

Fast Hierarchy Construction for Dense Subgraphs

Ahmet Erdem Saryüce
Sandia National Laboratories
Livermore, CA, USA
asariyu@sandia.gov

Ali Pinar
Sandia National Laboratories
Livermore, CA, USA
apinar@sandia.gov

ABSTRACT

Discovering dense subgraphs and understanding the relations among them is a fundamental problem in graph mining. We want to not only identify dense subgraphs, but also build a hierarchy among them (e.g., larger but sparser subgraphs formed by two smaller dense subgraphs). Peeling algorithms (k -core, k -truss, and nucleus decomposition) have been effective to locate many dense subgraphs. However, constructing a hierarchical representation of density structure, even correctly computing the connected k -cores and k -trusses, have been mostly overlooked. Keeping track of connected components during peeling requires an additional traversal operation, which is as expensive as the peeling process. In this paper, we start with a thorough survey and point to nuances in problem formulations that lead to significant differences in runtimes. We then propose efficient and generic algorithms to construct the hierarchy of dense subgraphs for k -core, k -truss, or any nucleus decomposition. Our algorithms leverage the disjoint-set forest data structure to efficiently construct the hierarchy during traversal. Furthermore, we introduce a new idea to avoid traversal. We construct the subgraphs while visiting neighborhoods in the peeling process, and build the relations to previously constructed subgraphs. We also consider an existing idea to find the k -core hierarchy and adapt for our objectives efficiently. Experiments on different types of large scale real-world networks show significant speedups over naive algorithms and existing alternatives. Our algorithms also outperform the hypothetical limits of any possible traversal-based solution.

1. INTRODUCTION

Graphs are used to model relationships in many applications such as sociology, the WWW, cybersecurity, bioinformatics, and infrastructure. Although the real-world graphs are sparse ($|E| \ll |V|^2$), vertex neighborhoods are dense [23]. Clustering coefficients [50], and transitivity [49] of real-world networks are also high and suggest the micro-scale dense structures. Literature is abundant with the benefits of dense subgraph discovery for various applications [30, 22]. Examples include finding communities in web [29, 13], and social networks [26], detecting spam groups in web [20], discovering migration patterns in stock market [14], improving soft-

ware understanding by analyzing static structure of large-scale software systems [55], analyzing gene co-expression networks [54], finding DNA motifs [16], quantifying the significance of proteins [53] and discovering molecular complexes [4] in protein interaction networks, identifying real-time stories in microblogging websites [3], and improving the throughput of social-networking sites [21].

k -core [43, 34], k -truss [39, 10, 56, 47, 44, 57, 26], and their generic variant for larger cliques, nucleus decomposition [42], are deterministic algorithms which are effective and efficient solutions to find dense subgraphs and creating hierarchical relations among them. They also known as *peeling algorithms* due to their iterative nature to reach the densest parts of the graph. Hierarchy has been shown to be a central organizing principle of complex networks, which is useful to relate communities of a graph and can offer insight into many network phenomena [9]. Peeling algorithms do not aim to find a single optimum dense subgraph, but rather gives many dense subgraphs with varying sizes and densities, and *hierarchy* among them, if supported by a post-processing traversal step [34, 42].

1.1 Problem, Misconception and Challenges

We focus on undirected, unattributed graphs. Hierarchy of dense subgraphs is represented as the tree structure where each node is a subgraph, each edge shows a containment relation, and the root node is the entire graph. The aim is to efficiently find the hierarchy by using peeling algorithms.

Misconception in the literature: Recent studies on peeling algorithms has interestingly overlooked the connectivity condition of k -cores and k -trusses. In the original definition of k -core, Seidman states that k -core is the maximal and connected subgraph where any vertex has at least degree k [43]. However, almost all the recent papers on k -core algorithms [8, 18, 19, 6, 36, 32, 28, 52, 51, 32] did not mention that k -core is a connected subgraph although they cite Seidman's seminal work [43]. On the k -truss side, the idea is introduced independently by Saito *et al.* [39] (as k -dense), Cohen [10] (as k -truss), Zhang and Parthasarathy [56] (as triangle k -core), and Verma and Butenko [47] (as k -community). They all define k -truss as a subgraph where any edge is involved in at least k triangles. Regarding the connectivity, Cohen [10], and Verma and Butenko [47] defined the k -truss as a single component subgraph, while others [39, 56] ignored the connectivity. In practice, overlooking the connectedness limits the contributions of most previous work regarding the performance and semantic aspects. More details are given in Section 3.

Finding k -cores requires traversal on the graph after the peeling process, where maximum k -core values of vertices are found. It is same for k -truss and nucleus decompositions

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>. For any use beyond those covered by this license, obtain permission by emailing info@vlldb.org.

Proceedings of the VLDB Endowment, Vol. 10, No. 3
Copyright 2016 VLDB Endowment 2150-8097/16/11.

Table 1: Speedups with our best algorithms for each decomposition. Starred columns (*) show lower bounds, when the other algorithm did not finish in 2 days or did a partial work. Best k -truss and (3,4) algorithms are significantly faster than alternatives, and also more efficient than the hypothetically best possible algorithm (HYPO) that does traversal to find the hierarchy.

	k -core		k -truss		(3,4) nucleus	
	NAIVE	HYPO	NAIVE	TCP* _[26]	HYPO	NAIVE*
Stanford3	25.50x	1.10x	12.58x	3.41x	1.48x	1321.89x
twit-hb	27.89x	1.33x	16.24x	3.27x	1.78x	38.96x
uk-2005	58.02x	1.68x	90.50x	11.07x	1.24x	1.98x

where the traversal is done on higher order structures. Constructing the hierarchy is only possible after that. However, it is not easy to track nested structure of subgraphs during a single traversal over entire graph. Traversing k -cores is cheap by a simple breadth-first search (BFS) in $O(|E|)$ time. When it comes to k -truss and higher order peeling algorithms, however, traversal becomes much costly due to the larger clique connectivity constraints.

1.2 Contributions

Motivated by the challenging cost of traversals and hierarchy construction, we focus on efficient algorithms to find the k -cores, k -trusses or any nuclei in general. Our contributions are as follows:

- **Thorough literature review:** We provide a detailed review of literature on peeling algorithms to point the misconception about k -core and k -truss definitions. We highlight the implications of these misunderstandings on the proposed solutions. We also stress the lack of understanding on the hierarchy construction and show that it is as expensive as the peeling process.
- **Hierarchy construction by disjoint-set forest:** We propose to use disjoint-set forest data structure (DF) to track the disconnected substructures that appear in the same node of the hierarchy tree. Disjoint-set forest is incorporated into the hierarchy tree by selectively processing the subgraphs in a particular order. We show that our algorithm is generic, i.e., works for *any* peeling algorithm.
- **Avoiding traversal:** We introduce a new idea to build the hierarchy without traversal. In the peeling process, we construct the subgraphs while visiting neighborhoods and bookkeep the relations to previously constructed subgraphs. Applying a lightweight post-processing operation to those tracked relations gives us all the hierarchy, and it works for *any* peeling algorithm.
- **Experimental evaluation:** All the algorithms we proposed are implemented for k -core, k -truss and (3,4)-nucleus decompositions, in which peeling is done on triangles and the four-clique involvements. Furthermore, we bring out an idea from Matula and Beck’s work [34], and adapt and implement it for our needs to solve the k -core hierarchy problem more efficiently. Table 1 gives a summary of the speedups we get for each decomposition. Our k -core hierarchy algorithm adaptation outperforms naive baseline by 58 times on uk-2005 graph. The best k -truss and (3,4) algorithms are significantly faster than alternatives. They also beat the hypothetically best possible algorithm (HYPO) that does traversal to find hierarchy. It is a striking result to show the benefit of our traversal avoiding idea.

Table 2: Summary of notations

Symbol	Description
K_r	r -clique; complete graph of r vertices
$\omega_s(u)$	K_s -degree of u ; number of s -cliques containing u
$\omega_{r,s}(H)$	$\min\{\omega_s(u) : u \in H\}$; min K_s -degree of a K_r in H
H^u	max k -(r, s) nucleus associated with the K_r u
$\lambda_s(u)$	$\omega_{r,s}(H^u)$; max k -(r, s) number of the K_r u
$\lambda_{r,s}(H)$	$\min\{\lambda_s(u) : u \in H\}$; min λ_s of a K_r in graph H
$T_{r,s}$	sub-(r, s) nucleus; maximal union of K_r s of same λ_s

2. PRELIMINARIES

This section presents building blocks for our work.

2.1 Nucleus decomposition

Let G be an undirected and simple graph. We start by quoting the Definitions 1 and 2 from [42]. We use K_r to denote an r -clique.

DEFINITION 1. Let $r < s$ be positive integers and \mathcal{S} be a set of K_s s in G .

- $K_r(\mathcal{S})$ is the set of K_r s contained in some $C \in \mathcal{S}$.
- The number of $C \in \mathcal{S}$ containing $u \in K_r(\mathcal{S})$ is the K_s -degree of u .
- Two K_r s u, u' are K_s -connected if there exists a sequence $u = u_1, u_2, \dots, u_k = u'$ in $K_r(\mathcal{S})$ such that for each i , some K_s $C \in \mathcal{S}$ contains $u_i \cup u_{i+1}$.

These definitions are generalizations of the standard vertex degree and connectedness. Indeed, setting $r = 1$ and $s = 2$ (so \mathcal{S} is a set of edges) yields exactly that. The main definition is as follows.

DEFINITION 2. Let k, r , and s be positive integers such that $r < s$. A k -(r, s) nucleus is a maximal union \mathcal{S} of K_s s such that:

- The K_s -degree of any $u \in K_r(\mathcal{S})$ is at least k .
- Any $u, u' \in K_r(\mathcal{S})$ are K_s -connected.

Figure 1: 2-(2,3) and 2-(2,4) nuclei on left and right.



Figure 1 gives an example for 2-(2,3) and 2-(2,4) nucleus. For $r = 1, s = 2$, a k -(1,2) nucleus is a maximal (induced) connected subgraph with minimum vertex degree k . This is exactly k -core [43]. Setting $r = 2, s = 3$ gives maximal subgraphs where every edge participates in at least k triangles, and edges are triangle-connected. This is almost the same definition of k -dense [39], k -truss [10], triangle-cores [56] and k -community [47]. The difference is on the connectivity condition; [10, 47] defines k -truss and k -community as a connected component, whereas [39, 56] do not mention connectedness, implicitly allowing disconnected subgraphs. The k -truss community defined by Huang *et al.* [26] is the same as the k -(2,3) nucleus: both require any pair of edges to be triangle-connected. More details can be found in Section 3. In the rest of the paper, we will be generic and present all our findings for the k -(r, s) nucleus decomposition, which subsumes k -core and k -truss definitions.

For an r -clique K_r u , $\omega_s(u)$ denotes the K_s -degree of u . For a subgraph $H \subseteq G$, $\omega_{r,s}(H)$ is defined as the minimum K_s -degree of a K_r in H , i.e., $\omega_{r,s}(H) = \min\{\omega_s(u) : u \in H\}$.

DEFINITION 3. The **maximum k - (r, s) nucleus** associated with a K_r u , denoted by H^u , is the k - (r, s) nucleus that contains u and has the largest $k = \omega_{r,s}(H^u)$ (i.e., $\nexists H$ s.t. $u \in H \wedge H$ is an l - (r, s) nucleus $\wedge l > k$).

The maximum k - (r, s) number of the r -clique u , denoted by $\lambda_s(u)$, is defined as $\lambda_s(u) = \omega_{r,s}(H^u)$.

Throughout the paper, $\lambda_s(u)$ implies that u is a K_r . We also abuse the notation as $\lambda(u)$ when r and s are obvious. The maximum k - $(1, 2)$ nucleus is same as the maximum k -core, defined in [41]. For a vertex, v , $\lambda_2(v)$ is also equal to the maximum k -core number of v [41], or core number of v [8, 36]. Likewise, for an edge e , $\lambda_3(e)$ is previously defined as the *trussness* of an edge in [26].

Building the (r, s) nucleus decomposition of a graph G is finding the λ_s of all K_r s in G and building the k - (r, s) nuclei for all k . The following corollary shows that given the λ_s of all K_r s, all (r, s) nuclei of G can be found.

COROLLARY 1. Given $\lambda_s(v)$ for all K_r $v \in G$ and assuming $\lambda_s(u) = k$ for a K_r u , the maximum k - (r, s) nucleus of u , denoted by H^u , consists of u as well as any K_r v that has $\lambda_s(v) \geq k$ and is reachable from u via a path P of K_s s such that $\forall C \in P, \lambda_{r,s}(C) \geq k$, where $\lambda_{r,s}(C) = \min\{\lambda_s(u) : u \in C$ (u is a K_r) $\}$.

Corollary 1 is easy to see for the k -core case, when $r = 1, s = 2$. All the traversed vertices are in H^u due to maximality property of k -cores, and all the vertices in H^u are traversed due to the connectivity condition, both mentioned in Definition 2. For the maximum k - $(2, 3)$ nucleus, we can also see the equality by Definition 2. For all edges e , $\lambda_3(e) \geq k$ satisfies the first condition, and the path of triangles, which does not contain any edge whose λ_3 is less than k , implies the second condition of Definition 2.

COROLLARY 2. H^u can be found by traversing G starting at the K_r u and including each K_r v to H^u if

- $\lambda_s(v) \geq k$
- \exists a K_s C s.t. $v \in C \wedge \lambda_{r,s}(C) \geq k$.

Repeating this traversal for all K_r s $u \in G$ gives all the k - (r, s) nuclei of G .

Traversal is trivial for k - $(1, 2)$ nucleus (k -core): include every vertex with greater or equal λ_2 . For $r = 2, s = 3$, maximum k - $(2, 3)$ nucleus is found by doing traversal on edges. Assuming the λ_3 value of the initial edge is k , the next edge in the traversal should have the λ_3 value $\geq k$; should be in the same triangle; and all the edges of this triangle should have λ_3 s greater-than or equal to k . Similar for $r = 3, s = 4$; traversal is done on triangles, and neighborhood condition is on the containment of triangles in four-cliques.

In summary, (r, s) nucleus decomposition problem has two phases: (1) peeling process which finds the λ_s values of K_r s, (2) traversal on the graph to find all the k - (r, s) nuclei. For $r = 1, s = 2$ case, the algorithm for finding λ_2 of vertices is based on the following property, as stated in [34]: to find the vertices with the λ_2 of k , all vertices of degree less than k and their adjacent edges are recursively deleted. For first phase, we provide the generic peeling algorithm in Alg. 1, which has been introduced in our earlier work [42], and for the second phase, we give the generic traversal algorithm in Alg. 2, which is basically the implementation of Corollary 2. The final algorithm, outlined in Alg. 3 combines the two.

Lastly, we define sub- (r, s) nucleus and strong K_s -connectedness to find the K_r s with same λ_s values. We will use them to efficiently locate all the k - (r, s) nuclei of given graph.

Algorithm 1: SET- $\lambda(G, r, s)$

```

1 Enumerate all  $K_r$ s in  $G(V, E)$ 
2 For every  $K_r$   $u$ , set  $\omega_s(u)$  as the number of  $K_s$ s containing  $u$ 
3 Mark every  $K_r$  as unprocessed
4 for each unprocessed  $K_r$   $u$  with minimum  $\omega_s(u)$  do
5    $\lambda_s(u) = \omega_s(u)$ ,  $max\lambda = \lambda_s(u)$ 
6   Find set  $S$  of  $K_s$ s containing  $u$ 
7   for each  $C \in S$  do
8     if any  $K_r$   $v \in C$  is processed then continue
9     for each  $K_r$   $v \in C, v \neq u$  do
10      if  $\omega_s(v) > \omega_s(u)$  then  $\omega_s(v) = \omega_s(v) - 1$ 
11   Mark  $u$  as processed
12 return array  $\lambda_s(\cdot)$  and  $max\lambda$ 

```

Algorithm 2: TRAVERSAL($G, r, s, \lambda(\cdot), max\lambda$)

```

1 for each  $k \in [1, max\lambda]$  do
2   visited $[v] = false, \forall K_r v \in G$ 
3   for each  $K_r u \in G$  do
4     if  $\lambda_s(u) = k$  and not visited $[u]$  then
5        $Q \leftarrow$  empty queue,  $Q.push(u)$ 
6        $S \leftarrow$  empty set,  $S.add(u)$ 
7       visited $[u] \leftarrow true$ 
8       while not  $Q.empty()$  do
9          $u \leftarrow Q.pop()$ 
10        for each  $K_r v$  s.t.
11           $(u \cup v \subset C) \wedge (C \text{ is a } K_s) \wedge (\lambda_{r,s}(C) \geq k)$  do
12          if not visited $[v]$  then
13             $Q.push(v), S.add(v)$ 
14            visited $[v] \leftarrow true$ 
15        report  $S$  // output the  $k$ - $(r, s)$  nucleus

```

Algorithm 3: NUCLEUSDECOMPOSITION(G, r, s)

```

1  $\lambda_s(\cdot), max\lambda \leftarrow$  SET- $\lambda(G, r, s)$  // Finding  $\lambda_s$  of  $K_r$ s
2 TRAVERSAL( $G, r, s, \lambda(\cdot), max\lambda$ ) // Finding all  $(r, s)$  nuclei

```

DEFINITION 4. Two K_r s u, u' with $\lambda_s(u) = \lambda_s(u')$ are **strongly K_s -connected** if there exists a sequence $u = u_1, u_2, \dots, u_k = u'$ such that:

- $\forall i, \lambda_s(u_i) = \lambda_s(u)$
- $\forall i, \exists C$ s.t.
 - C is a K_s
 - $(u_i \cup u_{i+1}) \subset C$
 - $\lambda_{r,s}(C) = \lambda_s(u)$.

DEFINITION 5. **sub- (r, s) nucleus**, denoted by $T_{r,s}$, is a maximal union of K_r s s.t. $\forall K_r$ pair $u, v \subset S$,

- $\lambda_s(u) = \lambda_s(v)$
- u and v are strongly K_s -connected

The sub- $(1, 2)$ nucleus is defined as the subcore in [41, 40]. All the notations are given in Table 2.

2.2 Disjoint sets problem

Disjoint-set data structure, also known as union-find, keeps disjoint dynamic sets, and maintains upon the operations that modifies the sets [11]. Each set has a representative. There are two main operations: UNION(x, y) merges the dynamic sets with ids x and y , and creates a new set, or just merge one of the sets into the other. FIND(x) returns the representative of the set which contains x .

Disjoint-set forest is introduced with two powerful heuristics [45]. In the disjoint-set forest, each set is a tree, each

Algorithm 4: DISJOINT-SET FOREST

```
LINK( $x, y$ ):           //  $x$  and  $y$  are nodes in the tree
if  $x.rank > y.rank$  then  $y.parent \leftarrow x$ 
else
   $x.parent \leftarrow y$ 
  if  $x.rank = y.rank$  then  $y.rank \leftarrow y.rank + 1$ 
FIND( $x$ ):
 $S \leftarrow$  empty set
while  $x.parent$  is not null do  $x \leftarrow x.parent, S.add(x)$ 
for each  $u \in S$  do  $u.parent \leftarrow x$ 
return  $x$ 
UNION( $x, y$ ): LINK(FIND( $x$ ), FIND( $y$ ))
```

node in the tree is an element of the set, and the root of each tree is the identifier of that set. To keep the trees flat, two heuristics are used that complement each other. First is union-by-rank, which merges the shorter tree under the longer one. Second heuristic is path-compression that makes each node on the find path point directly to the root. Time complexity with union-by-rank and path-compression heuristics is $O((m+n)\log^*n)$, where \log^*n is the inverse Ackermann function which is almost linear [45]. Pseudocode for FIND and UNION operations are given in Alg. 4.

3. LITERATURE AND MISCONCEPTIONS

In this section, we present a detailed review of related work on peeling algorithms. We point some misconceptions about the definitions and the consequences. Our focus is on peeling algorithms and their output, so we limit our scope to k -core and k -truss decompositions and their generalizations. Detailed literature review of dense subgraph discovery can be found in [22, 30].

3.1 k -core decomposition

The very first definition of a k -core related concept is given by Erdős and Hajnal [15] in 1966. They defined the degeneracy as the largest maximum core number of a vertex in the graph. Matula introduced the min-max theorem [35] for the same thing, highlighting the relationship between degree orderings of the graph and the minimum degree of any subgraph, and its applications to graph coloring problem. Degeneracy number has been rediscovered numerous times in the context of graph orientations and is alternately called the coloring number [33], and linkage [17].

First definition of the k -core subgraph is given by Seidman [43] for social networks analysis, and also by Matula and Beck [34], as k -linkage, for clustering and graph coloring applications, in the same year of 1983. Seidman [43] introduced the core collapse sequence, also known as degeneracy ordering of vertices, as an important graph feature. He states that k -cores are good *seedbeds* that can be used to find further dense substructures. Though, there is no algorithm in [43] on how to find the k -cores. Matula and Beck [34], on the other hand, gives algorithms for finding λ_2 values of vertices, and also finding all the k -cores of a graph (and their hierarchy) by using these λ_2 values, because there can be multiple k -cores for same k value. Both papers defined the k -core subgraph as follows:

“A **connected** and maximal subgraph H is k -core (k -linkage) if every vertex in H has at least degree k .” [43, 34]

The *connectedness* is an important detail in this definition because it requires a post-processing traversal operation on vertices to locate all the k -cores of the graph. Figure 2 shows

this. There are two 3-cores in the graph, and there is no way to distinguish them at the end of the peeling process by just looking at the λ values of vertices.

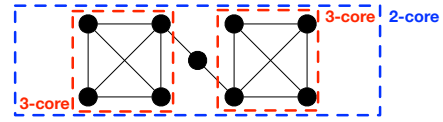


Figure 2: Multiple 3-cores

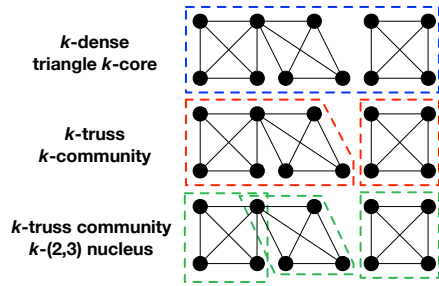
Batagelj and Zaversnik introduced an efficient implementation that uses bucket data structure to find the λ values of vertices [5]. They defined the k -core as a not necessarily connected subgraph, in contrast to previous work they cited [43, 34]. With this assumption, they claimed that their implementation finds all the k -cores of the graph.

Finding the relationships between k -cores of a graph has gained a lot of interest. Nested structure of k -cores reveals a hierarchy, and it has been shown to be useful for visualization [1] and understanding the underlying structure of complex networks arising in many domains. Carmi *et al.* [7] and Alvarez-Hamelin *et al.* [2] investigated the k -core hierarchy of internet topology at autonomous systems (AS) level. Healy *et al.* [25] compared the k -core hierarchies of real-world graphs in different domains and some generative models.

Given the practical benefit and efficiency of k -core decomposition, there has been a lot of recent work to adapt k -core algorithms for different data types or setups. Out of memory computation is an important topic for many graph analytic problems that deal with massive graphs not fitting in memory. Cheng *et al.* [8] introduced the first external-memory algorithm. Wen *et al.* [51] and Khaouid *et al.* [28] provided further improvements in this direction. Regarding the different type of graphs, Giatsidis *et al.* adapted the k -core decomposition for weighted [19] and directed [18] graphs. To handle the dynamic nature of real-world data, Sariyuce *et al.* [41] introduced the first streaming algorithms to maintain k -core decomposition of graphs upon edge insertions and removals. They recently improved these algorithms further by leveraging the information beyond 2-hop [40]. Li *et al.* [32] also proposed incremental algorithms for the same problem. More recently, Wu *et al.* [52] approached dynamic data from a different angle, and adapted k -cores for temporal graphs where possibly multiple interactions between entities occur at different times. Motivated by the incomplete and uncertain nature of the real network data, O’Brien and Sullivan [36] proposed new methods to locally estimate core numbers (λ values) of vertices when entire graph is not known, and Bonchi *et al.* [6] showed how to efficiently do the k -core decomposition on uncertain graphs, which has existence probabilities on the edges.

One common oversight in all those recent work (except [41, 42]) is that they ignore the **connectivity** of k -cores. This oversight does not change their results, but limit their contributions: they adapt/improve the peeling part of k -core decomposition, which finds the λ s of vertices, not the entire k -core decomposition which also needs traversal to locate all the (connected) k -cores. Considering the external memory k -core decomposition algorithms [8, 28, 51], existing works only focused on how to compute the λ values of vertices. However, the additional traversal operation in external memory is not taken into consideration which is at least as expensive as finding λ values. Finding the

Figure 3: k -dense [39] (triangle k -core [56]), k -truss [10] (k -community [47]) and k -truss community [26] (k - $(2,3)$ nucleus [42]) on the same graph for $k=2$. Each subgraph given by the corresponding algorithm is shown in dashed.



(connected) k -cores and constructing the hierarchy among them efficiently in the external memory computation model is not a trivial problem and will limit the performance of proposed algorithms for finding k -core subgraphs and constructing the hierarchy. Similar argument can be considered for weighted [19], probabilistic [6], and temporal [52] k -core decompositions, all of which have some kind of threshold-based adaptations on weights, probabilities and timestamps, respectively. On the other hand, connectedness definition is semantically unclear for some existing works like the directed graph core decomposition [18]. It is only defined that in- and out-degrees of vertices can be considered to find two λ values, but traversal semantic is not defined for finding subgraphs or constructing the hierarchy. One can think about building the hierarchy by considering the edges from lower level k -cores to higher level ones, or the opposite. To remedy those misconceptions, we focus on the efficient computation of traversal part for k -core decomposition and its higher-order variants.

3.2 k -truss decomposition

k -truss decomposition is inspired by the k -core and can be thought as the same peeling problem in a higher level that deals with triangles. It is independently introduced, with subtle differences, by several researchers. Chronologically, the idea is first proposed by Saito *et al.* [39], to the best of our knowledge, in 2006:

“ k -dense is a subgraph S if each adjacent vertex pair in S has more than or equal to $(k-2)$ common adjacent vertices in S .”

In other words, each edge in S should be involved in at least $k-2$ triangles. Nothing is mentioned about the connectedness of the vertices and edges, which implies that a k -dense subgraph might have multiple components. Saito *et al.* argue that k -dense is a good compromise between easy to compute k -cores and high quality k -cliques, and it is useful to detect communities in social networks. In 2008, Jonathan Cohen introduced the k -truss as a better model for cohesive subgraphs in social networks [10], which became the most popular naming in the literature:

“ k -truss is a one-component subgraph such that each edge is reinforced by at least $k-2$ pairs of edges making a triangle with that edge.”

In 2012, Zhang and Parthasarathy [56] proposed a new definition for visualization purposes:

“triangle k -core is a subgraph that each edge is contained within at least k triangles in the subgraph.”

Again there was no reference to the connectedness, implying multiple components can be observed in a triangle k -core. In the same year, Verma and Butenko [47] introduced the following:

“ k -community is a connected subgraph if every edge is involved in at least k triangles.”

The subtle difference between those papers is the connectedness issue. k -dense [39] and triangle k -core [56] definitions allow the subgraph to be disconnected whereas the k -truss [10] and k -community [47] are defined to be connected. All of these works only provided algorithms to find the λ_3 values of edges. k -dense and triangle k -cores can be found this way since they can be disconnected. However, finding the k -truss and k -community subgraphs requires a post-processing traversal operation, which increases the cost. As a stronger alternative to the k -truss, Huang *et al.* [26] introduced the k -truss community. The only difference is that each edge pair in a k -truss community is directly or transitively triangle-connected, where two edges should reside in the same triangle to be triangle-connected. The generic k - (r, s) nucleus, proposed by Sariyuce *et al.* [42], for $r = 2, s = 3$ gives the exact same definition. This brings a stronger condition on the connectivity structure, and shown to result in denser subgraphs than the classical k -truss definition [26]. However, it has an extra overhead of post-processing traversal operation that requires to visit triangles, which is more expensive than the traditional traversal. Authors devised TCP index, a tree structure at each vertex, to remedy this issue [26]. Figure 3 highlights the difference between those definitions on a simple example.

k -truss decomposition serves as a better alternative to the k -core. For most applications that k -core is useful for, k -truss decomposition performs better. Gregori *et al.* [24] investigated the structure of internet AS-level topologies by looking at the k -dense subgraphs, similar to Carmi *et al.* [7] and Alvarez-Hamelin *et al.* [2] who used k -core for same purpose. Orsini *et al.* [37] also investigated the evolution of k -dense subgraphs in AS-level topologies. It has been also used to understand the global organization of clusters in complex networks [44]. Colomer-de-Simon *et al.* used the hierarchy of k -dense subgraphs to visualize real-world networks, as Healy *et al.* [25] used the k -cores for the same objective.

Proven strength of k -truss decomposition drew further interest for adapting to different data types and setups, similar to the k -core literature. Wang and Cheng introduced external memory algorithms [48] and more improvements are provided by Zhao and Tung [57] for visualization purposes. More recently, Huang *et al.* [27] introduced probabilistic truss decomposition for uncertain graphs.

Similar to the k -core case, overlooking the connectivity constraints limits the contributions in the k -truss literature as well. For example, external memory k -truss decomposition [48] would be more expensive and require more intricate algorithms if it is done to find connected subgraphs by doing the traversal in external memory model. We believe that our algorithms for efficiently finding the k -trusses and constructing the hierarchy will be helpful to deal with this issue.

3.3 Generalizations

Given the similarity between k -core and k -truss decompositions, people have been interested in unified schemes to generalize the peeling process for a broader set of graph substructures.

Saito *et al.* pointed a possible direction of generalization in their pioneering work [39], where they defined k -dense subgraphs. Their proposal is as follows:

“Subgraph S is a h -level k -dense community if the vertices in every h -clique of S is adjacent to at least $h-k$ common vertices.” [39]

In other words, h -level k -dense community is the set of h -cliques where each h -clique is contained in at least $h-k$ number of $(h+1)$ -cliques. Note that, there is no connectivity constraint in the definition. h -level k -dense community subsumes the disconnected k -core, which contains multiple k -cores, for $h=1$. For $h=2$, it is their k -dense definition [39]. They claimed that h -level k -dense communities for $h>2$ are more or less same with $h=2$ and incurs higher computation cost. So they did not dive into more algorithmic details and stick with $h=2$.

Sariyuce *et al.* [42] introduced a broader definition to unify the existing proposals, which can be found by a generic peeling algorithm. As explained in Section 2.1, their definition subsumes k -core and k -truss community [26] concepts. It is also more generic than h -level k -dense community of [39], since (1) it allows to look for involvement of cliques whose size can differ by more than one, (2) enforces a stronger connectivity constraint to get denser subgraphs. h -level k -dense community can be expressed as the $k-(r, r+1)$ nucleus which does not have any connectivity constraint (k is actually $h-k$ and it does not matter). Well-defined theoretical notion of $k-(r, s)$ nucleus enables to provide a unified algorithm to find all the nuclei in graph, as explained in Section 2.1.

Sariyuce *et al.* [42] also analyzed the time and space complexity of (r, s) -nucleus decomposition. For the first phase, they report that finding λ values of nuclei (Alg. 1) requires $O(RT_r(G) + \sum_v \omega_r(v)d(v)^{s-r})$ time with $O(|K_r(G)|)$ space, where $RT_r(G)$ is K_r enumeration time, and second part is searching each K_s that a K_r is involved in ($\omega_r(v)$ is the number of K_r s containing vertex v , $d(v)$ is the degree of v , and $|K_r(G)|$ is the number of K_r s in G). For the second phase, traversal on the entire graph needs to access each K_r and examine all the K_s s it is involved. Its time complexity is the same as the second part of first phase: $O(\sum_v \omega_r(v)d(v)^{s-r})$ which also gives the total time complexity.

4. ALGORITHMS

In this part, we first highlight the challenging aspects of the traversal phase, then introduce two algorithms for faster computation of (r, s) -nucleus decomposition to meet those challenges.

4.1 Challenges of traversal

As mentioned in the previous section, time complexity of the traversal algorithm for (r, s) nuclei is $O(\sum_v \omega_r(v)d(v)^{s-r})$. However, designing an algorithm that constructs the hierarchy with this complexity is challenging. In [42], it is stated that finding the nuclei in the reverse order of λ is better since it enables to discover previously found components, thus avoiding repetition. No further details are given, though. This actually corresponds to finding all $T_{r,s}$ (sub- (r, s) nuclei of Definition 5), connected K_r s with the same λ value,

and incorporating the relations among them. But, keeping track of all the $T_{r,s}$ in a structured way is hard. Figure 4 shows a case for k -core ($r=1, s=2$). Traversal algorithm needs to understand the relation between $T_{1,2}$ s of equal λ that are not directly connected. For instance, $T_{r,s}$ A and E are in the same 2-core, but the traditional BFS will find 3 other $T_{r,s}$ (F, D, G) between those two. During the traversal operation, there is a need to detect each $k-(r, s)$ nucleus, determine containment relations and construct the hierarchy. One solution that can be thought is to construct the - expectedly smaller - supergraph which takes all the $T_{r,s}$ as vertices and their connections as edges. Then, repetitive traversals can be performed on this supergraph to find each $k-(r, s)$ nucleus and the hierarchy. However, it is not guaranteed to get a significantly smaller supergraph which can be leveraged for repetitive traversal. The $T_{r,s}$ structure of real-world networks, which are investigated in Section 5, also verify this concern. It is clear that there is a need for a lightweight algorithm/data structure that can be used on-the-fly, so that all the $k-(r, s)$ nuclei can be discovered with the hierarchy during the traversal algorithm.

The other challenge with the traversal algorithm is the high computational cost for $r \geq 2$ cases. Consider the $(2, 3)$ case. We need to traverse on edges, and determine the adjacencies of edges by looking at their common triangles. At each step, it requires to find the common neighbors of the vertices on extremities (of the edge), check whether each adjacent edge is previously visited, and push to queue if not. As explained at the end of Section 3.3, complexity becomes $O(3 * |\Delta|)$. Cost is getting much higher if we look for $(3, 4)$ nuclei, which is shown to give denser subgraphs with more detailed hierarchy. Ideally we are looking to completely avoid the costly traversal operation.

4.2 Disjoint-set forest idea

We propose to use disjoint-set forest data structure (DF) to track the disjoint $T_{r,s}$ (of equal λ), and construct the hierarchy where $T_{r,s}$ with smaller λ is on the upper side, and greater λ is on the lower side. DF has been used to track connected components of a graph and fits perfectly to our problem where we need to find the connected components at multiple levels.

DF-TRAVERSAL algorithm, outlined in Alg. 5, is used to replace the naive TRAVERSAL (Alg. 2) in NUCLEUSDECOMPOSITION (Alg. 3). Basically it finds all the $T_{r,s}$ in the decreasing order of λ . We construct the *hierarchy-skeleton* tree by using $T_{r,s}$. Each node in the hierarchy-skeleton is a $T_{r,s}$. We define *subnucleus* struct to represent a $T_{r,s}$. It consists of λ , **rank**, **parent** and **root** fields. λ field is the $\lambda(v)$ for $v \in T_{r,s}$, **rank** is the height of the node in the hierarchy-

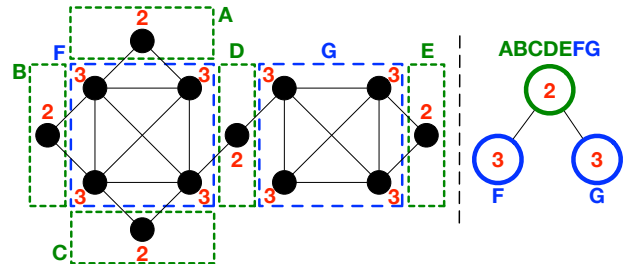


Figure 4: Example of $T_{1,2}$ s for $\lambda = 2$ and $\lambda = 3$. Hierarchy tree is shown on the right with participating $T_{1,2}$ s. λ s are shown in red. Traversal algorithm needs to infer that, for instance, components A and E are in the same 2-core.

skeleton, `parent` is a pointer to the parent node and `root` is a pointer to the root node of the hierarchy-skeleton. Default values for `parent` and `root` are null, and `rank` is 0. Figure 5 shows an example hierarchy-skeleton obtained by Alg. 5. Thin edges show the disjoint-set forests consisting of $T_{r,s}$ of equal λ value. The hierarchy of all (r, s) nuclei, the output we are interested, can be obtained by using the hierarchy-skeleton easily: we just take the child-parent links for which the λ values are different.

In the DF-TRAVERSAL algorithm, we store the `subnuclei` in `hrc` list (line 1) which also represents the hierarchy-skeleton. K_r s in a $T_{r,s}$ are stored by inverse-indices; `comp` keeps `subnucleus` index of each K_r in `hrc` (line 2). We also use `visited` to keep track of traversed K_r s (line 3). Main idea is to find each $T_{r,s}$ in decreasing order of λ (in lines 4-7). We construct the hierarchy-skeleton in a bottom-up manner this way and it lets us to use DF to find the representative $T_{r,s}$, i.e., the greatest ancestor, at any time. At each iteration we find an un-visited K_r with the λ value in order (line 6) and find its $T_{r,s}$ by SUBNUCLEUS algorithm (line 7), which also updates the hierarchy-skeleton.

Algorithm 5: DF-TRAVERSAL($G, r, s, \lambda(\cdot), max\lambda$)

```

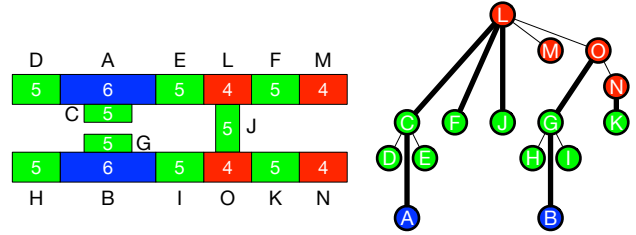
1 hrc ← list of subnuclei in a tree structure
2 comp(.) ∨  $K_r \in G$  // subnucleus id for each  $K_r$ 
3 visited[v] = false // lazy init
4 for each  $k \in [max\lambda, 1]$  do
5   for each  $K_r, u \in G$  do
6     if  $\lambda(u) = k$  and not visited[u] then
7       SUBNUCLEUS( $u, G, r, s, \lambda(\cdot), visited, comp, hrc$ )
8 root ← subnucleus, with  $\lambda = 0$ 
9 for each  $s \in hrc$  do
10  if s.parent is null then s.parent ← root
11 hrc.add(root) and REPORT ALL THE NUCLEI BY hrc, comp
```

Algorithm 6: SUBNUCLEUS($u, G, r, s, \lambda(\cdot), visited, comp, hrc$)

```

1 sn ← subnucleus, with  $\lambda = \lambda(u)$ , hrc.add(sn)
2 comp(u) ← sn,  $k \leftarrow \lambda(u)$ 
3 marked[v] = false // lazy init
4 merge ← list of subnuclei, merge.add(sn)
5 Q ← empty queue, Q.push(u)
6 visited[u] ← true
7 while not Q.empty() do
8   u ← Q.pop(), comp(u) ← sn
9   for each  $K_r, v$  s.t.
10    ( $u \cup v \subset C$ ) ∧ ( $C$  is a  $K_s$ ) ∧ ( $\lambda_{r,s}(C) = k$ ) do
11     if  $\lambda(v) = k$  then
12       if not visited[v] then
13         Q.push(v), visited[v] ← true
14         comp(v) ← sn
15     else
16       s ← comp(v) //  $\lambda(v) > k$ 
17       if not marked[s] then
18         marked[s] = true
19         s ← FIND-R(s)
20         if not marked[s] ∧  $s \neq sn$  then
21           if hrc(s). $\lambda > k$  then
22             hrc(s).parent ← hrc(s).root ← sn
23           else merge.add(s) // hrc(s). $\lambda = k$ 
24           marked[s] = true
25 for each  $m, n \in merge$  do
26  UNION-R(m, n)
```

Figure 5: A graph shown with $T_{r,s}$ regions on the left and the corresponding hierarchy-skeleton on the right. λ values of $T_{r,s}$ are the white numbers. Thin edges are disjoint-set forests.



SUBNUCLEUS (Alg. 6) starts by creating a `subnucleus`, with λ of the K_r of interest. We will store the discovered K_r s in this `subnucleus` (by inverse indices). We put this `subnucleus` into `hrc` (line 1) and assign its `comp` id (line 2). We use `marked` (line 3) to mark the adjacent `subnuclei` encountered during traversal so that unnecessary computation is avoided in lines 16-23. We do traversal (lines 7-23) by using a queue. At each step of the traversal, we process the next K_r in the queue. First, we assign its `comp` id as the new `subnucleus` (line 8) and then visit the adjacent K_r s residing in same K_s s, in which $\min \lambda$ of K_r is equal to $\lambda(v)$ (line 9). This is exactly the condition for $T_{r,s}$, given in Definition 5. For each adjacent K_r , its λ is either equal (line 10) or greater (line 14) by definition of $T_{r,s}$. If it is equal and not `visited` before, we visit it, put into queue (line 12) and also store in the current `subnucleus` (line 13). Otherwise, we find an adjacent `subnucleus` s with greater λ (line 15), that is already in hierarchy-skeleton, and can update the hierarchy-skeleton (lines 18-23) unless we had encountered s before (line 16).

Location of the `subnucleus` s in the hierarchy-skeleton is important. If it is parentless, we can just make it a child of the current `subnucleus` we build. If not, it means `subnucleus` s is a part of a larger structure and we should relate our current `subnucleus` to the representative of this large structure, which is the greatest ancestor of s that is guaranteed to have greater or equal λ (by line 9). So, hierarchy-skeleton update starts by finding the greatest ancestor of the s in line 18. FIND-R procedure is defined in Alg. 7. Its difference from FIND of Alg. 4 is that we use `root` field, not `parent`. `root` of a node implies its greatest ancestor in the hierarchy-skeleton, i.e., either it is the greatest ancestor or a few steps of `root` iterations would find the greatest ancestor. `parent` of a node, on the other hand, represents the links in hierarchy-skeleton, and not modified in FIND-R. After finding the root and making sure that it is not processed before (line 19), we can merge the current `subnucleus` to the hierarchy-skeleton. If the root has greater λ , we make it a child of our current `subnucleus` (line 21), by assigning both `root` and `parent` fields. Otherwise, we defer merging to the end (line 22), where we merge all `subnuclei` with equal λ by UNION-R operations (lines 24-25), defined in Alg. 7. UNION-R is slightly different than UNION of Alg. 4 in that it uses FIND-R instead of FIND and sets the `root` field of child node to the parent (in LINK-R).

Figure 5 displays the resulting hierarchy-skeleton for the $T_{r,s}$ regions shown on the left. We process $T_{r,s}$ in alphabetical order, which also conforms with decreasing order of λ . Consider the $T_{r,s}$ O, which is found and processed last. O finds the adjacent $T_{r,s}$ s I, J and K, during lines 9-23 of Alg. 6. All have greater λ values, so we focus on lines 14-

23. Greatest ancestor of I is G, and we make G child of O (line 21) since its λ is greater. Greatest ancestors of J and K are L and N, respectively, and they have equal λ values. So, we merge L and N with O in lines 24-25. Say we merge O and N first and O becomes **parent** of N since its rank is higher. Then, we merge O and L. Their ranks are equal and we arbitrarily choose L as the **parent**.

After traversing all the $T_{r,s}$, we create a root **subnucleus** to represent entire graph and make it parent to all the parentless nodes (lines 8-11 in Alg. 5). Time complexity of DF-TRAVERSAL does not change, i.e., $O(\sum_v \omega_r(v)d(v)^{s-r})$. Additional space is required by the auxiliary data structures in DF-TRAVERSAL. **hrc** needs $4 \cdot |T_{r,s}|$ (for four fields), and **comp** and **visited** requires $|K_r|$ space each. In addition, SUBNUCLEUS might require at most $2 \cdot |T_{r,s}|$ for **marked** and **merge**, and at most $|K_r|$ for Q , but reaching those upper bounds in practice is quite unlikely. Overall, additional space requirement of DF-TRAVERSAL at any instant is between $4 \cdot |T_{r,s}| + 2 \cdot |K_r|$ and $6 \cdot |T_{r,s}| + 3 \cdot |K_r|$. An upper bound for $|T_{r,s}|$ can be given as $|K_r|$, when each K_r is assumed to a **subnucleus**, but this case is also quite unlikely as we show in Section 5.

4.3 Avoiding traversal

All the $T_{r,s}$ can be detected without doing a traversal. A K_r is said to be processed, if its λ is assigned. During the peeling

Algorithm 7: NEW DISJOINT-SET FOREST

```

LINK-R( $x, y$ ):           //  $x$  and  $y$  are nodes in the tree
if  $x.\text{rank} > y.\text{rank}$  then  $y.\text{parent} \leftarrow x, y.\text{root} \leftarrow x$ 
else
   $x.\text{parent} \leftarrow y, x.\text{root} \leftarrow y$ 
  if  $x.\text{rank} = y.\text{rank}$  then  $y.\text{rank} \leftarrow y.\text{rank} + 1$ 
FIND-R( $x$ ):
 $S \leftarrow$  empty set
while  $x.\text{root}$  is not null do  $x \leftarrow x.\text{root}, S.\text{add}(x)$ 
for each  $u \in S$  do  $u.\text{root} \leftarrow x$ 
return  $x$ 
UNION-R( $x, y$ ): LINK-R(FIND-R( $x$ ), FIND-R( $y$ ))

```

Algorithm 8: FASTNUCLEUSDECOMPOSITION(G, r, s)

```

1 Enumerate all  $K_r$ s, mark them unprocessed, find their  $\omega_s$ 
2  $\text{hrc} \leftarrow$  list of subnuclei in a tree structure
3  $\text{comp}(v) = -1 \forall K_r v \in G$  // subnucleus id for each  $K_r$ 
4  $\text{ADJ} \leftarrow$  list of adjacent subnucleus pairs
5 for each unprocessed  $K_r u$  with minimum  $\omega_s(u)$  do
6    $\lambda(u) = \omega_s(u), \text{hrc}.\text{max}\lambda = \lambda(u)$ 
7    $sn \leftarrow$  subnucleus with  $\lambda = \lambda(u)$ 
8   Find set  $S$  of  $K_s$ s containing  $u$ 
9   for each  $C \in S$  do
10    if all  $K_r v \in C$  is unprocessed then
11     for each  $K_r v \in C, v \neq u$  do
12      if  $\omega_s(v) > \omega_s(u)$  then  $\omega_s(v) = \omega_s(v) - 1$ 
13    else
14      $D \leftarrow$  set of processed  $K_r$ s  $\subset C$ 
15      $w \leftarrow v \in D$  with the smallest  $\lambda$ 
16     if  $\lambda(w) = \lambda(u)$  then
17      if  $\text{comp}(u) = -1$  then  $\text{comp}(u) = \text{comp}(w)$ 
18      else UNION-R( $\text{comp}(u), \text{comp}(w)$ )
19     else  $\text{ADJ}.\text{add}(\text{comp}(u), \text{comp}(w))$  //  $\lambda(w) < \lambda(u)$ 
20   if  $\text{comp}(u) = -1$  then  $\text{comp}(u) \leftarrow sn, \text{hrc}.\text{add}(sn)$ 
21   Update all  $(-1, \cdot) \in \text{ADJ}$  with  $(\text{comp}(u), \cdot)$ 
22 BUILDHIERARCHY( $\text{ADJ}, \text{hrc}$ ) // postprocessing
23 create subnucleus root with  $\lambda = 0$ , tie to all existing roots
24  $\text{hrc}.\text{add}(\text{root})$  and REPORT ALL THE NUCLEI BY  $\text{hrc}, \text{comp}$ 

```

process, neighborhood of each K_r is examined, see lines 6-7 of Alg. 1, but processed neighbors are ignored (line 8). We leverage those ignored neighbors to construct the $T_{r,s}$. We introduce FASTNUCLEUSDECOMPOSITION algorithm (Alg. 8) to detect $T_{r,s}$ early in the peeling process so that costly traversal operation is not needed anymore.

At each iteration of the peeling process, a K_r with the minimum ω is selected and ω of its unprocessed neighbors are decremented. No information about the surrounding processed K_r s is used. If we check the processed neighbors, we can infer some connectivity information and use it towards constructing all the $T_{r,s}$ as well as the hierarchy-skeleton. For example, assume we are doing k -core decomposition and a vertex u with degree d is selected. We assign $\lambda(u) = d$ and check the unprocessed neighbors of u to decrement their degree, if greater than d . We can also examine the processed neighbors. λ of any processed neighbor is guaranteed to be less than or equal to d , by definition. Say v is a neighbor with $\lambda(v) = d$. Then, we can say that u and v are in the same $T_{r,s}$. Say w is another neighbor with $\lambda(w) < d$. Then, we can infer that maximum $d-(1,2)$ nucleus of w contains u , and $T_{r,s}$ of u is an ancestor of $T_{r,s}$ of w in the hierarchy-skeleton. Leveraging these pairwise relations enables us to find all the $T_{r,s}$ and construct the hierarchy-skeleton.

An important thing to note is that, it is not always possible to detect the $T_{r,s}$ of a K_r by only looking at the processed neighbors. Consider k -core decomposition on a star graph, for which all vertices has $\lambda = 1$. Center vertex is processed in the last two steps of peeling, so it is not possible to infer two connected vertices with equal λ until that time. We find *non-maximal* $T_{r,s}$ s (denoted as $T_{r,s}^*$) and combine them by using the disjoint-set forest algorithm. The difference from the DF-TRAVERSAL algorithm is that our hierarchy-skeleton will have more nodes because of non-maximal $T_{r,s}$.

Colored lines 13-22 in Alg. 8 implements our ideas. For each K_s we encountered (line 9), processed neighbors are explored starting from line 13. Note that, there is no need to check every adjacent and processed K_r in the same K_s , since the relations among them are already checked in previous steps. It is enough to find and process the $K_r w$ with minimum λ , as in line 15. If w has an equal λ value (line 16), we need to either put our K_r of interest to the **subnucleus** of w (line 17) or merge to the **subnucleus** of w by UNION-R operation (line 18). In FASTNUCLEUSDECOMPOSITION algorithm, we only build disjoint-set forests during the peeling process (until line 22). If $\lambda(w)$ happens to have

Algorithm 9: BUILDHIERARCHY(ADJ, hrc)

```

1  $\text{binned\_ADJ} \leftarrow \text{hrc}.\text{max}\lambda$  # of empty lists // to bin ADJ
2 for each  $(s, t) \in \text{ADJ}$  do //  $\text{hrc}(s).\lambda > \text{hrc}(t).\lambda$ 
3    $\text{binned\_ADJ}[\text{hrc}(t).\lambda].\text{add}(s, t)$ 
4 for each list  $l \in \text{binned\_ADJ}$  in  $[\text{hrc}.\text{max}\lambda \dots 1]$  order do
5    $\text{merge} \leftarrow$  empty list of subnucleus pairs
6   for each  $(s, t) \in l$  do
7      $s \leftarrow \text{FIND-R}(s), t \leftarrow \text{FIND-R}(t)$ 
8     if  $s \neq t$  then
9       if  $\text{hrc}(s).\lambda > \text{hrc}(t).\lambda$  then
10         $\text{hrc}(s).\text{parent} = \text{hrc}(s).\text{root} = t$ 
11       else  $\text{merge}.\text{add}(s, t)$ 
12 for each  $(s, t) \in \text{merge}$  do
13   UNION-R( $s, t$ )

```

Table 3: Statistics for the real-world graphs. Largest graph in the dataset has more than 37M edges. $|\Delta|$ is the number of triangles, $|K_4|$ is the number of four-cliques. Ratios of s -cliques to r -cliques are shown in columns 6,7,8, for $s \leq 4$ and $s - r = 1$. Columns 9 to 14 show sub- (r, s) nuclei numbers ($|T_{r,s}|$) and non-maximal sub- (r, s) nuclei numbers ($|T_{r,s}^*|$), artifact of Alg. 8, for the (r, s) values we interested. Last two columns are the number of connections from $T_{r,s}^*$ s with higher λ values to the ones with lower λ values.

	$ V $	$ E $	$ \Delta $	$ K_4 $	$\frac{ E }{ V }$	$\frac{ \Delta }{ E }$	$\frac{ K_4 }{ \Delta }$	$ T_{1,2} $	$ T_{1,2}^* $	$ T_{2,3} $	$ T_{2,3}^* $	$ T_{3,4} $	$ T_{3,4}^* $	$ c_{\downarrow}(T_{2,3}^*) $	$ c_{\downarrow}(T_{3,4}^*) $
skitter(SK)	1.7M	11.1M	28.8M	148.8M	6.54	2.59	5.17	1.1M	1.2M	2.1M	2.7M	1.4M	2M	8.0 M	8.2 M
Berkeley13(BE)	22.9K	852.4K	5.3M	26.6M	37.22	6.30	4.96	8.6K	9.3K	106.3K	137.7K	206.0K	261.1K	1.4 M	1.9 M
MIT(MIT)	6.4K	251.2K	2.3M	13.7M	39.24	9.44	5.77	2.7K	2.8K	27.3K	34.6K	77.6K	105.7K	509.1 K	1.3 M
Stanford3(ST)	11.6K	568.3K	5.8M	37.1M	49.05	10.27	6.37	4.7K	4.9K	56.8K	72.5K	185.2K	255.1K	1.2 M	2.9 M
Texas84(TX)	36.4K	1.6M	11.2M	70.7M	43.74	7.03	6.33	13.5K	14.9K	210.0K	266K	395.5K	498.9K	2.6 M	3.7 M
twit-hb(TW)	456.6K	12.5M	83.0M	429.7M	27.39	6.63	5.18	341.4K	350K	2.9M	3.3M	5.6M	7.3M	27.2 M	59.5 M
Google(GO)	916.4K	4.3M	13.4M	39.9M	4.71	3.10	2.98	408.7K	508.3K	386.9K	568.4K	251.4K	382.8K	982.8 K	814.7 K
uk-2005(UK)	129.6K	11.7M	837.9M	52.2B	90.60	71.35	62.36	1.5K	1.7K	837	837	836	836	0	0
wiki-0611(WK)	3.1M	37.0M	88.8M	162.9M	11.76	2.40	1.83	2.4M	2.5M	7.9M	9.7M	6.7M	9.3M	45.1 M	45.0 M

a smaller value (line 19), we put the pair of subnuclei to a list (ADJ), which will be used to build the hierarchy-skeleton after the peeling. We do not process the relations between subnuclei of different λ right away for two reasons: (1) subnucleus of the K_r of interest might not be assigned yet ($\text{comp}(u)$ is -1), (2) order of processing subnuclei relations is crucial to build the hierarchy-skeleton correctly and efficiently. Regarding (1), we take care of the K_r s not belonging to a subnucleus in lines 20 and 21. For (2), we have the BUILDHIERARCHY function (line 22), defined in Alg. 9.

In BUILDHIERARCHY, we create $\text{max}\lambda$ number of bins to distribute the subnuclei pairs based on the smaller λ of the subnucleus pair. The reason is same with our reverse order discovery of subnuclei in DF-TRAVERSAL (Alg. 5): we construct the hierarchy-skeleton in a bottom-up manner and it enables us to use disjoint-set forest algorithm to locate the k - (r, s) nuclei that we need. Distribution is done in lines 2-3. Then, we just process the binned list (binned_ADJ) in reverse order of λ values (line 4). We do the same operations to build the hierarchy-skeleton: lines 7-11 of BUILDHIERARCHY and lines 18-22 of SUBNUCLEUS algorithm (Alg. 6) are almost same. Once we finish each list in binned_ADJ, we union the accumulated subnuclei of equal λ values (lines 12-13), as we did in lines 24-25 of SUBNUCLEUS algorithm. Finally, in FASTNUCLEUSDECOMPOSITION we create a subnucleus to represent entire graph, make it parent to all parentless subnuclei, and give the hierarchy.

Avoiding traversal does not change the time complexity of overall algorithm, since the peeling part was already taking more time. Auxiliary data structures in FASTNUCLEUSDECOMPOSITION requires additional space, though. hrc needs $4 \cdot |T_{r,s}^*|$ in which subnuclei are not necessarily maximal, and comp needs $|K_r|$. ADJ structure corresponds to the connections from $T_{r,s}^*$ s with higher λ values to the ones with lower λ values, which we denote as $c_{\downarrow}(T_{r,s}^*)$. The upper bound for $|c_{\downarrow}(T_{r,s}^*)|$ is $\binom{s}{r}|K_s|$, when each K_r is assumed to be a $T_{r,s}^*$ and their λ values are adversary (see the end of Section 3.3 for details). However, it is unlikely as we show in Section 5. binned_ADJ in BUILDHIERARCHY is just an ordered version of ADJ, and needs the same amount of space; $|c_{\downarrow}(T_{r,s}^*)|$. Lastly, merge in BUILDHIERARCHY might require another $|c_{\downarrow}(T_{r,s}^*)|$ at most, but it is quite unlikely. Overall, additional space requirement of FASTNUCLEUSDECOMPOSITION at any instant is $4 \cdot |T_{r,s}^*| + 2 \cdot |c_{\downarrow}(T_{r,s}^*)| + |K_r|$ and additional $|c_{\downarrow}(T_{r,s}^*)|$ might be needed. Section 5 gives more details on $T_{r,s}^*$ structure of real-world networks and their impact to the memory cost.

5. EXPERIMENTS

We evaluated our algorithms on different types of real-world networks, obtained from SNAP [31], Network Repository [38] and UF Sparse Matrix Collection [12]. Our dataset includes an internet topology network (skitter), facebook friendship networks of some universities (Berkeley13, MIT, Stanford3, Texas84) [46], follower network of Twitter users tweeted about Higgs boson-like particle discovery (twit-hb), web networks (Google, uk-2005) and network of wikipedia pages (wiki-0611). We ignore the directions for directed graphs. Important statistics of the networks are given in Table 3. $|K_s|/|K_r|$ ratio gives an estimate for the k - (r, s) nucleus decomposition runtime, as explained at the end of Section 3.3, and we put them in columns 6-8 to show the challenging and diverse characteristics of the networks in our dataset. Note that most of the networks have relatively high edge density in the realm of real-world networks and it makes the computation more expensive. We also included graphs with various $|K_s|/|K_r|$ ratios to diversify our dataset. Last eight columns are shown to explain the runtime and memory costs of our algorithms. All the algorithms are implemented in C++ and compiled using gcc 5.2.0 at -O2 optimization level. All experiments are done on a Linux operating system running on a machine with Intel Xeon Haswell E5-2698 2.30 GHz processor with 128 GB of RAM.

Nucleus decomposition has been shown to give denser subgraphs and more detailed hierarchies for $s - r = 1$ cases, for fixed s [42]. We implemented and tested our algorithms for $s - r = 1$ cases where $s \leq 4$: giving us (1, 2), (2, 3) and (3, 4) nucleus decompositions. (1, 2)-nucleus decomposition is same as the k -core decomposition [43] and (2, 3) corresponds to k -truss community finding [26] (stronger definition of k -truss decomposition [10]). We consider the k - (r, s) nucleus as a set of K_r s. In our algorithms, we find the k - (r, s) nuclei for all k values and determine the hierarchy tree among those nuclei. We report the total time of peeling and traversal (or post-processing) that takes the graph as input and gives all the nuclei with an hierarchy.

5.1 k-core decomposition (or (1, 2) nuclei)

Peeling phase of the k -core decomposition, which finds λ values of vertices, is well studied, and most efficient implementation is used in our work. As mentioned before, traversal part, however, is mostly overlooked and there is no true algorithm to serve as a baseline.

In their pioneering work, Matula and Beck [34] introduced a high-level algorithm, named LCPS (Level Component Pri-

Table 4: k -core decomposition results. HYPO is the hypothetical limit for the best possible traversal based algorithm. NAIVE is Alg. 3, DFT is the one using Alg. 5, FND is the Alg. 8 and LCPS is our adaptation from [34]. Right-most column is the runtimes of fastest algorithm and rest are its speedups over other algorithms.

(1,2)-nuclei	speedups with respect to				time (s)
	HYPO	NAIVE	DFT	FND	
skitter	0.41x	5.94x	2.79x	1.33x	1.94
Berkeley13	0.82x	19.74x	1.68x	2.77x	0.05
MIT	0.82x	17.51x	1.74x	2.80x	0.01
Stanford3	0.78x	25.50x	1.76x	2.19x	0.03
Texas84	0.78x	23.01x	1.61x	2.57x	0.11
twit-hb	0.80x	27.89x	1.94x	2.60x	1.18
Google	0.15x	3.45x	0.40x	0.27x	3.83
uk-2005	0.86x	58.02x	2.61x	3.32x	0.17
wiki-0611	0.52x	10.09x	1.96x	1.41x	7.79
avg	0.66x	21.24x	1.83x	2.14x	best

ority Search), to detect k -cores and the hierarchy among them. LCPS algorithm traverses the vertices based on λ values in a selective order and outputs them with interspersed brackets. The vertices enclosed by paired brackets at depth $k + 1$ are the vertices of a k -core in G . The traversal can be started from any vertex, and neighbors are discovered and put in a priority queue along with their λ values. At each step, vertex with the maximum λ is chosen from the queue, and processed. Closed/open brackets are interspersed according to the λ values of consecutively processed vertices. Authors argue that time complexity of LCPS is $O(|E|)$, but an implementation may not always be possible owing to the difficulty of maintaining an appropriate priority queue [34].

We adapt and implement the LCPS algorithm for our objectives. We alleviate the problem of maintaining an appropriate priority queue by using the bucket data structure, known by the bucket sort [11]. During the traversal, we place the discovered vertices in the bucket according to their λ values, and find the one with the maximum λ value in $O(1)$ time. We also adapted the placing brackets part to have the tree hierarchy, where each node is a $T_{1,2}$. If the λ value of vertex of interest is equal to the one in the previous step, we stay in the same node. If it is greater, we go down in the tree by creating a chain of new nodes, otherwise we climb up in the tree. Number of these new nodes (or steps to climb) is the difference between λ s of current and previous vertex. As a result, we create a tree hierarchy with as many levels as the max λ in the graph.

Table 4 gives the total k -core decomposition runtimes for experimented algorithms; the fastest one shown on the right-most column and its speedups over other algorithms given on the corresponding columns. NAIVE is the base nucleus decomposition algorithm, given in Alg. 3, and uses the naive traversal to find all k -cores. DFT is the algorithm that uses DF-TRAVERSAL (Alg. 5) and FND is the FASTNUCLEUSDECOMPOSITION algorithm (Alg. 8). LCPS denotes the runtime of our LCPS adaptation. Note that we are not aware of any other adaptation/implementation of the LCPS algorithm in [34]. We also have a hypothetical runtime, shown as HYPO, which is the peeling time plus a BFS traversal on the entire graph. HYPO is a limit for the most efficient traversal (plus regular peeling) that can be done on a given graph. Although it does not compute the hierarchy and find k -cores, it shows the best that can be done with a traversal-based nucleus decomposition algorithm. Peeling phases of HYPO, NAIVE, DFT, and LCPS are same.

Results show that our LCPS adaptation outperforms other

alternatives. Detailed timings suggest that traversal time is almost same with peeling, satisfying the linear complexity in the number of edges. On average, LCPS is 21.24x faster than NAIVE algorithm. Speedups over DFT and FND are 1.83x and 2.14x. Our LCPS adaptation also runs in only 66% time of the theoretical limit of HYPO, where the overhead is due to the bucket data structure.

5.2 (2, 3) nuclei

Table 5 gives the runtime comparison for entire (2, 3) nucleus decomposition, where rightmost column is the fastest algorithm and other columns are its speedups over other algorithms. To start with, we state that our peeling implementation for (2, 3) nucleus decomposition is quite efficient; for instance [48] computes λ values of skitter graph in 281 secs where we can do it in 74 secs. We also tested wiki-Talk graph just to compare with [26], and we get $41/13 = 3.15$ times faster results despite our less powerful CPU. Apart from NAIVE, DFT and FND algorithms, we give HYPO as the bound for the best possible traversal-based nucleus decomposition algorithm that traverses graph over edges and triangles. In addition, we implemented the TCP index construction algorithm introduced by Huang *et al.* [26]. In that work, authors devise TCP index on vertices which enable fast traversal to answer max-(2, 3) nucleus queries on vertices. TCP index is a tree structure at each vertex, which is actually the maximum spanning forest of each ego-network. TCP index is constructed after the peeling process to be ready for incoming queries. Note that our implementation of TCP index construction is also quite efficient that is 1.8x faster than [26] for the dataset in their work. In Table 5, we show the time only for peeling plus TCP index construction. Entire graph still needs to be traversed using TCP to find all the (2, 3) nuclei and the hierarchy among them. Peeling phases of HYPO, NAIVE, TCP and DFT are same.

Before checking the fastest, we look at the DFT algorithm, which is the second best. It is 2.4x and 122.4x faster than TCP and NAIVE algorithms, respectively. Note that, TCP time is *even before the traversal*. In 40% of the time that TCP spends for being ready to answer queries, DFT is able to give all the answers, i.e, (2, 3) nuclei and the hierarchy. This shows the benefit of using disjoint-set forests to detect the nuclei and construct the hierarchy. One of our initial objectives was to meet the computational complexity of traversal phase by true design and implementation. DFT does that by keeping the traversal time close to the peeling time; traversal takes only 23% more time than peeling on average. Figure 6 shows the comparison, where we normalize runtimes with respect to total DFT. Also, DFT performs only 34% slower than the hypothetical bound (HYPO).

The fastest algorithm is FND. It is 215x faster than NAIVE on average and gives 1729 speedup on Google. It outperforms TCP by 4.3x on average and up to 11x (on uk-2005). FND also works 1.76x faster than DFT and Figure 6 shows the detailed comparison. For DFT, peeling is Alg. 1 and postprocessing is the traversal time (Alg. 5). For FND, postprocessing is the BUILDHIERARCHY algorithm (Alg. 9) and peeling is the rest in Alg. 8. FND is able to keep the total time comparable to the peeling of DFT; it is only 29% slower! This is exactly the benefit of avoiding traversal by early discovery of non-maximal $T_{2,3}$. On uk-2005 graph, benefit of avoiding traversal is most apparent; DFT performs worst compared to FND. Because there are only 837

Table 5: (2,3) and (3,4) nucleus decomposition results. Right-most column for each is the runtimes of the fastest algorithm and rest are its speedup over other algorithms. Starred numbers (*) show lower bounds, when the algorithm did not finish in 2 days. HYPO is the hypothetical limit for the best possible traversal based algorithm. For (2,3), it is done on edges and triangle connections, and for (3,4) it is on triangles and four-clique connections. NAIVE is Alg. 3 and TCP* is the indexing algorithm proposed by [26] and does not even include (2,3) nuclei finding. DFT is the one using Alg. 5 and FND is the Alg. 8.

	(2,3)-nuclei					(3,4)-nuclei			
	speedups with respect to				time (s)	speedups with respect to			time (s)
	HYPO	NAIVE	TCP*[26]	DFT	FND	HYPO	NAIVE	DFT	FND
skitter	1.54x	45.62x	3.69x	2.04x	91.3	1.78x	163.91*x	1.96x	1054.2
Berkeley13	1.13x	10.76x	3.41x	1.40x	7.3	1.42x	1812.47*x	1.52x	95.3
MIT	1.10x	11.68x	3.45x	1.33x	2.8	1.45x	3848.02*x	1.53x	44.9
Stanford3	1.10x	12.58x	3.41x	1.36x	7.8	1.48x	1321.89*x	1.58x	130.7
Texas84	1.10x	14.41x	3.35x	1.37x	16.8	1.49x	679.06*x	1.57x	254.5
twit-hb	1.33x	16.24x	3.27x	1.49x	255.5	1.78x	38.96*x	1.81x	4434.9
Google	1.31x	1729.93x	3.90x	1.59x	13.0	1.35x	1083.02*x	1.43x	159.6
uk-2005	1.68x	90.50x	11.07x	3.64x	562.5	1.24x	1.98*x	1.98*x	87329.6
wiki-0611	1.53x	7.06x	3.37x	1.63x	584.9	1.82x	23.00*x	1.91x	7513.5
avg	1.31x	215.4x	4.32x	1.76x	best	1.53x	> 996.92x	> 1.70 x	best

$T_{2,3}$ s which means no work is done by disjoint-set forest and all the time is spent on traversal (see Figure 6). As mentioned in Section 4, extended peeling in FND finds non-maximal $T_{2,3}$, and the count might be much larger than the number of $T_{2,3}$ s. However, it is not the case. Last 6 columns of Table 3 show that there is no significant difference; on average non-maximal $T_{2,3}$ s are 24% more than $T_{2,3}$ s.

DFT and FND algorithms require additional memory, as mentioned at the end of Section 4.2. The upper bound of additional space for DFT is $6 \cdot |T_{2,3}| + 3 \cdot |E|$, and takes at most ~ 650 MB for any graph in our dataset, where `int` is used to store each number (four bytes). Although the upper bound of $|T_{2,3}|$ is $|E|$, we observed that it is only 13.1% of $|E|$ on average, which can be calculated from Table 3. Regarding the FND, upper bound is given as $4 \cdot |T_{2,3}^*| + 3 \cdot |c_{\downarrow}(T_{2,3}^*)| + |E|$ (see Section 4.3) and it takes at most ~ 1.4 GB for the experimented graphs (two `ints` for each connection in $c_{\downarrow}(T_{2,3}^*)$). We observed that uk-2005 graph have so many edges with no triangles that results in isolated $T_{2,3}^*$ s (and $T_{3,4}^*$ s), so $|c_{\downarrow}(T_{2,3}^*)| = |c_{\downarrow}(T_{3,4}^*)| = 0$. For the rest, we see that $|c_{\downarrow}(T_{2,3}^*)|$ is only 8.7% of the upper bound ($\binom{3}{2}|\Delta|$), on average.

Lastly, we check FND vs. HYPO. FND is faster on all instances, and 1.31x on average, meaning that it outperforms **any possible** traversal-based algorithm for (2,3) nucleus decomposition. This is so important to show the strength of

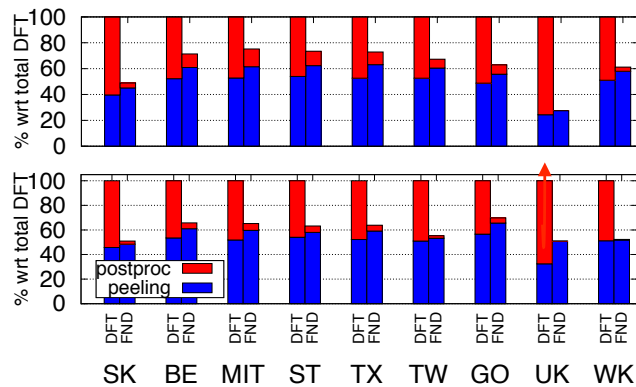


Figure 6: (2,3) [top] and (3,4) [bottom] nucleus decomposition comparison for DFT (the one using Alg. 5 for traversal) and FND (shown in Alg.8). Two main results: (1) Traversal part of DFT is close to the peeling part, (2) FND is able to keep the total runtime comparable to the peeling part of DFT

- (1) detecting the non-maximal $T_{2,3}$ early to avoid traversal,
- (2) using disjoint-set forest data structure to handle on-the-fly computation.

5.3 (3,4) nuclei

Table 5 and Figure 6 gives the comparison of all algorithms on (3,4)-nucleus decomposition, which has been shown to give subgraphs with highest density and most detailed hierarchy [42]. Results are similar to the (2,3) case: FND is the best and speedups are sharper. NAIVE algorithm could not be completed in 2 days for any graph. DFT also could not finish its computation on uk-2005 graph in that time. Regarding the memory requirement, we see that DFT and FND take at most ~ 10 GB and ~ 3.5 GB for the graphs in our dataset. Also, $|T_{3,4}|$ and $|c_{\downarrow}(T_{3,4}^*)|$ are far from the upper bounds, only of 3.9% and 2.5% on average. Overall, FND outperforms all others. Figure 6 shows that total time of FND is only 21% more than the peeling time of DFT and only 2% slower on wiki-0611. Most significantly, FND is 1.53x faster than hypothetical limit (HYPO) of any possible traversal-based (3,4) nucleus decomposition algorithm.

6. CONCLUSION

In this work, we focused on computing the hierarchy of dense subgraphs given by peeling algorithms. We first provided a detailed review of previous work and pointed misconceptions about k -core and k -truss decompositions. We proposed two generic algorithms to compute *any* nucleus decomposition. Our idea to leverage disjoint-set forest data structure for hierarchy computation works well in practice. We further improved the performance by detecting the subgraphs during peeling process to avoid traversal. We also adapted and implemented an existing idea for k -core decomposition hierarchy, and showed its benefit. Overall, our algorithms significantly outperformed the existing alternatives.

There are two open questions that might be worth to look at. First is about the analysis. Nested structures given by the resulting hierarchy only show the k -(r, s) nuclei. Instead looking at the $T_{r,s}$, which are many more than the k -(r, s) nuclei, might reveal more insight about networks. This actually corresponds to the hierarchy-skeleton structure that our algorithms produce. Second is regarding the performance. We believe that adapting the existing parallel peeling algorithms for the hierarchy computation can be helpful.

7. ACKNOWLEDGMENTS

This work was funded by the Applied Mathematics Program at the U.S. Department of Energy and Sandia's Laboratory Directed Research & Development (LDRD) program. Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

8. REFERENCES

- [1] J. I. Alvarez-Hamelin, A. Barrat, and A. Vespignani. Large scale networks fingerprinting and visualization using the k-core decomposition. In *NIPS*, pages 41–50, 2006.
- [2] J. I. Alvarez-Hamelin, L. Dall'Asta, A. Barrat, and A. Vespignani. K-core decomposition of Internet graphs: hierarchies, self-similarity and measurement biases. *Networks and Heterogeneous Media*, 3(2):371–293, 2008.
- [3] A. Angel, N. Sarkas, N. Koudas, and D. Srivastava. Dense subgraph maintenance under streaming edge weight updates for real-time story identification. *PVLDB*, 5(6):574–585, Feb. 2012.
- [4] G. D. Bader and C. W. Hogue. An automated method for finding molecular complexes in large protein interaction networks. *BMC Bioinformatics*, 4(1):1–27, 2003.
- [5] V. Batagelj and M. Zaversnik. An $o(m)$ algorithm for cores decomposition of networks. Technical Report cs/0310049, Arxiv, 2003.
- [6] F. Bonchi, F. Gullo, A. Kaltenbrunner, and Y. Volkovich. Core decomposition of uncertain graphs. In *KDD*, pages 1316–1325, 2014.
- [7] S. Carmi, S. Havlin, S. Kirkpatrick, Y. Shavitt, and E. Shir. A model of internet topology using k-shell decomposition. *PNAS*, 104(27):11150–11154, 2007.
- [8] J. Cheng, Y. Ke, S. Chu, and M. T. Ozsu. Efficient core decomposition in massive networks. In *ICDE*, pages 51–62, 2011.
- [9] A. Clauset, C. Moore, and M. E. J. Newman. Hierarchical structure and the prediction of missing links in networks. *Nature*, 453:98–101, 2008.
- [10] J. Cohen. Trusses: Cohesive subgraphs for social network analysis. National Security Agency Technical Report, 2008.
- [11] T. H. Cormen, C. Stein, R. L. Rivest, and C. E. Leiserson. *Introduction to Algorithms*. McGraw-Hill, 2001.
- [12] T. A. Davis and Y. Hu. The university of florida sparse matrix collection. *ACM Trans. Math. Softw.*, 38(1):1:1–1:25, 2011.
- [13] Y. Dourisboure, F. Geraci, and M. Pellegrini. Extraction and classification of dense communities in the web. In *WWW*, pages 461–470, 2007.
- [14] X. Du, R. Jin, L. Ding, V. E. Lee, and J. H. T. Jr. Migration motif: a spatial - temporal pattern mining approach for financial markets. In *KDD*, pages 1135–1144, 2009.
- [15] P. Erdős and A. Hajnal. On chromatic number of graphs and set-systems. *Acta Mathematica Hungarica*, 17:61–99, 1966.
- [16] E. Fratkin, B. T. Naughton, D. L. Brutlag, and S. Batzoglou. Motifcut: regulatory motifs finding with maximum density subgraphs. In *ISMB*, pages 156–157, 2006.
- [17] E. C. Freuder. A sufficient condition for backtrack-free search. *J. ACM*, 29(1):24–32, Jan. 1982.
- [18] C. Giatsidis, D. M. Thilikos, and M. Vazirgiannis. Evaluating cooperation in communities with the k-core structure. In *ASONAM*, pages 87–93, 2011.
- [19] C. Giatsidis, D. M. Thilikos, and M. Vazirgiannis. D-cores: measuring collaboration of directed graphs based on degeneracy. *Knowl. Inf. Syst.*, 35(2):311–343, 2013.
- [20] D. Gibson, R. Kumar, and A. Tomkins. Discovering large dense subgraphs in massive graphs. In *VLDB*, pages 721–732, 2005.
- [21] A. Gionis, F. Junqueira, V. Leroy, M. Serafini, and I. Weber. Piggybacking on social networks. *PVLDB*, 6(6):409–420, 2013.
- [22] A. Gionis and C. E. Tsourakakis. Dense subgraph discovery: Tutorial. In *KDD*, pages 2313–2314, 2015.
- [23] D. F. Gleich and C. Seshadhri. Vertex neighborhoods, low conductance cuts, and good seeds for local community methods. In *KDD*, pages 597–605, 2012.
- [24] E. Gregori, L. Lenzini, and C. Orsini. k-dense communities in the internet as-level topology. In *COMSNETS*, pages 1–10, 2011.
- [25] J. Healy, J. Janssen, E. Milios, and W. Aiello. *Characterization of Graphs Using Degree Cores*, pages 137–148. Springer Berlin Heidelberg, 2008.
- [26] X. Huang, H. Cheng, L. Qin, W. Tian, and J. X. Yu. Querying k-truss community in large and dynamic graphs. In *SIGMOD*, pages 1311–1322, 2014.
- [27] X. Huang, W. Lu, and L. V. Lakshmanan. Truss decomposition of probabilistic graphs: Semantics and algorithms. In *SIGMOD*, pages 77–90, 2016.
- [28] W. Khaouid, M. Barsky, S. Venkatesh, and A. Thomo. K-core decomposition of large networks on a single PC. *PVLDB*, 9(1):13–23, 2015.
- [29] R. Kumar, P. Raghavan, S. Rajagopalan, and A. Tomkins. Trawling the web for emerging cyber-communities. In *WWW*, pages 1481–1493, 1999.
- [30] V. E. Lee, N. Ruan, R. Jin, and C. Aggarwal. A survey of algorithms for dense subgraph discovery. In *Managing and Mining Graph Data*, volume 40. 2010.
- [31] J. Leskovec and A. Krevl. SNAP Datasets, June 2014.
- [32] R. Li, J. X. Yu, and R. Mao. Efficient core maintenance in large dynamic graphs. *IEEE TKDE*, 26(10):2453–2465, 2014.
- [33] D. Lick and A. White. k-degenerate graphs. *Canadian Journal of Mathematics*, 22:1082–1096, 1970.
- [34] D. Matula and L. Beck. Smallest-last ordering and clustering and graph coloring algorithms. *JACM*, 30(3):417–427, 1983.
- [35] D. W. Matula. A min-max theorem for graphs with application to graph coloring. *SIAM Review*, 10(4):481–482, 1968.
- [36] M. P. O'Brien and B. D. Sullivan. Locally estimating core numbers. In *ICDM*, pages 460–469, 2014.
- [37] C. Orsini, E. Gregori, L. Lenzini, and D. Krioukov. Evolution of the internet k-dense structure. *IEEE/ACM Trans. Netw.*, 22(6):1769–1780, 2014.
- [38] R. Rossi and N. Ahmed. The network data repository with interactive graph analysis and visualization. In *AAAI*, pages 4292–4293, 2015.
- [39] K. Saito and T. Yamada. Extracting communities from complex networks by the k-dense method. In *ICDMW*, 2006.
- [40] A. E. Sariyüce, B. Gedik, G. Jacques-Silva, K.-L. Wu, and Ü. V. Çatalyürek. Incremental k-core decomposition: algorithms and evaluation. *VLDB Journal*, pages 1–23, 2016.
- [41] A. E. Sariyüce, B. Gedik, G. Jacques-Silva, K.-L. Wu, and Ü. V. Çatalyürek. Streaming algorithms for k-core decomposition. *Proc. VLDB Endow.*, 6(6):433–444, Apr. 2013.
- [42] A. E. Sariyüce, C. Seshadhri, A. Pinar, and Ü. V. Çatalyürek. Finding the hierarchy of dense subgraphs using nucleus decompositions. In *WWW*, pages 927–937, 2015.
- [43] S. B. Seidman. Network structure and minimum degree. *Social Networks*, 5(3):269–287, 1983.
- [44] P. Simon, M. Serrano, M. Beiro, J. Alvarez-Hamelin, and M. Boguna. Deciphering the global organization of clustering in real complex networks. *Scientific Reports*, 3(2517), 2013.
- [45] R. E. Tarjan. Efficiency of a good but not linear set union algorithm. *JACM*, 22(2):215–225, Apr. 1975.
- [46] A. L. Traud, P. J. Mucha, and M. A. Porter. Social structure of facebook networks. *Physica A: Statistical Mechanics and its Applications*, 391(16):4165 – 4180, 2012.
- [47] A. Verma and S. Butenko. Network clustering via clique relaxations: A community based approach. In *DIMACS Graph Part. and Clustering*, pages 129–140, 2012.
- [48] J. Wang and J. Cheng. Truss decomposition in massive networks. *PVLDB*, 5(9):812–823, 2012.
- [49] S. Wasserman and K. Faust. *Social Network Analysis: Methods and Applications*. Cambridge Univ Press, 1994.
- [50] D. Watts and S. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393:440–442, 1998.
- [51] D. Wen, L. Qin, Y. Zhang, X. Lin, and J. Yu. I/o efficient core graph decomposition at web scale. In *ICDE*, pages 133–144, 2016.
- [52] H. Wu, J. Cheng, Y. Lu, Y. Ke, Y. Huang, D. Yan, and H. Wu. Core decomposition in large temporal graphs. In *Big Data, IEEE International Conf. on*, pages 649–658, 2015.
- [53] S. Wuchty and E. Almaas. Peeling the yeast protein network. *PROTEOMICS*, 5(2):444–449, 2005.
- [54] B. Zhang and S. Horvath. A general framework for weighted gene co-expression network analysis. *Statistical Applications in Genetics and Molecular Biology*, 4(1):Article 17+, 2005.
- [55] H. Zhang, H. Zhao, W. Cai, J. Liu, and W. Zhou. Using the k-core decomposition to analyze the static structure of large-scale software systems. *J. of Supercomputing*, 53(2):352–369, 2009.
- [56] Y. Zhang and S. Parthasarathy. Extracting analyzing and visualizing triangle k-core motifs within networks. In *ICDE*, pages 1049–1060, 2012.
- [57] F. Zhao and A. Tung. Large scale cohesive subgraphs disc. for social network visual analysis. In *PVLDB*, pages 85–96, 2013.