

The Core Decomposition of Networks: theory, algorithms and applications

(The VLDB Journal, 2019)

Fragkiskos D. Malliaros*, Christos Giatsidis[†], Apostolos N. Papadopoulos[‡],
and Michalis Vazirgiannis[§]

Abstract

The *core decomposition* of networks has attracted significant attention due to its numerous applications in real-life problems. Simply stated, the core decomposition of a network (graph) assigns to each graph node v , an integer number $c(v)$ (the core number), capturing how well v is connected with respect to its neighbors. This concept is strongly related to the concept of *graph degeneracy*, which has a long history in Graph Theory. Although the core decomposition concept is extremely simple, there is an enormous interest in the topic from diverse application domains, mainly because it can be used to analyze a network in a simple and concise manner by quantifying the significance of graph nodes. Therefore, there exists a respectable number of research works that either propose efficient algorithmic techniques under different settings and graph types or apply the concept to another problem or scientific area. Based on this large interest in the topic, in this survey, we perform an in-depth discussion of core decomposition, focusing mainly on: *i*) the basic theory and fundamental concepts, *ii*) the algorithmic techniques proposed for computing it efficiently under different settings, and *iii*) the applications that can benefit significantly from it.

Keywords: Core decomposition, graph mining, graph degeneracy, graph theory, algorithms

*Center for Visual Computing, CentraleSupélec, University of Paris-Saclay and Inria Saclay, France. Email: fragkiskos.malliaros@centralesupelec.fr

[†]Computer Science Laboratory, École Polytechnique, France. Email: xristosakamad@gmail.com

[‡]School of Informatics, Aristotle University of Thessaloniki, Greece. Email: papadopo@csd.auth.gr

[§]Computer Science Laboratory, École Polytechnique, France. Email: mvazirg@lix.polytechnique.fr

Contents

1	Introduction	4
2	Fundamental Concepts	8
2.1	Simple Graphs	8
2.2	Extensions to Other Types of Graphs	10
2.2.1	Directed Graphs	10
2.2.2	Weighted Graphs	10
2.2.3	Signed Graphs	11
2.2.4	Bipartite Graphs	11
2.2.5	Dynamic Graphs	11
2.2.6	Temporal Graphs	12
2.2.7	Probabilistic Graphs	12
2.2.8	Multilayer Graphs	12
2.2.9	Hidden Graphs	13
2.3	Generalized Core Decomposition	13
2.4	Extensions of the Core Decomposition	14
2.4.1	Truss Decomposition	14
2.4.2	Density-friendly Core Decomposition	15
2.4.3	Peak Decomposition	15
2.4.4	Nucleus Decomposition	16
2.4.5	Distance-based Core Decomposition	16
2.4.6	k -tip Decomposition	18
2.4.7	Radius-Bounded Decomposition	19
3	Algorithmic Techniques	21
3.1	In-Memory Computation	21
3.2	Disk-Resident Graphs	21
3.3	Core Decomposition in Dynamic Environments	24
3.4	Local Computation of Core Numbers	26
3.5	Parallel and Distributed Techniques	28
3.5.1	Parallel Computation	28
3.5.2	Distributed Computation	30
3.6	Probabilistic Core Decomposition	32
4	Representative Application Domains	33
4.1	Network Modeling and Analysis	33
4.2	Temporal Evolution	34
4.3	Anomaly Detection	34
4.4	Detection of Influential Spreaders	35
4.5	Network Visualization	36
4.6	Communities and Dense Subgraphs	38
4.7	Text Analytics	39
4.8	The Anchored k -Core Problem and Engagement Dynamics in Social Graphs	39
4.9	Graph Similarity	40

4.10 Physics, Biology and Ecology	41
4.11 Neuroscience	41
5 Conclusions and Further Research	42

1 Introduction

Graph management and *graph mining* are two important research areas with a plethora of significant practical applications [3, 36]. The main reason for this is the fact that graphs are ubiquitous and, therefore, their efficient management and mining is necessary to guarantee fast and meaningful knowledge discovery. Research on graph processing and mining was boosted by the need to explore and analyze massive graphs in Big Data Analytics tasks, aiming at computational efficiency and high scalability.

A *network* or *graph* (we will use the terms interchangeably) is denoted by $G(V, E)$, where V is the set of nodes or vertices and E is the set of edges or links. We will follow the trend in the literature and use the symbol n for the number of nodes ($n = |V|$) and the symbol m for the number of edges ($m = |E|$). The number of neighbors of a node $u \in V$ plays a central role in general, and it will be denoted by $\deg(u)$.

Figure 1(a) presents a simple graph $G(V, E)$ with $n = 8$ nodes and $m = 12$ edges. Based on the degree definition, $\deg(v_1) = 2$, $\deg(v_4) = 2$ whereas node v_3 has the highest degree $\deg(v_3) = 5$. Therefore, node v_8 has the smallest degree and node v_3 the highest.

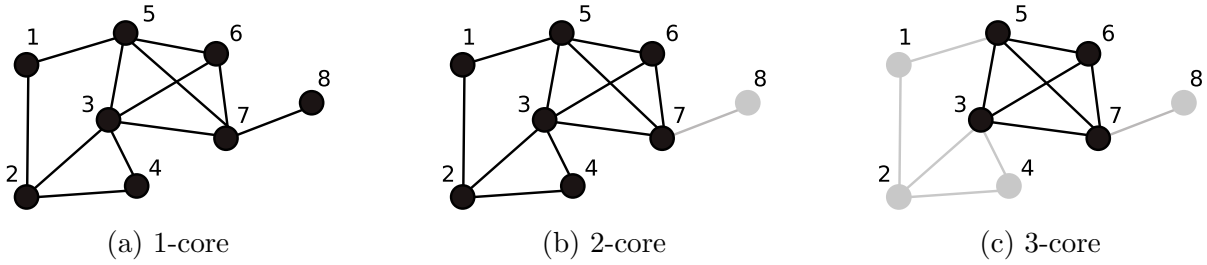


Figure 1: A small graph G with $n = 8$ nodes and $m = 12$ edges and the corresponding 1-core, 2-core and 3-core of G .

In many modern applications graphs are first class citizens. For example:

- The Web [80] can be modeled as a graph, where nodes represent web pages and edges represent hyperlinks among them.
- In a social network [2], nodes may represent users and edges may provide information regarding friendship relationships.
- In protein analysis, a protein-protein interaction network (PPI for short) [122] may be represented as a graph, where nodes represent different proteins and edges capture the interaction between proteins during a specific experiment.
- Graphs may also represent interactions among different types of nodes. For example, purchase information may be captured by using a bipartite graph, where a set of nodes (e.g., customers) interact with another set of nodes (e.g., products). In this case, a link between a customer c and a product p captures the fact that p was purchased by c .

These are only a few examples of applications that model elements and the interactions among them by using a graph structure. In these cases, it is important to rely on the structural properties of the interactions among the elements in order to discover useful and meaningful patterns.

Exploring and analyzing massive complex networks involves the execution of (usually) computationally intensive tasks, aiming at uncovering the network structure and detecting the presence of useful patterns

that could be proven significant. Some important graph mining tasks involve: reachability queries, graph partitioning, graph clustering, classification of graph nodes, predicting network evolution, discovering dense subgraphs, detecting influential spreaders.

In many cases we are searching for graph nodes that are considered “central” with respect to a specific problem at hand. Therefore, the concept of *node importance* is crucial in network analysis, since it is expected that among the nodes of a massive network, only a small fraction is of *high significance*. Evidently, one should first determine a method to quantify this significance (importance), since this concept is highly related to the application under consideration. For example, if we assume that we are interested in nodes with a large degree (simply because the number of links is important), then evidently, importance is directly measured by counting the number of neighbors of every node (a concept known as *degree centrality*). As another example, we may define importance in relation to the number of triangles each node participates in; the higher the number of triangles, the higher the node importance. Another popular measure of importance is the total number of shortest paths passing through a specific node (also known as *betweenness centrality* [50]). Also, one can quantify node importance using the concept of *random walks* and applying techniques similar to PageRank [27]. In such a case, the importance of a node is represented by the probability that this node will be visited by a random walker.

Based on the previous discussion, we realize that there is a plethora of different ways to define importance. However, it turns out that the concept of *core decomposition* can be used efficiently and effectively to quantify node importance in many different domains, thus, avoiding the use of more complex and computationally intensive algorithmic techniques. To be precise, the core decomposition of a simple graph G can be computed in linear time with respect to the number of edges of G , if the computation is done in main memory. Simply put, the k -core of a graph G is the maximal induced subgraph G_k , where the number of neighbors of every node u in G_k is at least k . The *core number* of a node u ($c(u)$) is defined as the maximum value of k such that u is contained in G_k . Figure 1 illustrates the 1-core, the 2-core and the 3-core of a toy graph. Based on the definition of the core number, it holds that $c(v_8) = 1$, $c(v_1) = 2$ and $c(v_6) = 3$.

Roadmap. Based on the fact that the core decomposition concept has numerous applications in diverse domains, in this survey we cover the topic as thoroughly as possible, by presenting the *concepts*, the *algorithmic techniques* used and also the fundamental *applications* that base their main results on the core number (or variations) of graph nodes. This survey was inspired by four tutorials presented by the authors [101] in the following conferences: *i*) 7th IEEE/ACM International Conference on Social Network Analysis and Mining (ASONAM) 2015, *ii*) 19th EDBT/ICDT Joint Conference 2016, *iii*) IEEE International Conference on Data Mining (ICDM) 2016 and *iv*) The European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECMLPKDD) 2017. To the best of our knowledge, this is the first comprehensive survey of the area. Existing works like [30, 23] cover only specific parts of the problem and they do not provide an in-depth discussion of the algorithms and the associated application domains.

Table 1 presents the timeline of the main research works related to core decomposition that are cited in this survey. We classify the related research contributions into two basic groups: *i*) concepts and algorithmic techniques and *ii*) representative application domains. In addition, the main topics covered in this survey are summarized in four categories (graph types, decomposition types, algorithms discussed and applications) as shown in Table 2. For each topic there are relevant references to the main text, to help the reader focus on the most relevant sections. The rest of this article is organized as follows:

- In Section 2 we present some fundamental concepts in addition to the basic definition of the core decomposition and some of its most widely used extensions for specific graph types. In particular,

Table 1: Core decomposition timeline.

Concepts and Algorithmic Techniques	Year	Representative Application Domains
coloring number [45]	1968	
degeneracy [94]	1970	
width [51]	1982	
	1983	dense subgraphs [133]
linkage [78]	1996	
	1999	adjacency matrix visualization [16]
	2000	web analysis [80]
generalized cores [17]	2002	
main-memory computation [18]	2003	
	2005	network analysis [8], visualization [6]
	2006	complex networks [41], fingerprinting and visualization [7]
	2007	internet topology [29]
truss decomposition [34]	2008	internet evolution [165], network analysis [69, 87, 9], neuroscience [67]
	2009	cell biology [99]
	2010	influential spreaders [79]
D-cores [62], disk-based computation [32], distributed computation [109], anchored k -core [20]	2011	communities [63], influence [28], neuroscience [147]
triangle cores (truss) [150, 166], weighted networks [54]	2012	visualizing triangle k -core [166]
dynamic cores [127], distributed computation [110], weighted networks [43]	2013	influential spreaders [168], engagement [103, 55], network analysis [1], neuroscience [135]
local estimation [114], uncertain graphs [24], S-cores [60]	2014	clustering [61], neuroscience [125], influential spreaders [119, 95]
temporal graphs [157], disk-based and in-memory [77], parallel computation [134], density-friendly decomposition [141]	2015	keyword extraction [124], influential communities [92], earthquake networks [68], cell biology [73, 44], neuroscience [156, 21], textual event detection [106], engagement [104]
incremental computation [126], disk-based computation [65, 153], distributed computation [12], uncertain truss [72]	2016	cell biology [42, 120], graph mining [136], network analysis [148], neuroscience [85], network decycling and dismantling [161], influential spreaders [102], keyword extraction [143], summarization [144]
hidden cores [139], multilayer graphs [53], parallel computation [152, 74], density-friendly decomposition [38], k -peak decomposition [66], (k, r) -core [164], core unravelling [162], truss in uncertain graphs [170], dynamic graphs [167], α - β -core [40]	2017	distributed clustering [33], community search [4, 46], engagement [163], ecology [58, 57], influential spreaders [5], keyword extraction [107]
distributed computation [71], probabilistic cores [121], radius-bounded cores [151], bipartite peeling [128], temporal graphs [52], local truss and nucleus [129]	2018	community search [90, 91], biology [56], ecology [112], core resilience [86], metabolic networks [49], patterns and anomaly detection [137], graph similarity [113]
disk-based computation [154], distance-based cores [25]	2019	network decycling and dismantling [132], statistical mechanics [111]

Table 2: Main topics covered.

GRAPH TYPES	DECOMPOSITIONS	ALGORITHMS	APPLICATIONS
simple (§2.1, §4)	simple core (§1, §3.1, §3.2)	baseline (§3.1)	network modeling (§4.1)
directed (§2.2.1, §2.2.3)	weighted core (§2.2.2, §4.4, §4.10)	linear (§3.1)	network analysis (§4.1)
weighted (§2.2.2, §4.7, §4.10)	D -core (§2.2.1)	disk-based (§3.2)	network evolution (§4.2)
signed (§2.2.3)	S -core (§2.2.3)	dynamic (§3.3)	anomaly detection (§4.3)
bipartite (§2.2.4)	$(\alpha-\beta)$ -cores (§2.2.4)	local computation (§3.4)	influential spreaders (§4.4)
dynamic (§2.2.5, §3.3)	$(k-\eta)$ -cores (§2.2.7)	parallel (§3.5.1)	network visualization (§4.5)
temporal (§2.2.6, §3.3, §4.2, §5)	generalized cores (§2.3)	distributed (§3.5.2)	communities (§4.6)
uncertain (§2.2.7, §3.6)	truss (§2.4.1, §4.4, §4.7)	probabilistic (§3.6)	dense subgraphs (§4.6)
multilayered (§2.2.8)	density-friendly (§2.4.2)		text analytics (§4.7)
hidden (§2.2.9)	peak (§2.4.3)		engagement dynamics (§4.8)
	nucleus (§2.4.4)		graph similarity (§4.9)
	distance-based (§2.4.5)		physics (§4.10)
	tip and wing (§2.4.6)		biology (§4.10)
	radius-bounded (§2.4.7)		ecology (§4.10)
			neuroscience (§4.11)

we will focus on simple, directed, weighted, signed, probabilistic, temporal, multilayer and hidden graphs which are frequently used in modern network-based applications.

- Section 3 offers a thorough overview of some of the the algorithmic techniques required for the core decomposition computation in different settings. In particular, we cover main memory computation, disk-based computation, local core number computation, parallel/distributed core decomposition and core decomposition in probabilistic graphs. Since it is not possible to cover every single algorithm in its full extend, we have selected a representative set of algorithms which we describe in detail and for the rest we offer the corresponding links to related research.
- Next, Section 4 elaborates on diverse application domains that benefit significantly from the core decomposition concept. More specifically, we present the application of the core decomposition concept in domains such as: network analysis, temporal evolution, influence maximization, dense subgraph discovery, community detection, text mining, biological network analysis and neuroscience. Evidently, the list is not exhaustive but we believe that these domains cover the vast majority of research performed in the area. We note that the research contributions presented in this section reuse existing techniques related to different versions of the k -core concept presented in the two previous sections and they apply them in different scientific areas. The purpose of this section is to

point out that the k -core decomposition is useful in many different disciplines.

- Finally, Section 5 concludes this survey and discusses briefly open problems and future research directions in areas like network representation learning, influence maximization and machine learning.

2 Fundamental Concepts

In this section, we cover the most important concepts related to core decomposition. In particular, we formally present the basic definitions and the most important properties which can be used by applications for performing more complex network analysis tasks. Table 3 illustrates some frequently used symbols and the corresponding interpretations.

Table 3: Frequently used symbols.

Symbol	Interpretation
G	a graph
V	set of vertices (or nodes) of G
E	set of edges of G
n	number of nodes ($n = V $)
m	number of edges ($m = E $)
u, v	some vertices of G
$N(u)$	set of direct neighbors of vertex u
$N_d(u)$	set of neighbors of u at a distance at most d
$\deg(u)$	degree of node u (number of incident edges)
$\deg_S(u)$	degree of node u in subgraph S
$c(u)$	core number of node u
$\delta^*(G)$	the degeneracy of graph G

2.1 Simple Graphs

Let $G(V, E)$ denote an *undirected* and *unweighted* graph, where V is the set of nodes and E is the set of edges. The k -core decomposition of G is a threshold-based hierarchical decomposition of G into nested subgraphs. The basic idea is that a threshold k is set on the degree of each node; nodes that do not satisfy the threshold, are excluded from the process. The following definitions provide some basic knowledge regarding the concepts around core decomposition.

Definition 1 (k -shell subgraph) *The k -shell is the subgraph of G defined by the nodes that belong to the k -core but not to the $(k + 1)$ -core.*

Definition 2 (k -core subgraph) *Let H be a subgraph of G , i.e., $H \subseteq G$. H is defined to be the k -core of G , denoted by G_k , if it is a maximal subgraph of G in which all nodes have degree at least k .*

Definition 3 (graph degeneracy $\delta^*(G)$) *The degeneracy [94] $\delta^*(G)$ of a graph G is defined as the maximum k for which graph G contains a non-empty k -core subgraph.*

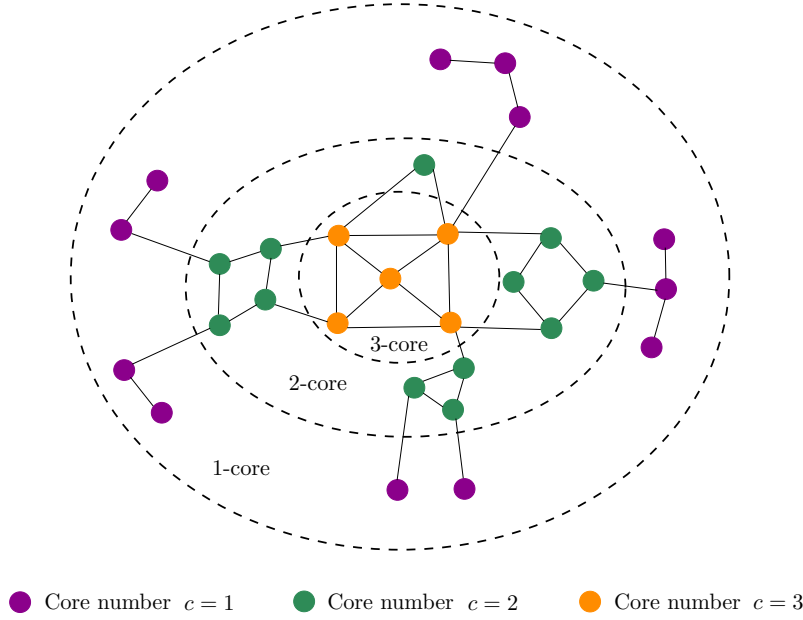


Figure 2: Example of the k -core decomposition.

Definition 4 (core number) A node v has core number $c(v) = k$, if it belongs to a k -core but not to the $(k + 1)$ -core.

Based on the above definitions, it is evident that if all the nodes of the graph have degree at least one, i.e., $\deg(v) \geq 1, \forall v \in V$, then the 1-core subgraph corresponds to the whole graph, i.e., $G_1 \equiv G$. Furthermore, assuming that $G_i, i = 0, 1, 2, \dots, \delta^*(G)$ is the i -core of G , then the k -core subgraphs are nested. Formally:

$$G_0 \supseteq G_1 \supseteq G_2 \supseteq \dots \supseteq G_{\delta^*(G)} \quad (1)$$

Typically, the subgraph $G_{\delta^*(G)}$ is called *the maximal k -core subgraph* of G . Figure 2 depicts an example of a graph and the corresponding k -core decomposition. As we observe, the degeneracy of this graph is $\delta^*(G) = 3$; therefore, the decomposition creates three nested k -core subgraphs, with the 3-core being the maximal one. An attempt to create a higher order core subgraph (i.e., the 4-core of the graph) would result in an empty subgraph, since the removal of one of the nodes belonging to the 3-core will force the removal of the remaining nodes. The nested structure of the k -core subgraphs is indicated by the dashed lines shown in Figure 2. Furthermore, the color of the nodes indicate the core number $c(u)$ of each node u .

It is important to note that the k -core subgraphs are not necessarily connected. As an example, consider the graph shown in Figure 3(a). The graph is composed of two cliques (complete subgraphs) of size four that are connected by a node x with a degree of 2. Evidently, the graph is a 2-core, since the degree of each node is at least 2. The transition from the 2-core to the 3-core of G will eliminate node x , since $\deg(x) = 2$. The remaining nodes constitute the 3-core of the graph, which evidently is disconnected as shown in Figure 3(b).

The concept of degeneracy in graphs, as defined above, is also known as *width* [51] and *linkage* [78]. It is also related to the *coloring number* α of a graph [45], which is defined as the least k for which there is an ordering \prec of the graph nodes, such that for every $v \in V$, the number of adjacent nodes $w \prec v$ is less than k .

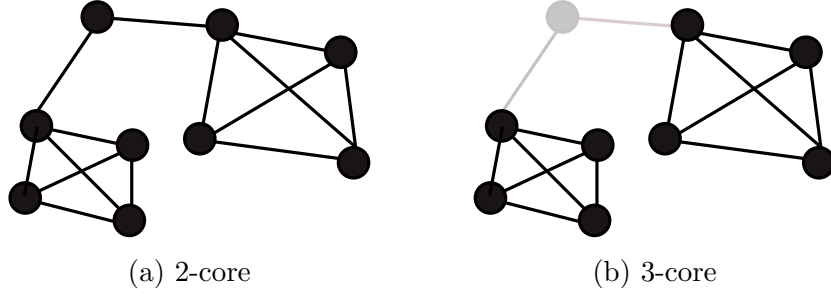


Figure 3: Example of a disconnected k -core subgraph. The 2-core is shown in (a). The removal of the node with degree 2 leads to (b) which depicts the 3-core (and also the maximum core of the graph). The 3-core is disconnected.

2.2 Extensions to Other Types of Graphs

The k -core decomposition described above considers that graphs are unweighted and undirected. However, many real-world networks carry rich semantics, as expressed by more complex graph types. To that end, there exist research efforts towards meaningful extensions of the k -core decomposition to other types of graphs. In most of the cases, these extensions pose additional challenges to the efficient computation of the decomposition as well.

2.2.1 Directed Graphs

Directed graphs or digraphs [13] are characterized by rich semantics in comparison to simple graphs, simply because edge direction is important. In a directed graph the degree of a node u may refer to the number of incoming links ($\deg_{in}(u)$) or to the number of outgoing links ($\deg_{out}(u)$). These are also known as the in-degree and the out-degree respectively.

Giatsidis et al. [64, 62] introduced D -cores, an extension of the k -core structure to *directed graphs*. In this case, the notion of (k, ℓ) -core is used to represent subgraphs in which all nodes have in-degree at least k and out-degree at least ℓ respectively.

2.2.2 Weighted Graphs

A weighted graph is characterized by the existence of weights on the graph edges. Each edge e is associated with a weight $w(e)$ that may represent the cost of the edge, or the strength of the link between the participating nodes, or any other type of quantification, depending on the application. Computing the core decomposition in a weighted graph is significantly harder than the computation in a simple graph, mainly because there is no easily derived bound on the core number of a node.

Definition 5 (Weighted k -core) Let $G(V, E, w)$ be a weighted graph with $w : E \rightarrow \mathbb{R}^+$ denoting a weight function that assigns weights on the edges of the graph. Let also $\deg_w(u)$ be the weighted degree of node u in (G, w) , i.e., $\deg_{G,w}(u) = \sum_{e \in N(u)} w(e)$. If (H, w') is a subgraph of G where w' is the restriction of w to the edges of H , we define $\Delta(H, w') = \min\{\deg_{H,w'}(u) | u \in V_H\}$. Moreover, the weighted degeneracy of G is defined as $\delta^*(G, w) = \max\{\Delta(H, w') | (H, w') \text{ is subgraph of } (G, w)\}$. Then, the k -core of (G, w) is the maximum-size subgraph (H, w') of (G, w) where $\Delta(H, w) \geq k$.

In [63, 54, 43], the authors propose efficient algorithms for computing the core decomposition in weighted graphs.

2.2.3 Signed Graphs

In [60], an extension of the k -core decomposition for *signed networks* was proposed. Signed networks [81, 82, 88] are used to capture the notion of positive and negative interactions among the nodes of a graph (e.g., trust/distrust, friend/foe relationships). Examples of such networks include the trust networks that can be produced by product review websites like Epinions¹ and the voting election network between the administrators of Wikipedia².

Let $G(V, E, w)$ be a signed directed graph, where $w : E \rightarrow \{+, -\}$ is a labeling of E , assigning either a positive or a negative sign on the edges of G . Given a vertex v of G , we denote by $\deg_{in}^+(v, G)$ (resp. $\deg_{out}^+(v, G)$) the positive in-degree (respectively positive outdegree) of v in G , i.e., the number of positive-signed edges tailing (respectively heading) on v . The negative in- and out-degrees of the vertices of G are defined in a similar manner and are denoted by $\deg_{in}^-(v, G)$ and $\deg_{out}^-(v, G)$.

Definition 6 (Signed (ℓ^t, k^s) -core) Let $G(V, E, w)$ be a signed graph. Let also $s, t \in \{+, -\}$ and $k, \ell \in \mathbb{N}$. We define the (ℓ^t, k^s) -core of G as the maximum size subgraph H of G where, for each vertex v of H , it holds that $\deg_{in}^s(v, H) \geq k$ and $\deg_{out}^t(v, H) \geq \ell$.

2.2.4 Bipartite Graphs

A graph $G(V_h, V_a, E_b)$ is called *bipartite* if the node set V can be partitioned into two disjoint sets V_h and V_a , where $V = V_h \cup V_a$, such that every edge $e \in E_b$ connects a node of V_h to a node of V_a , i.e., $e = (i, j) \in E \Rightarrow i \in V_h$ and $j \in V_a$. In other words, there are no edges between nodes of the same partition. A common approach to analyze bipartite networks is to project them into weighted or unweighted unipartite ones. Nevertheless, this simplistic approach has several drawbacks – with the major one being the fact that, a node with degree d in the original bipartite network will result in a d -clique in the projected one. Thus, hub nodes will dominate the maximal k -core subgraph produced by the core decomposition on the projected unipartite network. In an approach to generalize the concept of core decomposition in bipartite graphs under the context of recommender systems, where partitions correspond to users and items, Ding et al. [40] defined the concept of α - β -core, which constitutes a direct generalization of the traditional core decomposition.

Definition 7 (α - β -core) Let $G(V_h, V_a, E_b)$ be a bipartite graph and let $H(V'_h, V'_a, (V'_h, V'_a)|E_b)$ be a subgraph induced by $V'_h \subseteq V_h$ and $V'_a \subseteq V_a$. We call S as a candidate α - β -core, if $\forall u \in V_h$, $\deg_H(u) \geq \alpha$ and $\forall v \in V_a$, $\deg_H(v) \geq \beta$. A candidate α - β -core is a α - β -core if H is a maximal subgraph with respect to both α and β .

2.2.5 Dynamic Graphs

A *dynamic graph* is characterized by changes performed on the set of nodes and/or the set of edges. These changes may correspond to insertion or deletions of edges. Changes in the graph may have an impact on the core numbers of nodes. In the worst case, an insertion of a single edge may change all core numbers of nodes. On the other extreme, the insertion may cause no changes at all. The typical case is that an edge

¹www.epinions.com

²www.wikipedia.org

insertion (or deletion) will have an impact on some core numbers. Therefore, the challenge is to be able to monitor the core numbers of all nodes by applying only a few computations, avoiding the re-computation of the core decomposition.

In Section 3 we will discuss in detail, among others, a particular algorithm for monitoring the core decomposition in a dynamic graph reported in [127], where insertions and deletions of edges may be applied at any time. The main characteristic of that algorithm is that it detects the minimal set of nodes that must be checked for changes in the core numbers, thus reducing the overall processing cost significantly.

2.2.6 Temporal Graphs

A *temporal graph* is a special case of a dynamic graph. Two nodes u and v may be connected by an edge at multiple time instances or intervals. We may assume that each edge is annotated with a timestamp, denoting the time instance of the specific interaction. Also, two nodes may be linked for a specific time interval $[t_i, t_j]$ defined by two time instances t_i, t_j , where $t_i \leq t_j$.

In [157], the definition of the core decomposition is adapted to the case of temporal graphs. The concept of (k, h) -core is defined, where as usual k represents the degree of a node and h represents the number of multiple temporal edges between two vertices. Given a temporal graph G , the (k, h) -core of G is the largest subgraph $H_{(k,h)}$ of G such that every vertex u in $H_{(k,h)}$ has at least k direct neighbors, and there are at least h temporal edges between u and each one of its neighbors in $H_{(k,h)}$.

2.2.7 Probabilistic Graphs

A special category of graphs, includes graphs that introduce *uncertainty* with respect to the existence of nodes and edges [75]. For example, an edge e between nodes u and v may exist or not. The existence of an edge depends on several factors, mainly on the particular application under consideration. For example, in a social network where an edge corresponds to a message exchange between two users, the message will be sent with some probability (i.e., it is not sure that user u will send a message to user v). As another example, consider a protein-protein interaction network, where each node corresponds to a protein and each edge denotes interactions among proteins. In this case, we may realize that proteins u and v interact in 70% of the cases, which means that the edge $e_{u,v}$ will be present in the graph with a probability of 0.7. Also, uncertainty may be introduced on purpose for privacy reasons.

Computing the core decomposition of an uncertain graph is not trivial. One approach could be to transform the uncertain graph into a weighted graph, where the weight $w(e)$ of an edge e is inversely proportional to the existential probability $p(e)$ of the edge. However, this simplistic approach has severe drawbacks, since the meaning of the probability is lost during this transformation and it does not give any insight regarding the importance of the computed cores. To attack the problem, Bonchi et al. [24] proposed a core decomposition methodology for uncertain graphs. The problem was also studied later in [121]. The algorithm of [24] will be covered in detail in the following section that covers the major algorithmic techniques.

2.2.8 Multilayer Graphs

Usually, we assume that graph nodes are of the same type and also graph edges represent the same relationship among nodes. However, in many cases this simple view of the network may not represent reality. People interact in many different ways. For example, two persons may be friends in real life, but

also may be friends in a social network, may collaborate in a research project or may work in the same company. These are different relationship types that may be present.

In its simplest form, a *multilayer graph* (a.k.a. *multidimensional graph*) $G(V, \mathcal{E})$ is composed of a set of nodes V and a set of edge subsets $\mathcal{E} = E_1 \cup E_2 \cup \dots \cup E_\ell$, where ℓ is the total number of layers and E_j contains the set of edges present in the j -th layer of G . The first algorithm for computing the core decomposition of a multilayer graph is reported in [53].

Definition 8 (Multilayer k -core) *Given a multilayer graph $G(V, \mathcal{E})$ and an ℓ -dimensional integer vector \mathbf{k} , the multilayer \mathbf{k} -core of G is a maximal subgraph $G[C] = (C \subseteq V, E[C])$ such that $\mu(C, \ell) \geq k_\ell, \forall \ell$, where $\mu(C, \ell)$ is the minimum degree of a node in layer ℓ . Vector \mathbf{k} is referred to as coreness vector of $G[C]$.*

The authors of [53] not only provide a novel definition for the core numbers on a multilayer graph but also show that this definition has some nice properties regarding the density of the k -core subgraphs which are in sync with the core decomposition concept in simple graphs.

2.2.9 Hidden Graphs

Conventional graphs are characterized by the fact that both the set of vertices V and the set of edges E are known in advance, and are organized in such a way to enable efficient execution of basic tasks. Usually, the adjacency lists representation is being used, which is a good compromise between space requirements and computational efficiency. However, a concept that recently has started to gain significant interest is that of *hidden graphs* [11]. In contrast to conventional graphs, a hidden graph is defined as $G(V, f())$, where V is the set of vertices and $f()$ is a function $V \times V \rightarrow \{0, 1\}$ which takes as an input two vertex identifiers and returns true or false if the edge exists or not respectively. Therefore, in a hidden graph the edge set E is not given explicitly and it is inferred by using the function $f()$.

Hidden graphs constitute an interesting tool and a promising alternative to conventional graphs, since there is no need to represent the edges explicitly. This enables the analysis of different graph types that are implicitly produced by changing the function $f()$. Note that the total number of possible graphs that can be produced for the same set of vertices equals $2^{\binom{n}{2}}$, where $n = |V|$ is the number of vertices. It is evident, that the materialization of all possible graphs is not an option, especially when n is large. Therefore, hidden graphs is a tempting alternative to model relationships among a set of entities. On the other hand, there are significant challenges posed, since the existence of an edge must be verified by evaluating the function $f()$, which is costly in general.

Motivated by recent developments in the area [140, 159] for detecting the top- k nodes with the highest degrees in bipartite graphs, in [139], an algorithm is proposed to discover if a hidden graph contains a k -core subgraph or not, by applying as few edge probing queries as possible.

2.3 Generalized Core Decomposition

The k -core decomposition was initially introduced for the degree property of the nodes in a graph. A natural inquiry would be “why do we focus on node degrees?”. An obvious answer to this question would be that the degree of a node is relatively easily computed and it is a very simple concept. However, the degree is not the only node property that could be applied in this framework. Batagelj and Zaveršnik proposed the notion of *generalized cores* [17], which extends cores from degree to other node properties. In fact, any node property can be used potentially to define a different kind of core decomposition, where the concept of the core is associated with the node property under consideration.

Definition 9 (Generalized Cores or p -cores) Let $G = (V, E)$ be a graph and let $w : E \rightarrow \mathbb{R}$ be a function assigning values (or weights) to the edges of the graph. A node property function $p()$ that assigns real values on graph G , is defined as $p(v, C)$, where $v \in V$ and $C \subseteq V$. Then, a subgraph $H(V_H, E_H)$ induced by the set $V_H \subseteq V$ is called a p -core at level $t \in \mathbb{R}$ if and only if (i) $\forall v \in V_H : t \leq p(v, V_H)$ and (ii) V_H is a maximal set.

Recall that, a function $p()$ is called *monotone* if and only if the following property holds:

$$C_1 \subseteq C_2 \Rightarrow p(v, C_1) \leq p(v, C_2), \forall v \in V. \quad (2)$$

In [17] it was shown that for a monotone function p , the p -core at level t of the decomposition can be found by successively removing nodes with value of p less than t – as has been already described for the k -core decomposition. Furthermore, the subgraphs corresponding to the cores are nested, i.e., $t_1 < t_2 \Rightarrow H_{t_2} \subseteq H_{t_1}$. In fact, if we consider that function p corresponds to the degree of a node, i.e., $p(v, C) = d_v^C$, where d_v^C is the degree of node v in subgraph C , this function is monotone. Also, many other functions on the nodes v of the graph including the in-degree and out-degree, the weighted degree (i.e., sum of weights of the adjacent edges) and the number of cycles of length l that pass through node v , have been proven to be monotone; thus, the same procedure can be used to extract the corresponding p -cores.

2.4 Extensions of the Core Decomposition

In this section, we describe various extensions of the core decomposition, covering among others the notion of truss decomposition – a particular type of generalized cores based on the property of triangles.

2.4.1 Truss Decomposition

The K -truss decomposition extends the notion of k -core using triangles, i.e., cycle subgraphs of length three [34, 150, 166]. Let $G(V, E)$ be an undirected graph. We define as a *triangle* Δ_{uvw} a cycle subgraph of nodes $u, v, w \in V$. Additionally, the set of triangles of G is denoted by Δ_G . The *support* of an edge $e = (u, v) \in E$ is defined as $sup(e, G) = |\{\Delta_{uvw} : \Delta_{uvw} \in \Delta_G\}|$ and expresses the number of triangles that contain edge e .

Given an undirected graph G , the K -truss, $K \geq 2$, denoted by $T_K(V_{T_K}, E_{T_K})$, is defined as the largest subgraph of G , where every edge is contained in at least $K - 2$ triangles within the subgraph, i.e., $\forall e \in E_{T_K}, sup(e, T_K) \geq K - 2$. Based on that, the *truss number* of an edge $e \in G$ is defined as $t_{edge}(e) = \max\{K : e \in E_{T_K}\}$. Thus, if $t_{edge}(e) = K$, then $e \in E_{T_K}$ but $e \notin E_{T_{K+1}}$. We use K_{\max} to denote the maximum truss number of any edge $e \in E$. The K -class of a graph $G(V, E)$ is defined as $\Phi_K = \{e : e \in E, t_{edge}(e) = K\}$. Figure 4 shows an example graph and its K -classes.

Based on the above definitions, we can now introduce the concept of K -truss decomposition.

Definition 10 (K -truss decomposition) In a graph $G(V, E)$, the K -truss decomposition is defined as the task of finding the K -truss subgraphs of G , for all $2 \leq K \leq K_{\max}$. That is, the K -truss can be obtained by the union of all edges that have truss number at least K , i.e., $E_{T_K} = \bigcup_{j \geq K} \Phi_j$.

Since the K -truss decomposition is defined based on the number of triangles – a more “strict” criterion compared to the one of degree – it can intuitively be considered as the “core” of the k -core subgraph.

Lastly, we mention here that the concept of K -truss decomposition has recently been extended to the case of probabilistic (or uncertain) graphs [72, 170].

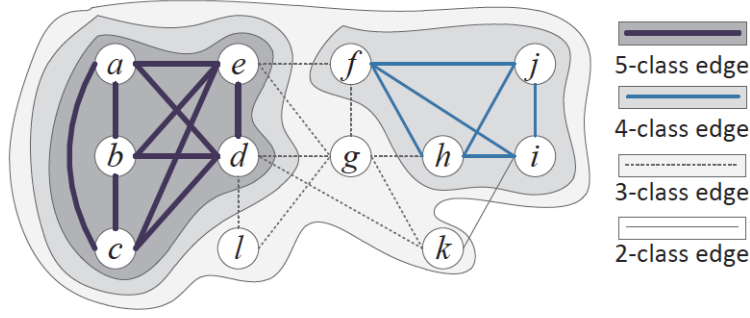


Figure 4: Example of a graph and its K -classes, $2 \leq K \leq 5$. The figure is courtesy of Wang and Cheng [150]. ©2012 VLDB Endowment.

2.4.2 Density-friendly Core Decomposition

One of the drawbacks of the k -core decomposition is that the nested k -core subgraphs do not satisfy a natural density property – simply defined as the ratio between the number of edges and nodes of the subgraph. In other words, the maximal k -core subgraph is not necessarily the densest subgraph of the graph. Based on this observation, Tatti and Gionis [141] introduced the concept of *density-friendly* graph decomposition, where *i*) the density of the inner core subgraphs given by the decomposition is higher than the density of the outer ones, and *ii*) the most inner subgraph will correspond to the densest subgraph. Furthermore, the authors of [141] have shown that the locally-dense decomposition can be computed in polynomial time. Note that, more recently, Danisch et al. [38] proposed a scalable algorithm for computing such a decomposition, based on convex programming.

2.4.3 Peak Decomposition

Another drawback of the k -core decomposition defined earlier has to do with the fact that it is computed globally; if the graph contains distinct regions of different densities, the sparser among these regions might be neglected by the decomposition. To deal with this issue, the authors of [66] have proposed the notion of k -peak decomposition, which aims at finding the centers of distinct regions in the graph – viewing the global structure of the graph as a set of regions, each one resembling a mountain with a central peak.

More precisely, given a graph G , the k -contour can be defined as follows:

Definition 11 (k -contour) *Given a graph $G(V, E)$, a k -contour is the induced subgraph over the maximal set of nodes, such that *i*) the k -contour does not contain nodes from a higher contour (i.e., values higher than k), and *ii*) each node in the k -contour has at least k links to other nodes in the k -contour.*

Based on that, we can define the *peak number* of a node as the value k such that the node belongs to a k -contour. Then, a k -peak decomposition of a graph G is defined as the assignment of each node to exactly one contour. Figure 5 shows an example graph and the corresponding k -core and k -peak decomposition.

Definition 12 (k -peak) *Given a graph G , a k -peak is the induced subgraph of the union of j -contours, $\forall j \leq k$.*

Lastly, as shown in the paper [66], similar to the case of k -core decomposition, the k -peak decomposition is also unique, i.e., each node has a single unique peak number.

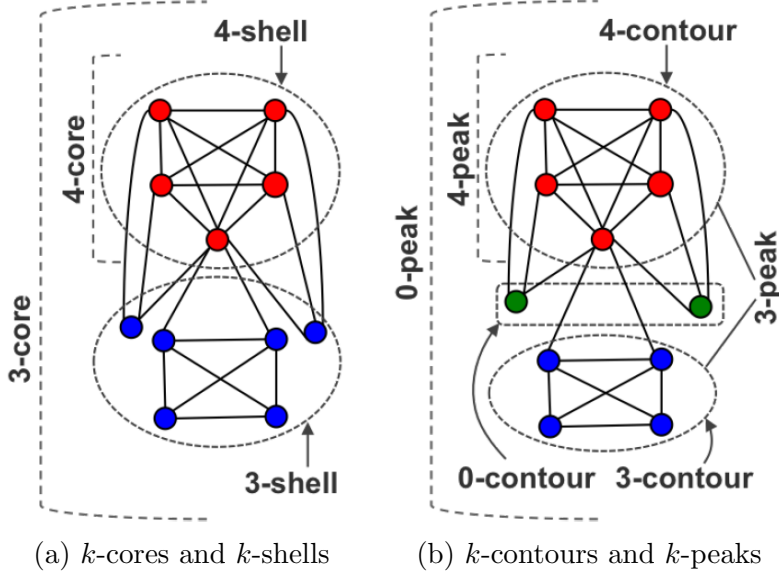


Figure 5: Example of the k -core and k -peak decomposition. The peak number of a node is at most its core number. The figure is courtesy of Govindan et al. [66]. ©2017 International World Wide Web Conference Committee.

2.4.4 Nucleus Decomposition

Another extension of the core decomposition is the *nucleus decomposition* [130]. The basic motivation here comes from the problem of dense subgraph detection, where the goal is to identify structures of dense subgraphs within a large graph and to understand how those structures are related to each other. In particular, we are interested in extracting a global, hierarchical representation of many dense subgraphs.

To this direction, the authors of [130] defined the notion of *nuclei* in a graph: an (r, s) -nucleus, for fixed and small positive integers $r < s$, is defined as a maximal subgraph where every r -clique (i.e., complete graph of r nodes) is part of many s -cliques. Furthermore, nuclei subgraphs that do not contain one another, cannot share an r -clique. Based on that, for various values of r and s ($r < s$), it can be shown that the (r, s) -nuclei form an hierarchical decomposition of the graph – where the density of the nuclei is increasing as we move towards the leaves of the decomposition. In practice, the authors of [130] have observed that the $(3, 4)$ -nuclei provide the most interesting decomposition of real-world graphs. Figure 6 depicts an example of the hierarchical structure of $(3, 4)$ -nuclei decomposition in a snapshot of the Facebook graph composed by 88K nodes. Each node of the structure corresponds to a $(3, 4)$ -nucleus, and the tree edges indicate containment. More generally, an ancestor nucleus contains all descendant nuclei. The figure also shows the scale and densities of the various nuclei subgraphs.

2.4.5 Distance-based Core Decomposition

In the basic formulation of the k -core decomposition, the k -core subgraph is defined as the maximal subgraph where every node has at least k neighbors within the subgraph. Bonchi et al. [25] introduced a generalization of the k -core decomposition, taking into account the shortest path distance among nodes. In particular, they proposed the concept of *distance-generalized* core decomposition, referring to as the (k, h) -core, i.e., the maximal subgraph in which every node has at least k neighbors within the subgraph at distance smaller or equal to h .

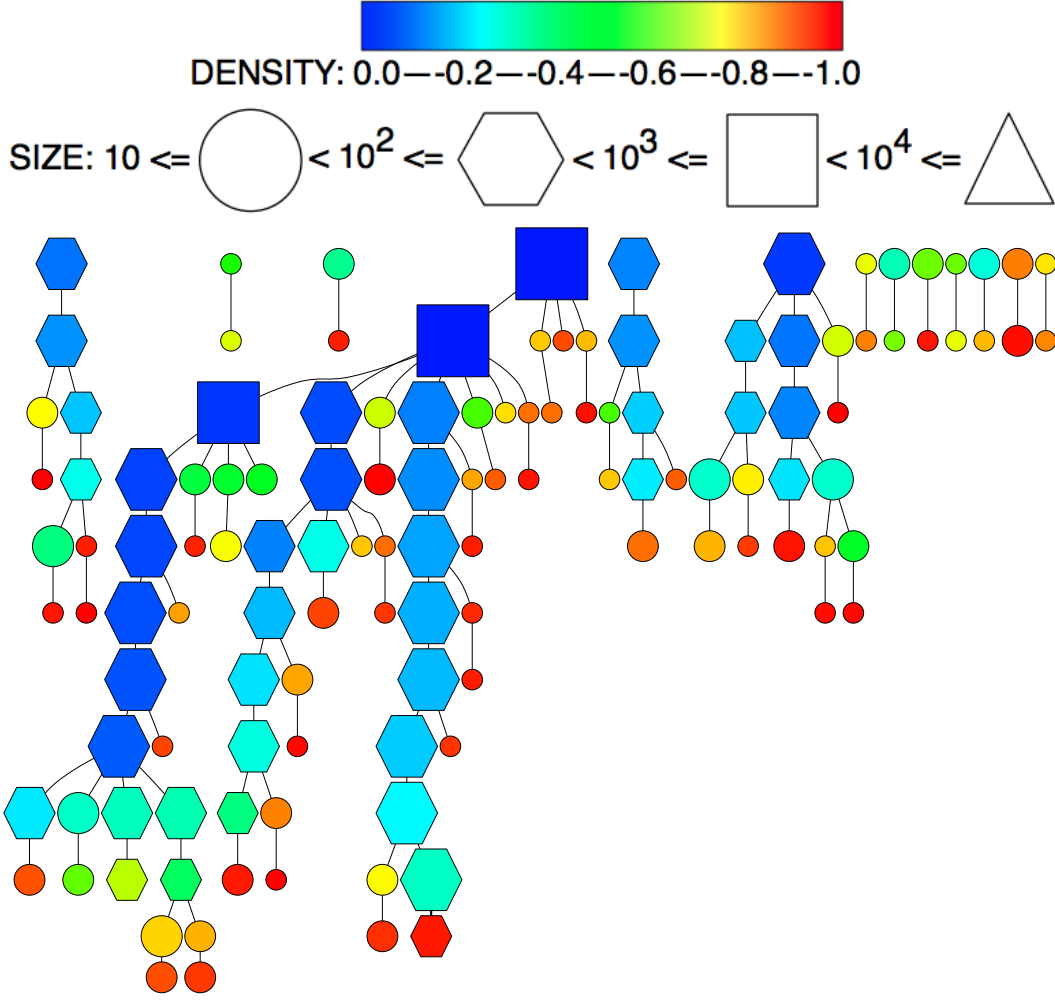


Figure 6: (3,4)-nuclei subgraphs for a snapshot of the Facebook graph [130]. Branching depicts the different regions in the graph. The figure is courtesy of Sarıyüce et al. [130]. ©2015 International World Wide Web Conference Committee (IW3C2).

Definition 13 (h -degree of a node) Let $G(V, E)$ be an undirected graph, and let $G[S] = (S, E[S])$ be the subgraph induced by S . Given a positive integer $h \in \mathbb{N}^+$, the h -neighborhood of a node $v \in S$ in $G[S]$ is defined as $N_{G[S]}(v, h) = \{u \in S \mid u \neq v, \deg_{G[S]}(u, v) \leq h\}$, where $d_{G[S]}(u, v)$ is the shortest-path distance between u and v computed on $G[S]$. The h -degree of a node v is defined as the size of its h -neighborhood, i.e., $\deg_{G[S]}^h(v) = |N_{G[S]}(v, h)|$.

Definition 14 ((k, h) -core) Given a threshold integer $h \in \mathbb{N}^+$ and an integer $k \geq 0$, the (k, h) -core (or h -neighborhood k -core) of a graph G is a maximal subgraph $G[C_k](C_k, E[C_k])$ in which all nodes have h -degree at least k , i.e., $\deg_{G[C_k]}^h(v) \geq k$.

Note that, the traditional k -core decomposition is an instance of the (k, h) -decomposition, where the degree of a node v corresponds to the number of nodes in the graph which have distance ≤ 1 from v , i.e., the size of the 1-neighborhood of v . Figure 7 shows an example of the distance-based core decomposition. Contrary to the traditional k -core decomposition which can also be interpreted as the $(k, 1)$ -core (Fig. 7 (left)), the $(k, 2)$ -core decomposition (Fig. 7 (right)) is able to detect groups of nodes (nodes 4 to 13) that form a denser and more well-structured region.

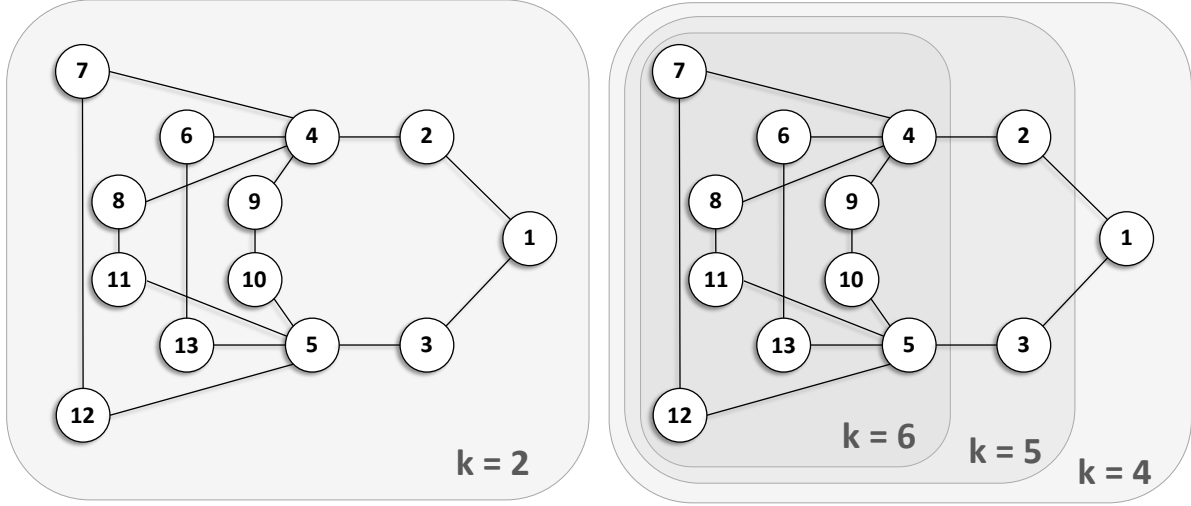


Figure 7: Example of (k, h) -core decomposition. (Left) The traditional k -core decomposition ($k = 2$), which corresponds to the $(k, 1)$ -core under the distance-generalized setting. (Right) The $(k, 2)$ -core decomposition. The figure is courtesy of Bonchi et al. [25]. ©2019 Association for Computing Machinery.

In a similar spirit, Zhang et al. [164] extended the notion of core decomposition to capture pairwise similarities among nodes, based on node attributes. In particular, the concept of (k, r) -core was introduced – being able to capture both graph structural constraints as well as node similarity constraints. Let $G(V, E, A)$ be an attributed graph, where A denotes the attributes of the nodes. Let $\text{sim}(u, v)$ be the similarity of two nodes $u, v \in V$, based on their attribute values (e.g., user profile information and interests in social networks). Given a positive integer k and a subgraph S , the *structural constraint* refers to the traditional degree-based requirement of core decomposition, i.e., $\deg(v, S) \geq k \forall v \in S$. In addition, for a node v and a similarity threshold r , $DP(v, S)$ denotes the number of nodes in S that are dissimilar to v (based on the threshold r). Then, we say that subgraph S satisfies the *similarity constraint* if $DP(v, S) = 0$ for each node $v \in S$. Based on the structural and similarity constraints, we can define the notion of (k, r) -core.

Definition 15 ((k, r) -core) We say that a subgraph S is a (k, r) -core if S satisfies both the structural and similarity constraints.

2.4.6 k -tip Decomposition

Motivated by the task of dense subgraph detection in bipartite networks, Sariyüce and Pinar [128] introduced the concept of bipartite *graph peeling* to detect dense subgraphs and the relationships among them. The main idea is to rely on *higher-order structural motifs* [19] that are able to capture cohesiveness in bipartite graphs. More precisely, the authors of [128] have used a particular motif, called *butterfly* subgraph, which corresponds to a $(2, 2)$ -biclique (bipartite clique with 2 nodes at each partition) – the overall goal is to discover bipartite subgraphs including many butterfly structures and to construct relations among them. Then, two types of decomposition can be defined: the *tip decomposition* and the *wing decomposition*.

Definition 16 (k -tip decomposition) A bipartite subgraph $H(U, V, E) \subset G$ induced on U , is a k -tip if and only if: i) each node $u \in U$ takes part in at least k butterflies; ii) each node pair $(u, v) \in U$ is connected by a series of butterflies; iii) H is maximal, i.e., there is no other k -tip that subsumes H .

Note that, two nodes $u, w \in U$ are connected by a series of butterfly subgraphs, if there exists a sequence of nodes $u = v_1, v_2, \dots, v_k = w$ such that some butterfly contains v_i and v_{i+1} , for each i .

Definition 17 (Tip number $\theta(u)$) *The tip number $\theta(u)$ of a node $u \in U$ is the largest value t such that there exists a t -tip that contains u . Then, the tip decomposition of a bipartite graph $G(U, V, E)$ is to find the tip numbers of nodes in U .*

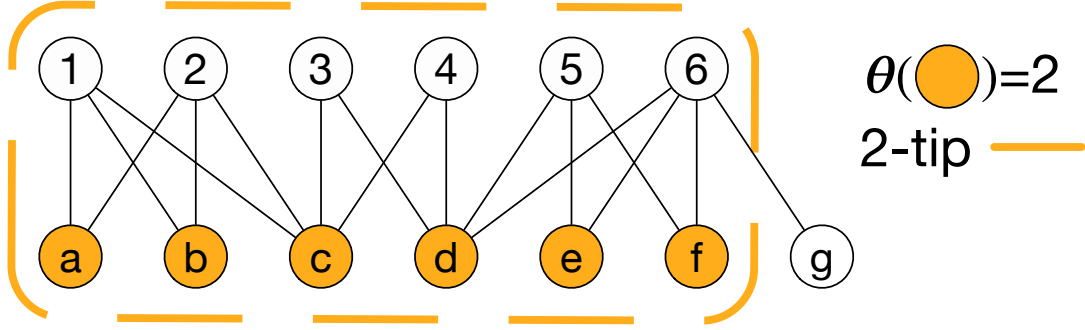


Figure 8: Example of k -tip decomposition. Nodes a, b, e , and f participate at two butterfly subgraphs, while nodes c and d at three. Notice that, nodes c and d cannot have a tip number of 3 since their induced subgraph has just one butterfly. Therefore, nodes a to f form a 2-tip, and their tip number will be $\theta(a - f) = 2$. The figure is courtesy of Saryüce and Pinar [128]. ©2018 Association for Computing Machinery.

Figure 8 gives an example graph and its k -tip decomposition. Notice that, the k -tip decomposition does not allow node overlaps, which might be the case in many real-world bipartite networks (e.g., author-paper collaboration networks). To allow node overlaps, the authors of [128] have introduced the concept of k -wing decomposition, where the focus is on the edges of the network instead of the nodes.

Definition 18 (k -wing decomposition) *A bipartite subgraph $H(U, V, E) \subset G$ induced on U , is a k -tip if and only if: i) each edge $(u, v) \in E$ participates in at least k butterflies; ii) each edge pair $(u_1, v_1), (u_2, v_2) \in E$ is connected by a series of butterflies; iii) H is maximal, i.e., there is no other k -wing that subsumes H .*

Definition 19 (Wing number $\psi(e)$) *The wing number $\psi(e)$ of an edge $e \in E$ is the largest value s such that there exists a s -wing that contains e . Then, the wing decomposition of a bipartite graph $G(U, V, E)$ is to find the wing numbers of edges in a graph G .*

2.4.7 Radius-Bounded Decomposition

Motivated by the fact that *geo-social networks* (a.k.a. *geo-spatial social networks* or *location-based social networks* (LBSNs)) have emerged recently [14], the work in [151] studies the adaptation of the core decomposition process in cases where additional spatial information is available. More precisely, let $G(V, E)$ be a geo-social graph, where each vertex $v \in V$ is associated with location information $(v.x, v.y)$, which corresponds to the position of v in the two-dimensional space. Given a set of vertices, we can now define the notion of *minimum covering circle* (MCC).

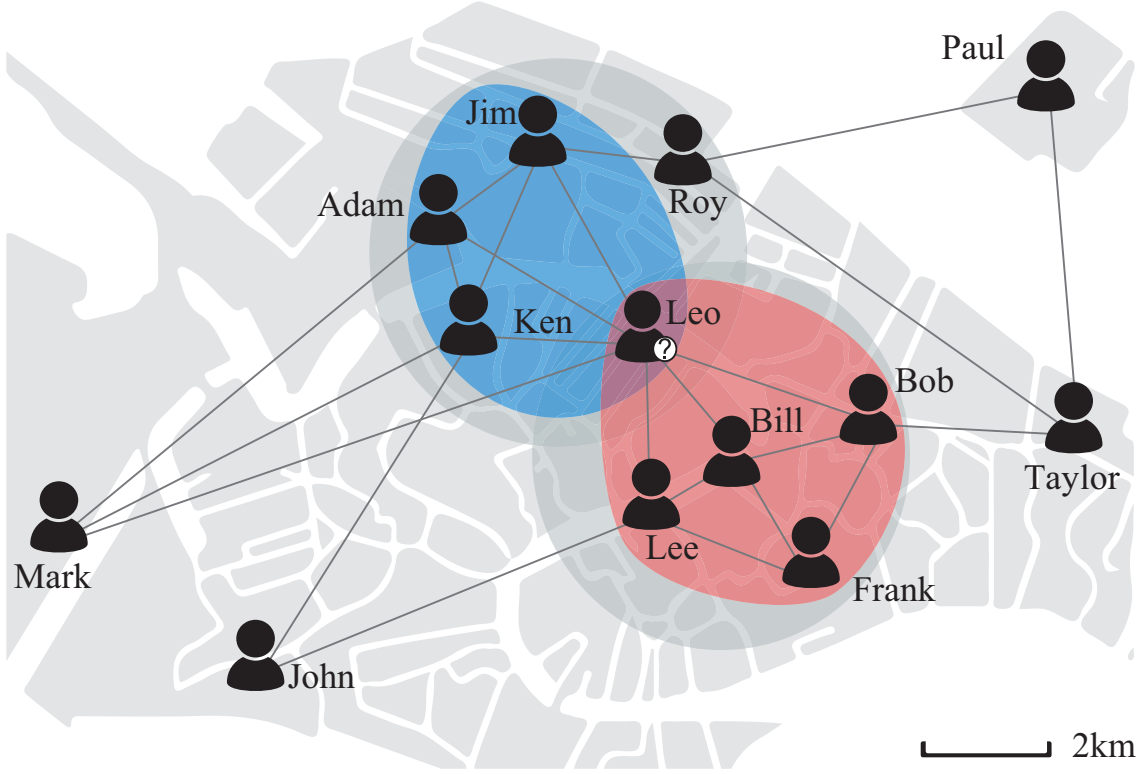


Figure 9: Example of a geo-social network, where edges represent friendship relationships among users and locations on the map correspond to home locations of users. The figure is courtesy of Wang et al. [151]. ©2018 IEEE.

Definition 20 (Minimum Covering Circle (MCC)) *Given a set of vertices $S \subseteq V$, the minimum covering circle of S is the circle which encloses all the vertices $v \in S$ with the smallest radius. The vertices which lie on the boundary of an MCC are called boundary vertices.*

Based on the definition of MCC, we can now define the concept of *radius-bounded k -core*, which extends k -core decomposition to geo-social graphs.

Definition 21 (Radius-bounded k -core) *Given a geo-social network $G(V, E)$, a vertex $w \in V$, a positive integer k and a query radius r , subgraph $G_k^r \subseteq G$ is defined as a radius-bounded k -core if the following constraints are satisfied:*

1. *Connectivity constraint: $G_k^r \subseteq G$ is a connected subgraph that contains w ;*
2. *Social constraint: for all vertices v of G_k^r , $\deg_{G_k^r}(v) \geq k$;*
3. *Spatial constraint: the MCC of the nodes belonging to G_k^r has a radius $r' \leq r$;*
4. *Maximality constraint: there exists no other radius-bounded k -core $G_k^{r'} \supsetneq G_k^r$ satisfying constraints (1), (2) and (3).*

Figure 9 shows an example of a geo-social network, where edges correspond to friendship relationships among users and spatial locations on the map correspond to home locations of users. Setting $r = 3$ and $k = 3$, there are two radius-bounded k -core subgraphs recommended to the user named Leo, as illustrated by the two shadowed areas, i.e., $\{\text{Leo}, \text{Ken}, \text{Jim}, \text{Adam}\}$ and $\{\text{Leo}, \text{Bill}, \text{Frank}, \text{Bob}, \text{Lee}\}$.

3 Algorithmic Techniques

In this section, we focus on algorithmic techniques for the computation of the k -core decomposition in different settings. The design of an algorithm for core decomposition depends on many diverse factors such as the type of the graph (simple, directed, signed, weighted, probabilistic), the hardware infrastructure (memory-based, disk-based, parallel, distributed), the type of the output (exact, approximate), just to name a few. The goal of this section is to provide details for representative algorithmic techniques that solve the k -core decomposition problem under different settings. Since there is a plethora of different algorithmic techniques we are going to discuss the most representative ones. In addition, we note that the techniques presented in this section are generic, meaning that potentially they can be applied to many different types of graphs reported in Section 2.

3.1 In-Memory Computation

The computation of the k -core decomposition of a graph can be performed through a simple process that is based on the following rationale: to extract the k -core subgraph, all nodes with degree less than k and their adjacent edges should be recursively deleted [133]. That way, beginning with $k = 1$, the algorithm removes all the nodes (and the incident edges) with degree equal to or less than k , until no such node remains in the graph. Also notice that, removing edges that are incident to a node may cause reductions to the degree of neighboring nodes; the degree of some nodes may become at most k , and thus, they should also be removed at this step of the algorithm. When all remaining nodes have degree $d_v > k$, k is increased by one and the process is repeated until no more remaining nodes are left in the graph.

A straight-forward implementation of this algorithm requires a priority queue to store the nodes, prioritized by their degree. The removal of a node requires in the worst case the deletion of $n - 1$ edges which translates to the execution of $n - 1$ decrease-key operations in the priority queue. Since each edge is examined exactly once, the worst case complexity of the naive algorithms becomes $\mathcal{O}(m \cdot \log n)$.

However, the problem can be solved in linear time, by using bin-sorting as it was demonstrated in [18] by Batagelj and Zaveršnik. However, the same idea was applied by Matula and Beck [105]. All nodes are maintained sorted with respect to their degrees by using a comparison-free sorting algorithm which maintains separate bins for each degree value. Clearly, for a graph with n nodes, the minimum degree of a node is 1 (assuming no isolated nodes exist) and the maximum is $n - 1$. Thus, by keeping an in-memory array of all possible degree values and keeping track of bin boundaries, each edge deletion can be handled in $\mathcal{O}(1)$ time, resulting in a total complexity of $\mathcal{O}(m + n)$. Note that, maintaining the k -core decomposition of a graph is equivalent of keeping the core number $c(u_i), \forall u_i \in V$.

3.2 Disk-Resident Graphs

A natural extension to the previous algorithm is based on the fact that most of the interesting real-world networks are too large to fit in main memory. Therefore, we need efficient algorithmic techniques for providing the core decomposition in cases where the graph is stored in secondary storage (i.e., on disk).

The first algorithm (EMCORE) to attack the problem in secondary storage was reported in [32]. EMCORE assumes that the graph resides on disk and it performs the following steps: *i*), graph partitioning, *ii*) core number estimation and *iii*) recursive top-down processing.

Graph partitioning. The purpose of this step is to decompose the graph into small subgraphs, so that core number computation can be performed in each subgraph separately. Evidently, each of these subgraphs must fit in main memory.

The algorithm scans the input graph only once and partitions the vertex set V into a set of p mutually disjoint vertex subsets $U = \{U_1, \dots, U_p\}$, where $V = \bigcup_i U_i$ and for any i, j , $U_i \cap U_j = \emptyset$. The graph partitioning algorithm starts reading nodes from the disk-resident graph and maintains in memory as many nodes as possible. Let U_{mem} denote the current memory-resident part of the graph. For any examined node u , if u is not connected to any of the vertex subsets currently in U_{mem} , a new partition is created and u becomes the only member of it. Otherwise, u is assigned to the vertex subset to which u has the most connections. In case we reach the block limits, a new block is flushed to the disk. Also, if the memory capacity is reached, the largest partition of the memory-resident part U_{mem} is flushed to the disk. This process is continuous until all nodes have been scanned.

Core number estimation. For each subgraph determined by the partitioning process, the upper bound of the core number of each node is determined. This bound becomes tighter during the course of the algorithm. Initially, the upper bound $\psi(u)$ of the core number of each node u , is set to the degree of u , i.e., $\psi(u) = \deg(u)$.

The refinement of the upper bound is performed by using the following observation: for a vertex u let Z denote the subset of the neighbors of u such that their upper bound is strictly less than $\psi(u)$. Then, for any nonempty subset $Z' \subseteq Z$, if $\deg(u) - |Z'| < \psi(u)$, then the upper bound $\psi(u)$ can be tightened as $\psi(u) = \max\{\deg(u) - |Z'|, \psi_{max}(Z')\}$, where $\psi_{max}(Z')$ denotes the maximum upper bound found in Z' .

The algorithm is executed every time a block needs to be flushed to disk during the partitioning process described previously. Note that in order to determine the subset Z' , there is no need to generate all possible subsets of Z which would lead to checking $2^{|Z|}$ subsets. Instead, nodes in Z are sorted in decreasing order of $\psi(u)$ and keep the first nodes in the order that minimize the value of $f(Z')$. This results in a much more efficient computation, since we only need to check at most $|Z|$ subsets.

Recursive top-down processing. Based on the upper bound of the core number, the k -classes are recursively computed for a convenient value of k . The value of k is determined so that the relevant subgraph can fit in main memory. More specifically, a range of values $[k_l, k_u]$ is determined, where $k_l \leq k_u$. The target is to determine the value of k_l , given the value of k_u and b , which is the maximum number of blocks that can be accommodated in main memory. The value of b is simply $\lfloor M/B \rfloor$, where M is the total memory capacity and B is the size of the disk block.

Let $\Psi_{k_l}^{k_u}$ denote the subset of nodes with core upper bound estimation in the range $[k_l, k_u]$, i.e., $\Psi_{k_l}^{k_u} = \{u : u \in V, k_l \leq \psi(u) \leq k_u\}$. At each recursive step, the algorithm constructs a subgraph that is relevant for the computation of the exact core numbers of the nodes $v \in \Psi_{k_l}^{k_u}$.

Let K be the set of values of k such that the nodes in the set $\Psi_k^{k_u}$ are distributed in at most b different vertex sets of U . Formally, the set K is computed as follows:

$$K = \{k : 1 < k \leq k_u, |U_x : U_x \in U, U_x \cap \Psi_k^{k_u} \neq \emptyset| \leq b\}$$

Based on this, the value of k_l is set to $\min\{k : k \in K\}$ if $K \neq \emptyset$ or it is set to k_u otherwise. This way, at each recursive step, the algorithm loads as many parts of the graph as possible into main memory. Since both k_l and k_u are determined, the algorithm proceeds with the computation of the core numbers of all nodes in $\Psi_{k_l}^{k_u}$ by loading in main memory the corresponding subgraphs. Then, the core number of each node u currently in main memory is refined. Before the execution of the next recursive step, the nodes that have their core number refined are removed from main memory (together with the corresponding edges) after *depositing* some important bookkeeping information that may be needed by the next recursive

step.

In conclusion, EMCORE manages to compute the core number of all nodes without the requirement that the graph fits in main memory. This fact enables the computation of the core decomposition of large disk-resident graphs. The algorithm performs $\mathcal{O}(k_{max})$ iterations over a graph G , where k_{max} is the maximum core number of G .

A limitation of EMCORE is that it may require a significant number of I/O operations to detect the appropriate partitions. In [65] a space-efficient algorithm is proposed (NIMBLECORE), that provides accurate estimates of the core numbers by using $\mathcal{O}(n)$ space for graphs with power-law degree distribution and $\mathcal{O}(n \log d_{max})$ space for arbitrary graphs, where d_{max} is the maximum node degree. Another implementation of EMCORE has been reported in [77], which is based on GraphChi [84]. In the same work, the classic code decomposition algorithm [18] is implemented using the Webgraph [22] framework. Webgraph is a graph compression framework that provides efficient access to a compressed graph.

The limitations of EMCORE are highlighted in [153] especially concerning the amount of main memory it might need. The authors note that the worst case memory complexity is that of $\mathcal{O}(n+m)$ and provide a better bounded algorithm for computing the degeneracy of a graph while maintaining most of it stored in disk. Their approach relies on maintaining only the vertices of the graph in memory with some additional properties and having the edges stored in disk. They start with the basic property of locality for the core number of a vertex that is also used in [109] for the design of a distributed k-core algorithm.

The locality property states that the core number $core(v)$ of a node $v \in V$ can be calculated based on the following recursive function :

$$core(v) = \max k \text{ s.t. } |\{u \in N(v) | core(u) \geq k\}| \geq k \quad (3)$$

The intuition of this function is that one may start with an upper bound of the core number for each node and keep updating it based on the core numbers of its' neighbors until it converges. The upper bound can be the degree of the node as the core number cannot be higher than that.

The authors of [153] start with this as their first (out of three) and most naive algorithm (SemiCore) and calculate the core numbers while having to load the list of neighbors from disk (where the edges of the graph are stored). They further optimize the I/O requests (from the disk) of SemiCore by noticing that not all nodes need to be checked at each iteration. In fact if the neighbors of a node v have not been updated in iteration i then v does not need to be updated with the aforementioned equation in iteration $i + 1$. By maintaining the range of node ids that have been updated (with a v_{min} and v_{max} value) they present SemiCore+ which does less requests as it skips the neighbor loading for nodes outside the (v_{min}, v_{max}) range. The final improvement (SemiCore*) is introduced by maintaining for each node v the number of neighbors u with $core(u) \geq core(v)$. This information remains in memory and when they calculate the core of v they update the count for each u . Then on each iteration if the count is larger than $core(v)$ the node is skipped as it already satisfies the minimum requirement for having its current core value (i.e., the number of neighbors with greater or equal core number are enough to support the current core number of v).

The authors of [153] present again their improvements in their subsequent work of [154] where they also tackle the problem of degeneracy ordering. In this problem, we may consider a topological (acyclic) ordering of the graph where the nodes are ordered arbitrarily and the original non-directed edges are given a direction from the earliest node to the latest. We can define the degeneracy of the ordering as the maximum out-degree and the degeneracy of the graph can be shown to be the minimum degeneracy of any ordering [48]. An ordering whose degeneracy is equal to the original degeneracy of the non-direct graph is an degenerate ordering. This task of ordering a graph falls outside the scope of this survey but

the specific ordering can be useful in other algorithmic problems (e.g., finding shortest paths in a graph).

In [154] they utilize the algorithms first introduced in [153] to build degenerate orderings and maintain when the graph is updated. Given an arbitrary original ordering, their solution to build the degenerate ordering relies on calculating the core value of each node with SemiCore* and then add the nodes that have a degree lower than the graph degeneracy at the end of the ordering while removing them from the graph.

3.3 Core Decomposition in Dynamic Environments

The previous techniques discussed so far mainly assume that the input graph is available in the very beginning of the execution. However, many modern applications are characterized by frequent updates, meaning that the structure of the graph may change over time by insertions/deletions of nodes or edges. Evidently, if the graph changes then one can reevaluate the core numbers by running the core decomposition algorithm again. Although this approach will provide the correct results, it is expected that the performance will degrade, especially if updates are frequent.

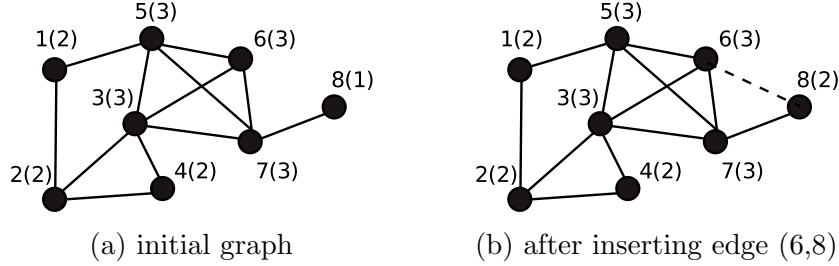


Figure 10: Changes performed to the core numbers after inserting a new edge linking nodes 6 and 8. Only the core number of node 8 needs to be updated.

As an example, consider the graph shown in Figure 10(a). The core numbers are shown in parentheses near each node identifier. Assume that a new edge between nodes 6 and 8 is inserted. The updated core numbers are shown in Figure 10(b). It is evident that only the core numbers of node 8 will change from 1 to 2. No other changes are required, since the core number of the rest of the nodes does not change.

A possible alternative is to maintain the previous core numbers and perform only incremental changes to the core numbers, based on the parts of the graph that have change. For example, the insertion of a new edge in the graph may impact the core numbers of specific nodes. The basic idea is to detect the set of affected nodes and recompute their core numbers without recomputing the core numbers for all the nodes. A similar approach can be followed when an edge is deleted from the graph. Performing incremental changes is expected to be much more efficient than re-executing the core decomposition algorithm from scratch. In this section, we discuss the methodology proposed in [127], which provides an incremental way to update the core numbers. The incremental algorithm is based on the following foundations:

1. If an edge is inserted to or removed from G , the core number of any node u can change by at most one.
2. If an edge (u, v) is inserted to or removed from G , where $c(u) < c(v)$, then $c(v)$ cannot change.
3. If an edge (u, v) is inserted into G , then all of the vertices whose core numbers have changed should

form a connected subgraph. Similarly, if an edge (u, v) is removed from G , then all the vertices whose core numbers have changed should form a connected subgraph.

4. If an edge (u, v) is inserted (removed) and $c(u) \leq c(v)$, then only the vertices w that have $c(w) = c(u)$ and are reachable from u via a path that consists of vertices with core numbers equal to $c(u)$, may have their core numbers incremented (decremented).

In [127], three different algorithms are studied: *i*) SUBCORE, which is based on the aforementioned foundations, *ii*) PURECORE, which applies some additional optimizations and *iii*) TRAVERSAL, which manages to reduce the number of examined nodes even further and shows the best overall performance. Based on the experimental results in [127], even the SUBCORE algorithm, which is the simplest among the three proposed, manages to update the core numbers up to 14,000 faster than the standard baseline algorithm which recomputes the core decomposition from scratch. Interestingly, a very similar set of foundations is also defined in [93] for the same topic. The two pieces of work are seemingly developed in parallel and there is no direct comparison of their proposed algorithms but the theorems appear to be equivalent.

The later work in [157] addresses [127] and [93] as being impractical for large graphs (in the aspect of the number of nodes). Dynamic graphs are referred to as temporal in [157] and the number of edges between two vertices may be more than one as a mean to describe an edge being in multiple instances in time (i.e., if an edge exists in t_1 and t_2 then it is counted as two edges). This transforms the problem of k -core decomposition to a (k, h) -core decomposition where the h is the number of edges a node shares with all of its neighbors and it is an additional degree/dimension in the core decomposition model. Effectively, by thresholding h (similar to k) one calculates cores as different temporal instances.

A proposal for improving the TRAVERSAL algorithm is presented in [167] where the authors focus on limiting the cardinality of the set of nodes need to be examined. While this work comes after [157], it does not compare against it but one may see it as a direct evolution of the technique proposed in [127]. The main point of [167] is that the TRAVERSAL algorithm may visit a very large number of edges to decide the set of affected nodes by the insertion of edges. The improvements of their solution rely on the following ideas:

- An ordering of the nodes by the sets of the k -core decomposition. (i.e., nodes belonging to core $k - 1$ are before nodes in k) and a further ordering in each core by the order the nodes were removed during the decomposition algorithm.
- Two degrees for each node u :
 - $deg(u)^+$: The number of nodes that are after u . This is essentially the number of nodes that can potentially support the increment of the current core number of u by being after u in the ordering (and can have their core number increased).
 - $deg(u)^*$: The number of nodes that are before u have the same core number as k and could potentially have their own core number increased. In order to determine if their core number may increase they use $deg(u)^* + deg(u)^+$.

The latter recursive check takes a crucial part in their algorithm as it creates a chain effect through the ordering in the visitation of the nodes and their exclusion from the set of nodes that will have their core number increased. The case of edge deletion follows the TRAVERSAL algorithm in logic.

In the scenario of low-memory environments, [153] -beyond the presentation of memory efficient decomposition algorithms- also tackles the topic of core maintenance. In fact, it is shown that a small

modification of the SEMICORE* is enough to handle the maintenance of the core number for the case of a single edge deletion. The basic intuition of this modification is that the main algorithm can be called on the current core values by only updating beforehand the count property of one of the nodes connected by the edge. The node with the lowest core number out of the two is selected or both if they have equal core numbers. The case of edge insertion is more complicated as the core number of a node v may increase and that can cascade to all nodes that are reachable from v on a path of nodes u where $core(u) = core(v)$. In a baseline attempt, the authors provide a modification of SEMICORE* where the neighbors of a node v are labeled for a possible update if they belong in the aforementioned path (i.e., if v is in the path and the core number has the expected value). This is further improved by modifying the original “count” logic (of SEMICORE*) to take into account the updated core value only. In this aspect, the path is restricted by the count of neighbors (of a candidate node v) which have a core value equal or greater than the updated one.

Moreover, in [154] (presented in the “Disk-Resident Graphs” section), we also see solutions regarding maintenance in the degeneracy ordering task. This approach relies on maintaining two additional properties for each node upon insertion and deletion of an edge:

- *Level-value*: a partitioning of the nodes where each partition contains nodes that their removal order would not affect the final ordering. The neighbors of the nodes of each partition have to be in the same or higher level partition.
- *Upper-degree*: the number of neighbors that are in the same or higher level-value partition.

In other publications, two distributed algorithms are proposed for the calculation of the temporal core decomposition. The first is based on the Pregel model [100] and is similar to the classic k -core algorithm on Pregel but it is considered inefficient as it has high memory, communication and computation costs. A more efficient distributed algorithm is then presented based on a block-centric model for graph computations [158] (Blogel) which is elaborated in the **Distributed Computation** section.

In [52] cores in a temporal graph are considered to exist in temporal intervals Δ and are named span-cores. The authors then make the note that a span-core at k, Δ is contained in a span core k, Δ' if $k' \leq k$ & $\Delta' \subseteq \Delta$. Based on this, they define a maximal span-core as one where it is not contained in another span core (the previous condition cannot be satisfied by another core).

In order to figure out whether a graph of $\Delta = [t_s, t_e]$ contains a maximal span-core, the authors prove that one only needs the core numbers from $\Delta' = [t_s - 1, t_e]$ and $\Delta'' = [t_s, t_e + 1]$. This eventually motivates the authors to start from larger temporal spans instead of calculating the core decomposition at each time-instance t of the graph. In this manner, the authors do not consider every core decomposition at each t as a maximal one.

3.4 Local Computation of Core Numbers

The main characteristic of the aforementioned techniques is that the core number of a node is determined by taking into account the whole graph. Moreover, the output in all cases is the core number of every node $u \in V$. An interesting alternative is to try to estimate the core number of a node u by considering only the neighborhood of u (e.g., the 1-hop, 2-hop 3-hop neighborhood, etc). The main advantage of such an approach is that we do not need to consider the whole graph, which potentially leads to more efficient computation. Also, in many cases we need to compute the core number of a small subset of nodes. However, we expect that the core number determined by such a *local computation* may not match the core number determined if the whole graph is taken into account.

The first work towards local core number estimation was reported in [114]. For a node u let $N_d(u)$ denote the set of nodes at a *distance* at most d from u . In its simplest form, the distance can be the shortest path distance, which, in the case of unweighted graphs, translates to the minimum number of hops between two nodes. Let G_d^u denote the graph induced by the d -neighborhood of u . One possible approach is to compute the core number of node u in the induced subgraph G_d^u . Let $c_d(u)$ denote the core number of node u computed in the induced subgraph G_d^u . By increasing d it is expected that a more accurate estimate of the core number may be achieved, since a larger induced subgraph is being used. Evidently, if d is large enough so that $G_d^u = G$, then $c_d(u) = c(u)$, meaning that the core number will be accurately computed. It is not hard to prove that always $c_d(u) \leq c(u)$, i.e., the core number computed in the induced subgraph G_d^u is a lower bound on the exact core number $c(u)$ computed in the whole graph.

In the sequel, we elaborate on a more sophisticated estimator which, in contrast to the previous one, uses some additional foundations to provide tighter bounds for the core number. Let u the node we are interested in. In case $d = 0$, the only available bound that we may use is the degree of u , since $c(u) \leq \deg(u)$. By setting $d = 1$, we may use additional information related to the 1-hop neighborhood of u . For example, if the core numbers of the direct neighbors of u are known, then the value of $c(u)$ can be computed accurately, by utilizing the following rationale:

- A node u belongs to the k -core of G , if and only if u has at least k direct neighbors in the k -core.
- Let u_1, u_2, \dots be the direct neighbors of u with known core numbers. Then, it holds that:

$$c(u) = \max_{1 \leq i \leq \deg(u)} (\min(c(u_i), \deg(u) - i + 1)) \quad (4)$$

Evidently, the assumption that the core numbers of all neighbors of u are known, is quite restrictive. However, the rationale of the previous idea is very useful in deriving an upper bound of the core number of a node, based on upper bounds of the core numbers of its neighbors. In fact, Equation 4 is valid for any $f(u) \geq c(u)$, $\forall u \in V$. Thus, we can derive the following recurrence, where u_1, u_2, \dots are the direct neighbors of u ordered in increasing order based on $\hat{c}_{d-1}()$, i.e., $\hat{c}_{d-1}(u_i) \leq \hat{c}_{d-1}(u_{i+1})$:

$$\hat{c}_d(u) = \begin{cases} \max_{1 \leq i \leq \deg(u)} (\min(\hat{c}_{d-1}(u_i), \deg(u) - i + 1)), & d > 0 \\ 0, & d = 0 \end{cases}$$

Evidently, $\hat{c}_d(u)$ is an upper bound for $c(u)$.

Based on the previous discussion, the exact core number $c(u)$ of a node u is bounded by the two values $c_d(u)$ and $\hat{c}_d(u)$:

$$c_d(u) \leq c(u) \leq \hat{c}_d(u), \forall d \geq 0$$

Based on the experimental results given in [114], both estimators provide satisfactory results for $d = 2$ for different input graphs. The accuracy of the estimators is quantified by considering the percentage of the total number of nodes for which the estimators give the exact core number. In most of the cases, the propagating estimator $\hat{c}_d(u)$ is more efficient than the induced estimator $c_d(u)$ in terms of accuracy as defined previously. In fact, the propagating estimator manages to achieve an accuracy between 80% and 90%, in the majority of the experiments performed, for $d = 2$. As a concluding remark, in cases where the core number of specific nodes needs to be computed, the methodology described is an effective alternative to the complete core decomposition process since it is more efficient with respect to runtime and provides satisfactory accuracy results.

Truss and nucleus decomposition. In a more recent work, Sariyüce et al. [129] capitalized on the relationship between core number and h -index [98], in order to propose efficient local algorithms for truss and nucleus decomposition with convergence guarantees. Lu et al. [98] have introduced an alternative formulation for the k -core decomposition that considers local information, utilizing the concept of h -index which is widely used in scientometrics. In particular, the degree of the nodes are used as the initial core number estimates; then, each node updates its estimates as the h -index of the neighbor’s core number estimates. At the end of the process, it was shown that the estimates converge to the core numbers. As we have presented in Section 2, the nucleous decomposition is a framework that generalizes the k -core and k -truss decompositions, using higher order structures to find dense subgraphs. The authors of [129] generalized the work by Lu et al. for any nucleus decomposition – providing also theoretical quarantess about the convergence properties. In addition, the proposed algorithms are highly parallel due to the fact that they operate based on local computations.

3.5 Parallel and Distributed Techniques

The main feature of the algorithmic techniques discussed so far is that they work in a centralized environment. However, mining massive graphs in a centralized manner does not provide scalable solutions. A very promising alternative is the exploitation of *multiple resources* to attack the problem. Towards this goal, there are two different research directions based on architectural assumptions: *i*) solving the problem in a shared-memory *multi-core machine* (the *parallel* case) and *ii*) solving the problem in a *cluster of machines* (the *distributed* case).

3.5.1 Parallel Computation

Parallel computation of core decomposition in multi-core processors was first investigated in [134], where the PARK algorithm was proposed. PARK was designed to work efficiently in multi-core processors where locality of reference is very important. In contrast to other techniques, PARK carefully reduces the number of random memory accesses performed. Note that random memory accesses may invalidate the caches in a multi-core system leading to significant performance degradation.

PARK uses three data structures: *Core*, *Curr* and *Next*. *Core* is an array of size n , initialized to the degrees of the nodes, i.e., initially $Core(u) = deg(u)$ for any node u . Recall that the degree of a node serves as an upper bound for its core number. *Core* is updated continuously during the course of the algorithm and its final values correspond to the core numbers of the nodes. *Curr* contains the set of nodes to be processed during the current iteration, whereas *Next* contains the nodes to be processed in the next iteration.

The processing of a node u involves accessing the direct neighbors of u and decreasing their degree if they have not already been processed. Computation is performed in different levels. While processing level l , all nodes that belong to the l -shell are processed. This is performed in two steps, SCAN and LOOP:

- During the SCAN phase, the array *Core* is scanned and all relevant nodes belonging to the l -shell are collected in the set *Curr*. Formally: $Curr = \{u : deg(u) = l\}$.
- The LOOP phase is executed in sublevels. In each sublevel, all nodes in the set *Curr* are processed. While processing node u , if a neighbor v is moved to *Curr*, which means that $deg(v) = 1$, then node v is added to the set *Next*. At the end of the sublevel, the contents of *Next* are transferred to the set *Curr* to be processed in the next sublevel.

In the previous discussion, the description of PARK is based on a centralized environment. In the sequel we provide the changes need to be applied in order to use PARK in a parallel (shared-memory) setting. First, we assume that processing will be performed by a set of threads T . Let t denote the total number of threads in T . Graph nodes are split among the t threads, meaning that each thread handles roughly n/t nodes. However, since the data structures used may need to be written by multiple threads concurrently, race conditions may appear and they must be handled carefully to avoid inconsistencies. To facilitate effective parallel processing, some modifications must be applied. Updates performed to the *Curr* set must be atomic, which means that write operations must be protected due to race conditions. For example, adding a new node id to *Curr* requires an `atomicIncrement()`. Similar atomic operations are required in several other cases, to protect the consistency of the data structures.

In addition, we need to invoke `fork()` to create multiple instances that will execute in parallel and also `join()` to wait for the threads to finish before we report the core numbers back to the caller of PARK. Moreover, there is a need to include synchronization calls. The invocation of `synchronize()` sets a barrier which must be reached by all running threads before code execution can continue.

To get an idea of the performance improvement that PARK achieves, for the Friendster graph [89] which contains 65 million nodes and 1.8 billion edges, PARK needs roughly 160 seconds to compute the core numbers whereas the centralized algorithm requires almost 1000 seconds on a machine with 8 physical cores. Similar behavior is observed in the majority of the datasets used for experimental evaluation.

Based on the main idea of PARK, a more scalable algorithm (PKC) has been reported in [74]. The main motivation for the design of PKC is that PARK does not scale well when the number of running threads is increased beyond 32. This behavior implies that PARK cannot provide the required efficiency when the number of physical cores is increased. On the other hand, PKC manages to scale well by increasing the level of parallelism. The first modification applied on PARK is based on the observation that arrays *Curr* and *Next* are not need to be maintained. On the other hand, PKC utilizes only the arrays *Core* and *Buff* and applies additional optimizations as follows:

- To process nodes at level l , the array *Core* is scanned and if the degree of a node is equal to l , it is appended to *Buff*. Then, nodes in *Buff* are processed until *Buff* becomes empty. A node u of *Buff* is processed by checking its neighbors. The current core value of u 's neighbors are examined and if the current core number is larger than l , the core number of a neighbor is decreased. During this process, new nodes may be added to *Buff*. The iteration is complete when all nodes contained in the l -core have been determined and the array *Buff* is empty. The algorithm proceeds by setting $l \leftarrow l + 1$ and it terminates when all the nodes are examined.
- Based on the power-law degree distribution in real-world networks, it is observed that the vast majority of nodes have a small core number. For example, in a typical case, more than 90% of the nodes have a relatively small core number. Thus, 10% of the nodes are classified as high-degree nodes. For these high-degree nodes, we need to check the core numbers of their neighbors, although we know that most of them have a smaller core number. PKC exploits this observation by creating a new *Core* array as soon as a significant percentage of the nodes has been processed.

Based on performance evaluation results reported in [74], PKC is significantly more efficient and scalable in comparison to PARK and thus, it is more suitable for large levels of parallelism.

A recent research work reported in [152] studies the parallel k -core computation in dynamic graphs. An important contribution of that work is the definition of the *superior edge set*, which is composed by edges that may change the core number of nodes by at most one. This observation enables the processing of these edges in parallel. In addition, the authors provide sufficient conditions to enable the identification

of nodes whose core numbers will change when a superior edge set is inserted or deleted. These techniques are able to increase the level of parallelism significantly, reducing the number of iterations required to determine the final values of the core numbers.

3.5.2 Distributed Computation

Although the computation of the core decomposition using parallelism is very attractive compared to the centralized approach, still there are significant limitations. The use of shared memory may become a bottleneck by increasing the number of parallel resources in the system. Anyway, the level of parallelism may increase up to a point using a shared-memory architecture. In addition, if the size of the graph grows significantly, we may face storage problems since the graph may not be accommodated in main memory. One may argue that in such a case, disk-based techniques, like the EMCORE algorithm which was covered in a previous section, could be applied. The problem with the algorithms that utilize secondary storage is that they perform a large number of passes over the data. As the graph grows larger the number of passes is expected to increase leading to inefficient computation.

During the last fifteen years we have witnessed a tremendous progress in data-driven distributed computing. In addition to the numerous ad-hoc solutions, several unified distributed platforms like MapReduce [39], Hadoop [155] and Spark [160] appeared and paved the way for nowadays cluster computing. These platforms are based on a cluster of shared-nothing machines (usually of commodity hardware) and they are able to execute complex data mining and machine learning algorithms over massive datasets efficiently. For the remaining of this section we will focus on a distributed core decomposition algorithm that was initially designed for a cluster of machines without any specific organization. The only requirement is that the processors may communicate by exchanging messages through the interconnect (usually a high-speed LAN). However, the algorithm was later adapted in order to be applicable in the Spark distributed engine.

The first distributed core decomposition algorithm was reported in [109, 110]. In the general case, each processing unit is responsible for multiple graph nodes. To simplify the presentation we will adopt the *one-to-one* scenario, meaning that we assume that each graph node corresponds to a single processing unit (a processor or a core). This model, resembles the Pregel’s [100] “*think as a vertex*” point of view. In other words, to design an algorithm we should take the point of view of a graph node u and try to provide the answer based on the information collected from the neighbors of u in an iterative manner. Each graph node u maintains the following information:

- $core(u)$: This is the currently most accurate estimate for the core number of u , which is initialized to the degree of u .
- $est[u_1, \dots, u_l]$: This is an array (or hashmap) storing the current estimates for the core numbers of u ’s neighbors. More specifically, $est[u_j]$ is the most up-to-date estimate of the core number of u_j known by u . Initially, $est[u_j] \leftarrow \infty$.
- $changed(u)$: This is a boolean flag that is set to **true** whenever the value of $core(u)$ changes. This attribute is initialized to false.

The computation for a node u begins by sending a message to all its neighbors. Since initially there is no accurate estimate for $c(u)$ (the core number of u), we may use the node degree as an upper bound. Node u prepares a message $msg[u, deg(u)]$ containing the node identifier of u and the current best estimate of the core number. This message is transmitted to all neighbors of u .

Assume now that u receives a message $msg[u_i, k]$ from node u_i (i.e., the i -th neighbor of u). From this message, we know that k is the current best estimate for $core(u_i)$. Therefore, if $k < est[u_i]$, this means that the previous estimate that node u knows about u_i is larger than the last received, and it should be updated by setting $est[u_i] \leftarrow k$. Once node u receives such a message, a new estimate for $core(u)$ may be produced. Let h denote the newly estimated value of the core number based on the information contained in $est[u_1, \dots, u_l]$. If $h < core(u)$, then $core(u) \leftarrow h$ and also $changed(u) \leftarrow true$, which means that a new estimate for $core(u)$ is available.

The main part of the algorithm is composed of three phases, depending on the event being handled. Initially, and before the start of execution an initialization step is applied. In this step, the only valid estimate for the core number of the degree of u . Therefore, the message send by node u to all its neighbors simply contains the degree of u . Every time a new message is received by u from one of its neighbors u_i , an update is performed in case the new value received has an impact on the estimation the $core(u)$. Node u sends periodically (every Δt time instances) the value of $core(u)$ to its neighbors, in case $core(u)$ has changed. The algorithm terminates when no change is performed on any core number. If this happens, then the core number estimate for each node equals its actual core number. To be able to achieve this convergence several techniques may be applied, e.g., centralized, decentralized, barrier synchronization.

Another more efficient technique is to execute the algorithm for a fixed number of rounds. The main motivation for this alternative is that after the first few rounds the core number estimates are quite close to the actual core numbers. Therefore, there is an efficiency vs. accuracy trade-off since on one hand the number of rounds is fixed but on the other hand the core numbers reported may not be 100% accurate. We have performed some experiments to test the accuracy of the algorithm by using a fixed number of iterations. Table 4 presents some representative results. More specifically, the table shows the percentage of nodes that have the correct core number after 20 rounds (iterations). We observe, that the accuracy is adequate for most realistic scenarios taking into account that an exact computation would require a significant number of rounds (shown in column **max num of iters**). All datasets are available at [89].

Table 4: Percentage of nodes with correct core numbers after 20 iterations (Accuracy@20).

Graph	Max Num of Iters	Accuracy@20
Orkut	191	88.76%
LiveJournal	99	98.49%
Web-Stanford	538	90.2%
Enron Email	28	99.7%

In [157], the Pregel model is deemed inefficient for large temporal graphs and an alternative is proposed based on the Bogel model [158]. This approach partitions the graph into blocks V_b that are accompanied with the information of which vertices are also connected to V_b (V_b^+). The core numbers in each V_b are computed and then the core numbers of the V_b^+ vertices are used to update their respective block (V_b). The degree of a node in V_b takes into account also V_b^+ . The main intuition of the update is that the vertices in V_b^+ with a core number lower than the minimum degree of V_b will not contribute to the core number of vertices in the future. Those vertices are removed from V_b^+ and core numbers are recomputed in V_b recursively.

An interesting technique reported in [12], involves the distributed core decomposition computation in dynamic graphs. In the previous section, we have discussed a similar topic which was related to

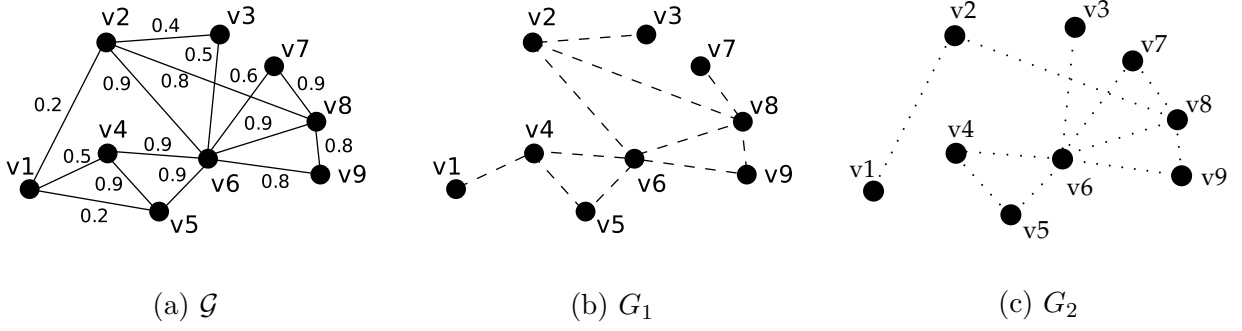


Figure 11: A probabilistic graph \mathcal{G} and two possible instances G_1 and G_2 . The numbers near the edges denote existential probabilities.

parallel processing of dynamic graphs. The proposed distributed algorithms exploit the graph topology to compute all k -cores and maintain them under frequent insertions and deletions of edges. The main idea is that when an edge is inserted or deleted, a limited number of core numbers is updated. The proposed technique is implemented by using the AKKA framework³ and the experimental evaluation has shown that the algorithm is efficient when applied in large dynamic graphs.

3.6 Probabilistic Core Decomposition

The aforementioned algorithmic techniques operate on *certain* graphs, meaning that graph nodes and edges are present with certainty (they always exist). However, as stated in Section 2, many applications require some kind of *uncertainty* associated with nodes or edges. In such a case, the graph becomes *uncertain* or *probabilistic*, meaning that the edge (or node) will be present in the network with some probability. For the following discussion, we will assume that graph nodes are certain (i.e., they exist at all times) whereas uncertainty is associated only with edges.

In the sequel we describe the algorithm reported in [24]. Let $\mathcal{G} = (V, E, p)$ be an uncertain (a.k.a probabilistic) graph, where $p : E \rightarrow (0, 1]$ is a function that assigns probabilities to the edges of the graph. A widely used approach to analyze uncertain graphs is the one of *possible worlds*, where each possible world constitutes a deterministic realization of \mathcal{G} . According to this model, an uncertain graph \mathcal{G} is interpreted as a set $\{G = (V, E_G)\}_{E_G \subseteq E}$ of $2^{|E|}$ possible deterministic graphs [123, 115, 116]. Let $G \sqsubseteq \mathcal{G}$ indicate that G is a possible world of \mathcal{G} . Then, the probability that $G = (V, E_G)$ is observed as a possible world of \mathcal{G} (assuming independence of edge existence) is given by the following formula:

$$\Pr[G|\mathcal{G}] = \prod_{e \in E_G} p(e) \prod_{e \in E \setminus E_G} (1 - p(e)) \quad (5)$$

As an example, consider the probabilistic graph \mathcal{G} shown in Figure 11(a). Two possible instances of \mathcal{G} are given in Figure 11(b) and 11(c). High-probability edges are expected to show up more frequently in instances of \mathcal{G} in comparison to low-probability edges. In this example, G_1 and G_2 are two of the *possible worlds* that can be produced by using \mathcal{G} as a template.

One of the novelties of the approach proposed in [24] is that the degree of a node is expressed by using probabilistic arguments. First, we need an expression of the probability that the degree of a node u is more than k . Note that this is a natural concept taking into account that the degree of a node is

³<https://akka.io>

in general different in different instances of the probabilistic graph \mathcal{G} . This probability is expressed as follows:

$$\Pr[\deg(u) \geq k] = \sum_{G \in \mathcal{G}_u^k} \Pr[G] \quad (6)$$

where \mathcal{G}_u^k is the set of instances of \mathcal{G} where u has a degree at least k . Next, we introduce the η -degree of a node u , denoted as $\eta\deg(u)$ and defined as follows:

$$\eta\deg(u) = \max\{k \leq |N(u)| \mid \Pr[\deg(u) \geq k] \geq \eta\} \quad (7)$$

More specifically, the $\eta\deg(u)$ is the maximum k for which the probability that the degree of u is more than k , is more than η . Recall that $|N(u)|$ is the number of direct neighbors of u in the probabilistic graph \mathcal{G} . Based on the variables k and η , the concept of (k, η) -core is defined: the (k, η) -core of a probabilistic graph \mathcal{G} is a maximal subgraph $\mathcal{H}(V_{\mathcal{H}}, E_{\mathcal{H}}, p)$ such that the probability that each vertex $u \in V_{\mathcal{H}}$ has degree no less than k in \mathcal{H} is greater than or equal to η , i.e., $\forall u \in V: \Pr[\deg_{\mathcal{H}}(u) \geq k] \geq \eta$.

The *probabilistic core decomposition* problem is defined as follows: Given an uncertain graph \mathcal{G} and a probability threshold $\eta \in [0, 1]$, find the (k, η) -core decomposition of \mathcal{G} , that is the set of all (k, η) -cores of \mathcal{G} .

In contrast to the standard core decomposition approach, degree computation is substituted by η -degree computation in the probabilistic case. More specifically, all η -degrees are computed at the beginning of the processing and η -degrees of neighboring nodes are updated accordingly. However, while the degree computation is straightforward in the deterministic case, computing η -degrees are far from trivial. In [24], interesting techniques are developed for computing η -degrees efficiently, based on *dynamic programming*.

Probabilistic core decomposition was also studied recently in [121], where (k, θ) -cores were proposed. In contrast to (k, η) -cores, (k, θ) -cores capture the likelihood of a node to be a k -core member in different instances of an uncertain graph \mathcal{G} . Given k and the probability threshold θ , the technique of [121] detects nodes such that with probability at least θ are included in the k -core in possible worlds.

4 Representative Application Domains

The core decomposition concept, despite its simplicity, has been applied successfully in many different scientific disciplines. It turns out that the core number of network nodes plays different roles, depending on the context being used and the type of the network applied. In this section, we present the most representative use cases and the associated results obtained.

4.1 Network Modeling and Analysis

The k -core decomposition and its extensions have been extensively used in several applications. Seidman [133] was the first that proposed to use the tool of k -core decomposition in social network analysis, as an easy to compute and effective way to extract dense subgraphs. Later, other studies in large scale real networks followed [69, 63], including the analysis of Microsoft Instant Messenger (MSN) [87] and Facebook [146] social graphs. In a similar way, the decomposition has been applied to study [8, 9] and model [29] the Internet graph.

Furthermore, several theoretical studies about the structure of real networks have been presented from the *statistical physics* community [41, 148]. In [70], it was shown that the k -core plays a central role for the modeling of real world graphs and their percolation properties. Based on that, a graph generation

model was introduced, and the properties of the generated graphs has been compared against a variety of real networks.

On a conceptually opposite side, some works have studied how the concept of core decomposition can be used to study the resilience or robustness of a network. Adiga and Vullikanti [1] examined the robustness of the top (i.e., maximal) cores under sampling and in pertrubed (i.e., noisy) networks. The authors of [161] have used the k -core as a heuristic tool in the process of graph decycling and dismantling. Their COREHD algorithm (which utilizes the core decomposition) provides an efficient manner for finding a minimal set of nodes which if removed will decycle (and then dismantle) a graph. Their work is further improved by the recent work in [132] which improves the heuristic by selecting nodes of high degree in the k -core with low degree neighbors. In a similar spirit, the authors of [86] proposed metrics for measuring core resilience in order to characterize the robustness of the core structure of a network when nodes and edges are removed.

4.2 Temporal Evolution

An interesting application of temporal cores is found in the aforementioned work of [52] (Dynamic Graphs) for the detection of anomalous contacts in social networks. Besides the extensions of the k -core structure into temporal graphs, the evolution of degeneracy has also been studied as a property through time (e.g. like density, diameter etc.). In [165] the authors use k -core to study the evolution of the Internet through time and discover that the majority of new nodes are added to the periphery of the graph while the size of the maximal k -core is quite stable through time.

4.3 Anomaly Detection

Shin et al. [136, 137] examined the properties of the k -core decomposition in a wide range of real-world networks. Their main observations include a set of empirical patterns that hold across several real-world graphs and can further be used to detect anomalies. The main observation, called the *mirror pattern*, indicates that the core number of the nodes of a graph has a strong positive correlation with the degree (which essentially represents an upper bound for the core number).

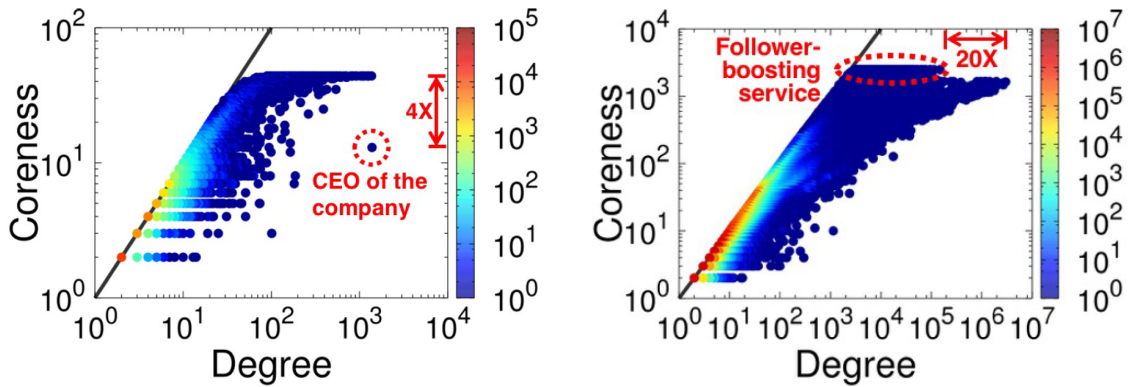


Figure 12: The *mirror pattern* of an Email communication graph (left) and a snapshot of the Twitter network (right). Each plot depicts the correlation between the degree and core number of the nodes of the graph. The figure is courtesy of Shin et al. [136]. ©2016 IEEE.

Figure 12 depicts the mirror pattern of two real-world graphs: an Email communication network and a snapshot of the Twitter graph. As one can observe, there is strong positive correlation between node

degree and core number. For example, in the case of the Email network, Spearman’s rank correlation coefficient has value $\rho = 0.99$, while in Twitter network the correlation coefficient is $\rho = 0.95$. Intuitively, the mirror pattern implies that nodes with high core number have also the tendency to have high degree and vice versa.

Nevertheless, one may observe that some nodes deviate from this “ideal” behavior; as the authors mention, these nodes correspond to two different types of anomalies: “lonerstars” (i.e., nodes mostly connected to “loners”) and “lockstep behavior” (i.e., a group of similarly behaving nodes). In the Email network of Figure 12 (left), the marked node has the highest degree but relatively low core number; this node corresponds to a secondary email account of the former CEO of the company, which was used only to receive emails.

In the case of the Twitter network, the nodes with the highest core number in Figure 12 (right) have been marked. Those nodes, have relatively low degrees and slightly deviate from the mirror pattern. Taking a closer look on the corresponding Twitter accounts, the authors noticed that at least 78% of those nodes were directly involved in a “Follower-Boosting” service – thus, can be annotated as anomalies.

In order to conduct this analysis, they devised a novel approach for estimating the core number in huge graphs (CORE-D, Algorithm). CORE-D is based on estimating the number of triangles by using the work of [145] for sampling and estimating the triangle count in large networks. The degeneracy of the graph can be estimated based on coefficients (w_0, w_1) which are computed from real data using linear regression.

4.4 Detection of Influential Spreaders

Detecting influential spreaders is an important topic for understanding how information diffuses in social networks. An intuitive notion in this domain is that individuals with high connectivity would contribute more in the diffusion process. This would naturally lead to metrics like betweenness centrality to be utilized for the identification of good spreaders (e.g., [117]).

The work of Kitsak et al. [79] finds a contradiction to that – as being highly connected in a network is not sufficient. It is pointed out that the quality of the connections are also important (i.e., the neighbors must be also well connected) and that the k -core is better at finding highly influential nodes. The observation that the core number of a node is a good predictor of its spreading capabilities, formed a new line of research in the area of influence spreading. A similar study that was conducted in [119], including PageRank as well as additional influence spreading metrics, lead to the same conclusion in favour of k -core.

Naturally, several improvements and extensions have been made to this approach (e.g., [98, 118, 97, 168]). One such improvement ranks nodes of a network by the sum of the core numbers of its neighbors [95]. In [28], the authors apply k -core (k -shell) on data from Twitter and notice that the skew in the degrees of the nodes creates a unnecessary number of cores (thousands) with most of the nodes existing in the lowered ranked ones (up to 4). To limit the number of cores, instead of mapping k connection to the k -th level, they assign $2^k - 1$ to k . In order to provide a more sophisticated decomposition, weights can be introduced in the graph. The authors of [5] provide a weighting scheme that represents the interaction among nodes and apply a weighted version of the k -core algorithm. This weight is specific to the nature of the graph (e.g. based on retweets in Twitter).

Other extensions of the k -core decomposition can be utilized as well. A prime example is the use of k -truss (triangle-based) [102]. This more restrictive version of graph degeneracy provides a smaller and more refined set of nodes in the maximal k -truss subgraph (which is a subset of the maximal k -core). The k -truss structure also captures the cohesiveness of the graph.

The work of [92] utilizes the k -core as a preprocessing step for an algorithm that extracts influential communities. A basic assumption in that work is that we can weight the graphs with some influential metric (e.g., Pagerank). Then the k -influential community is defined as an induced subgraph H^k of G where:

- H^k is connected.
- Each node in H^k has a degree of k .
- There is no other subgraph that satisfies the other two criteria, is not a subgraph of H^k and has a minimum weight (among its nodes) lower than H^k .

Algorithm 1: TOPCOM (G, W, r, k)

Input: $G(V, E)$, W , r , and k
Result: The top- r k -influential communities

```

1  $G_0 \leftarrow G, i \leftarrow 0;$ 
2 while  $G_i$  contains a  $k$ -core do
3   Compute the maximal  $k$ -core  $C^k(G_i)$ ;
4   Let  $H^k(i)$  be the maximal connected component
     of  $C^k(G_i)$  with the smallest influence value;
5   Let  $u$  be the smallest-weight node in  $H^k(i)$ ;
6   Delete  $u$ ;
7   Let  $G_{i+1}$  be a subgraph of  $C^k(G_i)$  after
     deleting  $u$ ;
8    $i \leftarrow i + 1;$ 
9 if  $i \geq r$  then
10  return  $H^k(i1), \dots, H^k(ir)$ 
11 else
12  return  $H^k(i1), \dots, H^k(0)$ 
```

While there is a much more efficient algorithm presented in that work, we present here in Algorithm 1 the naive basic version of computing the top- r k -connected communities as it is more intuitive. The authors note that this is better than k -truss as it includes the influence of each community (the minimum node-weight) but there is no direct comparison.

The work of [137] uses k -core decomposition for detecting influential spreaders as well (besides anomaly detection). In their approach (CORE-S Algorithm) for this application, they apply “vanilla” k -core decomposition and rank the potential spreaders with their eigen-centrality within the core (instead of the global graph). They compare their approach against k -core, k -truss, and eigen-centrality presenting better results than them both in efficiency and accuracy.

4.5 Network Visualization

The nested decomposition of k -core organizes vertices efficiently into groups for visual analysis as well. One of the earliest pieces of work for k -core based graph visualization focused on the presentation of the graph’s adjacency matrix [16]. The main idea is to reorder the vertices in rows and columns by their core number.

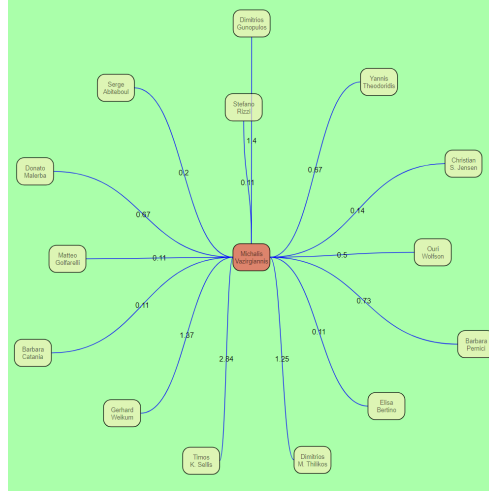


Figure 13: Example of k -core-based ego network.

In general, the k -core has been used to display examples of results from the analysis of real world graphs by focusing in the most dense cores. In [63] the authors display dense cliques of collaboration found in the academia with fractional (weighted) cores. Triangle cores are also used on publication data in [166] as well as the Wikipedia graph and protein to protein interaction networks among others to display examples of discovered cliques. A similar concept to the triangle cores, the m -core is defined based on the number of triangles an edge (instead of a vertex) belongs to [169]. Based on the m -coreness of an edge, a vertex will belong to an m -core if at least one of its endpoint vertices belongs to it. The m -core is utilized in [35] on an internet graph as well as on the E. Coli metabolic network to support that real world networks are organized with mechanisms that are based on local instead of global properties.

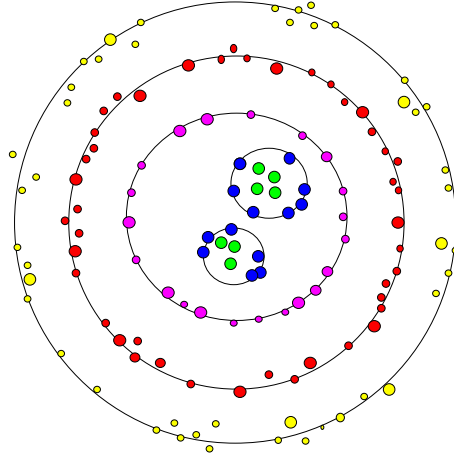


Figure 14: Example of the model for graph visualization with k -cores of [6, 7].

A variety of system and software applications have been developed to provide visualizations of graphs either at their entirety or at specific sub-graphs. Focusing on those that utilize graph degeneracy, Gephi [15] is a popular graph visualization tool that includes -among many others- k -core as a vertex positioning algorithm to organize vertices in concentric nested circles that are equal to the number of cores while the position of the a vertex in each circle is random. In a work with a specific focus of evaluation of individuals,

the authors of [59], present a tool for selecting and displaying the ego network of researchers in the graph of academic collaborations where only the colleagues of at least equal core number are included while the rest are hidden (example shown in Figure 13).

Finally, [6, 7] offers a well-developed tool for degeneracy based graph visualization. An mock-up of what it produces can be found in Fig 14. The nodes are organized in a nested manner that indicates their core number while the degree of the vertices and their proximity among them is also taken into account. Specifically the degree is represented by the size of the vertex while the proximity is displayed by positioning them in relevant proximity at the nested circles.

4.6 Communities and Dense Subgraphs

In a recent work on community detection [61], the authors built upon the properties of the decomposition to speed-up the execution time of computationally intensive graph clustering algorithms, such as *spectral clustering*. In particular, they have proposed CORECLUSTER, an efficient clustering framework that can be used along with any known graph clustering algorithm. The approach capitalizes on processing the graph in a hierarchical way guided by its k -core decomposition. Nodes are clustered in an incremental manner that preserve the clustering structure of the graph, while making the execution of the chosen clustering algorithm much faster due to the smaller size of the graph’s partitions onto which the algorithm operates.

In addition to the community detection problem, another interesting problem, *community search* has been addressed recently. In community search, given a node v the aim is to detect the most appropriate community for v . This problem was introduced in [138] and it attracted significant interest. In most of the cases, graph nodes have specific characteristics (i.e., attributes) that must be taken into account during community search. This concept was introduced in [47], where node attributes play a key role towards supporting a new query termed *attributed community query* (ACQ). In fact, one of the properties that a subset S of nodes must satisfy to qualify as a potential answer is *structure cohesiveness*, which essentially is the k -core property (i.e., for each node $u \in S$, $\deg_S(u) \geq k$). The structure cohesiveness property is used also in [46], where the focus is to support *spatial-aware communities* (SAC).

In [91] the concept of k -core is combined with the concept of *skyline*, in order to spot communities that are not dominated. This technique is applied to graphs with d numeric attributes per node. Therefore, each node may be seen as a point in the d -dimensional space. The skyline concept was introduced in the data management community by [26]. In fact, the skyline concept is based on the concept of *maximal vectors* which was studied before in Computational Geometry [83]. The dominance property states that a point p_1 dominates another point p_2 , if p_1 is as good as p_2 in every dimension, and p_1 is strictly better than p_2 in at least one of them. The points that are not dominated, form the *skyline set*. In the case of communities, a community c_1 dominates another community c_2 , if c_1 is at least as good as c_2 in every dimension and strictly better than c_2 in at least one dimension. Comparisons are performed based on an aggregation function $f_i()$ that takes into account the i -th dimension of all nodes of the community.

An interesting application of the k -core concept is the detection of *persistent communities* in networks that change over time, reported in [90]. The main idea behind that work is that a persistent community should exist during a specified period of time. One of the contributions of [90] is the definition of the (θ, τ) -persistent k -core. The challenging part in this concept is that the problem of finding the largest (θ, τ) -persistent k -core is NP-hard. The authors propose algorithmic techniques to solve the problem efficiently.

Another “famous” application of the k -core decomposition is the identification of dense subgraphs; Andersen and Chellapilla [10] were based on this to propose solutions with approximation guarantees for variants of the densest subgraph problem. In a similar spirit, variants of the community detection problem

has been addressed utilizing the properties of k -core decomposition, including local community detection techniques [37] and the influential community search problem [92, 4] where the notion of influence is defined as the minimum weight of the nodes in that community.

4.7 Text Analytics

Recently, the concept of k -core decomposition has been also applied in information networks used to represent textual information. Models for graph construction from textual data can be found at [108]. In short, one may consider elements of text (n-grams, single terms etc.) as vertices and as edges the co-occurrence of those elements. The edges can be enriched with additional properties depending on the application. One common approach assigns weights to the edges based on the frequency of term co-occurrence withing a fixed window - where a sliding window is assumed over the text.

A variety of graph based techniques have been utilized on those graphs. For example, [108] utilizes Pagerank and [96] uses HITS for ranking words in text based on their ranking in the corresponding graph. This, in both cases, aims at of the task of keyword extraction. Naturally, the ranking provided from degeneracy has been utilized as well for keyword extraction [124]. As the authors note in their work, beyond better results than the he other aforementioned approaches, the maximal k -core automatically decides the number of keywords in contrast with other methodologies where a fixed number is selected. Later on, Tixier et al. [143] further refined the concept of core (and truss) decomposition for the task of keyword extraction. Overall, the additional benefit of flexibility on the properties of the graph motivated the utilization of advanced approaches and applications of the k -core in text mining.

One such example is found in the later work of detecting events in Twitter streams [106]. The graph of words is build there from tweets in time-windows of fixed size and the detection of the event is based on thresholding the vertex/term weighted degree in the maximal core (i.e. the degrees define the event appearance). Another example [144], utilizes the k -core and k -truss structure for text visualization and summarization as an online real-time application. Moreover, the degeneracy approach has found application on keyword extraction from multiparty conversations [107]. In all cases there is a benefit from the efficient calculation of the k -core which allows almost instant results.

4.8 The Anchored k -Core Problem and Engagement Dynamics in Social Graphs

A common behavior of users in social networks is that their decisions are influenced by that of their neighbors, exhibiting the so-called *positive network effect*: assuming some notion of utility or gain per individual, this is increasing with the number of friends that behave in a certain way. For example, it has been empirically observed that users are more likely to engage to the activities of a social network if their friends do so. Based on that, how can we design or modify a social network in order to maximize the engagement of its users?

Assume that all users in a community are initially engaged and each individual has two strategies: to remain engaged in the activities of the community or to drop out [103]. An individual will remain engaged if at least k of his/her friends are engaged (i.e., degree constraint). A user with less than k engaged friends will decide to drop out, and his/her decision might spread over the network forming a cascade of departures (i.e., other individuals might drop out too). When the collapse stops, the remaining engaged users correspond to the k -core subgraph. That way, the size of k -core can be used to measure the overall engagement of the social network. In the related literature, many empirical studies have used the core number of nodes or the size of the maximal k -core subgraph to characterize the engagement properties of individual nodes or even the engagement characteristics of the whole graph [103, 55, 104].

Based on the previously described model of user engagement in social networks, Bhawalkar et al. [20] introduced the *anchored k -core* problem, which examines how to prevent unravelling on the network: we aim at retaining (anchoring) some individuals, so as to maximize the number of users that will remain engaged when the unraveling stops (i.e., the size of the maximal k -core subgraph). Once a node v in G is *anchored*, it is always retained by the k -core decomposition regardless of its degree (i.e., it is never removed by the decomposition) [162, 163].

Definition 22 (Anchored k -core subgraph) *Given an undirected graph G and a vertex set $A \subset G$, the anchored k -core subgraph, denoted by $C_k(G_A)$, is the corresponding k -core of G with vertices in A anchored.*

Definition 23 (Anchored k -core problem) *Given an undirected graph G , a degree constraint k and a budget b , the anchored k -core problem aims at finding a set A of b nodes, such that the size of the resulting anchored k -core, $C_k(G_A)$, is maximized.*

If we have a set A of anchor nodes, then we can directly use the linear time algorithm presented in Section 3 to compute $C_k(G_A)$. However, finding the optimal A is a computational difficult problem; it has been shown that, when $k \geq 3$, the anchored k -core problem is NP-hard [20]. Zhang et al. [162] have proposed an efficient heuristic algorithm, called OLAK, to deal with the complexity constraints of the anchored k -core problem.

4.9 Graph Similarity

In a very recent work [113], the hierarchy of the core decomposition is utilized to provide a general framework for computing similarity metrics among graphs.

Graph similarity is an upcoming topic in the domains of computational biology, chemistry and natural language processing. Simply put, when computing similarities among graphs basic structures (e.g. trees, cycles) are compared between graphs in a local or global level with graph kernels. The aforementioned work [113], contributes by utilizing existing kernels at equivalent core levels between graphs in order to compare the structures at similar levels of connectivity.

Algorithm 2: GRAPHSIMKERNEL(G, G')

Input: A pair of graphs G and G'

Result: Result of the kernel function, val

```

1  $val = 0$ 
2  $\delta_{min} = \min(\delta(G), \delta(G'))$ 
3 Let  $C_i, C'_i$  be the  $i$ -core of  $G, G'$  for  $i = 0, \dots, \delta_{min}$ 
4 for  $i = \delta_{min}$  to 0 do
5    $val = val + \text{kernel}(C_i, C'_i)$ 
```

Algorithm 2 is straightforward as it accumulates the similarities of a base-kernel along with the corresponding cores between two graphs. Despite its simplicity in implementation, it outperforms the utilized baselines in comparison with the result if they (the baselines) were used without the framework. The evaluation is performed on classification tasks of real world graphs from a variety of domains.

4.10 Physics, Biology and Ecology

Many interactions in real organisms are modeled as networks. Biological networks have been studied significantly by many different perspectives. One of these, is related to the discovery of core/periphery structure of biological networks. In [99] this idea is studied for protein-protein interactions networks (PPI). It turns out that in addition to the discovery of interesting correlations between connectivity and biological properties, the core/periphery structures help to reveal the existence of multiple levels of protein expression dynamics. Moreover, as reported in [73], residues belonging to inner cores are more conserved than those at the periphery of the network and also it seems that these groups are functionally and structurally critical. Another important result was reported in [44] which studied the relation between the core numbers of proteins and mutation rates. It turns out that the mutation rates for the interior cores is lower.

In addition to homogeneous networks that appear in biology, bipartite networks are also quite frequent: gene-protein, host-pathogen, predator-prey. In [56], two new visualization types are proposed that exploit the structural properties of these networks to improve readability. The basis of these methods is the core decomposition of the bipartite graph.

Other Biology-related works that use the concept of core decomposition include: [42], which applies the concept in weighted biological networks, [120] that performs protein complex prediction and [49] which uses core decomposition in plant metabolic networks.

The core decomposition has been effectively applied to other branches of Biology, such as Ecology. For example, in [57] an application is presented to plot bipartite ecological networks. Also, in [58], the authors study different techniques to identifying the species for which the networks are most vulnerable to cascade extinctions. It turns out that the core decomposition concept sheds light on the understanding of the robustness properties in mutualistic networks. Also, similar ideas can be found in [112], which studies the concept of structural collapse in mutualistic ecosystems. More specifically, based on the authors, it was shown that “when species located at the maximum k -core of the network go extinct as a consequence of sufficiently weak interaction strengths, the system will reach the *tipping point* of its collapse”.

In a surprising application of the classic k -core decomposition, [111] finds that when that modeling viscosity in materials as a network of contacts leads to the k -core being the tool for explaining a phenomenon referred as jamming (increment of viscosity with increment of density). In fact the network of contacts show a behavior similar to an Erdos-Renyi graph model in the aspect of k -core percolation.

4.11 Neuroscience

The study and the understanding of the brain is an ongoing adventure. It turns out the core decomposition concept plays an important role in this study. One of the first works that applied the concept was published in [67]. In that work, the authors performed an in-depth analysis of the brain functional network which is composed of parts of the brain that are functionally interconnected in a dense manner. The main result of this study is that regions of the brain that belong to the structural core, share high degree, strength, and betweenness centrality, and they operate as hubs linking other major structural modules.

In a similar line, the work in [147] demonstrates that parts of the brain with high connectivity (i.e., brain hubs) form a so-called “rich club”, which means that there is a tendency for high-degree nodes to be more densely connected among themselves than nodes of a lower degree. The result is that this “rich club” provides important information on the higher-level topology of the brain functional network.

Later, this idea was developed further in [135] in order to compare the brain organization of pigeons

and mammals. The main result of this study was that the pigeon *telencephalon*⁴ is organized in a similar manner to that of a mammal.

Additional research efforts in the area that use the core decomposition concept as a first class citizen include: [125] which studies the influence of wiring cost on the large-scale architecture of human cortical connectivity, [21] which examines the way the brain functional network reorganizes during cognition, [156] which links the concept of *cell assemblies* to that of *k*-core and studies a specific type of cell assembly called *k*-assembly and [85] which studies the hierarchical cortical organization of the human brain.

5 Conclusions and Further Research

The core decomposition of a graph is a concept that has been studied for many years and is applied in many different problems in diverse scientific areas – also displaying its value in real world applications. The main reason for its ubiquity lies in the fact that it provides an efficient manner for organizing the graph into hierarchical structures of increasing cohesiveness.

In this survey article, we have provided a thorough review on existing approaches for applications and algorithmic techniques concerning degeneracy-based graph decompositions. In the application domain, most of the works have focused primarily on real-world scale-free networks. A possible explanation is that, random graphs do not exhibit interesting degeneracy-based properties. In other words, there is no actual discrepancy of the core numbers of nodes in random graphs. Concerning algorithmic techniques, we have examined *k*-core (and its variants) being applied to a multitude of computational models covering many scenarios. One perhaps issue, we notice that in most cases the approaches are tailored by utilizing specialized structures which makes difficult to utilize the same model in a more generic graph analysis scenario (we elaborate on this below).

We have organized this review by three major axes *i*) graph types, *ii*) algorithmic techniques, and *iii*) applications. We note that many publications could be mentioned in all of these axes as the graph type is usually defined in a specific scenario for “not simple” graphs which in turn requires re-definition (or extension of the definition) of degeneracy and consecutively the re-design of the algorithm. For this reason we place each piece work into one of these axes based on its major focus.

To the extend of our knowledge, this survey is the first in attempting to cover degeneracy to this extend. Two notable attempts in similar reviews exist in the general literature as book chapters found in “Encyclopedia of Social Network Analysis and Mining” [23] and in “Cohesive Subgraph Computation over Large Sparse Graphs” [31] both covering to a lesser extend a subset of the main topics presented in this manuscript.

Although core decomposition has been covered by many different perspectives, still there is room for more work in the area. Next, we discuss briefly possible topics of interest for future research.

Algorithmic Aspects. A variety of techniques have been studied for *k*-core decomposition (in memory, streaming, in parallel), but most of them focus on a single dimension of decomposition based on a function upon properties of the node i.e., the degree, the number of triangles etc. As graph models become more elaborate (direction, label, timestamp), the algorithms are redesigned for a specific computational model. While this is expected (one designs an algorithm specific to the needs of application and core extension), there is a lack for a uniform model regardless of the dimensions across which the core is computed. One of the original works for Generalized Cores [17] attempted at providing such a general model but this

⁴The telencephalon is the most highly developed and anterior part of the forebrain, composed mainly of the cerebral hemispheres (<https://en.wikipedia.org/wiki/Cerebrum>).

would not be applicable in e.g., temporal cores as the additional dimension can be used in a different manner than the degree (which lead to a variety of works specifically for temporal cores).

Through out this review, the reader may notice that different models for graph decomposition try to solve the same problem in principle for computing the decomposition -while they study different graph properties. Nevertheless, a variety of algorithms has been introduced that solve the same problem (with similar approaches). As such, it is expected in future research for "generalized cores" to re-appear for a greater extension of dimensions and computational models.

Furthermore, there is a lot of potential for graph computational models in general in distributed environments due to the advances in this domain both in academia but in the industry as well combined with need for scalable algorithms on fast evolving data. While we cover here cases for distributed, parallel and streaming computation for large graphs, these approaches are yet to be adopted in general in real world applications and for now are specific implementations. One major factor for this is due to the fact that they require specific data structures and designs that do not generalize in usability by other algorithms e.g., the model used in incremental k -core decomposition [?] is very efficient but it might not be very efficient in other graph analysis tasks like eigen-centrality computation or triangle counting. As such, it is hard to adopt a solution if it requires a great cost for transformation in order to use other graph mining algorithms.

Core Decomposition in Machine Learning. While we saw here the potential of k -core decomposition even in graphs that are extracted from non-graph data (e.g., in graphs of words), it seems that it has not been explored in too many cases in the domain of Machine Learning or Data Science in general. While the topic is beyond the scope of this review, there are many cases where a graph structure is assumed as representation of the data in some form (e.g., manifold based decomposition [142], spectral clustering on none graph data through affinity kernels [149]).

Despite that graph structures are used quite frequently through data transformations and affinity kernels, the k -core structure of those graphs has not even been examined for its properties and potentials (e.g., feature extraction). This is somewhat, surprising as the k -core algorithm is quite efficient and -as it has been seen in this review- it offers the potential of finding out-liers or anomalies in real-world graphs. Of course, other types of data (e.g., points in the Euclidean space) will not have the same properties as the graphs that are usually studied with degeneracy. Nevertheless, at the minimum a study is lacking on the properties of "other" graphs.

Representation Learning on Graphs. A recent work (core2vec [131]) showed the potential of using k -core decomposition in feature engineering with deep learning techniques. In general this work lies in the domain of representing nodes with latent features – as an embedding in a multidimensional space – that are learned automatically. Usually, this embedding represents similar vertices as vectors with a high similarity in the vector space model. The aforementioned work describes an alternative to using random walk during one of the phases of an existing technique, by biasing the walk with core-based information (focusing the walk in the same core). Although this can be considered a marginal modification, it offers motivation on working in this sub-domain of machine learning. We have seen throughout this review a variety of core models with different interpretations – capturing different (structural) properties of the vertices and subgraphs in general. While these properties can be considered as features, one could utilize them in a similar way to the core2vec model (e.g., as preprocessing techniques), in order to optimize graph embedding techniques.

Influence Maximization. The topic of influence maximization and the identification of influential spreaders has received a lot of attention in the domain of online social networks and recommender systems due to its economical applications but also due to the general emergence of interest in the study of information diffusion.

As we have seen in Section 4, the k -core has played a major role in this domain and further improvements did not involve only the extension of the k -core but also its combination with other methodologies. As only a few basic models have been studied, the exploration of the more elaborate ones remains still an open subject (e.g., in directed graphs). Moreover, any lateral property defined on this topic will always have the potential of being combined with the k -core structure in similar manners to the ones reviewed. Lastly, the k -core has mainly been used so far as a heuristic algorithm to identify nodes with good spreading properties. It remains open problem how to combine the properties of the core decomposition with the greedy algorithm by Kempe et al. [76], towards a scalable core-based influence maximization algorithm with theoretical guarantees.

iyuce2016iyuce2016

References

- [1] A. Adiga and A. K. S. Vullikanti. How robust is the core of a network? In H. Blockeel, K. Kersting, S. Nijssen, and F. Železný, editors, *Machine Learning and Knowledge Discovery in Databases*, pages 541–556, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg. (Cited on pages 6 and 34.)
- [2] C. C. Aggarwal, editor. *Social Network Data Analytics*. Springer, 2011. (Cited on page 4.)
- [3] C. C. Aggarwal and H. Wang. *Managing and Mining Graph Data*. Springer, 2010. (Cited on page 4.)
- [4] E. Akbas and P. Zhao. Truss-based community search: A truss-equivalence based indexing approach. *Proceedings of the VLDB Endowment*, 10(11):1298–1309, Aug. 2017. (Cited on pages 6 and 39.)
- [5] M. A. Al-garadi, K. D. Varathan, and S. D. Ravana. Identification of influential spreaders in online social networks using interaction weighted k -core decomposition method. *Physica A: Statistical Mechanics and its Applications*, 468:278–288, 2017. (Cited on pages 6 and 35.)
- [6] J. Alvarez-Hamelin, L. Dall’Asta, A. Barrat, and A. Vespignani. K -core decomposition: A tool for the visualization of large scale networks. *Adv. Neural Inf. Process. Syst.*, 18, 04 2005. (Cited on pages 6, 37, and 38.)
- [7] J. I. Alvarez-hamelin, A. Barrat, and A. Vespignani. Large scale networks fingerprinting and visualization using the k -core decomposition. In *NIPS ’06: Advances in Neural Information Processing Systems*, pages 41–50, 2006. (Cited on pages 6, 37, and 38.)
- [8] J. I. Alvarez-Hamelin, L. Dall’Asta, A. Barrat, and A. Vespignani. k -core decomposition: a tool for the analysis of large scale internet graphs, 2005. (Cited on pages 6 and 33.)
- [9] J. I. Alvarez-Hamelin, L. Dall’Asta, A. Barrat, and A. Vespignani. k -core decomposition of internet graphs: Hierarchies, self-similarity and measurement biases. *NHM*, 3(2):371, 2008. (Cited on pages 6 and 33.)
- [10] R. Andersen and K. Chellapilla. Finding dense subgraphs with size bounds. In *WAW*, pages 25–37, 2009. (Cited on page 38.)

- [11] D. Angluin and J. Chen. Learning a hidden graph using $o(\log n)$ queries per edge. *J. Comput. Syst. Sci.*, 74(4):546–556, June 2008. (Cited on page 13.)
- [12] S. Aridhi, M. Brugnara, A. Montresor, and Y. Velegrakis. Distributed k -core decomposition and maintenance in large dynamic graphs. In *Proceedings of the 10th ACM International Conference on Distributed and Event-based Systems*, DEBS '16, pages 161–168, New York, NY, USA, 2016. ACM. (Cited on pages 6 and 31.)
- [13] J. Bang-Jensen and G. Z. Gutin. *Digraphs: Theory, Algorithms and Applications*. Springer Publishing Company, Incorporated, 2nd edition, 2008. (Cited on page 10.)
- [14] J. Bao, Y. Zheng, D. Wilkie, and M. Mokbel. Recommendations in location-based social networks: a survey. *GeoInformatica*, 19(3):525–565, Jul 2015. (Cited on page 19.)
- [15] M. Bastian, S. Heymann, M. Jacomy, et al. Gephi: an open source software for exploring and manipulating networks. *Icwsn*, 8(2009):361–362, 2009. (Cited on page 37.)
- [16] V. Batagelj, A. Mrvar, and M. Zaversnik. Partitioning approach to visualization of large graphs. In *International Symposium on Graph Drawing*, pages 90–97. Springer, 1999. (Cited on pages 6 and 36.)
- [17] V. Batagelj and M. Zaversnik. Generalized cores. *CoRR*, cs.DS/0202039, 2002. (Cited on pages 6, 13, 14, and 42.)
- [18] V. Batagelj and M. Zaversnik. An $o(m)$ algorithm for cores decomposition of networks, 2003. cite arxiv:cs/0310049. (Cited on pages 6, 21, and 23.)
- [19] A. R. Benson, D. F. Gleich, and J. Leskovec. Higher-order organization of complex networks. *Science*, 353(6295):163–166, 2016. (Cited on page 18.)
- [20] K. Bhawalkar, J. Kleinberg, K. Lewi, T. Roughgarden, and A. Sharma. Preventing unraveling in social networks: the anchored k -core problem. In *ICALP '11: Proceedings of the 39th International Colloquium Conference on Automata, Languages, and Programming*, pages 440–451, 2011. (Cited on pages 6 and 40.)
- [21] M. Bola and B. Sabel. Dynamic reorganization of brain functional networks during cognition. *NeuroImage*, 114, 03 2015. (Cited on pages 6 and 42.)
- [22] P. Boldi and S. Vigna. The webgraph framework i: Compression techniques. In *Proceedings of the 13th International Conference on World Wide Web*, WWW '04, pages 595–602, New York, NY, USA, 2004. ACM. (Cited on page 23.)
- [23] F. Bonchi, F. Gullo, and A. Kaltenbrunner. *Core Decomposition of Massive, Information-Rich Graphs*, pages 1–11. Springer New York, New York, NY, 2017. (Cited on pages 5 and 42.)
- [24] F. Bonchi, F. Gullo, A. Kaltenbrunner, and Y. Volkovich. Core decomposition of uncertain graphs. In *KDD*, pages 1316–1325, 2014. (Cited on pages 6, 12, 32, and 33.)
- [25] F. Bonchi, A. Khan, and L. Severini. Distance-generalized core decomposition. In *Proceedings of the 2019 ACM SIGMOD International Conference on Management of Data*, 2019. (Cited on pages 6, 16, and 18.)

- [26] S. Borzsony, D. Kossmann, and K. Stocker. The skyline operator. In *Proceedings 17th International Conference on Data Engineering*, pages 421–430, April 2001. (Cited on page 38.)
- [27] S. Brin and L. Page. The anatomy of a large-scale hypertextual web search engine. In *Proceedings of the Seventh International Conference on World Wide Web 7, WWW7*, pages 107–117, Amsterdam, The Netherlands, The Netherlands, 1998. Elsevier Science Publishers B. V. (Cited on page 5.)
- [28] P. Brown and J. Feng. Measuring user influence on twitter using modified k-shell decomposition. In *The Social Mobile Web*, volume WS-11-02 of *AAAI Workshops*. AAAI, 2011. (Cited on pages 6 and 35.)
- [29] 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. (Cited on pages 6 and 33.)
- [30] L. Chang and L. Qin. *Cohesive Subgraph Computation over Large Sparse Graphs*. Springer, 2018. (Cited on page 5.)
- [31] Q.-L. Chang, Lijun. *Minimum Degree-Based Core Decomposition*, pages 21–39. Springer Series in the Data Sciences, 2018. (Cited on page 42.)
- [32] J. Cheng, Y. Ke, S. Chu, and M. T. Oszu. Efficient core decomposition in massive networks. In *ICDE*, pages 51–62, 2011. (Cited on pages 6 and 21.)
- [33] S.-T. Cheng, Y.-C. Chen, and M.-S. Tsai. Using k -core decomposition to find cluster centers for k -means algorithm in graphx on spark. In *Proceedings of the 8-th International Conference on Cloud Computing, GRIDs, and Virtualization*, pages 93–98, 2017. (Cited on page 6.)
- [34] J. Cohen. Trusses: Cohesive subgraphs for social network analysis. *National Security Agency Technical Report*, 2008. (Cited on pages 6 and 14.)
- [35] P. Colomer-de Simón, M. A. Serrano, M. G. Beiró, J. I. Alvarez-Hamelin, and M. Boguná. Deciphering the global organization of clustering in real complex networks. *Scientific reports*, 3:2517, 2013. (Cited on page 37.)
- [36] D. J. Cook and L. B. Holder. *Mining Graph Data*. John Wiley & Sons, 2006. (Cited on page 4.)
- [37] W. Cui, Y. Xiao, H. Wang, and W. Wang. Local search of communities in large graphs. In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data*, pages 991–1002, 2014. (Cited on page 39.)
- [38] M. Danisch, T.-H. H. Chan, and M. Sozio. Large scale density-friendly graph decomposition via convex programming. In *Proceedings of the 26th International Conference on World Wide Web, WWW '17*, pages 233–242, 2017. (Cited on pages 6 and 15.)
- [39] J. Dean and S. Ghemawat. Mapreduce: Simplified data processing on large clusters. In *Proceedings of the 6th Conference on Symposium on Operating Systems Design & Implementation - Volume 6, OSDI'04*, pages 10–10, Berkeley, CA, USA, 2004. USENIX Association. (Cited on page 30.)
- [40] D. Ding, H. Li, Z. Huang, and N. Mamoulis. Efficient fault-tolerant group recommendation using alpha-beta-core. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, CIKM '17*, pages 2047–2050, New York, NY, USA, 2017. ACM. (Cited on pages 6 and 11.)

- [41] S. N. Dorogovtsev, A. V. Goltsev, and J. F. F. Mendes. k -core organization of complex networks. *Physical Review Letters*, 96:040601, 2006. (Cited on pages 6 and 33.)
- [42] M. Eidsaa. *Core Decomposition Analysis of Weighted Biological Networks*. PhD thesis, NTNU, 2016. (Cited on pages 6 and 41.)
- [43] M. Eidsaa and E. Almaas. s -core network decomposition: A generalization of k -core analysis to weighted networks. *Phys. Rev. E*, 88:062819, 2013. (Cited on pages 6 and 11.)
- [44] A. I. Emerson, S. Andrews, I. Ahmed, T. K. Azis, and J. A. Malek. K -core decomposition of a protein domain co-occurrence network reveals lower cancer mutation rates for interior cores. *Journal of Clinical Bioinformatics*, 5(1):1, Mar 2015. (Cited on pages 6 and 41.)
- [45] P. Erdős and A. Hajnal. On chromatic number of graphs and set-systems. *Acta Mathematica Academiae Scientiarum Hungarica*, 17(1-2):61–99, 1966. (Cited on pages 6 and 9.)
- [46] Y. Fang, R. Cheng, X. Li, S. Luo, and J. Hu. Effective community search over large spatial graphs. *Proc. VLDB Endow.*, 10(6):709–720, Feb. 2017. (Cited on pages 6 and 38.)
- [47] Y. Fang, R. Cheng, S. Luo, and J. Hu. Effective community search for large attributed graphs. *Proc. VLDB Endow.*, 9(12):1233–1244, Aug. 2016. (Cited on page 38.)
- [48] M. Farach-Colton and M.-T. Tsai. Computing the degeneracy of large graphs. In *Latin American Symposium on Theoretical Informatics*, pages 250–260. Springer, 2014. (Cited on page 23.)
- [49] H. A. Filho, J. Machicao, and O. M. Bruno. A hierarchical model of metabolic machinery based on the k -core decomposition of plant metabolic networks. *PLOS ONE*, 13(5):1–15, 05 2018. (Cited on pages 6 and 41.)
- [50] L. C. Freeman. A set of measures of centrality based on betweenness. *Sociometry*, 40(1):35–41, 1977. (Cited on page 5.)
- [51] E. C. Freuder. A sufficient condition for backtrack-free search. *J. ACM*, 29(1):24–32, 1982. (Cited on pages 6 and 9.)
- [52] E. Galimberti, A. Barrat, F. Bonchi, C. Cattuto, and F. Gullo. Mining (maximal) span-cores from temporal networks. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pages 107–116. ACM, 2018. (Cited on pages 6, 26, and 34.)
- [53] E. Galimberti, F. Bonchi, and F. Gullo. Core decomposition and densest subgraph in multilayer networks. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, CIKM '17*, pages 1807–1816, New York, NY, USA, 2017. ACM. (Cited on pages 6 and 13.)
- [54] A. Garas, F. Schweitzer, and S. Havlin. A k -shell decomposition method for weighted networks. *New Journal of Physics*, 14(8), 2012. (Cited on pages 6 and 11.)
- [55] D. Garcia, P. Mavrodiev, and F. Schweitzer. Social resilience in online communities: The autopsy of friendster. In *COSN '13: Proceedings of the First ACM Conference on Online Social Networks*, pages 39–50, 2013. (Cited on pages 6 and 39.)
- [56] J. Garcia-Algarra, J. Pastor, M. L. Mouronte, and J. Galeano. A structural approach to disentangle the visualization of bipartite biological networks. *Complexity*, 2018:1–11, 02 2018. (Cited on pages 6 and 41.)

- [57] J. Garcia-Algarra, J. M. M. Pastor, M. L. Mouronte, and J. Galeano. Bipartgraph: An interactive application to plot bipartite ecological networks. *bioRxiv*, 2017. (Cited on pages 6 and 41.)
- [58] J. García-Algarra, J. Pastor, J. Iriondo, and J. Galeano. Ranking of critical species to preserve the functionality of mutualistic networks using the k-core decomposition. *PeerJ*, 5(e3321), 2017. (Cited on pages 6 and 41.)
- [59] C. Giatsidis, K. Berberich, D. M. Thilikos, and M. Vazirgiannis. Visual exploration of collaboration networks based on graph degeneracy. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1512–1515. ACM, 2012. (Cited on page 38.)
- [60] C. Giatsidis, B. Cautis, S. Maniu, D. M. Thilikos, and M. Vazirgiannis. Quantifying trust dynamics in signed graphs, the s-cores approach. In *Proceedings of the 2014 SIAM International Conference on Data Mining, Philadelphia, Pennsylvania, USA, April 24-26, 2014*, pages 668–676, 2014. (Cited on pages 6 and 11.)
- [61] C. Giatsidis, F. D. Malliaros, D. M. Thilikos, and M. Vazirgiannis. Corecluster: A degeneracy based graph clustering framework. In *AAAI ’14: Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, pages 44–50, 2014. (Cited on pages 6 and 38.)
- [62] C. Giatsidis, D. M. Thilikos, and M. Vazirgiannis. D-cores: Measuring collaboration of directed graphs based on degeneracy. In *ICDM ’11: Proceedings of the 11th IEEE International Conference on Data Mining*, pages 201–210, 2011. (Cited on pages 6 and 10.)
- [63] C. Giatsidis, D. M. Thilikos, and M. Vazirgiannis. Evaluating cooperation in communities with the k-core structure. In *ASONAM ’11: Proceedings of the International Conference on Advances in Social Networks Analysis and Mining*, pages 87–93, 2011. (Cited on pages 6, 11, 33, and 37.)
- [64] 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. (Cited on page 10.)
- [65] P. Govindan, S. Soundarajan, T. Eliassi-Rad, and C. Faloutsos. Nimblecore: A space-efficient external memory algorithm for estimating core numbers. In *ASONAM*, pages 207–214. IEEE Computer Society, 2016. (Cited on pages 6 and 23.)
- [66] P. Govindan, C. Wang, C. Xu, H. Duan, and S. Soundarajan. The k-peak decomposition: Mapping the global structure of graphs. In *Proceedings of the 26th International Conference on World Wide Web, WWW ’17*, pages 1441–1450, 2017. (Cited on pages 6, 15, and 16.)
- [67] P. Hagmann, L. Cammoun, X. Gigandet, R. Meuli, C. J. Honey, and O. Sporns. Mapping the structural core of human cerebral cortex. *PLOS Biology*, 6(7):e159, 2008. (Cited on pages 6 and 41.)
- [68] X. He, H. Zhao, W. Cai, G.-G. Li, and F.-D. Pei. Analyzing the structure of earthquake network by k-core decomposition. *Physica A: Statistical Mechanics and its Applications*, 421:34–43, 2015. (Cited on page 6.)
- [69] J. Healy, J. Janssen, E. Milios, and W. Aiello. Characterization of graphs using degree cores. In *WAW ’08: Algorithms and Models for the Web-Graph*, pages 137–148, 2008. (Cited on pages 6 and 33.)

- [70] L. Hébert-Dufresne, A. Allard, J.-G. Young, and L. J. Dubé. Percolation on random networks with arbitrary k-core structure. *Physical Review E*, 88(6):062820, 2013. (Cited on page 33.)
- [71] X. Hu, F. Liu, V. Srinivasan, and A. Thomo. k-core decomposition on giraph and graphchi. In L. Barolli, I. Woungang, and O. K. Hussain, editors, *Advances in Intelligent Networking and Collaborative Systems*, pages 274–284, Cham, 2018. Springer International Publishing. (Cited on page 6.)
- [72] X. Huang, W. Lu, and L. V. Lakshmanan. Truss decomposition of probabilistic graphs: Semantics and algorithms. In *Proceedings of the 2016 ACM SIGMOD International Conference on Management of Data*, SIGMOD '16, pages 77–90, 2016. (Cited on pages 6 and 14.)
- [73] A. E. Isaac and S. Sinha. Analysis of core–periphery organization in protein contact networks reveals groups of structurally and functionally critical residues. *Journal of Biosciences*, 40(4):683–699, Oct 2015. (Cited on pages 6 and 41.)
- [74] H. Kabir and K. Madduri. Parallel k-core decomposition on multicore platforms. In *IPDPS Workshops*, pages 1482–1491. IEEE Computer Society, 2017. (Cited on pages 6 and 29.)
- [75] V. Kassiano, A. Gounaris, A. N. Papadopoulos, and K. Tsichlas. Mining uncertain graphs: An overview. In T. Sellis and K. Oikonomou, editors, *Algorithmic Aspects of Cloud Computing*, pages 87–116, Cham, 2017. Springer International Publishing. (Cited on page 12.)
- [76] D. Kempe, J. Kleinberg, and E. Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '03, pages 137–146, 2003. (Cited on page 44.)
- [77] 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. (Cited on pages 6 and 23.)
- [78] L. M. Kirousis and D. M. Thilikos. The linkage of a graph. *SIAM J. Comput.*, 25(3):626–647, 1996. (Cited on pages 6 and 9.)
- [79] M. Kitsak, L. K. Gallos, S. Havlin, F. Liljerosand, L. Muchnik, H. E. Stanley, and H. A. Makse. Identification of influential spreaders in complex networks. *Nature Physics*, 2010. (Cited on pages 6 and 35.)
- [80] R. Kumar, P. Raghavan, S. Rajagopalan, D. Sivakumar, A. Tomkins, and E. Upfal. The web as a graph. In *PODS*, 2000. (Cited on pages 4 and 6.)
- [81] J. Kunegis, A. Lommatzsch, and C. Bauckhage. The slashdot zoo: Mining a social network with negative edges. In *Proceedings of the 18th International Conference on World Wide Web*, WWW '09, pages 741–750, New York, NY, USA, 2009. ACM. (Cited on page 11.)
- [82] J. Kunegis, S. Schmidt, A. Lommatzsch, J. Lerner, E. W. D. Luca, and S. Albayrak. Spectral analysis of signed graphs for clustering, prediction and visualization. In *SDM*, pages 559–570. SIAM, 2010. (Cited on page 11.)
- [83] H. T. Kung, F. Luccio, and F. P. Preparata. On finding the maxima of a set of vectors. *J. ACM*, 22(4):469–476, Oct. 1975. (Cited on page 38.)

- [84] A. Kyrola, G. Blelloch, and C. Guestrin. Graphchi: Large-scale graph computation on just a pc. In *Proceedings of the 10th USENIX Conference on Operating Systems Design and Implementation*, OSDI'12, pages 31–46, Berkeley, CA, USA, 2012. USENIX Association. (Cited on page 23.)
- [85] N. Lahav, B. Ksherim, E. Ben-Simon, A. Maron-Katz, R. Cohen, and S. Havlin. K-shell decomposition reveals hierarchical cortical organization of the human brain. *New Journal of Physics*, 18(8):083013, 2016. (Cited on pages 6 and 42.)
- [86] R. Laishram, A. E. Sariyüce, T. Eliassi-Rad, A. Pinar, and S. Soundarajan. Measuring and improving the core resilience of networks. In *Proceedings of the 2018 World Wide Web Conference*, WWW '18, pages 609–618, 2018. (Cited on pages 6 and 34.)
- [87] J. Leskovec and E. Horvitz. Planetary-scale views on a large instant-messaging network. In *WWW '08: Proceedings of the 17th International Conference on World Wide Web*, pages 915–924, 2008. (Cited on pages 6 and 33.)
- [88] J. Leskovec, D. Huttenlocher, and J. Kleinberg. Signed networks in social media. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, pages 1361–1370, New York, NY, USA, 2010. ACM. (Cited on page 11.)
- [89] J. Leskovec and A. Krevl. SNAP Datasets: Stanford large network dataset collection. <http://snap.stanford.edu/data>, June 2014. (Cited on pages 29 and 31.)
- [90] R. Li, J. Su, L. Qin, J. X. Yu, and Q. Dai. Persistent community search in temporal networks. In *2018 IEEE 34th International Conference on Data Engineering (ICDE)*, pages 797–808, April 2018. (Cited on pages 6 and 38.)
- [91] R.-H. Li, L. Qin, F. Ye, J. X. Yu, X. Xiao, N. Xiao, and Z. Zheng. Skyline community search in multi-valued networks. In *Proceedings of the 2018 ACM SIGMOD International Conference on Management of Data*, SIGMOD '18, pages 457–472, New York, NY, USA, 2018. ACM. (Cited on pages 6 and 38.)
- [92] R.-H. Li, L. Qin, J. X. Yu, and R. Mao. Influential community search in large networks. *Proceedings of the VLDB Endowment*, 8(5):509–520, 2015. (Cited on pages 6, 36, and 39.)
- [93] R.-H. Li, J. X. Yu, and R. Mao. Efficient core maintenance in large dynamic graphs. *IEEE Transactions on Knowledge and Data Engineering*, 26(10):2453–2465, 2014. (Cited on page 25.)
- [94] D. R. Lick and A. T. White. k-degenerate graphs. *Canadian Journal of Mathematics*, 22:1082–1096, 1970. (Cited on pages 6 and 8.)
- [95] J.-H. Lin, Q. Guo, W.-Z. Dong, L.-Y. Tang, and J.-G. Liu. Identifying the node spreading influence with largest k-core values. *Physics Letters A*, 378(45):3279–3284, 2014. (Cited on pages 6 and 35.)
- [96] M. Litvak and M. Last. Graph-based keyword extraction for single-document summarization. In *Proceedings of the workshop on Multi-source Multilingual Information Extraction and Summarization*, pages 17–24. Association for Computational Linguistics, 2008. (Cited on page 39.)
- [97] L. Lü, D. Chen, X.-L. Ren, Q.-M. Zhang, Y.-C. Zhang, and T. Zhou. Vital nodes identification in complex networks. *Physics Reports*, 650:1–63, 2016. (Cited on page 35.)

- [98] L. Lü, T. Zhou, Q.-M. Zhang, and H. E. Stanley. The h-index of a network node and its relation to degree and coreness. *Nature Communications*, 7:10168 EP –, 01 2016. (Cited on pages 28 and 35.)
- [99] F. Luo, B. Li, X.-F. Wan, and R. H. Scheuermann. Core and periphery structures in protein interaction networks. *BMC Bioinformatics*, 10(Suppl 4):s8, 2009. (Cited on pages 6 and 41.)
- [100] G. Malewicz, M. H. Austern, A. J. Bik, J. C. Dehnert, I. Horn, N. Leiser, and G. Czajkowski. Pregel: A system for large-scale graph processing. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of Data*, SIGMOD '10, pages 135–146, New York, NY, USA, 2010. ACM. (Cited on pages 26 and 30.)
- [101] F. D. Malliaros, A. N. Papadopoulos, and M. Vazirgiannis. Core decomposition in graphs: Concepts, algorithms and applications. In *EDBT*, pages 720–721. OpenProceedings.org, 2016. (Cited on page 5.)
- [102] F. D. Malliaros, M.-E. G. Rossi, and M. Vazirgiannis. Locating influential nodes in complex networks. *Scientific reports*, 6:19307, 2016. (Cited on pages 6 and 35.)
- [103] F. D. Malliaros and M. Vazirgiannis. To stay or not to stay: modeling engagement dynamics in social graphs. In *22nd ACM International Conference on Information and Knowledge Management, CIKM'13*, pages 469–478, 2013. (Cited on pages 6 and 39.)
- [104] F. D. Malliaros and M. Vazirgiannis. Vulnerability assessment in social networks under cascade-based node departures. *EPL (Europhysics Letters)*, 110(6):68006, 2015. (Cited on pages 6 and 39.)
- [105] D. W. Matula and L. L. Beck. Smallest-last ordering and clustering and graph coloring algorithms. *J. ACM*, 30(3):417–427, July 1983. (Cited on page 21.)
- [106] P. Meladianos, G. Nikolentzos, F. Rousseau, Y. Stavrakas, and M. Vazirgiannis. Degeneracy-based real-time sub-event detection in twitter stream. In *ICWSM*, pages 248–257, 2015. (Cited on pages 6 and 39.)
- [107] P. Meladianos, A. Tixier, I. Nikolentzos, and M. Vazirgiannis. Real-time keyword extraction from conversations. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, volume 2, pages 462–467, 2017. (Cited on pages 6 and 39.)
- [108] R. Mihalcea and P. Tarau. Texttrank: Bringing order into text. In *Proceedings of the 2004 conference on empirical methods in natural language processing*, 2004. (Cited on page 39.)
- [109] A. Montresor, F. De Pellegrini, and D. Miorandi. Distributed k-core decomposition. In *PODC*, pages 207–208, 2011. (Cited on pages 6, 23, and 30.)
- [110] A. Montresor, F. De Pellegrini, and D. Miorandi. Distributed k-core decomposition. *IEEE Transactions on Parallel and Distributed Systems*, 24(2):288–300, 2013. (Cited on pages 6 and 30.)
- [111] F. Morone, K. Burleson-Lesser, H. Vinutha, S. Sastry, and H. A. Makse. The jamming transition is a k-core percolation transition. *Physica A: Statistical Mechanics and its Applications*, 516:172–177, 2019. (Cited on pages 6 and 41.)
- [112] F. Morone, G. Ferraro, and H. A. Makse. The k-core as a predictor of structural collapse in mutualistic ecosystems. *Nature Physics*, 10 2018. (Cited on pages 6 and 41.)

- [113] G. Nikolentzos, P. Meladianos, S. Limnios, and M. Vazirgiannis. A degeneracy framework for graph similarity. In *IJCAI*, pages 2595–2601, 2018. (Cited on pages 6 and 40.)
- [114] M. P. O’Brien and B. D. Sullivan. Locally estimating core numbers. In *ICDM*, pages 460–469, 2014. (Cited on pages 6 and 27.)
- [115] P. Parchas, F. Gullo, D. Papadias, and F. Bonchi. The pursuit of a good possible world: Extracting representative instances of uncertain graphs. In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data*, pages 967–978, 2014. (Cited on page 32.)
- [116] P. Parchas, F. Gullo, D. Papadias, and F. Bonchi. Uncertain graph processing through representative instances. *ACM Trans. Database Syst.*, 40(3):20:1–20:39, 2015. (Cited on page 32.)
- [117] R. Pastor-Satorras and A. Vespignani. Epidemic spreading in scale-free networks. *Physical review letters*, 86(14):3200, 2001. (Cited on page 35.)
- [118] S. Pei and H. A. Makse. Spreading dynamics in complex networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2013(12):P12002, 2013. (Cited on page 35.)
- [119] S. Pei, L. Muchnik, J. S. Andrade Jr, Z. Zheng, and H. A. Makse. Searching for superspreaders of information in real-world social media. *Scientific reports*, 4:5547, 2014. (Cited on pages 6 and 35.)
- [120] M. Pellegrini, M. Baglioni, and F. Geraci. Protein complex prediction for large protein protein interaction networks with the core & peel method. *BMC Bioinformatics*, 17(12):372, Nov 2016. (Cited on pages 6 and 41.)
- [121] Y. Peng, Y. Zhang, W. Zhang, X. Lin, and L. Qin. Efficient probabilistic k-core computation on uncertain graphs. In *2018 IEEE 34th International Conference on Data Engineering (ICDE)*, pages 1192–1203, 2018. (Cited on pages 6, 12, and 33.)
- [122] E. M. Phizicky and S. Fields. Protein-protein interactions: methods for detection and analysis. *Microbiological reviews*, 59 1:94–123, 1995. (Cited on page 4.)
- [123] M. Potamias, F. Bonchi, A. Gionis, and G. Kollios. K-nearest neighbors in uncertain graphs. *Proceedings of the VLDB Endowment*, pages 997–1008, 2010. (Cited on page 32.)
- [124] F. Rousseau and M. Vazirgiannis. Main core retention on graph-of-words for single-document keyword extraction. In *ECIR ’15: Proceedings of the 37th European Conference on Information Retrieval*, pages 382–393, 2015. (Cited on pages 6 and 39.)
- [125] D. Samu, A. K. Seth, and T. Nowotny. Influence of wiring cost on the large-scale architecture of human cortical connectivity. *PLOS Computational Biology*, 10(4):1–24, 04 2014. (Cited on pages 6 and 42.)
- [126] A. E. Sarıyüce, B. Gedik, G. Jacques-Silva, K.-L. Wu, and Ü. V. Çatalyürek. Incremental k-core decomposition: algorithms and evaluation. *The VLDB Journal*, 25(3):425–447, Jun 2016. (Cited on page 6.)
- [127] A. E. Sarıyüce, B. Gedik, G. Jacques-Silva, K.-L. Wu, and U. V. Çatalyürek. Streaming algorithms for k-core decomposition. *Proceedings of the VLDB Endowment*, 6(6):433–444, Apr. 2013. (Cited on pages 6, 12, 24, and 25.)

- [128] A. E. Sariyüce and A. Pinar. Peeling bipartite networks for dense subgraph discovery. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining, WSDM*, pages 504–512, 2018. (Cited on pages 6, 18, and 19.)
- [129] A. E. Sariyüce, C. Seshadhri, and A. Pinar. Local algorithms for hierarchical dense subgraph discovery. *Proc. VLDB Endow.*, 12(1):43–56, Sept. 2018. (Cited on pages 6 and 28.)
- [130] A. E. Sariyüce, C. Seshadhri, A. Pinar, and U. V. Catalyurek. Finding the hierarchy of dense subgraphs using nucleus decompositions. In *Proceedings of the 24th International Conference on World Wide Web, WWW '15*, pages 927–937, 2015. (Cited on pages 16 and 17.)
- [131] S. Sarkar, A. Bhagwat, and A. Mukherjee. Core2vec: A core-preserving feature learning framework for networks. In *IEEE/ACM 2018 International Conference on Advances in Social Networks Analysis and Mining, ASONAM 2018*, pages 487–490, 2018. (Cited on page 43.)
- [132] C. Schmidt, H. D. Pfister, and L. Zdeborová. Minimal sets to destroy the k-core in random networks. *Physical Review E*, 99(2):022310, 2019. (Cited on pages 6 and 34.)
- [133] S. B. Seidman. Network Structure and Minimum Degree. *Social Networks*, 5:269–287, 1983. (Cited on pages 6, 21, and 33.)
- [134] N. Shailaja Dasari, D. Ranjan, and M. Zubair. Park: An efficient algorithm for k-core decomposition on multicore processors. *Proceedings - 2014 IEEE International Conference on Big Data, IEEE Big Data 2014*, pages 9–16, 01 2015. (Cited on pages 6 and 28.)
- [135] M. Shanahan, V. Bingman, T. Shimizu, M. Wild, and O. Güntürkün. Large-scale network organization in the avian forebrain: a connectivity matrix and theoretical analysis. *Frontiers in Computational Neuroscience*, 7:89, 2013. (Cited on pages 6 and 41.)
- [136] K. Shin, T. Eliassi-Rad, and C. Faloutsos. Corescope: Graph mining using k-core analysis - patterns, anomalies and algorithms. In *ICDM*, pages 469–478. IEEE, 2016. (Cited on pages 6 and 34.)
- [137] K. Shin, T. Eliassi-Rad, and C. Faloutsos. Patterns and anomalies in k-cores of real-world graphs with applications. *Knowl. Inf. Syst.*, 54(3):677–710, 2018. (Cited on pages 6, 34, and 36.)
- [138] M. Sozio and A. Gionis. The community-search problem and how to plan a successful cocktail party. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '10*, pages 939–948, 2010. Exported from <https://app.dimensions.ai> on 2019/04/27. (Cited on page 38.)
- [139] P. Strouthopoulos and A. N. Papadopoulos. Core discovery in hidden graphs. *CoRR (to appear in Data and Knowledge Engineering)*, abs/1712.02827, 2017. (Cited on pages 6 and 13.)
- [140] Y. Tao, C. Sheng, and J. Li. Finding maximum degrees in hidden bipartite graphs. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of Data, SIGMOD '10*, pages 891–902, New York, NY, USA, 2010. ACM. (Cited on page 13.)
- [141] N. Tatti and A. Gionis. Density-friendly graph decomposition. In *WWW*, pages 1089–1099, 2015. (Cited on pages 6 and 15.)
- [142] J. B. Tenenbaum, V. De Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *science*, 290(5500):2319–2323, 2000. (Cited on page 43.)

- [143] A. Tixier, F. D. Malliaros, and M. Vazirgiannis. A graph degeneracy-based approach to keyword extraction. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1860–1870. Association for Computational Linguistics, 2016. (Cited on pages 6 and 39.)
- [144] A. Tixier, K. Skianis, and M. Vazirgiannis. Gowvis: a web application for graph-of-words-based text visualization and summarization. *Proceedings of ACL-2016 System Demonstrations*, pages 151–156, 2016. (Cited on pages 6 and 39.)
- [145] C. E. Tsourakakis, U. Kang, G. L. Miller, and C. Faloutsos. Doulion: counting triangles in massive graphs with a coin. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 837–846. ACM, 2009. (Cited on page 35.)
- [146] J. Ugander, B. Karrer, L. Backstrom, and C. Marlow. The anatomy of the facebook social graph, 2011. cite arxiv:1111.4503Comment: 17 pages, 9 figures, 1 table. (Cited on page 33.)
- [147] M. P. van den Heuvel and O. Sporns. Rich-club organization of the human connectome. *Journal of Neuroscience*, 31(44):15775–15786, 2011. (Cited on pages 6 and 41.)
- [148] T. Verma, F. Russmann, N. Araújo, J. Nagler, and H. Herrmann. Emergence of core-peripheries in networks. *Nature Communications*, 7, 2016. (Cited on pages 6 and 33.)
- [149] U. Von Luxburg. A tutorial on spectral clustering. *Statistics and computing*, 17(4):395–416, 2007. (Cited on page 43.)
- [150] J. Wang and J. Cheng. Truss decomposition in massive networks. *Proceedings of the VLDB Endowment*, 5(9):812–823, May 2012. (Cited on pages 6, 14, and 15.)
- [151] K. Wang, X. Cao, X. Lin, W. Zhang, and L. Qin. Efficient computing of radius-bounded k-cores. In *2018 IEEE 34th International Conference on Data Engineering (ICDE)*, pages 233–244, 2018. (Cited on pages 6, 19, and 20.)
- [152] N. Wang, D. Yu, H. Jin, C. Qian, X. Xie, and Q. Hua. Parallel algorithm for core maintenance in dynamic graphs. In *2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)*, volume 00, pages 2366–2371, June 2017. (Cited on pages 6 and 29.)
- [153] D. Wen, L. Qin, Y. Zhang, X. Lin, and J. X. Yu. I/O efficient core graph decomposition at web scale. In *2016 IEEE 32nd International Conference on Data Engineering (ICDE)*, pages 133–144. IEEE, 2016. (Cited on pages 6, 23, 24, and 25.)
- [154] D. Wen, L. Qin, Y. Zhang, X. Lin, and J. X. Yu. I/o efficient core graph decomposition: Application to degeneracy ordering. *IEEE Transactions on Knowledge & Data Engineering*, 31(1):75–90, Jan. 2019. (Cited on pages 6, 23, 24, and 26.)
- [155] T. White. *Hadoop: The Definitive Guide*. O’Reilly, 4th edition, 2015. (Cited on page 30.)
- [156] C. I. Wood and I. V. Hicks. The minimal k-core problem for modeling k-assemblies. *The Journal of Mathematical Neuroscience (JMN)*, 5(1):14, Jul 2015. (Cited on pages 6 and 42.)
- [157] H. Wu, J. Cheng, Y. Lu, Y. Ke, Y. Huang, D. Yan, and H. Wu. Core decomposition in large temporal graphs. In *BigData*, pages 649–658. IEEE, 2015. (Cited on pages 6, 12, 25, and 31.)

- [158] D. Yan, J. Cheng, Y. Lu, and W. Ng. Blogel: A block-centric framework for distributed computation on real-world graphs. *Proceedings of the VLDB Endowment*, 7(14):1981–1992, 2014. (Cited on pages 26 and 31.)
- [159] M. L. Yiu, E. Lo, and J. Wang. Identifying the most connected vertices in hidden bipartite graphs using group testing. *IEEE Transactions on Knowledge & Data Engineering*, 25:2245–2256, 10 2013. (Cited on page 13.)
- [160] M. Zaharia, R. S. Xin, P. Wendell, T. Das, M. Armbrust, A. Dave, X. Meng, J. Rosen, S. Venkataraman, M. J. Franklin, A. Ghodsi, J. Gonzalez, S. Shenker, and I. Stoica. Apache spark: A unified engine for big data processing. *Commun. ACM*, 59(11):56–65, Oct. 2016. (Cited on page 30.)
- [161] L. Zdeborová, P. Zhang, and H.-J. Zhou. Fast and simple decycling and dismantling of networks. *Scientific reports*, 6:37954, 2016. (Cited on pages 6 and 34.)
- [162] F. Zhang, W. Zhang, Y. Zhang, L. Qin, and X. Lin. Olak: An efficient algorithm to prevent unraveling in social networks. *Proceedings of the VLDB Endowment*, 10(6):649–660, Feb. 2017. (Cited on pages 6 and 40.)
- [163] F. Zhang, Y. Zhang, L. Qin, W. Zhang, and X. Lin. Finding critical users for social network engagement: The collapsed k-core problem. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, pages 245–251, 2017. (Cited on pages 6 and 40.)
- [164] F. Zhang, Y. Zhang, L. Qin, W. Zhang, and X. Lin. When engagement meets similarity: Efficient (k,r)-core computation on social networks. *Proceedings of the VLDB Endowment*, 10(10):998–1009, June 2017. (Cited on pages 6 and 18.)
- [165] G.-Q. Zhang, G.-Q. Zhang, Q.-F. Yang, S.-Q. Cheng, and T. Zhou. Evolution of the Internet and its cores. *New Journal of Physics*, 10(12):123027+, Dec. 2008. (Cited on pages 6 and 34.)
- [166] Y. Zhang and S. Parthasarathy. Extracting analyzing and visualizing triangle k-core motifs within networks. In *ICDE '12: Proceedings of the 2012 IEEE 28th International Conference on Data Engineering*, pages 1049–1060, 2012. (Cited on pages 6, 14, and 37.)
- [167] Y. Zhang, J. X. Yu, Y. Zhang, and L. Qin. A fast order-based approach for core maintenance. In *2017 IEEE 33rd International Conference on Data Engineering (ICDE)*, pages 337–348, April 2017. (Cited on pages 6 and 25.)
- [168] R. Zhuo-Ming, L. Jian-Guo, S. Feng, H. Zhao-Long, and G. Qiang. Analysis of the spreading influence of the nodes with minimum k-shell value in complex networks. *Acta Physica Sinica*, 62(10):108902, 2013. (Cited on pages 6 and 35.)
- [169] V. Zlatić, D. Garlaschelli, and G. Caldarelli. Networks with arbitrary edge multiplicities. *EPL (Europhysics Letters)*, 97(2):28005, 2012. (Cited on page 37.)
- [170] Z. Zou and R. Zhu. Truss decomposition of uncertain graphs. *Knowledge and Information Systems*, 50(1):197–230, Jan. 2017. (Cited on pages 6 and 14.)