and $\delta = 30$ minutes counts only the motifs that take at least 30 minutes and at most 60 minutes to form. With this type of analysis, we find that certain question-and-answer patterns on Stack Overflow need at least 30 minutes to develop. We also see that in online private messaging, star patterns constructed by outgoing messages sent by one user tend to increase in frequency from time scales of 1 to 20 minutes before peaking and then declining in frequency.

All in all, our work defines a flexible notion of motifs in temporal networks and provides efficient algorithms for counting them. It enables new analyses in a variety of scientific domains and paves a new way for modeling dynamic complex systems.

2. RELATED WORK

Our work builds upon the rich literature on network motifs in static graphs, where these models have proved crucial to under-

standing the mechanisms driving complex systems [19] and to characterizing classes of static networks [25, 29]. Furthermore, motifs are critical for understanding the higher-order organizational patterns in networks [3, 4]. On the algorithmic side, a large amount of research has been devoted simply to counting triangles in undirected static graphs [14].

Prior definitions of temporal network motifs either do not account for edge ordering [30], only have heuristic counting algorithms [7], or assume temporal edges in a motif must be consecutive events for a node [13]. In the last case, the restrictive definition permits fast counting algorithms but misses important structures. For example, many related edges occurring in a short burst at a node would not be counted together. In contrast, δ -temporal motifs capture *every* occasion that edges form a particular pattern within the prescribed time window.

There are several studies on pattern formation in growing networks where one only considers the addition of edges to a static graph over time. In this context, motif-like patterns have been used to create evolution rules that govern the ways that networks develop [5, 24]. The way we consider ordering of temporal edges in our definition of δ -temporal motifs is similar in spirit. There are also several analyses on the formation of triangles in a variety of social networks [9, 12, 15]. In contrast, in the temporal graphs we study here, three nodes may form a triangle several times.

3. PRELIMINARIES

We now provide formal definitions of temporal graphs and δ -temporal motifs. In Section 4, we provide algorithms for counting the number of δ -temporal motifs in a given temporal graph.

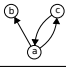
Temporal edges and graphs. We define a *temporal edge* to be a timestamped directed edge between an ordered pair of nodes. We call a collection of temporal edges a *temporal graph* (Fig. 1A). Formally, a temporal graph T on a node set V is a collection of tuples (u_i, v_i, t_i) , $i = 1, \dots, m$, where each u_i and v_i are elements of V and each t_i is a timestamp in \mathbb{R} . We refer to a specific (u_i, v_i, t_i) tuple as a *temporal edge*. There can be many temporal edges directed from u to v , and we refer to them as *edges between u and v* . We assume that the timestamps t_i are unique so that the tuples may be strictly ordered. This assumption makes the presentation of the definitions and algorithms clearer, but our methods can easily be adapted to the case when timestamps are not unique. When it is clear from context, we refer to a temporal edge as simply an *edge*. Finally, by ignoring timestamps and duplicate edges, the temporal graph induces a standard directed graph, which we call the *static graph* G of T with *static edges*, i.e., (u, v) is an edge in G if and only if there is some temporal edge (u, v, t) in T .

δ -temporal motifs and motif instances. We formalize δ -temporal motifs with the following definition.

Definition. A k -node, l -edge, δ -temporal motif is a sequence of l edges, $M = (u_1, v_1, t_1), (u_2, v_2, t_2) \dots, (u_l, v_l, t_l)$ that are time-ordered within a δ duration, i.e., $t_1 < t_2 < \dots < t_l$ and $t_l - t_1 \leq \delta$, such that the induced static graph from the edges is connected and has k nodes.

Note that with this definition, many edges between the same pair of nodes may occur in the motif M . Also, we note that the purpose of the timestamps is to induce an ordering on the edges. Fig. 1B illustrates a particular 3-node, 3-edge δ -temporal motif.

The above definition provides a template for a particular pattern, and we are interested in how many times a given pattern occurs in a dataset. Intuitively, a collection of edges in a given temporal graph is an instance of a δ -temporal motif M if it matches the same edge pattern and all of the edges occur in the right order within the



| | 15s (c, a) | 17s (a, c) | 25s (a, b) | 28s (a, c) | 30s (a, c) | 32s (c, a) | 35s (a, c) |
|----------------------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|
| counts[(a, b)] | 0 | 0 | 1 | 1 | 1 | 1 | 1 |
| counts[(a, c)] | 0 | 1 | 1 | 1 | 2 | 2 | 3 |
| counts[(c, a)] | 1 | 1 | 1 | 0 | 0 | 1 | 1 |
| counts[(a, b)(a, c)] | 0 | 0 | 0 | 1 | 2 | 2 | 3 |
| counts[(c, a)(a, c)] | 0 | 1 | 1 | 0 | 0 | 0 | 1 |
| counts[(a, b)(a, c)(c, a)] | 0 | 0 | 0 | 0 | 0 | 2 | 2 |
| start | 1 | 1 | 1 | 3 | 3 | 3 | 3 |
| end | 1 | 2 | 3 | 4 | 5 | 6 | 7 |

Figure 2: Example execution of Alg. 1 for counting instances of the δ -temporal motif M in Fig. 1. Each column shows the value of counters at the end of the for loop that processes temporal edges. Color indicates change in the variable: incremented (blue), decremented (red), incremented and decremented (purple), or no change (black). At the end of execution, counts[(a, b)(a, c)(c, a)] = 2 for the two instances of the temporal motif M with center node a . Here we only show the counters needed to count M ; in total, Alg. 1 maintains 39 total counters for this input edge sequence, 25 of which are non-zero.

δ time window (Fig. 1C). Formally, we say that any time-ordered sequence $S = (w_1, x_1, t'_1), \dots, (w_l, x_l, t'_l)$ of l unique edges is an instance of the motif $M = (u_1, v_1, t_1), \dots, (u_l, v_l, t_l)$ if

1. There exists a bijection f on the vertices such that $f(w_i) = u_i$ and $f(x_i) = v_i, i = 1, \dots, l$, and
2. the edges all occur within δ time, i.e., $t'_l - t'_1 \leq \delta$

A central goal of this work is to count the number of ordered subsets of edges from a temporal graph T that are instances of a particular motif. In other words, given a k -node, l -edge δ -temporal motif, we seek to find how many of the $l!(\binom{m}{l})$ ordered length- l sequences of edges in the temporal graph T are instances of the motif. A naive approach to this problem would be to simply enumerate all ordered subsets and then check if it is an instance of the motif. In modern datasets, the number of edges m is typically quite large (we analyze a dataset in Section 5 with over two billion edges), and this approach is impractical even for $l = 2$. In the following section, we discuss several faster algorithms for counting the number of instances of δ -temporal motifs in a temporal graph.

4. ALGORITHMS

We now present several algorithms for exactly counting the number of instances of δ -temporal motifs in a temporal graph. We first present a general counting algorithm in Section 4.1, which can count instances of any k -node, l -edge temporal motif faster than simply enumerating over all size- l ordered subsets of edges. This algorithm is optimal for counting 2-node temporal motifs in the sense that it is linear in the number of edges in the temporal graph. In Section 4.2, we provide faster, specialized algorithms for counting specific types of 3-node, 3-edge temporal motifs (Fig. 3).

4.1 General counting framework

We begin with a general framework for counting the number of instances of a k -node, l -edge temporal motif M . To start, consider H to be the static directed graph induced by the edges of M . A sequence of temporal edges S is an instance of M if and only if the static subgraph induced by edges in S is isomorphic to H , the ordering of the edges in S matches the order in M , and all the edges in S span a time window of at most δ time units. This leads to the following general algorithm for counting instances of M in a temporal graph T :

Algorithm 1: Algorithm for counting the number of instances of all possible l -edge δ -temporal motifs in an ordered sequence of temporal edges. We assume the keys of counts[.] are accessed in order of length.

Input: Sequence S' of edges $(e_1 = (u_1, v_1), t_1), \dots, (e_L, t_L)$ with $t_1 < \dots < t_L$, time window δ

Output: Number of instances of each l -edge δ -temporal motif M contained in the sequence

start $\leftarrow 1$, counts \leftarrow Counter(default = 0)

for end = 1, ..., L :

while $t_{\text{start}} + \delta < t_{\text{end}}$ **do**
| $\text{DecrementCounts}(e_{\text{start}})$, start += 1
| $\text{IncrementCounts}(e_{\text{end}})$

return counts

Procedure $\text{DecrementCounts}(e)$

counts[e] -= 1
for suffix **in** counts.keys of length $< l - 1$:
| counts[concat(e , suffix)] -= counts[suffix]

Procedure $\text{IncrementCounts}(e)$

for prefix **in** counts.keys.reverse() of length $< l$:
| counts[concat(prefix, e)] += counts[prefix]
 counts[e] += 1

1. Identify all instances H' of the static motif H induced by M within the static graph G induced by the temporal graph T (e.g., there are three instances of H induced by M in Fig. 1).
2. For each static motif instance H' , gather all temporal edges between pairs of nodes forming an edge in H' into an ordered sequence $S' = (u_1, v_1, t_1), \dots, (u_L, v_L, t_L)$.
3. Count the number of (potentially non-contiguous) subsequences of edges in S' occurring within δ time units that correspond to instances of M .

The first step can use known algorithms for enumerating motifs in static graphs [27], and the second step is a simple matter of fetching the appropriate temporal edges. To perform the third step efficiently, we develop a dynamic programming approach for counting the number of subsequences (instances of motif M) that match a particular pattern within a larger sequence (S'). The key idea is that, as we stream through an input sequence of edges, the count of a given length- l pattern (i.e., motif) with a given final edge is computed from the current count of the length- $(l - 1)$ prefix of the pattern. Inductively, we maintain auxiliary counters of all of the prefixes of the pattern (motif). Second, we also require that all edges in the motif be at most δ time apart. Thus, we use the notion of a moving time window such that any two edges in the time window are at most δ time apart. The auxiliary counters now keep track of only the subsequences occurring within the current time window. Last, it is important to note that the algorithm only counts the number of instances of motifs rather than enumerating them.

Alg. 1 counts all possible l -edge motifs that occur in a given sequence of edges. The data structure counts[.] maintains auxiliary counts of all (ordered) patterns of length at most l . Specifically, counts[$e_1 \dots e_r$] is the number of times the subsequence [$e_1 \dots e_r$] occurs in the current time window (if $r < l$) or the number of times the subsequence has occurred within all time windows of length δ (if $r = l$). We also assume the keys of counts[.] are accessed in order of length. Moving the time window forward by adding a new edge into the window, all edges ($e = (u, v), t$) farther than δ time from the new edge are removed from the window and the appropriate counts are decremented (the $\text{DecrementCounts}()$ method). First,

the single edge counts ($|e|$) are updated. Based on these updates, length-2 subsequences formed with e as its first edge are updated and so on, up through length- $(l - 1)$ subsequences. On the other hand, when an edge e is added to the window, similar updates take place, but in reverse order, from longest to shortest subsequences, in order to increment counts in subsequences where e is the last edge (the *IncrementCounts()* method). Importantly, length- l subsequence counts are incremented in this step but never decremented. As the time window moves from the beginning to the end of the sequence of edges, the algorithm accumulates counts of all length- l subsequences in all possible time windows of length δ .

Fig. 2 shows the execution of the Alg. 1 for a particular sequence of edges. Note that the figure only displays values of counts $[\cdot]$ for contiguous subsequences of the motif M , but the algorithm keeps counts for other subsequences as well. In general, there are $O(l^2)$ contiguous subsequences of an l -edge motif M , and there are $O(|H|^l)$ total keys in counts $[\cdot]$, where $|H|$ is the number of edges in the static subgraph H induced by M , in order to count all l -edge motifs in the sequence (i.e., not just motif M).

We now analyze the complexity of the overall 3-step algorithm. We assume that the temporal graph T has edges sorted by timestamps, which is reasonable if edges are logged in their order of occurrence, and we pre-process T in linear time such that we can access the sorted list of all edges between u and v in $O(1)$ time. Constructing the time-sorted sequence S' in step 2 of the algorithm then takes $O(\log(|H|)|S'|)$ time. Each edge inputted to Alg. 1 is processed exactly twice: once to increment counts when it enters the time window and once to decrement counts when it exits the time window. As presented in Alg. 1, each update changes $O(|H|^l)$ counters resulting in an overall complexity of $O(|H|^l|S'|)$. However, one could modify Alg. 1 to only update counts for contiguous subsequences of the sequence M , which would change $O(l^2)$ counters and have overall complexity $O(l^2|S'|)$. We are typically only interested in small constant values of $|H|$ and l (for our experiments in Section 5, $|H| \leq 3$ and $l = 3$), in which case the running time is linear in the size of the input to the algorithm, i.e., $O(|S'|)$.

In the remainder of this section we analyze our 3-step algorithm with respect to different types of motifs (2-node, stars, and triangles) and argue benefits as well as deficiencies of the proposed framework. We show that for 2-node motifs, our general counting framework takes time linear in the total number of edges m . Since all the input data needs to be examined for computing exact counts, this means the algorithm is optimal for 2-node motifs. However, we also show that for star and triangle motifs the algorithm is not optimal, which then motivates us to design faster algorithms in Sec. 4.2.

General algorithm for 2-node motifs. We first show how to map 2-node motifs to the framework described above. Any induced graph H of a 2-node δ -temporal motif is either a single or a bidirectional edge. In either case, it is straightforward to enumerate over all instances of H in the static graph. This leads to the following procedure: (1) for each pair of nodes u and v for which there is at least one edge, gather and sort the edges in either direction between u and v ; (2) call Alg. 1 with these edges. The obtain the total motif count the counts from each call to Alg. 1 are then summed together.

We only need to input each edge to Alg. 1 once, and under the assumption that we can access the sorted directed edges from one node to another in $O(1)$ time, the merging of edges into sorted order takes linear time. Therefore, the total running time is $O(2^l m)$, which is linear in the number of temporal edges m . We are mostly interested in small patterns, i.e., cases when l is a small constant. Thus, this methodology is optimal (linear in the input, m) for counting 2-node δ -temporal motif instances.

General algorithm for star motifs. Next, we consider k -node, l -




Figure 3: All 2-node and 3-node, 3-edge δ -temporal motifs. The green background highlights the four 2-node motifs (bottom left) and the grey background highlights the eight triangles. The 24 other motifs are stars. We index the 36 motifs $M_{i,j}$ by 6 rows and 6 columns. The first edge in each motif is from the green to the orange node. The second edge is the same along each row, and the third edge is the same along each column.

edge star motifs M , whose induced static graph H consists of a center node and $k - 1$ neighbors, where edges may occur in either direction between the center node and a neighbor node. For example, in the top left corner of Fig. 3, $M_{1,1}$ is a star motif with all edges pointing toward the center node. In such motifs, the induced static graph H contains at most $2(k - 1) = 2k - 2$ static edges—one incoming and outgoing edge from the center node to each neighbor node. We have the following method for counting the number of instances of k -node, l -edge star motifs: (1) for each node u in the static graph and for each unique set of $k - 1$ neighbors, gather and sort the edges in either direction between u and the neighbors; (2) count the number of instances of M using Alg. 1. The counts from each call to Alg. 1 are summed over all center nodes.

The major drawback of this approach is that we have to loop over each size- $(k - 1)$ neighbor set. This can be prohibitively expensive even when $k = 3$ if the center node has large degree. In Section 4.2, we shall design an algorithm that avoids this issue for the case when the star motif has $l = 3$ edges and $k = 3$.

General algorithm for triangle motifs. In triangle motifs, the induced graph H consists of 3 nodes and at least one directed edge between any pair of nodes (see Fig. 3 for all eight of the 3-edge triangle motifs). The induced static graph H of M contains at least three and at most six static edges. A straightforward algorithm for counting l -edge triangle motifs in a temporal graph T is:

1. Use a fast static graph triangle enumeration algorithm to find all triangles in the static graph G induced by T [14].
2. For each triangle (u, v, w) , merge all temporal edges from each pair of nodes to get a time-sorted list of edges. Use Alg. 1 to count the number of instances of M .

This approach is problematic as the edges between a pair of nodes may participate in many triangles. Fig. 4 shows a worst-case example for the motif $M = (w, u, t'_1), (w, v, t'_2), (u, v, t'_3)$ with $\delta = \infty$. In this case, the timestamps are ordered by their index. There are $m - 2n$ edges between u and v , and each of these edges forms an instance of M with every w_i . Thus, the overall worst-case running time of the algorithm is $O(\text{TriEnum} + m\tau)$, where TriEnum is the time to enumerate the number of triangles τ in the static graph. In




Figure 4: Worst-case example for counting triangular motifs with Alg. 1.


the following section, we devise an algorithm that significantly reduces the dependency on τ from linear to sub-linear (specifically, $\sqrt{\tau}$) when there are $l = 3$ edges.

4.2 Faster algorithms

The general counting algorithm from the previous subsection counts the number of instances of any k -node, l -edge δ -temporal motif, and is also optimal for 2-node motifs. However, the computational cost may be expensive for other motifs such as stars and triangles. We now develop specialized algorithms that count certain motif classes faster. Specifically, we design faster algorithms for counting all 3-node, 3-edge star and triangle motifs (Fig. 3 illustrates these motifs). Our algorithm for stars is linear in the input size, so it is optimal up to constant factors.

Fast algorithm for 3-node, 3-edge stars. With 3-node, 3-edge star motifs, the key drawback of using the previous algorithmic approach would be that we would have to loop over all pairs of neighbors given a center node. Instead, we will count all instances of star motifs for a given center node in just a single pass over the edges adjacent to the center node.

We use a dynamic programming approach for counting star motifs. First, note that every temporal edge in a star with center u is defined by (1) a neighbor node, (2) a direction of the edge (outward from or inward to u), and (3) the timestamp. With this insight we then notice that there are 3 classes of star motifs on 3 nodes and 3 edges:



where each class has $2^3 = 8$ motifs for each of the possible directions on the three edges.

Now, suppose we process the time-ordered sequence of edges containing the center node u . We maintain the following counters when processing an edge with timestamp t_j :

- $\text{pre_sum}[\text{dir1}, \text{dir2}]$ is the number of sequentially ordered pairs of edges in $[t_j - \delta, t_j)$ where the first edge points in direction dir1 and the second edge points in direction dir2
- $\text{post_sum}[\text{dir1}, \text{dir2}]$ is the analogous counter for the time window $(t_j, t_j + \delta]$.
- $\text{mid_sum}[\text{dir1}, \text{dir2}]$ is the number of pairs of edges where the first edge is in direction dir1 and occurred at time $t < t_j$ and the second edge is in direction dir2 and occurred at time $t' > t_j$ such that $t' - t \leq \delta$.

If we are currently processing an edge, the “pre” class gets $\text{pre_sum}[\text{dir1}, \text{dir2}]$ new motif instances for any choice of directions dir1 and dir2 (specifying the first two edge directions) and the current edge serves as the third edge in the motif (hence specifying the third edge direction). Similar updates are made with the $\text{post_sum}[\cdot, \cdot]$ and $\text{mid_sum}[\cdot, \cdot]$ counters, where the current edge serves as the first or second edge in the motif, respectively.

In order for our algorithm to be efficient, we must quickly update our counters. To aid in this, we introduce two additional counters:

Algorithm 2: Algorithmic framework for faster counting of 3-node, 3-edge star and triangle temporal motifs. The fast star counting method (Alg. 3) and triangle counting method (Alg. 4) implement different versions of the $\text{Push}()$, $\text{Pop}()$, and $\text{ProcessCurrent}()$ subroutines.

Input: Sequence of edges $(e_1 = (u_1, v_1), t_1), \dots, (e_L, t_L)$ with $t_1 < \dots < t_L$, time window δ

Initialize counters pre_nodes , post_nodes , mid_sum , pre_sum , and post_sum ; $\text{start} \leftarrow 1$, $\text{end} \leftarrow 1$

for $j = 1, \dots, L$:

while $t_{\text{start}} + \delta < t_j$ **do**

$\text{Pop}(\text{pre_nodes}, \text{pre_sum}, e_{\text{start}})$, $\text{start} += 1$

while $t_{\text{end}} \leq t_j + \delta$ **do**

$\text{Push}(\text{post_nodes}, \text{post_sum}, e_{\text{end}})$, $\text{end} += 1$

$\text{Pop}(\text{post_nodes}, \text{post_sum}, e_j)$

$\text{ProcessCurrent}(e_j)$

$\text{Push}(\text{pre_nodes}, \text{pre_sum}, e_j)$

- $\text{pre_nodes}[\text{dir}, v_i]$ is the number of times node v_i has appeared in an edge with u with direction dir in the time window $[t_j - \delta, t_j)$
- $\text{post_nodes}[\text{dir}, v_i]$ is the analogous counter but for the time window $(t_j, t_j + \delta]$.

Following the ideas of Alg. 1, it is easy to update these counters when we process a new edge. Consequently, $\text{pre_sum}[\cdot, \cdot]$, $\text{post_sum}[\cdot, \cdot]$, and $\text{mid_sum}[\cdot, \cdot]$ can be maintained when processing an edge with just a few simple primitives:

- $\text{Push}()$ and $\text{Pop}()$ update the counts for $\text{pre_nodes}[\cdot, \cdot]$, $\text{post_nodes}[\cdot, \cdot]$, $\text{pre_sum}[\cdot, \cdot]$ and $\text{post_sum}[\cdot, \cdot]$ when edges enter and leave the time windows $[t_j - \delta, t_j)$ and $(t_j, t_j + \delta]$.
- $\text{ProcessCurrent}()$ updates motif counts involving the current edge and updates the counter $\text{mid_sum}[\cdot, \cdot]$.

We describe the general procedure in Alg. 2, which will also serve as the basis for our fast triangle counting procedure, and Alg. 3 implements the subroutines $\text{Push}()$, $\text{Pop}()$, and $\text{ProcessCurrent}()$ for counting instances of 3-node, 3-edge star motifs. The $\text{count_pre}[\cdot, \cdot, \cdot]$, $\text{count_post}[\cdot, \cdot, \cdot]$, and $\text{count_mid}[\cdot, \cdot, \cdot]$ counters in Alg. 3 maintain the counts of the three different classes of stars described above.

Finally, we note that our counting scheme incorrectly includes instances of 2-node motifs such as $M = (u, v_i, t_1), (u, v_i, t_2), (u, v_i, t_3)$, but we can use the efficient 2-node motif counting algorithm to account for this. Putting everything together, we have the following procedure:

1. For each node u in the temporal graph T , get a time-ordered list of all edges containing u .
2. Use Algs. 2 and 3 to count star motif instances.
3. For each neighbor v of a star center u , subtract the 2-node motif counts using Alg. 1.

If the m edges of T are time-sorted, the first step can be done in linear time. The second and third steps run in linear time in the input size. Each edge is used in steps 2 and 3 exactly twice: once for each end point as the center node. Thus, the overall complexity of the algorithm is $O(m)$, which is optimal up to constant factors.

Fast algorithm for 3-edge triangle motifs. While our fast star counting routine relied on counting motif instances for all edges adjacent to a given *node*, our fast triangle algorithm is based on counting instances for all edges adjacent to a given *pair of nodes*. Specifically, given a pair of nodes u and v and a list of common neighbors w_1, \dots, w_d , we count the number of motif instances for triangles (w_i, u, v) . Given all of the edges between these three

Algorithm 3: Implementation of Alg. 2 subroutines for efficiently counting instances of 3-node, 3-edge star motifs. Temporal edges are specified by a neighbor nbr , a direction dir (incoming or outgoing), and a timestamp. The “:” notation represents a selection of all indices in an array.

```

Initialize counters  $\text{count\_pre}, \text{count\_post}, \text{count\_mid}$ 
Procedure  $\text{Push}(\text{node\_count}, \text{sum}, e = (\text{nbr}, \text{dir}))$ 
┌  $\text{sum}[:, \text{dir}] += \text{node\_count}[:, \text{nbr}]$ 
└  $\text{node\_count}[\text{dir}, \text{nbr}] += 1$ 
Procedure  $\text{Pop}(\text{node\_count}, \text{sum}, e = (\text{nbr}, \text{dir}))$ 
┌  $\text{node\_count}[\text{dir}, \text{nbr}] -= 1$ 
└  $\text{sum}[\text{dir}, :] -= \text{node\_count}[:, \text{nbr}]$ 
Procedure  $\text{ProcessCurrent}(e = (\text{nbr}, \text{dir}))$ 
┌  $\text{mid\_sum}[:, \text{dir}] -= \text{pre\_nodes}[:, \text{nbr}]$ 
├  $\text{count\_pre}[:, :, \text{dir}] += \text{pre\_sum}[:, :]$ 
├  $\text{count\_post}[\text{dir}, :, :] += \text{post\_sum}[:, :]$ 
├  $\text{count\_mid}[:, \text{dir}, :] += \text{mid\_sum}[:, :]$ 
└  $\text{mid\_sum}[\text{dir}, :] += \text{post\_nodes}[:, \text{nbr}]$ 
return  $\text{count\_pre}, \text{count\_post}, \text{count\_mid}$ 

```

nodes, the counting procedures are nearly identical to the case of stars. We use the same general counting method (Alg. 2), but the behavior of the subroutines $\text{Push}()$, $\text{Pop}()$, and $\text{ProcessCurrent}()$ depends on whether or not the edge is between u and v .

These methods are implemented in Alg. 4. The input is a list of edges adjacent to a given pair of neighbors u and v , where each edge consists of four pieces of information: (1) a neighbor node nbr , (2) an indicator of whether or not the node nbr connects to node u or node v , (3) the direction dir of the edge, and (4) the timestamp. The node counters ($\text{pre_nodes}[:, \cdot, \cdot]$ and $\text{post_nodes}[:, \cdot, \cdot]$) in Alg. 4 have an extra dimension compared to Alg. 3 to indicate whether the counts correspond to edges containing node u or node v (denoted by “uorv”). Similarly, the sum counters ($\text{pre_sum}[:, \cdot, \cdot]$, $\text{mid_sum}[:, \cdot, \cdot]$ and $\text{post_sum}[:, \cdot, \cdot]$) have an extra dimension to denote if the first edge is incident on node u or node v .

Recall that the problem with counting triangle motifs by the general framework in Alg. 1 is that a pair of nodes with many edges might have to be counted for many triangles in the graph. However, with Alg. 4, we can simultaneously count all triangles adjacent to a given pair of nodes. What remains is that we must assign each triangle in the static graph to a pair of nodes. Here, we propose to assign each triangle to the pair of nodes in that triangle containing the largest number of edges, which is sketched in Alg. 5. Alg. 5 aims to process as many triangles as possible for pairs of nodes with many edges. The following theorem says that this is faster than simply counting for each triangle (described in Section 4.1). Specifically, we reduce $O(m\tau)$ complexity to $O(m\sqrt{\tau})$.

Theorem. *In the worse case, Alg. 5 runs in time $O(\text{TriEnum} + m\sqrt{\tau})$, where TriEnum is the time to enumerate all triangles in the static graph G , m is the total number of temporal edges, and τ is the number of static triangles in G .*

Proof. Let σ_i be the number of edges between the i th pair of nodes with at least one edge, and let $p_i \geq 1$ be the number of times that edges on this pair are used in a call to Alg. 4 by Alg. 5. Since Alg. 4 runs in linear time in the number of edges in its input, the total running time is on the order of $\sum_i \sigma_i p_i$.

The σ_i are fixed, and we wish to find the values of p_i that maximize the summation. Without loss of generality, assume that the σ_i are in decreasing order so that the most number of edges between a pair of nodes is σ_1 . Consequently, $p_i \leq i$. Note that each triangle contributes to at most a constant repeat processing of edges for a given pair of nodes. Hence, $\sum_i p_i \leq c\tau$ for some constant c . The

Algorithm 4: Implementation of Alg. 2 subroutines for counting 3-edge triangle motifs containing a specified pair of nodes u and v . Temporal edges are specified by a neighbor nbr , a direction dir (incoming or outgoing), an indicator “uorv” denoting if the edge connects to u or v , and a timestamp. The “:” notation represents a selection of all indices in an array.

```

Initialize counter  $\text{count}$ 
Procedure  $\text{Push}(\text{node\_count}, \text{sum}, e = (\text{nbr}, \text{dir}, \text{uorv}))$ 
┌ if  $\text{nbr} \in \{u, v\}$  then return
├  $\text{sum}[1\text{-uorv}, :, \text{dir}] += \text{node\_count}[1\text{-uorv}, :, \text{nbr}]$ 
└  $\text{node\_count}[\text{uorv}, \text{dir}, \text{nbr}] += 1$ 
Procedure  $\text{Pop}(\text{node\_count}, \text{sum}, e = (\text{nbr}, \text{dir}, \text{uorv}))$ 
┌ if  $\text{nbr} \in \{u, v\}$  then return
├  $\text{node\_count}[\text{uorv}, \text{dir}, \text{nbr}] -= 1$ 
└  $\text{sum}[\text{uorv}, \text{dir}, :] -= \text{node\_count}[1\text{-uorv}, :, \text{nbr}]$ 
Procedure  $\text{ProcessCurrent}(e = (\text{nbr}, \text{dir}, \text{uorv}))$ 
┌ if  $\text{nbr} \notin \{u, v\}$  then
├  $\text{mid\_sum}[1\text{-uorv}, :, \text{dir}] -= \text{pre\_nodes}[1\text{-uorv}, :, \text{nbr}]$ 
├  $\text{mid\_sum}[\text{uorv}, \text{dir}, :] += \text{post\_nodes}[1\text{-uorv}, :, \text{nbr}]$ 
└ else
├  $\text{utov} = (\text{nbr} == u) \text{ XOR } \text{dir}$ 
├ for  $0 \leq i, j, k \leq 1$  :
├ ┌  $\text{count}[i, j, k] += \text{mid\_sum}[j \text{ XOR } \text{utov}, i, k]$ 
├ └  $\text{post\_sum}[i \text{ XOR } \text{utov}, j, 1 - k]$ 
├ └  $\text{pre\_sum}[k \text{ XOR } \text{utov}, 1 - i, 1 - j]$ 
└ end
/* count key map to Fig. 3: */
/*  $[0, 0, 0] \mapsto M_{1,3}, [0, 0, 1] \mapsto M_{1,4}, [0, 1, 0] \mapsto M_{2,3}$  */
/*  $[0, 1, 1] \mapsto M_{2,4}, [1, 0, 0] \mapsto M_{3,5}, [1, 0, 1] \mapsto M_{3,6}$  */
/*  $[1, 1, 0] \mapsto M_{4,5}, [1, 1, 1] \mapsto M_{4,6}$  */
return  $\text{count}$ 

```

Algorithm 5: Sketch of fast algorithm for counting the number of 3-edge δ -temporal triangle motifs in a temporal graph T .

```

Enumerate all triangles in the undirected static graph  $G$  of  $T$ 
 $\sigma \leftarrow$  number of temporal edges on each static edge  $e$  in  $G$ 
foreach static triangle  $\Delta = (e_1, e_2, e_3)$  in  $G$  do
┌  $e_{\max} = \arg \max_{e \in \Delta} \{\sigma_e\}$ 
├ foreach  $e$  in  $\Delta$  do Add  $e_{\max}$  to edge set  $\ell_e$ 
└ /*  $e' \in \ell_e$  if  $e \in \Delta$  and  $\Delta$  assigned to  $e'$  */
foreach temporal edge  $(e = (u, v), t)$  in time-sorted  $T$  do
┌ foreach  $e' \in \ell_e$  do Append  $(e, t)$  to temporal-edge list  $a_{e'}$ 
└ /*  $(e, t) \in a_{e'}$  if  $e \in \Delta$  and  $\Delta$  assigned to  $e'$  */
foreach undirected edge  $e$  in  $G$  do
└ Update counts using Alg. 4 with input  $a_e$ 

```

summation $\sum_i \sigma_i p_i$ is maximized when $p_1 = 1, p_2 = 2$, and so on up to some index $j = O(\sqrt{\tau})$ for which $\sum_{i=1}^j p_j = \sum_{i=1}^j i = c\tau$. Now given that the p_j are fixed and the σ_i are ordered, the summation is maximized when $\sigma_1 = \sigma_2 = \dots = \sigma_j = m/j$. In this case, $\sum_i \sigma_i p_i = \sum_{i=1}^j (m/j)(1 + i) = O(m\sqrt{\tau})$. \square

5. EXPERIMENTS

Next, we use our algorithms to reveal patterns in a variety of temporal network datasets. We find that the number of instances of various δ -temporal motifs reveal basic mechanisms of the networks. Datasets and implementations of our algorithms are available at <http://snap.stanford.edu/temporal-motifs>.

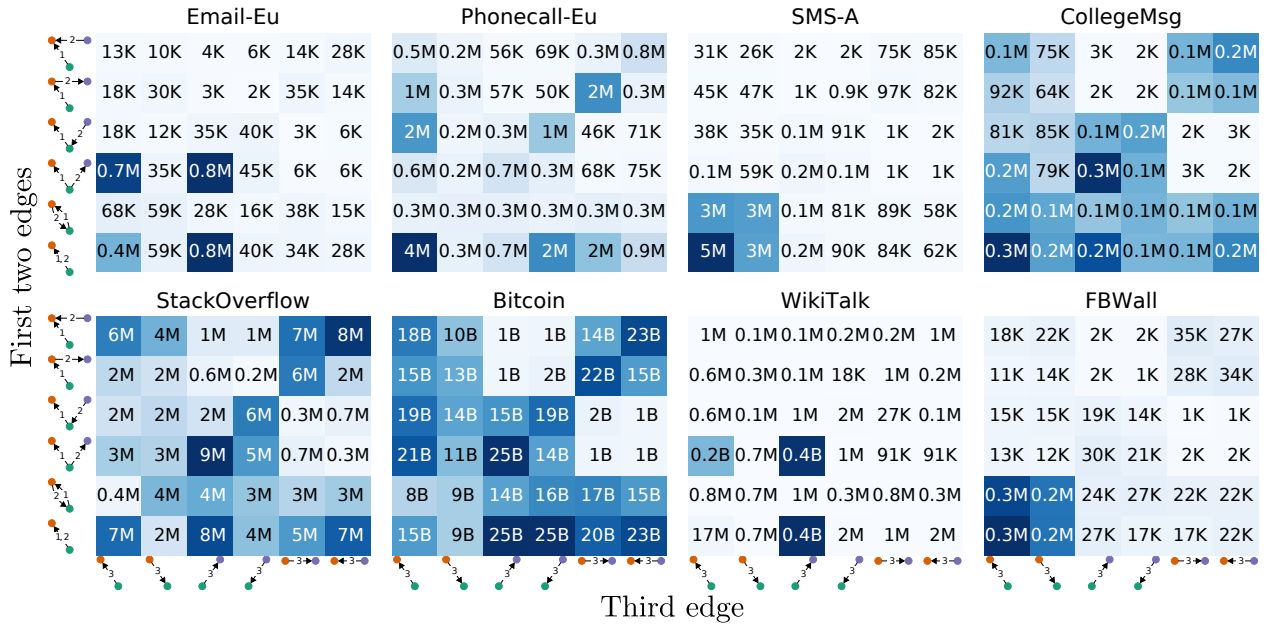


Figure 5: Counts of instances of all 2- and 3-node, 3-edge δ -temporal motifs with $\delta = 1$ hour. For each dataset, counts in the i th row and j th column is the number of instances of motif $M_{i,j}$ (see Fig. 3); this motif is the union of the two edges in the row label and the edge in the column label. For example, there are 0.7 million instances of motif $M_{4,1}$ in the EMAIL-EU dataset. The color for the count of motif $M_{i,j}$ indicates the fraction over all $M_{i,j}$ on a linear scale—darker blue means a higher count.

Table 1: Summary statistics of datasets.

| dataset | # nodes | # static edges | # edges | time span (days) |
|---------------|---------|----------------|---------|------------------|
| EMAIL-EU | 986 | 2.49K | 332K | 803 |
| PHONCALL-EU | 1.05M | 2.74M | 8.55M | 7 |
| SMS-A | 44.1K | 67.2K | 545K | 338 |
| COLLEGEMSG | 1.90K | 20.3K | 59.8K | 193 |
| STACKOVERFLOW | 2.58M | 34.9M | 47.9M | 2774 |
| BITCOIN | 24.6M | 88.9M | 123M | 1811 |
| FBWALL | 45.8K | 264K | 856K | 1560 |
| WIKITALK | 1.09M | 3.13M | 6.10M | 2277 |
| PHONCALL-ME | 18.7M | 360M | 2.04B | 364 |
| SMS-ME | 6.94M | 51.5M | 800M | 89 |

5.1 Data

We gathered a variety of datasets in order to study the patterns of δ -temporal motifs in several domains. The datasets are described below and summary statistics are in Table 1. The time resolution of the edges in all datasets is one second.

EMAIL-EU. This dataset is a collection of emails between members of a European research institution [17]. An edge (u, v, t) signifies that person u sent person v an email at time t .

PHONCALL-EU. This dataset was constructed from telephone call records for a major European service provider. An edge (u, v, t) signifies that person u called person v starting at time t .

SMS-A. Short messaging service (SMS) is a texting service provided on mobile phones. In this dataset, an edge (u, v, t) means that person u sent an SMS message to person v at time t [28].

COLLEGEMSG. This dataset is comprised of private messages sent on an online social network at the University of California, Irvine [21]. Users could search the network for others and then initiate conversation based on profile information. An edge (u, v, t) means that user u sent a private message to user v at time t .

STACKOVERFLOW. On stack exchange web sites, users post questions and receive answers from other users, and users may comment on both questions and answers. We derive a temporal network by creating an edge (u, v, t) if, at time t , user u : (1) posts an answer to user v 's question, (2) comments on user v 's question, or (3) comments on user v 's answer. We formed the temporal network from the entirety of Stack Overflow's history up to March 6, 2016.

BITCOIN. Bitcoin is a decentralized digital currency and payment system. This dataset consists of all payments made up to October 19, 2014 [11]. Nodes in the network correspond to Bitcoin addresses, and an individual may have several addresses. An edge (u, v, t) signifies that bitcoin was transferred from address u to address v at time t .

FBWALL. The edges of this dataset are wall posts between users on the social network Facebook located in the New Orleans region [26]. Any friend of a given user can see all posts on that user's wall, so communication is public among friends. An edge (u, v, t) means that user u posted on user v 's wall at time t .

WIKITALK. This dataset represents edits on user talk pages on Wikipedia [16]. An edge (u, v, t) signifies that user u edited user v 's talk page at time t .

PHONCALL-ME and SMS-ME. This dataset is constructed from phone call and SMS records of a large telecommunications service provider in the Middle East. An edge (u, v, t) in PHONCALL-ME means that user u initiated a call to user v at time t . An edge (u, v, t) in SMS-ME means that user u sent an SMS message to user v at time t . We use these networks for scalability experiments in Section 5.3.

5.2 Empirical observations of motif counts

We first examine the distribution of 2- and 3-node, 3-edge motif instance counts from 8 of the datasets described in Section 5.1 with $\delta = 1$ hour (Fig. 5). We choose 1 hour for the time window as this is close to the median time for a node to take part in




Figure 6: Fraction of all 2 and 3-node, 3-edge δ -temporal motif counts that correspond to two groups of motifs ($\delta = 1$ hour). Motifs on the left capture “blocking” behavior, common in SMS messaging and Facebook wall posting, and motifs on the right exhibit “non-blocking” behavior, common in email.




Figure 7: Distribution of switching behavior amongst the non-blocking motifs. Switching is least common on Stack Overflow and most common in email.

three edges in most of our datasets. We make a few empirical observations uniquely available due to temporal motifs and provide possible explanations for these observations.

Blocking communication. If an individual typically waits for a reply from one individual before proceeding to communicate with another individual, we consider it a *blocking* form of communication. A typical conversation between two individuals characterized by fast exchanges happening back and forth is blocking as it requires complete attention of both individuals. We capture this behavior in the “blocking motifs” $M_{5,1}$, $M_{5,2}$ and $M_{6,2}$, which contain 3 edges between two nodes with at least one edge in either direction (Fig. 6, left). However, if the reply doesn’t arrive soon, we might expect the individual to communicate with others without waiting for a reply from the first individual. This is a non-blocking form of communication and is captured by the “non-blocking motifs” $M_{4,1}$, $M_{4,3}$ and $M_{6,3}$ having edges originating from the same source but directed to different destinations (Fig. 6, right)

The fractions of counts corresponding to the blocking and non-blocking motifs out of the counts for all 36 motifs in Fig. 3 uncover several interesting characteristics in communication networks ($\delta = 1$ hour; see Fig. 6). In FBWALL and SMS-A, blocking communication is vastly more common, while in EMAIL-EU non-blocking communication is prevalent. Email is not a dynamic method of communication and replies within an hour are rare. Thus, we would expect non-blocking behavior. Interestingly, the COLLEGEMSG dataset shows both behaviors as we might expect individuals to engage in multiple conversations simultaneously. In complete contrast, the PHONECALL-EU dataset shows neither behavior. A simple explanation is that a single edge (a phone call) captures an entire conversation and hence blocking behavior does not emerge.

Cost of switching. Amongst the non-blocking motifs discussed above, $M_{4,1}$ captures two consecutive switches between pairs of nodes whereas $M_{4,3}$ and $M_{6,3}$ each have a single switch (Fig. 7, right). Prevalence of $M_{4,1}$ indicates a lower cost of switching targets, whereas prevalence of the other two motifs are indicative of a higher cost. We observe in Fig. 7 that the ratio of 2-switch to 1-




Figure 8: Fraction of 3-edge δ -temporal triangle motif counts ($\delta = 1$ hour) corresponding to cyclic triangles (right) in the four datasets for which this fraction is the largest. BITCOIN has a much higher fraction compared to all other datasets.




Figure 9: Percentage of explained variance of relative counts of collections of datasets plotted as a function of the number of principal components. In datasets from the same domain, 90% of variance is explained with fewer components.

switch motif counts is the least in STACKOVERFLOW, followed by WIKITALK, COLLEGEMSG and then EMAIL-EU. On Stack Overflow and Wikipedia talk pages, there is a high cost to switch targets because of peer engagement and depth of discussion. On the other hand, in the COLLEGEMSG dataset there is a lesser cost to switch because it lacks depth of discussion within the time frame of $\delta = 1$ hour. Finally, in EMAIL-EU, there is almost no peer engagement and cost of switching is negligible.

Cycles in BITCOIN. Of the eight 3-edge triangle motifs, $M_{2,4}$ and $M_{3,5}$ are cyclic, i.e., the target of each edge serves as the source of another edge. We observe in Fig. 8 that the fraction of triangles that are cyclic is much higher in BITCOIN compared to any other dataset. This can be attributed to the transactional nature of BITCOIN where the total amount of bitcoin is limited. Since remittance (outgoing edges) is typically associated with earnings (incoming edges), we should expect cyclic behavior.

Datasets from the same domain have similar counts. Static graphs from similar domains tend to have similar motif count distributions [18, 25, 29]. Here, we find similar results in temporal networks. We formed two collections of datasets from similar domains. First, we took subsets of the EMAIL-EU dataset corresponding to email communication within four different departments at the institution. Second, we constructed temporal graphs from the stack exchange communities Math Overflow, Super User, and Ask Ubuntu to study in conjunction with the STACKOVERFLOW dataset. We form count distributions by normalizing the counts of the 36 different motifs in Fig. 5. For datasets from a similar domain, we expect that if the count distributions are similar, then most of the variance is captured by a few principal components. To compare, we use four datasets from dissimilar domains (EMAIL-EU, PHONECALL-EU, SMS-A, WIKITALK). Fig. 9 shows that to explain 90% variance, EMAIL-EU subnetworks need just one principal component, stack exchange networks need two, and the dissimilar networks need three.

Motif counts at varying time scales. We now explore how motif counts change at different time scales. For the STACKOVERFLOW dataset we counted the number of instances of 2- and 3-node, 3-edge δ -temporal motifs for $\delta = 60, 300, 1800,$ and 3600 seconds

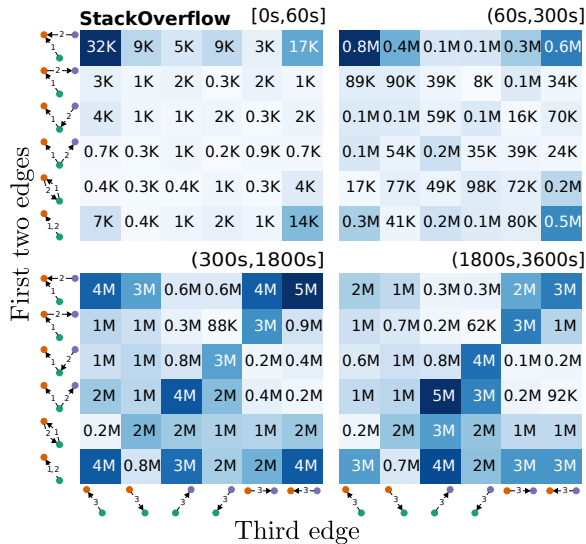


Figure 10: Counts for all 2- and 3-node, 3-edge δ -temporal motifs in four time intervals for the STACKOVERFLOW dataset. For each interval, the count in the i th row and j th column is the number of instances of motif $M_{i,j}$ (see Fig. 3).




Figure 11: Counts over various time scales for the motifs representing a node sending 3 outgoing messages to 1 or 2 neighbors in the COLLEGMMSG dataset.

(Fig. 10). These counts determine the number of motifs that completed in the intervals $[0, 60]$, $(60, 300]$, $(300, 1800s]$, and $(1800, 3600]$ seconds (e.g., subtracting 60 second counts from 300 second counts gives the interval $(60, 300]$). Observations at smaller timescales reveal phenomenon which start to get eclipsed at larger timescales. For instance, on short time scales, motif $M_{1,1}$ (Fig. 10, top-left corner) is quite common. We suspect this arises from multiple, quick comments on the original question, so the original poster receives many incoming edges. At larger time scales, this behavior is still frequent but relatively less so. Now let us compare counts for $M_{1,5}$, $M_{1,6}$, $M_{2,5}$, $M_{2,6}$ (the four in the top right corner) with counts for $M_{3,3}$, $M_{3,4}$, $M_{4,3}$, $M_{4,4}$ (the four in the center). The former counts likely correspond to conversations with the original poster while the latter are constructed by the same user interacting with multiple questions. Between 300 and 1800 seconds (5 to 30 minutes), the former counts are relatively more common while the latter counts only become more common after 1800 seconds. A possible explanation is that the typical length of discussions on a post is about 30 minutes, and later on, users answer other questions.

Next, we examine messaging behavior in the COLLEGMMSG dataset at fine-grained time intervals. We counted the number of motifs consisting of a single node sending three outgoing messages

Table 2: Time to count the eight 3-edge δ -temporal triangle motifs ($\delta = 3600$) using the general counting method (Alg. 1) and the fast counting method (Alg. 5).

| dataset | # static triangles | time, Alg. 1 (seconds) | time, Alg. 5 (seconds) | speedup |
|---------------|--------------------|------------------------|------------------------|---------|
| WIKITALK | 8.11M | 51.1 | 26.6 | 1.92x |
| BITCOIN | 73.1M | 27.3K | 483 | 56.5x |
| SMS-ME | 78.0M | 2.54K | 1.11K | 2.28x |
| STACKOVERFLOW | 114M | 783 | 606 | 1.29x |
| PHONECALL-ME | 667M | 12.2K | 8.59K | 1.42x |




Figure 12: Time to count 3-edge motifs on the first k temporal million edges in the PHONECALL-ME as a function of k .

to one or two neighbors (motifs $M_{6,1}$, $M_{6,3}$, $M_{4,1}$, and $M_{4,3}$ in the time bins $[10(i-1), 10i)$ seconds, $i = 1, \dots, 500$ (Fig. 11)). We first notice that at small time scales, the motif consisting of three edges to a single neighbor ($M_{6,1}$) occurs frequently. This pattern could emerge from a succession of quick introductory messages. Overall, motif counts increase from roughly 1 minute to 20 minutes and then decline. Interestingly, after 5 minutes, counts for the three motifs with one switch in the target ($M_{6,1}$, $M_{6,3}$, and $M_{4,3}$) grow at a faster rate than the counts for the motif with two switches ($M_{4,1}$). As mentioned above, this pattern could emerge from a tendency to send several messages in one conversation before switching to a conversation with another friend.

5.3 Algorithm scalability

Finally, we performed scalability experiments of our algorithms. All algorithms were implemented in C++, and all experiments ran using a single thread of a 2.4GHz Intel Xeon E7-4870 processor. We did not measure the time to load datasets into memory, but our timings include all pre-processing time needed by the algorithms (e.g., the triangle counting algorithms first find triangles in the static graph). We emphasize that our implementation is single threaded, and the methods can be sped up with a parallel algorithm.

First, we used both the general counting method (Alg. 1) and the fast counting method (Alg. 5) to count the number of all eight 3-edge δ -temporal triangle motifs in our datasets ($\delta = 1$ hour). Table 2 reports the running times of the algorithms for all datasets with at least one million triangles in the static graph. For all of these datasets, our fast temporal triangle counting algorithm provides significant performance gains over the general counting method, ranging between a 1.29x and a 56.5x speedup. The gains of the fast algorithm are the largest for BITCOIN, which is due to some pairs of nodes having many edges between them and also participating in many triangles.

Second, we measured the time to count various 3-edge δ -temporal motifs in our largest dataset, PHONECALL-ME. Specif-

ically, we measured the time to compute (1) 2-node motifs, (2) 3-node stars, and (3) triangles on the first k million edges in the dataset for $k = 250, 500, \dots, 2000$ (Fig. 12). The time to compute the 2-node, 3-edge motifs and the 3-node, 3-edge stars scales linearly, as expected from our algorithm analysis. The time to count triangle motifs grows superlinearly and becomes the dominant cost when there is a large number of edges. For practical purposes, the running times are quite modest. With two billion edges, our methods take less than 3.5 hours to complete (executing sequentially).

6. DISCUSSION

We have developed δ -temporal network motifs as a tool for analyzing temporal networks. We introduced a general framework for counting instances of any temporal motif as well as faster algorithms for certain classes of motifs and found that motif counts reveal key structural patterns in a variety of temporal network datasets. Our work opens a number of avenues for additional research. First, our fast algorithms are designed for 3-node, 3-edge star and triangle motifs. We expect that the same general techniques can be used to count more complex temporal motifs. Next, it is important to note that our fast algorithms only *count* the number of instances of motifs rather than *enumerate* the instances. This concept has also been used to accelerate static motif counting [22]. Temporal motif enumeration algorithms provide an additional algorithmic design challenge. There is also a host of theoretical questions in this area for lower bounds on temporal motif counting. Finally, motif counts can also be measured with respect to a null model [13, 19]. Such analysis may yield additional discoveries. Importantly, our algorithms will speed up such computations, which use raw counts from many random instances of a generative null model.

Acknowledgements. We thank Moses Charikar for valuable discussion. This research has been supported in part by NSF IIS-1149837, ARO MURI, DARPA SIMPLEX and NGS2, Boeing, Bosch, Huawei, Lightspeed, SAP, Tencent, Volkswagen, Stanford Data Science Initiative, and a Stanford Graduate Fellowship.

7. REFERENCES

- [1] M. Araujo, S. Papadimitriou, S. Günnemann, C. Faloutsos, P. Basu, A. Swami, E. E. Papalexakis, and D. Koutra. Com2: fast automatic discovery of temporal (‘comet’) communities. In *PAKDD*, 2014.
- [2] A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.
- [3] A. R. Benson, D. F. Gleich, and J. Leskovec. Tensor spectral clustering for partitioning higher-order network structures. In *SDM*, 2015.
- [4] A. R. Benson, D. F. Gleich, and J. Leskovec. Higher-order organization of complex networks. *Science*, 353(6295):163–166, 2016.
- [5] M. Berlingerio, F. Bonchi, B. Bringmann, and A. Gionis. Mining graph evolution rules. In *ECML PKDD*, 2009.
- [6] D. M. Dunlavy, T. G. Kolda, and E. Acar. Temporal link prediction using matrix and tensor factorizations. *TKDD*, 5(2):10, 2011.
- [7] S. Gurukar, S. Ranu, and B. Ravindran. Commit: A scalable approach to mining communication motifs from dynamic networks. In *SIGMOD*, 2015.
- [8] P. Holme and J. Saramäki. Temporal networks. *Physics Reports*, 519(3):97–125, 2012.
- [9] H. Huang, J. Tang, S. Wu, L. Liu, et al. Mining triadic closure patterns in social networks. In *WWW*, 2014.
- [10] A. Z. Jacobs, S. F. Way, J. Ugander, and A. Clauset. Assembling thefacebook: Using heterogeneity to understand online social network assembly. In *Web Science*, 2015.
- [11] D. Kondor, M. Pósfai, I. Csabai, and G. Vattay. Do the rich get richer? an empirical analysis of the bitcoin transaction network. *PLOS ONE*, 9(2):e86197, 2014.
- [12] G. Kossinets and D. J. Watts. Empirical analysis of an evolving social network. *Science*, 311(5757):88–90, 2006.
- [13] L. Kovanen, M. Karsai, K. Kaski, J. Kertész, and J. Saramäki. Temporal motifs in time-dependent networks. *JSTAT*, 2011(11):P11005, 2011.
- [14] M. Latapy. Main-memory triangle computations for very large (sparse (power-law)) graphs. *Theoretical Computer Science*, 407(1):458–473, 2008.
- [15] J. Leskovec, D. Huttenlocher, and J. Kleinberg. Signed networks in social media. In *CHI*, 2010.
- [16] J. Leskovec, D. P. Huttenlocher, and J. M. Kleinberg. Governance in social media: A case study of the wikipedia promotion process. In *ICWSM*, 2010.
- [17] J. Leskovec, J. Kleinberg, and C. Faloutsos. Graph evolution: Densification and shrinking diameters. *TKDD*, 1(1):2, 2007.
- [18] R. Milo, S. Itzkovitz, N. Kashtan, R. Levitt, S. Shen-Orr, I. Ayzenshtat, M. Sheffer, and U. Alon. Superfamilies of evolved and designed networks. *Science*, 303(5663):1538–1542, 2004.
- [19] R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon. Network motifs: simple building blocks of complex networks. *Science*, 298(5594):824–827, 2002.
- [20] M. E. Newman. The structure and function of complex networks. *SIAM Review*, 45(2):167–256, 2003.
- [21] P. Panzarasa, T. Opsahl, and K. M. Carley. Patterns and dynamics of users’ behavior and interaction: Network analysis of an online community. *JASIST*, 2009.
- [22] A. Pinar, C. Seshadhri, and V. Vishal. ESCAPE: Efficiently Counting All 5-Vertex Subgraphs. *arXiv: 1610.09411*, 2016.
- [23] C. Tantipathananandh, T. Berger-Wolf, and D. Kempe. A framework for community identification in dynamic social networks. In *KDD*, 2007.
- [24] J. Ugander, L. Backstrom, and J. Kleinberg. Subgraph frequencies: Mapping the empirical and extremal geography of large graph collections. In *WWW*, 2013.
- [25] A. Vazquez, R. Dobrin, D. Sergi, J.-P. Eckmann, Z. Oltvai, and A.-L. Barabási. The topological relationship between the large-scale attributes and local interaction patterns of complex networks. *PNAS*, 101(52):17940–17945, 2004.
- [26] B. Viswanath, A. Mislove, M. Cha, and K. P. Gummadi. On the evolution of user interaction in Facebook. In *WOSN*, 2009.
- [27] S. Wernicke and F. Rasche. Fanmod: a tool for fast network motif detection. *Bioinformatics*, 22(9):1152–1153, 2006.
- [28] Y. Wu, C. Zhou, J. Xiao, J. Kurths, and H. J. Schellnhuber. Evidence for a bimodal distribution in human communication. *PNAS*, 107(44):18803–18808, 2010.
- [29] Ö. N. Yaveroglu, N. Malod-Dognin, D. Davis, Z. Levnajic, V. Janjic, R. Karapandza, A. Stojmirovic, and N. Przulj. Revealing the hidden language of complex networks. *Scientific Reports*, 4, 2014.
- [30] Q. Zhao, Y. Tian, Q. He, N. Oliver, R. Jin, and W.-C. Lee. Communication motifs: a tool to characterize social communications. In *CIKM*, 2010.