



Bounds and algorithms for graph trusses

*Paul Burkhardt*¹ *Vance Faber*² *David G. Harris*³

¹U.S. National Security Agency, Ft. Meade, MD 20755

²IDA/Center for Computing Sciences, Bowie MD 20707

³Department of Computer Science, University of Maryland, College Park, MD 20742

Abstract

The k -truss, introduced by Cohen (2005), is a graph where every edge is incident to at least k triangles. This is a relaxation of the clique. It has proved to be a useful tool in identifying cohesive subnetworks in a variety of real-world graphs. Despite its simplicity and its utility, the combinatorial and algorithmic aspects of trusses have not been thoroughly explored.

We provide nearly-tight bounds on the edge counts of k -trusses. We also give two improved algorithms for finding trusses in large-scale graphs. First, we present a simplified and faster algorithm, based on approach discussed in Wang & Cheng (2012). Second, we present a theoretical algorithm based on fast matrix multiplication; this converts a triangle-generation algorithm of Björklund et al. (2014) into a dynamic data structure.

Submitted: November 2018	Reviewed: June 2019	Revised: July 2019	Accepted: January 2020
	Final: April 2020	Published: April 2020	
Article type: Regular Paper		Communicated by: X. He	

1 Introduction

In a number of contexts, a group of interacting agents can be represented in terms of an undirected graph $G = (V, E)$. For example, in a social network, the vertices may represent people with an edge if they know each other. One basic task is to find a *cohesive subnetwork* of G : a maximal subgraph whose vertices are “highly connected” [17]. This may represent a discrete community in the overall network, or another type of subgroup with a high degree of mutual relationship. We emphasize that since we are ultimately trying to understand a non-mathematical property of G , we cannot give an exact definition of a cohesive subnetwork.

A number of graph-theoretic structures can be used to find cohesive subnetworks in G . A clique is the most highly connected substructure. An alternate choice, suggested by [17], is the k -core, which is defined as a maximal connected subgraph in which each vertex has degree at least k .

Cohen [6, 7] proposed a stronger heuristic based on triangle counts called the *truss*. Formally, a k -truss is defined to be a graph in which every edge is incident to at least k triangles and which has no isolated vertices. Note that a $(k + 2)$ -clique is a k -truss.¹ A k -truss-component of G is defined to be a maximal edge set $L \subseteq E$ such that the edge-induced subgraph $G(L)$ is a connected k -truss.

The k -truss has been rediscovered and renamed several times. The earliest example was its definition as a k -dense core [16], which was motivated by the goal of detecting dense communities where the k -core proved to be too coarse. It was also defined as a *triangle k -core* in [23] and used as a motif exemplar in graphs. Other names include *k -community* [20] and *k -brace* [18].

The task of determining the truss-components of G is called *truss decomposition*. The truss-components of G can be derived from an associated hypergraph H which is defined as follows: the vertex set of H is the edge set of G , and the edge set of H is the set of triangles of G . Each k -core of H corresponds to a k -truss-component of G .

The *trussness* of an edge e of G , denoted $\tau(e)$, is defined to be the maximal value k such that e is in a k -truss-component of G . Equivalently, $\tau(e)$ is the coreness of the edge e regarded as a node of H . Truss decomposition algorithms typically first compute $\tau(e)$ for all edges e . The k -truss-components (for any value of k) can then be found by depth-first search of G restricted to edges e with $\tau(e) \geq k$.

An appealing feature of trusses is that truss decomposition algorithms are relatively practical, making them feasible for large graphs. As a starting point, Cohen’s original algorithm [7] was essentially an adaptation of a graph core decomposition algorithm of Matula & Beck [14] to the hypergraph H . This was improved by Wang & Cheng [21] and Huang et al. [12] by avoiding explicit generation of H . These algorithms compute truss decomposition in $O(m\alpha(G))$ time and $O(m)$ memory, where $\alpha(G)$ denotes the arboricity of G . Note that $m\alpha(G)$ is at most $O(m^{3/2})$.

¹Cohen defined the k -truss as being a connected graph such that every edge is incident to at least $k - 2$ triangles. This was presumably chosen so that a k -clique is a k -truss.

The k -truss-component structure of graphs has become a key tool for pattern mining and community detection in a variety of scientific and social network studies, e.g. [18, 4]. Additionally, the truss serves as a fast filter for finding cliques since a $(k + 2)$ -clique is a k -truss. In these applications, the parameter k represents the anticipated size of a meaningful cohesive subnetwork. Typically, k may be much smaller than n . In this setting, it is often useful to compute a *truncated* truss decomposition up to some chosen parameter k_{trunc} . This entails computing $\tau(e)$ for edges e with $\tau(e) \leq k_{\text{trunc}}$; other edges e record that $\tau(e) > k_{\text{trunc}}$ but do not record the precise value. For instance, Verma et al. [19] discussed an algorithm for sparse graphs with runtime $O(mk_{\text{trunc}}\Delta)$, where Δ is the maximum degree of G .

1.1 Our contributions and overview

Despite its simplicity and its use for understanding real-world networks, the truss has seen little formal analysis, either from a combinatorial or algorithmic point of view. We address these gaps in this paper.

Our main combinatorial subject of investigation is the minimum number of edges in a k -truss. Section 2 gives asymptotically tight bounds for this quantity. Section 3 analyzes a stricter notion of *critical k -truss*, which is a k -truss none of whose subgraphs are themselves k -trusses. We summarize the main results of these sections as follows:

Theorem 1 *The minimum number of edges in a connected k -truss on n vertices, is $n(1 + k/2) + \Theta(k^2)$. The minimum number of edges in a critical k -truss on n vertices is $nk/2 + \Theta(n + k^2)$.*

Section 4 describes a new simple algorithm for truss decomposition. We analyze this algorithm in terms of a graph parameter we refer to as the *average degeneracy* $\bar{\delta}(G)$, defined as:

$$\bar{\delta}(G) = \frac{1}{m} \sum_{\substack{\text{edges} \\ e=(u,v)}} \min(\deg(u), \deg(v)).$$

We have $\bar{\delta}(G) \leq 2\alpha(G) \leq O(\sqrt{m})$. While $\alpha(G)$ is a local property of the graph, which can be influenced by a few high-degree nodes, the parameter $\bar{\delta}(G)$ is a global property and can be much smaller than $\alpha(G)$. To the best of our knowledge, this parameter has not been studied before. We show the following result:

Theorem 2 *There is an algorithm to compute the truss decomposition of G in $O(m\bar{\delta}(G)) \leq O(m\alpha(G)) \leq O(m^{3/2})$ time and $O(m)$ memory.*

This algorithm is inspired by Wang & Cheng [21] but uses much simpler and faster data structures; it should be practical for large-scale graphs.

Section 5 gives an alternative algorithm for truncated truss decomposition based on matrix multiplication. We present here a slightly simplified summary

in terms of the linear algebra constant ω , i.e. the value for which multiplication of $N \times N$ matrices can be performed in $N^{\omega+o(1)}$ time. (The current best estimate [9] is $\omega \approx 2.38$.) The algorithm is theoretically appealing, but the algorithm in Section 4 is more likely to be useful in practice.

Theorem 3 *There is an algorithm to compute the truncated truss decomposition up to any desired parameter k_{trunc} in $m^{\frac{2\omega}{\omega+1}+o(1)} k_{trunc}^{\frac{1}{\omega+1}}$ time and $m^{\frac{4}{\omega+1}+o(1)} + m^{1+o(1)} k_{trunc}$ memory.*

Fast matrix multiplication has been used for previous triangle counting and enumeration algorithms [1, 3]. For our decomposition algorithm, we turn these into dynamic data structures to maintain triangle lists as edges of G are removed. See Section 5 for more precise costing.

1.2 Truss combinatorics in the context of cohesive subnetworks

The truss is an interesting but somewhat obscure combinatorial object, and this is the primary reason for studying its combinatorial properties. In addition, there is an important connection between the extremal bounds for trusses and its use as a heuristic for cohesive subnetworks.

Intuitively, a cohesive subnetwork on n nodes should be highly connected. Thus the edge counts should be quite high, perhaps on the order of n^2 . By contrast, our results in Section 2 show an example of a connected k -truss with edge counts as low as $\Theta(kn)$. It may be surprising that a connected k -truss can be extremely sparse despite its use in finding cohesive networks. These can be viewed as pathological cases where the truss heuristic does a poor job at discovering the underlying graph structure.

The extremal example consists of a series of $(k+2)$ -cliques connected at vertices. This is clearly a collection of multiple distinct cohesive subnetworks. In particular, it contains many subgraphs which are themselves k -trusses. This extremal example motivates us to define a stricter notion of *critical k -truss* as a heuristic for finding cohesive subnetworks: namely, a collection of edges which is a k -truss but which contains no smaller k -truss.

It seems reasonable that this restriction might give a more robust heuristic for cohesive subnetworks. Yet, we will show in Section 3 that this has similar extremal examples. The additional restriction of criticality does not significantly increase the minimum edge count. These results suggest that, despite their usefulness for real-world graphs, both the k -truss and the critical k -truss can be fallible heuristics.

1.3 Notation

We let n denote the number of vertices and m the number of edges of a graph $G = (V, E)$. The neighborhood of a vertex $v \in V$ is the set $N(v) = \{u : (u, v) \in E\}$.

$E\}$ and $d(v) = |N(v)|$ is the degree of v . We also define $N^+(v) = N(v) \cup \{v\}$. For simplicity, we assume throughout that G has no isolated vertices and $m \geq n/2$.

We define a triangle to be a set of three vertices (v_1, v_2, v_3) where edges $(v_1, v_2), (v_2, v_3), (v_3, v_1)$ are all present in G . We also write (e_1, e_2, e_3) for this triangle, where edges e_1, e_2, e_3 are given by $e_1 = (v_1, v_2), e_2 = (v_2, v_3), e_3 = (v_3, v_1)$. We define $\Delta(e)$ and $\Delta(v)$ to be the number of triangles containing an edge e or vertex v respectively. We say edge e_1 and e_2 are neighbors if they share a vertex.

For an integer t we define $[t] = \{1, \dots, t\}$. We assume there is some fixed, but arbitrary, indexing of the vertices, and we define $ID(v) \in [n]$ to be the identifier of vertex v .

For a vertex subset $U \subseteq V$, we define the *vertex-induced subgraph* $G[U]$ to be the graph on vertex set U and edge set $\{(u, v) \in E : u, v \in U\}$. For an edge set $L \subseteq E$, we define the *edge-induced subgraph* $G(L)$ to be the graph on edge set L and vertex set $\{v \in V : (u, v) \in L \text{ for some } u \in V\}$.

The complete graph on n vertices (n -clique) is denoted by K_n .

We will analyze some algorithms in terms of the *degeneracy* of graph G , which we denote $\delta(G)$. See Appendix A for further definitions and properties.

2 Minimum edge counts for the k -truss

We begin by collecting a few simple observations on the vertex counts in a k -truss.

Observation 4 *Any vertex v in a k -truss G must have degree at least $k + 1$. Furthermore, the graph $G[N^+(v)]$ has at least $\binom{k+1}{2}$ triangles and $\binom{k+2}{2}$ edges.*

Proof: Let e be any edge on v . This edge e has at least k triangles, giving k other edges incident on v . Each of these $k + 1$ edges has at least k triangles in $G[N^+(v)]$; furthermore, each such triangle is counted by at most two edges incident on v . So $\Delta(v)$ is at least $k(k + 1)/2$. Finally, the number of edges in $G[N^+(v)]$ is precisely $d(v) + \Delta(v)$, which is at least $(k + 1) + \binom{k+1}{2} = \binom{k+2}{2}$. \square

Observation 5 *The minimum number of vertices in a k -truss is exactly $k + 2$.*

Proof: Each vertex must have degree at least $k + 1$, thus each k -truss contains at least $k + 2$ vertices. On the other hand, K_{k+2} is a k -truss with $k + 2$ vertices. \square

The properties can be used to bound the clustering coefficient, a common measure of graph density. Formally, the *clustering coefficient* of a vertex v is defined as $cc(v) = \Delta(v) / \binom{d(v)}{2}$. Thus, $cc(v)$ is a real number in the range $[0, 1]$.

Observation 6 *In a k -truss, each vertex v has clustering coefficient $cc(v) \geq \binom{k+1}{2} / \binom{d(v)}{2}$. In particular, if $d(v) = k + 1$, then $cc(v) = 1$ and $G[N^+(v)]$ is a $(k + 2)$ -clique.*

We also get a simple bound on the edge counts.

Observation 7 *If graph G has m edges, then $\tau(e) \leq \sqrt{2m + 1/4} - 3/2 \leq \sqrt{2m}$ for all edges e .*

Proof: Suppose that edge $e = (u, v)$ has $\tau(e) = k$. Let G' denote the k -truss component containing edge e . By Observation 4, $G'[N^+(v)]$ contains at least $\binom{k+2}{2}$ edges, and hence $\binom{k+2}{2} \leq m$. \square

Let us define $M_{n,k}$ to be the minimum number of edges for a *connected* k -truss on n vertices. Our main result in this section is to estimate this quantity $M_{n,k}$, showing the following tight bounds:

Theorem 8 *For every $k \geq 1$ and $n \geq k + 2$, we have*

$$(n - 1)(1 + k/2) \leq M_{n,k} \leq n(1 + k/2) + \Theta(k^2).$$

Furthermore, if $n \equiv 1 \pmod{k + 1}$, then $M_{n,k} = (n - 1)(1 + k/2)$.

As an immediate corollary of Theorem 8, we also get a tight bound on triangle counts:

Corollary 9 *A connected k -truss on n vertices must contain at least $\frac{(n-1)(k+2)k}{6}$ triangles.*

Proof: The graph has at least $m \geq (n - 1)(1 + k/2)$ edges. Each edge has at least k triangles; since each triangle is incident to three edges, this implies there are at least $mk/3$ triangles. \square

For the upper bound of Theorem 8, we use a construction based on vertex contraction. Namely, for a pair of graphs G, H , define $G * H$ to be the graph obtained by contracting an arbitrary vertex of G to an arbitrary vertex of H . The resulting graph $G * H$ has $|V(G)| + |V(H)| - 1$ vertices and $|E(G)| + |E(H)|$ edges.

Now when $n \equiv 1 \pmod{k + 1}$, we get the upper bound $M_{n,k} \leq (n - 1)(1 + k/2)$ by taking $G = A_1 * \dots * A_s$ where A_1, \dots, A_s are copies of K_{k+2} and $s = \frac{n-1}{k+1}$. This also shows that the bound of Corollary 9 is tight in this case.

A slightly modified construction shows the upper bound $M_{n,k} \leq n(1 + k/2) + O(k^2)$ for arbitrary values of n . Namely, let $n = s(k + 1) + r$ for integer r in the range $k + 2 \leq r < 2k + 3$, and consider $G = A_1 * \dots * A_s * B$, where A_1, \dots, A_s are copies of K_{k+2} and B is a copy of K_r . Noting that $r \leq 2k$, we calculate the edge count of G as $m = s\binom{k+2}{2} + \binom{r}{2} \leq n(1 + k/2) + O(k^2)$.

We now turn to prove the lower bound of Theorem 8. Let G be a connected k -truss. Since G is connected, it has a spanning tree T , which we may take to be a rooted tree (with an arbitrary root). For a triangle t of G , we say that t is *single-tree* if exactly one edge is in T , otherwise it is *double-tree*. (If all three edges were in T , this would be a cycle on T .) If an edge $e \in E - T$ participates in a double-tree triangle, we say that e is *double-tree-compatible* otherwise it is *double-tree-incompatible*.

Proposition 10 *Let (u, v, w) be a single-tree triangle where $(u, v) \in T$. Then either (u, w) or (v, w) is double-tree-incompatible.*

Proof: If $(u, w), (v, w)$ are both double-tree-compatible, this would imply that there are vertices x, y with $(u, x), (x, w), (v, y), (w, y) \in T$. Since $(u, w) \in E - T$, we must have $y \neq u$. But then u, x, w, y, v, u is a cycle on the tree T , which is a contradiction. \square

Our proof strategy will be to construct a function $F : T \times [k] \rightarrow E - T$, and then argue that F is 2-to-1. This will show that $E - T$ has cardinality at least $|T|k/2$, and so $|E| \geq |T| + |T|k/2 = (n - 1)(1 + k/2)$ which is the lower bound we need to show.

To define F , consider an edge $e = (x, y) \in T$, where y is a T -child of x . Arbitrarily select k triangles $t_{e,1}, \dots, t_{e,k}$ involving e . For $i = 1, \dots, k$, we define $F(e, i)$ as follows:

- If $t_{e,i}$ is double-tree, then $F(e, i)$ is the unique off-tree edge of $t_{e,i}$.
- If $t_{e,i}$ is single-tree and exactly one off-tree edge f of $t_{e,i}$ is double-tree-incompatible, then $F(e, i) = f$.
- If $t_{e,i}$ is single-tree and both off-tree edges of $t_{e,i}$ are double-tree-incompatible, then $F(e, i)$ is the off-tree edge of $t_{e,i}$ containing y .

In light of Proposition 10, this fully defines the function F . We refer to an edge e as a preimage of f if $F(e, i) = f$ for some index $i \in [k]$. Since any triangle is determined by two edges, such index i is uniquely determined by e and f .

Proposition 11 *Suppose that edge $f = (u, v) \in E - T$ is double-tree-incompatible, and (u, x) is a preimage of f where x is a T -child of u . Then edge (x, v) must be double-tree-compatible. Furthermore, f does not have a preimage (u, y) where y is a T -child of u distinct from x .*

Proof: For the first result, suppose that $F((u, x), i) = f$ and consider the triangle $t_{(u,x),i}$. Since f is double-tree-incompatible, necessarily $t_{(u,x),i}$ is single-tree. If the other edge (x, v) in this triangle were also double-tree-incompatible, then since x is a T -child of u we would have $F((u, x), i) = (x, v) \neq f$, a contradiction.

For the second result, suppose that (u, x) and (u, y) are preimages of f . By the argument in the preceding paragraph the edges (x, v) and (y, v) are both double-tree-compatible. So there are vertices r, s with $(x, r), (r, v), (y, s), (s, v) \in T$. One can then check that v, s, y, u, x, r, v is a cycle on T , a contradiction. \square

Proposition 12 *Every edge $f = (u, v) \in E - T$ has at most two preimages under F .*

Proof: Case I: f is double-tree-compatible. The only possible preimages to f would come from double-tree triangles in which f is the unique off-tree edge. There can only be a single such triangle; for, if (u, v, x) and (u, v, y) were

two such triangles, then u, x, v, y, u would be a cycle on T . So the only possible preimages of f are the two tree-edges in this double-tree triangle.

Case II: f is double-tree-incompatible. We first claim that f cannot have three preimages $(u, x), (u, y), (u, z)$. For, in this case, at least two vertices, say without loss of generality x, y , must be T -children of u . This is ruled out by Proposition 11.

So suppose that f has three preimages $(u, x), (u, y), (v, z)$. By Proposition 11, it cannot be that both x and y are T -children of u . So assume without loss of generality that y is the T -parent of u and x is a T -child of u .

By Proposition 11, the edge (x, v) is double-tree-compatible. So there is some vertex r such that $(x, r), (r, v) \in T$. Since u is the T -parent of y , this implies that x is the T -parent of r and r is the T -parent of v and v is the T -parent of z .

Since (v, z) is a preimage of f and v is the T -parent of z , by Proposition 11 the edge (z, u) must be double-tree-compatible. So we have $(z, s), (s, u) \in T$ for some vertex s . It can be seen that u, x, r, v, z, s, u is a cycle on T , a contradiction. \square

This completes the proof of the lower bound of Theorem 8.

3 Critical connectivity of the k -truss

We define a *critical k -truss* as follows:

Definition 13 (Critical k -truss) *Graph G is a critical k -truss if G is a k -truss, but $G(E')$ is not a k -truss for any non-empty $E' \subsetneq E$.*

It is clear that critical k -trusses are connected. The extremal graphs $A_1 * \dots * A_s$ for the upper bound in Theorem 8 are far from critical, as each subgraph A_i is a k -truss. Arguably, critical k -trusses are more relevant for community detection — if a graph contains smaller k -trusses, then it is a conglomeration of communities rather than a single cohesive subnetwork of its own.

We begin with some simple observations.

Observation 14 *There is no critical k -truss with exactly $k + 3$ vertices. In a critical k -truss with more than $k + 3$ vertices, every vertex has degree at least $k + 2$.*

Proof: By Observation 4, each vertex v has $d(v) \geq k + 1$. By Observation 6, if $d(v) = k + 1$, then the neighborhood of v is a copy of K_{k+2} , which is impossible in a critical k -truss. Furthermore, if $n = k + 3$ and $d(v) = k + 2$ for all vertices v , then G is a copy of K_{k+3} , and in particular G contains a copy of the k -truss K_{k+2} . \square

Observation 15 *The graph K_3 is the only critical 1-truss.*

Proof: Suppose that G is a critical 1-truss, and let e_1 be an edge of G . So e_1 is contained in some triangle with edges (e_1, e_2, e_3) . Then $G(\{e_1, e_2, e_3\})$ is a 1-truss. \square

Let us define $M_{n,k}^*$ to be the minimum number of edges in a critical n -node k -truss; we have $M_{n,k}^* = \infty$ if no such critical k -truss exists. To avoid edge cases covered by Observations 14 and 15, we assume that $k \geq 2$ and $n \geq k + 4$.

For $k = 2$, we can compute the precise value of $M_{n,k}^*$:

Theorem 16 *For all $n \geq 6$ we have $M_{n,2}^* = 3n - 6$.*

Proof: For the upper bound, consider the graph consisting of a cycle C of length $n - 2$, plus two new vertices x_1, x_2 with edges to every vertex in C . This is a critical 2-truss with n vertices and $3n - 6$ edges.

For the lower bound, let $G = (V, E)$ be a critical 2-truss with m edges and n vertices. Define \mathcal{V} to be the vector space of all functions from E to the finite field $GF(2)$, and for any triangle t of G we define χ_t to be the characteristic function of t , i.e. $\chi_t(e) = 1$ if and only if $e \in t$.

Select U to be any smallest set of triangles in G with the property that every edge is in at least two triangles of U . This is well-defined since G is a 2-truss. We claim that for every proper subset $W \subsetneq U$, the sum $\sum_{t \in W} \chi_t$ is not identically zero.

For, suppose it is, and let L denote the set of edges appearing in the triangles $t \in W$. For any edge $e \in L$, we then have $\sum_{t \in W} \chi_t(e) = 0$ in the field $GF(2)$. Since the sum is taken modulo two, there are at least two triangles t_1, t_2 in $G(L)$ containing e (by definition of L , there is at least one). Thus $G(L)$ is a 2-truss. Since G is a critical 2-truss, we must have $L = E$. Thus, every edge of G is covered by at least two triangles in W . This contradicts minimality of U .

We have shown that there is at most one linear dependency among the functions χ_t for $t \in U$ (namely, corresponding to $W = U$). A standard result (see e.g. [8]) is that the vector space \mathcal{V} has dimension $m - n + 1$. This implies that $|U| - 1 \leq m - n + 1$. On the other hand, every edge is in at least two triangles in U and so $3|U| \geq 2m$. Putting these inequalities together gives $m \geq 3n - 6$. \square

Our main result in this section is to estimate $M_{n,k}^*$, showing a result similar to Theorem 8. Specifically, we will show that

$$M_{n,k}^* = (\Theta(1) + k/2)n + \Theta(k^2).$$

(A more precise estimate is shown in Theorem 18.)

Lemma 1 *For any integers $k \geq 2, n \geq k + 4$ we have*

$$M_{n+2,k+2}^* \leq M_{n,k}^* + 2n, \quad M_{n+1,k+1}^* \leq M_{n,k}^* + n.$$

Proof: To show the first bound, let $G = (V, E)$ be a critical k -truss with n vertices and $m = M_{n,k}^*$ edges. Create a new graph G' , which has all the

vertices and edges of G , plus two new vertices x_1, x_2 . We add a new set F of edges connecting x_1, x_2 to the previous vertices, where F is chosen so that $G' = (V \cup \{x_1, x_2\}, E \cup F)$ has the properties that (i) G' is a $(k+2)$ -truss and (ii) F is inclusion-wise minimal with this property.

To show this is well-defined, we need to show that property (i) is satisfied when F is the set of all possible edges between the new vertices and the old ones. In this case, G' has no isolated vertices, as G is a k -truss and has none. Also, any edge $e \in E$ has k triangles from G and two new triangles in F . Finally, for each edge $e = (x_i, v)$ where $v \in G$, there is a triangle in G' for each neighbor of v in G . By Observation 14, this implies that $\Delta(e) \geq k+2$.

Clearly G' has $n+2$ vertices and has $m+|F| \leq m+2n$ edges. To show G' is critical, suppose there are edge subsets $E' \subseteq E, F' \subseteq F$ such that $G'(E' \cup F')$ is a $(k+2)$ -truss. Then $G(E')$ must be a k -truss, as removing the edges incident to x_1, x_2 can only remove 2 triangles per edge. Since G is critical, this implies $E' = \emptyset$ or $E' = E$. If $E' = \emptyset$, then $G(F')$ must be a k -truss. However, $G(F)$ itself has no triangles, and so we must have $F' = \emptyset$, i.e. $E' \cup F' = \emptyset$. On the other hand, if $E' = E$, then $G'(E \cup F')$ is a $(k+2)$ -truss; by definition of F , this implies that $F' = F$ and hence $E' \cup F' = E \cup F$.

The second bound is essentially identical, except that we add only a single vertex instead of two vertices. \square

Proposition 17 *For $n \geq k+4$, we have $M_{n,k}^* \leq n(k+1) - k^2/2 - 2k + 1/2$.*

Proof: Suppose first that k is even. From Theorem 16, we have $M_{i,2}^* = 3i - 6$ for any integer $i \geq 6$. By repeated applications of Lemma 1 we get

$$M_{i+2j,2+2j}^* \leq M_{i,2}^* + 2i + 2(i+2) + \cdots + 2(i+2j-2) \leq 3i - 6 + 2j(j+i-1).$$

Setting $j = (k-2)/2$ and $i = n+2-k$ gives $M_{n,k}^* \leq n(k+1) - k^2/2 - 2k$.

If k is odd, then by Lemma 1 we have $M_{n,k}^* \leq M_{n-1,k-1}^* + (n-1)$. Since $k-1$ is even, the argument of the preceding paragraph gives $M_{n-1,k-1}^* \leq (n-1)k - (k-1)^2/2 - 2(k-1)$, and so

$$M_{n,k}^* \leq (n-1)k - \frac{(k-1)^2}{2} - 2(k-1) + (n-1) = n(k+1) - \frac{k^2}{2} - 2k + 1/2.$$

\square

We are now ready for the main construction to show the upper bound on $M_{n,k}^*$. This uses a type of graph embedding in the torus; we describe the construction in more detail in Appendix B.

Lemma 2 *Suppose there exists a graph T embedded in the torus with r faces, where each edge appears in two distinct faces, and each face F has $s_F \geq 4$ edges. Let $g = \sum_F s_F$. Then for $k \geq 3$ there is a critical k -truss with $r(k-2) + g/2$ vertices and $r \binom{k-1}{2} + (k-1 + 1/2)g$ edges.*

Proof: For each face F , let us define $T(F)$ to be the corresponding subgraph of T . We form the graph G by starting with T . For each face F , we insert a

copy of K_{k-1} , which we denote by $C(F)$. We add an edge from every vertex of $C(F)$ to every vertex in $T(F)$, and we let $H(F)$ denote these edges. We also define $G(F) = C(F) \cup H(F) \cup T(F)$. See Figure 1.




Figure 1: This shows $G(F)$ for a single face F (here, the outside square), for $k = 4$. The triangle inside the square is $C(F)$. The square on the outside is $T(F)$. Each vertex of $C(F)$ is connected to each vertex of face $T(F)$, via an edge of $H(F)$.

We first compute the number of vertices and edges in G . First, since each edge of T appears in exactly two faces, T has $\frac{1}{2} \sum_F s_F = g/2$ edges. By Euler’s formula, T therefore has $g/2 - r$ vertices. Each face F of T gives $k - 1$ vertices and $\binom{k-1}{2}$ edges in $C(F)$ and $(k - 1)s_F$ edges in $H(F)$. So G has $(g/2 - r) + r(k - 1) = r(k - 2) + g/2$ vertices and has $g/2 + \sum_F (\binom{k-1}{2} + (k - 1)s_F) = r\binom{k-1}{2} + (k - 1 + 1/2)g$ edges, as we have claimed.

Let us check that G is a k -truss. For an edge e of T , the two corresponding faces include copies of K_{k-1} , so e has at least $2k - 2 \geq k$ triangles. An edge of $C(F)$ has $k - 3$ triangles within $C(F)$ and at least $s_F \geq 4$ triangles from vertices of $T(F)$, for a total of $k + 1$ triangles. An edge of $H(F)$ has 2 triangles in $T(F)$ and $k - 2$ triangles within $C(F)$, a total of k triangles.

We next need to show that G is critical. Suppose that $G(L)$ is a k -truss for $L \subseteq E$; we need to show that $L = E$ or $L = \emptyset$. We do this in four stages.

- (a) We first claim that, for every face F , either L contains all the edges in $H(F)$, or none of them. For, suppose L omits an edge $e = (u, v) \in H(F)$, where $u \in T(F)$ and $v \in C(F)$. Every other edge e' in $H(F)$ incident to u would then have at most $k - 1$ triangles in L , and so also $e' \notin L$. Thus, L contains no edges of $H(F)$ incident to u .

Next, consider a vertex u' adjacent to u in $T(F)$. Any edge $e' \in H(F)$ incident on u' can now have at most $k - 1$ triangles in L , since one of its triangles in $G(F)$ used an edge of $H(F)$ incident on u . Thus, L must omit all the edges of $H(F)$ incident on u' as well.

Continuing this way around the cycle $T(F)$, we see that $H(F) \cap L = \emptyset$.

- (b) We next claim that for every face F , if L omits any edge $e \in G(F)$ then it omits all the edges in $H(F)$. For, note that this edge e participates in

a triangle with some edge $e' \in H(F)$. Thus e' has at most $k - 1$ triangles in L and must be omitted. By part (a), this implies that all the edges in $H(F)$ are omitted.

- (c) We next claim that for every face F , either L contains all the edges in $G(F)$, or none of them. From parts (a) and (b), we see that if L omits any such edge, then it omits all the edges in $H(F)$. This implies that each edge $e \in C(F)$ has at most $k - 2$ triangles in L , and so $e \notin L$. Similarly, each edge $e \in T(F)$ has at most $k - 1$ triangles in L , coming from the graph $C(F')$ where F' is the other face touching e ; thus also $e \notin L$.
- (d) Finally, we claim that either $L = E$ or $L = \emptyset$. For, suppose that G omits an edge $e \in G(F)$ for some face F . By part (c), L omits every edge in $G(F)$. Now note that if F' touches F in T , then $G(F')$ omits an edge, namely, the common edge of F and F' . By part (c) this implies $G(F') \cap L = \emptyset$. Since T is connected, continuing this way around T we see that $L = \emptyset$. \square

We now get our final estimate for $M_{n,k}^*$:

Theorem 18 *For $k \geq 2$ and $n \geq k + 4$, we have*

$$n(k/2 + 1) \leq M_{n,k}^* \leq n(k/2 + 5/2 - 1/k) + O(k^2) \leq nk/2 + O(n + k^2).$$

Proof: The lower bound is an immediate consequence of Observation 14. Lemma 17 shows that $M_{n,k}^* \leq n(k + 1) - (k^2/2 + 2k - 1/2)$; when $n \leq 2k$, this quantity is $O(k^2)$ and we are done. Similarly, Theorem 16 already shows this result when $k = 2$.

So suppose that $k \geq 3$ and $n > 2k$, and we want to show the upper bound on $M_{n,k}^*$. We write $n = ik + j$, where $j \in \{0, \dots, k - 1\}$ and $i \geq 2$. As we show in Lemma 3, for these parameters there is a toroidal embedding satisfying the conditions of Lemma 2, whose faces consist of two $(j + 4)$ -cycles and $i - 2$ four-cycles.

The sum of edge counts for this embedding is given by $\sum_F s_F = 4(i - 2) + 2(j + 4) = 4i + 2j$. By Lemma 2, there is a critical k -truss with $ik + j = n$ vertices and $m = i \binom{k-1}{2} + (k - 1 + 1/2)(4i + 2j)$ edges. As $j \leq k - 1$, we have

$$m = \frac{(3k - 1)(k - 2)j + (k(k + 5) - 2)n}{2k} \leq (k/2 + 5/2 - 1/k)n + O(k^2).$$

\square

4 Practical truss decomposition algorithm

We now present Algorithm 2 to compute the trussness $\tau(e)$ of every edge e in a graph. Recall that $\tau(e)$ is the maximal value k such that e is in a k -truss and that, after this has been computed, we can easily compute the k -truss-components of G by discarding all edges e with $\tau(e) < k$ and finding the connected components of the resulting graph.

Our algorithm here is inspired by the Wang & Cheng [21] and Huang et al. [12] algorithms, but uses simpler data structures. To explain, let us provide a brief summary of their algorithms. At each stage, they find an edge with the fewest incident triangles in the remaining graph, remove this edge, and then update the triangle counts for all neighboring edges. If an edge e has k triangles when it is removed, then it has $\tau(e) = k + 1$. This process continues until all the edges have been removed from the graph.

While this is conceptually simple, it can be cumbersome to implement in practice. In particular, this requires relatively heavy-weight data structures to maintain the edges sorted in increasing order of triangle counts. For example, Wang & Cheng use a method of [2] based on a four-level hierarchy of associative arrays. (See [5] for related data structures.) While this could certainly be implemented, it is also clearly more complex than primitive data structures such as arrays.

The key idea of our new algorithm is that, instead of sorting the edges by triangle count, we only maintain an unordered list of edges whose triangle count is below a given threshold. We can afford to periodically re-scan the graph for edges with few triangles.

Let us begin by recalling the standard simple algorithm to *enumerate* the triangles in G :

Algorithm 1

```

1: for each edge  $e = (u, v) \in E$  with  $d(u) \leq d(v)$  do
2:   for all  $w \in N(u)$  do
3:     if  $(v, w) \in E$  and  $\text{ID}(w) > \max(\text{ID}(u), \text{ID}(v))$  then
4:       Output triangle  $(u, v, w)$ 
5:     end if
6:   end for
7: end for

```

The following bound is immediate:

Observation 19 *Algorithm 1 runs in $O(m\bar{\delta}(G))$ time and outputs every triangle exactly once.*

Instead of statically listing triangles, as in Algorithm 1, our truss decomposition Algorithm 2 keeps track of them as edges are removed from the graph. The main data structure is the array Δ , which stores the number of triangles in the residual graph containing any given edge e . We also use a sentinel value denoted \emptyset to indicate that edge e is no longer present in the residual graph. Other data structures include a stack S and a list L of edges which need to be processed.

Algorithm 2

```

1: Initialize an empty stack  $S$ , and initialize an edge-list  $L = E$ 
2: Using Algorithm 1, compute the triangle counts  $\Delta(e)$  for all  $e \in E$ 
3: for  $k = 1, \dots, \sqrt{2m}$  do
4:   for all edges  $e \in L$  do
5:     if  $\Delta(e) = k - 1$  then push  $e$  onto  $S$ 
6:     if  $\Delta(e) = \emptyset$  then remove  $e$  from  $L$ 
7:   end for
8:   while  $S$  is non-empty do
9:     Pop edge  $e$  from  $S$ . Let  $e = (u, v)$  such that  $d(u) \leq d(v)$ 
10:     $\Delta(e) \leftarrow \emptyset$ 
11:    for all  $w \in N(u)$  do
12:      if  $(v, w) \in E$  and  $\Delta(u, w) \neq \emptyset$  and  $\Delta(v, w) \neq \emptyset$  then
13:         $\Delta(u, w) \leftarrow \Delta(u, w) - 1$ . if  $\Delta(u, w) = k - 1$  then push  $(u, w)$  to  $S$ 
14:         $\Delta(v, w) \leftarrow \Delta(v, w) - 1$ . if  $\Delta(v, w) = k - 1$  then push  $(v, w)$  to  $S$ 
15:      end if
16:    end for
17:    Output  $\tau(e) = k - 1$ 
18:  end while
19: end for

```

We refer to each iteration of the loop at line (3) of Algorithm 2 as *round* k .

Let us first remark on the implementation of L . At first glance, it would appear to require a linked list, since we need to remove edges e from L in line (6). However, we only delete elements while we are iterating over the entire list, and so we can instead store L as a simple array. When we want to delete e from L , we just swap it to the end of the buffer instead.

We also note that round $k = 1$ can be simplified: for an edge e with $\Delta(e) = 0$, we can immediately output $\tau(e) = 0$ and we do not need to push e onto the stack. This optimization can be useful for graphs which have a relatively small number of triangles.

We now show Algorithm 2 has the claimed complexity and correctly computes the values $\tau(e)$. At any given point in the algorithm, we define the set of edges $e \in E$ with $\Delta(e) \neq \emptyset$ as the *residual edges* and denote them by R .

Proposition 20 *Any edge e gets added to S at most once over the entire lifetime of Algorithm 2.*

Proof: If e gets added to S in round k , then line (10) ensures that $\Delta(e) = \emptyset$ by the end of round k , so that e never gets added in subsequent rounds. Also, an edge e can be added to S at most once in a given round, since before adding e to S in line (13) or (14) we first decrement $\Delta(e)$. \square

Proposition 21 *Algorithm 2 maintains the following loop invariants on the data structures:*

- (a) *The array Δ correctly records triangle counts for the graph $G(R)$.*

(b) For every edge $e \in R$, either $\Delta(e) \geq k$ or $e \in S$.

(c) Every edge $e \in S$ satisfies $e \in R$ and $\Delta(e) < k$.

Proof: Since $R = E$ initially, these properties are satisfied at line (2). Now suppose these properties are satisfied at the end of round $k - 1$; we want to show they remain satisfied in round k as well.

From property (b), we know that $\Delta(e) \geq k - 1$ for all edges $e \in R$ at the beginning of round k . Lines (4) — (7) maintain the properties. Now, consider the state just before line (9), where we are processing some edge $e \in S$. Property (c) is maintained since e does not appear elsewhere in S . For property (a), since e gets removed from R at line (10), the counts in Δ must be updated for all triangles of R involving e . Line (13) — (14) are reached for each such triangle and $\Delta(f)$ is indeed properly adjusted for each edge f neighboring e . Finally for property (b), note that if any such edge f was placed into S , it would necessarily have $f \in R$ and $\Delta(f) = k - 1$. \square

Proposition 22 *An edge e is in R at the end of round k if and only if $\tau(e) \geq k$.*

Proof: By Proposition 21(b), every edge $f \in R$ has either $\Delta(f) \geq k$ or $f \in S$. Since S is empty after line (18), this implies that each $e \in R$ has at least k triangles in $G(R)$, and hence $\tau(e) \geq k$.

Conversely, suppose that some edge e with $\tau(e) \geq k$ gets removed from R before round $k + 1$. Let f be the first such edge removed and let E' denote the k -truss-component containing f . Consider the state before line (9) when f is removed. Since f is the first such removed edge, all the other edges in E' remain in R and so f still has at least k triangles in R . This contradicts Proposition 21(c). \square

Theorem 23 *Algorithm 2 correctly computes $\tau(e)$ for all edges e .*

Proof: Suppose that line (17) outputs $\tau(e) = k - 1$ for some edge e . By Proposition 21(c), this implies that e must have been in R at the beginning of round k , and hence by Proposition 22 we have $\tau(e) \geq k - 1$. On the other hand, e got removed from R at line (9), and so by Proposition 22 we have $\tau(e) < k$. Thus, $\tau(e)$ is indeed $k - 1$.

Note that the termination condition $k = \sqrt{2m}$ of line (3) follows from Observation 7. \square

Theorem 24 *Algorithm 2 runs in $O(m\bar{\delta}(G))$ time and $O(m)$ memory.*

Proof: The array Δ can be referenced or updated in $O(1)$ time. Bearing in mind our remarks about implementing L as an array, all operations on S and L take $O(1)$ time. Observation 19 shows that line (2) takes $O(m\bar{\delta}(G))$ time and $O(m)$ memory. The data structures Δ, S, L are indexed by edges, and so overall take $O(m)$ memory.

Next, let L_k denote the value of list L at the beginning of round k . By Proposition 22, L_k consists solely of edges with $\tau(e) \geq k - 1$. The runtime of

the loop at lines (4) — (7) at round k is linear in the length of L_k , and so the total work over all rounds is a constant factor times

$$\sum_k |L_k| \leq \sum_k |\{e \in E : \tau(e) \geq k - 1\}| \leq \sum_e (\tau(e) + 1).$$

For any edge $e = (u, v)$ with $\tau(e) = k$, we have $\Delta(e) \geq k$ so clearly $d(u) \geq k$ and $d(v) \geq k$. Thus $\tau(e) \leq k \leq \min(d(u), d(v))$, and so

$$\sum_k |L_k| \leq \sum_e 1 + \min(d(u), d(v)) \leq m + m\bar{\delta}(G) \leq O(m\bar{\delta}(G)).$$

Finally, consider lines (8) — (18). By Proposition 21(c), every edge $e = (u, v)$ appears at most once in S . The enumeration in line (11) takes $O(\min(d(u), d(v)))$ time. So again, the total work for this loop is $\sum_{(u,v) \in E} \min(d(u), d(v)) = m\bar{\delta}(G)$. \square

5 Truncated truss decomposition using matrix multiplication

We now develop a theoretically more efficient algorithm for truncated truss decomposition up to some given bound k_{trunc} . This allows us to compute the k -truss-components for any value $k \leq k_{\text{trunc}}$, by discarding edges with $\tau(e) < k$ and running depth-first search on the resulting graph.

The algorithm here, like Algorithm 2, is based on removing edges and updating triangle counts. The crux of the algorithm is the following observation. When an edge is removed, it must be incident on fewer than k triangles in the residual graph. If we could enumerate these triangles efficiently, then we could potentially update their edges in only $O(k)$ time.

There is a long history of fast matrix multiplication for triangle enumeration [1, 22, 3]. These prior algorithms are inherently static: they treat the graph G as a fixed input, and the output is the triangle count or list of triangles. To compute the values $\tau(e)$, by contrast, we must *dynamically* maintain the triangle lists as edges are removed. We develop an algorithm based on a methods of [11, 3] which reduce triangle enumeration to finding witnesses for boolean matrix multiplication.

Our algorithm uses multiplication of rectangular matrices; see [13] for further details. To measure the cost of this operation, we define the function $\Gamma : [0, 1] \rightarrow \mathbb{R}_+$ as:

$$\Gamma(s) = \inf\{p \in \mathbb{R} : \exists \text{ an } O(t^p)\text{-time algorithm for } t \times \lceil t^s \rceil \text{ by } \lceil t^s \rceil \times t \text{ matrix multiplication}\}$$

It is shown in [10] that $\Gamma(s) = 2$ for $s \leq 0.31$, and it is conjectured that $\Gamma(s) = 2$ for all $s \in [0, 1]$. The value $\Gamma(1)$ is also known as the *linear-algebra constant* ω . By standard reductions, there is a single randomized algorithm to multiply $t \times \lceil t^s \rceil$ by $\lceil t^s \rceil \times t$ matrices in time $t^{\Gamma(s)+o(1)}$ for all s .

5.1 Algorithm description

We begin by generating $L = 10k_{\text{trunc}} \log n$ random vertex sets X_1, \dots, X_L , where each vertex goes into each X_i independently with probability $q = \frac{1}{k_{\text{trunc}}}$. We also define $Y_v = \{\ell \in [L] : v \in X_\ell\}$ for each vertex v . The algorithm maintains two data structures corresponding to these sets, keeping track of the following information for each edge $e = (u, v)$:

1. The total triangle count $\Delta(e)$
2. $S(e, \ell) = \sum_{w \in N(u) \cap N(v) \cap X_\ell} \text{ID}(w)$ for each $\ell \in [L]$

An outline is shown in Algorithm 3. We provide more detail below on the implementation and runtimes of the steps. The proof of correctness is essentially the same as Theorem 23, so we do not provide it here.

Algorithm 3

- 1: Generate the vertex subsets $X_1, \dots, X_L \subseteq V$ and corresponding sets Y_v
 - 2: Compute $S(e, \ell)$ and $\Delta(e)$ for every edge e and every $\ell \in [L]$
 - 3: **for** $k = 1, \dots, k_{\text{trunc}}$ **do**
 - 4: **while** $\Delta(e) < k$ for some edge $e \in G$ **do**
 - 5: Output $\tau(e) = k - 1$
 - 6: Enumerate all the triangles involving e in the residual graph G
 - 7: Remove e from G and appropriately update S, Δ
 - 8: **end while**
 - 9: **end for**
 - 10: **for all** remaining edges in G **do** Output $\tau(e) \geq k_{\text{trunc}}$
-

We write $k_{\text{trunc}} = m^a$; since $\tau(e) \leq \sqrt{2m}$ for all edges e , we assume that $k_{\text{trunc}} \leq \sqrt{2m}$ and hence $a \in [0, \frac{1}{2} + o(1)]$. We also use the notation $\tilde{O}(x) = x \times \text{polylog}(x)$ for any quantity x .

Line (1) With standard sorting methods, this takes $\tilde{O}(Ln)$ time and memory in the worst case.

Line (2) We use a method based on [22] for this step. The calculation of $\Delta(e)$ is similar to $S(e, \ell)$ so we only show the latter.

We divide the vertices into two classes: the *heavy* vertices v (if $d(v) > m^{1-b}$) and the *light* vertices (if $d(v) \leq m^{1-b}$). Here, b is a parameter in the range $[a, 1]$ we will set later. We let H denote the set of heavy vertices, and observe that $|H| \leq 2m/m^{1-b} = 2m^b$.

We first compute the contribution to S coming from triangles with at least one light vertex. We do this by looping over light vertices and pairs of their incident edges. For each such triangle $t = (u, v, w)$, we update $S((u, v), \ell) \leftarrow S((u, v), \ell) + \text{ID}(w)$ for each $\ell \in Y_v$; we similarly update values $S((u, w), \ell)$ and $S((v, w), \ell)$. This simple algorithm has expected runtime $m^{2-b+o(1)}$.

We next consider the triangles with only heavy vertices. For each index ℓ we form matrices B and B' of dimensions $|H| \times |H \cap X_\ell|$ as follows:

$$B(v, w) = \begin{cases} 1 & \text{if } (v, w) \in E \\ 0 & \text{if } (v, w) \notin E \end{cases} \quad B'(v, w) = \begin{cases} \text{ID}(w) & \text{if } (v, w) \in E \\ 0 & \text{if } (v, w) \notin E \end{cases}$$

For an edge $e = (u, v)$ on heavy vertices u, v we compute the contribution to $S(e, \ell)$ from heavy vertices w as:

$$\sum_{\substack{w \in N(u) \cap N(v) \\ w \in H \cap X_\ell}} \text{ID}(w) = (B^\top B')_{u,v}.$$

This matrix multiplication $B^\top B'$ can be computed in $|H|^{\Gamma(\frac{\log |H \cap X_\ell|}{\log |H|}) + o(1)}$ time and $|H|^2 = O(m^{2b})$ memory. As $|H \cap X_\ell|$ is a binomial random variable with mean $q|H| \leq 2m^{b-a}$, the total expected time for this over all indices ℓ is $m^{a+b\Gamma(1-a/b)+o(1)}$.

Line (4) We use similar data structures to Algorithm 2 for this. For each round k , we check if $\Delta(e) < k$ for each edge e . Whenever we process an edge e in lines (4) — (9), we check if $\Delta(f) = k - 1$ for some neighboring edge f and, if so, add f to a stack.

Line (6) We use the following primary algorithm to enumerate triangles containing edge $e = (u, v)$: for each $\ell \in [L]$ with $S(e, \ell) \leq n$, let x_ℓ be the vertex with $\text{ID}(x_\ell) = S(e, \ell)$ and test if there are edges $(u, x_\ell), (v, x_\ell)$ in the residual graph. If so, then output a triangle (u, v, x_ℓ) .

Let us define A to be the set of vertices w where the edges (u, w) and (v, w) remain in the current residual graph, and let $W \subseteq A$ denote the set of all triangle vertices x_ℓ enumerated in the primary algorithm. Note that we also maintain the residual triangle count $\Delta(e) = |A|$. If $|W|$ is equal to the known value $\Delta(e)$, then we have found all the triangles containing e . If $|W| < \Delta(e)$, then use a second, slower, fallback option: we simply loop over all vertices in V .

We argue now that $|W| = \Delta(e)$ with high probability. Note first that the update sequence in Algorithm 3 is not affected by the randomness in the sets X_ℓ . So A can be regarded as a deterministic quantity. For an arbitrary vertex $w \in A$, observe that if $A \cap X_\ell = \{w\}$ for some index ℓ , then $S(e, \ell) = \text{ID}(w)$ and hence w will go into W . For each ℓ , we have $A \cap X_\ell = \{w\}$ with probability $q(1-q)^{|A|-1}$. As $|A| \leq k \leq k_{\text{trunc}}$ and $q = \frac{1}{k_{\text{trunc}}}$ and $L = 10k_{\text{trunc}} \log n$, this shows that

$$\Pr(w \notin W) \leq (1 - q(1 - q)^{k_{\text{trunc}}})^L \leq n^{-3.6}.$$

By a union bound over $w \in A$, this implies that $\Pr(W = A) \geq 1 - n^{-2.6}$. The primary algorithm of generating the set W takes $O(L) = O(k_{\text{trunc}} \log n)$ time. The fallback algorithm takes $O(n)$ time in the worst case, so it contributes expected runtime at most $n^{-2.6} \times O(n) = o(1)$.

Line (7) Suppose we are removing edge $e = (u, v)$, and we have enumerated all the triangles containing e in the residual graph. For each such triangle (u, v, w) , we update Δ by decrementing $\Delta(u, w)$ and $\Delta(v, w)$. We update S by

setting $S((u, w), \ell) \leftarrow S((u, w), \ell) - \text{ID}(v)$ for each $\ell \in Y_v$ and $S((v, w), \ell) \leftarrow S((v, w), \ell) - \text{ID}(u)$ for each $\ell \in Y_u$.

Since there are at most $k \leq k_{\text{trunc}}$ triangles, and since the sets Y_u and Y_v have expected size $O(\log n)$, the overall expected time for line (8) is $O(k_{\text{trunc}} \log n)$.

5.2 Putting it together

We can now calculate the overall complexity of the algorithm.

Theorem 25 *For any real number r satisfying $\Gamma\left(\frac{2-a(2+r)}{2-a}\right) \leq r \leq 2/a - 2$, Algorithm 3 can be implemented to run in $m^{\frac{2r+a}{r+1}+o(1)}$ expected time and $\tilde{O}(m^{\frac{4-2a}{r+1}} + m^{1+a})$ memory.*

Proof: The algorithm uses $\tilde{O}(m^{1+a})$ memory to store the arrays $S(e, \ell)$ and $\Delta(e)$. Line (1) takes $\tilde{O}(Ln)$ time, which is at most $m^{1+a+o(1)}$.

Lines (5) — (8) are executed at most once per edge, so their total expected time is at most $\tilde{O}(mk_{\text{trunc}}) \leq m^{1+a+o(1)}$. Since $r \geq \Gamma\left(\frac{2-a(2+r)}{2-a}\right) \geq 2$ and $a \leq \frac{1}{2} + o(1)$, this is at most $m^{\frac{2r+a}{r+1}+o(1)}$.

For line (2), we set $b = \frac{2-a}{r+1}$, which is in the range $a \leq b \leq 1$. This uses $\tilde{O}(m^{\frac{4-2a}{r+1}})$ memory. The runtime is $m^{2-b+o(1)} + m^{a+b\Gamma(1-a/b)+o(1)}$, which by our bound on r is at most $m^{\frac{2r+a}{r+1}+o(1)}$. \square

The current known bounds on $\Gamma(s)$ are highly non-linear functions of s . Thus, the parameters for Theorem 25 cannot be optimized in closed form. We can obtain slightly crude estimates in terms of ω or using analysis of Γ from [10]. We focus on the case where k_{trunc} is small; note that even $k_{\text{trunc}} = O(1)$ may be relevant for cohesive subnetworks of real-world graphs.

Theorem 26 *1. Algorithm 3 runs in $m^{\frac{2\omega+a}{\omega+1}+o(1)}$ time and $\tilde{O}(m^{\frac{4-2a}{\omega+1}} + m^{1+a})$ memory.*

2. If $\omega = 2$, Algorithm 3 runs in $m^{4/3+o(1)}k_{\text{trunc}}^{1/3}$ time and $\tilde{O}(m^{4/3} + mk_{\text{trunc}})$ memory.

3. For $k_{\text{trunc}} \leq m^{0.029}$, Algorithm 3 runs in $O(m^{1.4071}k_{\text{trunc}}^{0.0648})$ time and $O(m^{1.1860})$ memory.

Proof:

1. If $a \leq 0.4$, then apply Theorem 25 with $r = \omega$; note then that $2/a - 2 \geq 3 \geq r$ and $\Gamma\left(\frac{2-a(2+r)}{2-a}\right) \leq \omega = r$ as required. This runs in the claimed time and memory.

If $a > 0.4$, then apply Theorem 25 with $r = 2$. Note then that $2/a - 2 \geq 2 = r$ and $\Gamma\left(\frac{2-a(2+r)}{2-a}\right) \leq \Gamma(0.25) = 2 = r$. Algorithm 3 runs in $m^{\frac{4+a}{3}+o(1)}$ time and $\tilde{O}(m^{\frac{4-2a}{3}} + m^{1+a})$ memory. Since $\omega \geq 2$, this is also in the required range.

2. This follows immediately from the preceding paragraphs, setting $\omega = 2$.
3. From [10] we see that $\Gamma(0.95) \leq 2.333789$ and from [9] we see that $\omega = \Gamma(1) \leq 2.3728639$. As the function Γ is concave-up [13], this implies that when $0.95 \leq x \leq 1$ we have

$$\Gamma(x) \leq 2.333789 + 0.781498(x - 0.95) \quad (1)$$

Set $r = 14.43632 - \frac{110.419519}{9.15323-a}$ and let $s = \frac{2-a(2+r)}{2-a}$. We have $s \geq 0.95$ for $a \leq 0.029$, and from Eq. (1) we can verify that $\Gamma(s) \leq r \leq 2/a - 2$ for a in this range. Further calculus shows that then $\frac{2r+a}{r+1} \leq 1.40704 + 0.0648a$ so the runtime is $O(m^{1.4071} k_{\text{trunc}}^{0.0648})$. The memory is $\tilde{O}(m^{1+a} + m^{\frac{4-2a}{r+1}})$, which in this range is $O(m^{1.1860})$. \square

6 Acknowledgments

Thanks to Michael Murphy, Noah Streib, Lowell Adams, Tad White and Randy Dougherty for ideas and discussion. Thanks to the anonymous journal and conference reviewers for helpful suggestions and pointing us to some useful references.

A Properties of degeneracy and average degeneracy

The *degeneracy* of graph G , denoted $\delta(G)$, is the minimum value such that for every vertex set $U \subseteq V$, there is a vertex in $G[U]$ of degree at most $\delta(G)$. Many graph classes have bounded degeneracy, for example, planar graphs have degeneracy at most 5.

We remark that there is a connection between degeneracy and graph trussness:

Proposition 27 *For any edge e we have $\tau(e) \leq \delta(G) - 1$.*

Proof: Consider any k -truss-component E' , and let U denote the set of vertices with at least one edge in E' . Then $G[U]$ contains $G(E')$. By Observation 4, every vertex in $G(E')$ has degree at least $k + 1$. So $G[U]$ has minimum degree at least $k + 1$, which shows that $\delta(G) \geq k + 1$. \square

A number of previous triangle and trussness algorithms such as [5, 12] have analyzed runtime in terms of degeneracy as well as a closely related graph parameter known as *arboricity*, denoted $\alpha(G)$. This is also related to a parameter known as the *pseudo-arboricity* $\alpha^*(G)$. See [15] for definitions and background. The following are some standard and well-known bounds:

Proposition 28 *Consider a graph $G = (V, E)$ with $m = |E|$ edges.*

- G has an edge-orientation in which every vertex has out-degree at most $\alpha^*(G)$.
- G has an acyclic edge-orientation in which every vertex has out-degree at most $\delta(G)$.
- $\alpha(G) \leq \sqrt{2m}$.
- $\alpha^*(G) \leq \alpha(G) \leq \alpha^*(G) + 1$.
- $\alpha(G) \leq \delta(G) < 2\alpha^*(G)$.

Note that, in light of the last two bounds in Proposition 28, asymptotic runtime bounds are equivalent for α, α^* , and δ .

For any edge $e = (u, v)$, define $h(e) = \min(d(u), d(v))$. With this notation, we recall the definition of average degeneracy as $\bar{\delta}(G) = \frac{1}{m} \sum_{\text{edges } e} h(e)$. The following result shows some intuition behind the term ‘‘average degeneracy.’’ To state it informally, ‘‘most’’ edges of G have degeneracy ‘‘not much larger’’ than $\bar{\delta}(G)$.

Proposition 29 *Consider a graph $G = (V, E)$.*

1. *The average degeneracy $\bar{\delta}(G)$ satisfies $\bar{\delta}(G) \leq 2\alpha^*(G)$.*
2. *For any $x \in (0, 1)$, there is an edge-set $L \subseteq E$ with $|L| \geq (1 - x)|E|$ and $\delta(G(L)) \leq \bar{\delta}(G)/x$.*

Proof: The first result (in a slightly weaker form) was shown by [5]; for completeness, we provide a version of their proof here. By Proposition 28 there exists an orientation of E such that every vertex has out-degree at most $\alpha^*(G)$. Now compute:

$$\begin{aligned} \sum_{e \in E} h(e) &\leq \sum_{\substack{e=(u,v) \in E \\ e \text{ oriented to } v}} d(u) = \sum_{u \in V} \text{out-degree}(u) \times d(u) \\ &\leq \sum_{u \in V} \alpha^*(G) \times d(u) = 2m\alpha^*(G) \end{aligned}$$

For the second result, let L denote the set of edges e with $h(e) \leq s = \bar{\delta}(G)/x$. Since $\sum_e h(e) = m\bar{\delta}(G)$, we must have $|E - L|s \leq m\bar{\delta}(G)$, i.e. $|L| \geq (1 - x)|E|$.

Now consider a vertex set $U \subseteq V$ and let $G' = G(L)[U]$. Each edge $e \in G'$ has an endpoint $u \in U$ with $d(u) \leq s$. Thus, the total number of edges in G' is at most $\sum_{u \in U: d(u) \leq s} d(u) \leq |U|s$. Since this holds for all U , the graph $G(L)$ has degeneracy at most s . \square

B Constructions of toroidal graph embedding

Lemma 3 *For any integers $i \geq 0, t \geq 4$, there is a graph embedding in the torus whose faces consist of two t -cycles and i four-cycles, and where each edge is in two distinct faces.*

Proof: There are two cases depending on the parity of t , as shown in Figure 2.

Case I: $t = 2s$. We view each t -cycle as a rectangle with side lengths $s - 1$ and 1 . They are stacked vertically on top of i unit-length squares. Overall, we have one large rectangle with vertical sides of length $i + 2(s - 1)$ and horizontal sides of length 1 .

To get a torus, we first identify the top and bottom edges (marked X) to form a cylinder with circular circumference $i + 2(s - 1)$. We next identify the left side of the cylinder with a twisted version of the right side, namely, we rotate one of the sides by $s - 1$. Thus, the two edges marked Y are identified.

Case II: $t = 2s + 1$. We view each t -cycle as a trapezoid with base length 1 and side lengths $s - 1$ and s , joined to form a rectangle of height $2s - 1$. We put the i squares below them, and use a similar twisting process to join them into a torus.




Figure 2: Two t -cycles packed on top of i four-cycles. The case of even values of t is shown on the left side and the case of odd values of t on the right side.

The horizontal edges clearly have distinct faces. Also, every left edge is identified with a right edge from a distinct face. For example, when $t = 2s$, the top $s - 1$ left edges come from the top t -cycle and they are identified with the $s - 1$ right edges from the bottom t -cycle. \square

References

- [1] N. Alon, R. Yuster, and U. Zwick. Finding and counting given length cycles. *Algorithmica*, 17(3):209–223, 1997. doi:10.1007/BF02523189.
- [2] V. Batagelj and M. Zaversnik. Fast algorithms for determining (generalized) core groups in social networks. *Advanced in Data Analysis and Classification*, 5(2):129–145, 2011. doi:10.1007/s11634-010-0079-y.
- [3] A. Björklund, R. Pagh, V. V. Williams, and U. Zwick. Listing triangles. In *Proc. 41st International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 223–234, 2014. doi:10.1007/978-3-662-43948-7_19.
- [4] H. Cheng, X. Yan, and J. Han. Mining graph patterns. In *Frequent Pattern Mining*, pages 307–338. Springer, 2014. doi:10.1007/978-1-4419-6045-0_12.
- [5] N. Chiba and T. Nishizeki. Arboricity and subgraph listing algorithms. *SIAM Journal on Computing*, 14(1):210–223, 1985. doi:10.1137/0214017.
- [6] J. Cohen. Unpublished technical report. 2005.
- [7] J. Cohen. Graph twiddling in a MapReduce world. *Computing in Science and Engineering*, 11(4):29–41, 2009. doi:10.1109/MCSE.2009.120.
- [8] R. Diestel. *Graph Theory: Springer Graduate Text GTM 173*, volume 173. 2012.
- [9] F. L. Gall. Powers of tensors and fast matrix multiplication. In *Proc. 39th International Symposium on Symbolic and Algebraic Computation (IS-SAC)*, pages 296–303, 2014. doi:10.1145/2608628.2608664.
- [10] F. L. Gall and F. Urrutia. Improved rectangular matrix multiplication using powers of the Coppersmith-Winograd tensor. In *Proc. 29th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1029–1046, 2018. doi:10.1137/1.9781611975031.67.
- [11] L. Gkasienciec, M. Kowaluk, and A. Lingas. Faster multi-witnesses for boolean matrix multiplication. *Information Processing Letters*, 109(4):242–247, 2009. doi:10.1016/j.ipl.2008.10.012.
- [12] X. Huang, L. V. Lakshmanan, J. X. Yu, and H. Cheng. Approximate closest community search in networks. *Proc. VLDB Endowment*, 9(4):276–287, 2015.
- [13] G. Lotti and F. Romani. On the asymptotic complexity of rectangular matrix multiplication. *Theoretical Computer Science*, 23(2):171–185, 1983. doi:10.1016/0304-3975(83)90054-3.

- [14] D. W. Matula and L. L. Beck. Smallest-last ordering and clustering and graph coloring algorithms. *Journal of the ACM*, 30(3):417–427, 1983. doi:10.1145/2402.322385.
- [15] J.-C. Picard and M. Queyranne. A network flow solution to some nonlinear 0–1 programming problems, with applications to graph theory. *Networks*, 12(2):141–159, 1982. doi:10.1002/net.3230120206.
- [16] K. Saito and T. Yamada. Extracting communities from complex networks by the k -dense method. In *Proc. 6th annual IEEE International Conference on Data Mining Workshops (ICDM)*, pages 300–304, 2006. doi:10.1109/ICDMW.2006.76.
- [17] S. B. Seidman. Network structure and minimum degree. *Social Networks*, 5:269–287, 1983. doi:10.1016/0378-8733(83)90028-X.
- [18] J. Ugander, L. Backstrom, C. Marlow, and J. Kleinberg. Structural diversity in social contagion. *Proc. National Academy of Sciences*, 109(16):5962–5966, 2012. doi:10.1073/pnas.1116502109.
- [19] A. Verma, A. Buchanan, and S. Butenko. Solving the maximum clique and vertex coloring problems on very large sparse networks. *INFORMS Journal on Computing*, 27(1):164–177, 2015. doi:10.1287/ijoc.2014.0618.
- [20] A. Verma and S. Butenko. Network clustering via clique relaxations: A community based approach. *Graph Partitioning and Graph Clustering*, 588:129, 2013. doi:10.1090/conm/588/11702.
- [21] J. Wang and J. Cheng. Truss decomposition in massive networks. *Proc. VLDB Endowment*, 5(9):812–823, 2012. doi:10.14778/2311906.2311909.
- [22] R. Yuster and U. Zwick. Fast sparse matrix multiplication. *ACM Transactions on Algorithms*, 1(1):2–13, 2005. doi:10.1145/1077464.1077466.
- [23] Y. Zhang and S. Parthasarathy. Extracting analyzing and visualizing triangle k -core motifs within networks. In *Proc. 28th International Conference on Data Engineering (ICDE)*, pages 1049–1060, 2012. doi:10.1109/ICDE.2012.35.