An Experimental Evaluation of *k*-core Decomposition on Giraph and GraphChi

Xin Hu, Fangming Liu, Venkatesh Srinivasan, Alex Thomo

University of Victoria, Canada
{huxin,jasmine1,srinivas,thomo}@uvic.ca

Abstract. The analysis of characteristics of large-scale graphs has shown tremendous benefits in social networks, spam detection, epidemic disease control, analyzing software systems and so on. However, today, processing graph algorithms on massive datasets is not an easy task not only because of the large data volume, but also the complexity of the graph algorithm. Therefore, a number of large-scale processing platforms have been developed to tackle these problems. GraphChi is a popular system that is capable of executing massive graph datasets on a single PC. Some researchers claim that GraphChi has the same or even better performance, compared with distributed graph-analytics platforms such as the popular Apache Giraph. In this paper, we implement a well-optimized k-core decomposition algorithm on Giraph. Then we provide a comparison of the performance of running the k-core decomposition algorithm in Giraph and GraphChi using various graph datasets.

Keywords: vertex-centric model, graph theory, Apache Giraph, GraphChi

1 Introduction

Graphs are widely used as a data structure for representing relationships between different objects or people. In many applications it is of great benefit to discover graph structure and analyse it. For example, advertisement companies utilize social network structure to find targeted communities in order to spread their commercials [6]; by detecting densely connected networks among web links, an organization can facilitate combating spam [24]; as another example, people simulate infectious disease spread by using graph structure in order to be better prepared to stop an epidemic [4]; software engineers use graphs to extract and analyse large-scale software systems so that the complexity of the systems is tackled more effectively [31].

In all these applications, the graphs are grouped into different subgraphs where some are dense and the others are sparse. Intuitively, nodes in a dense network have close ties with each other. So detecting those dense subgraphs in real-life networks is an important task in graph analytics. There are already many early studies on finding dense components such as in [25], [15], and [7]. One popular study is the k-core decomposition which attempts to find all the maximal connected subgraphs of a graph so that all vertices within it have

at least k neighbors. In another dense graph searching study, [29] proposed a novel density measure that extracts optimal Quasi-Cliques and returns denser subgraphs with smaller diameters.

However today, graphs with millions of nodes and billions of edges are very common, and thus, doing analytics on large graph datasets is challenging due to the sheer size and complexity of graph computations. For example, in April 2017, the Facebook social network graph has over 1.97 billion monthly active users and more than 140 billion friendship connections, followed by Whatsapp and WeChat, with 1.2 billion users and 889 millions users, respectively [27].

Pregel [19], developed by Google in 2010, is a scalable and fault-tolerant platform that provides APIs for supporting large graph processing.

This model provides a vertex-centric computation model which enables users to only focus on programming itself without knowledge of the mechanisms behind it. The Pregel vertex-centric (VC) model has been implemented by several open source projects, for example, Apache Giraph is a framework designed to process iterative graph algorithms that can be parallelized across multiple commodity machines. Giraph became popular after careful engineering by Facebook researchers in 2012 to scale the computation of PageRank to a trillion-edge graph of user interactions using 200 machines [1].

Systems like Apache Giraph and Pregel require a distributed computing cluster to process large scale graph data quickly and effectively. Although distributed computing facilities such as cloud computing clusters are becoming more common and accessible, nevertheless, the question of how to process large scale graph data effectively without distributed commodity computing clusters is an interesting avenue for a data analyst who may need to analyze a large graph dataset but is unable to access a distributed computing cluster. GraphChi proposed by Kyrola and Guestrin [5] is a disk-based, vertex-centric system, which segments a large graph into different partitions. Then, a novel parallel sliding window algorithm is implemented to reduce random access to the data graph. Graphchi can process hundreds to thousands of graph vertices per second. GraphChi became popular, around the same time in 2012, as it made possible to perform intensive graph computations in a single PC in just under 59 minutes, whereas the distributed systems were taking 400 minutes using a cluster of about 1,000 computers (as reported by MIT Technology Review [23]). Since then, new versions of GraphChi and Apache Giraph have been released, where new ideas and optimizations have been implemented. Therefore, one needs to validate again the claims made several years ago. In [17], Lu and Thomo present a detailed evaluation of computing PageRank, shortest-paths, and weakly-connected-components on Giraph and GraphChi. In this work, we embark in computing k-core decomposition using a vertex-centric algorithm on Giraph and GraphChi. We adapt for this the algorithm of Montresor et. al. [22]. The latter was implemented and optimized for GraphChi by Khaouid et. al. in [12].

2 Graphs and Cores

We consider undirected and unweighted graphs. We denote a graph by $G = (V, E)$, where V denotes the set of n vertices and E denotes the set of m edges linking the vertices. The neighbors of a vertex v are denoted by $N_G(v)$, and the degree of this vertex is denoted by $d_G(v)$.

Let $K \subseteq V$ be a subset of vertices of a graph $G = (V, E)$. We have the following definitions.

Definition 1 *Graph $G(K) = (K, E_K)$, where $E_K = \{(u, v) \in E : u, v \in K\}$ is called* the subgraph of G induced by K.

Definition 2 *(k-core): Graph $G(K)$ is a k-core if and only if for each $v \in G(K)$, $d_{G(K)}(v) \geq k$, and $G(K)$ is a maximum size subgraph with this property.*

The process of finding the k-core of a graph is to recursively prune all vertices that have degree less than k until converging to a subgraph in which all the vertices have degrees greater or equal to k.

Definition 3 *(Coreness): A vertex has a coreness of k if it is in the k-core but not in the $k+1$-core.*

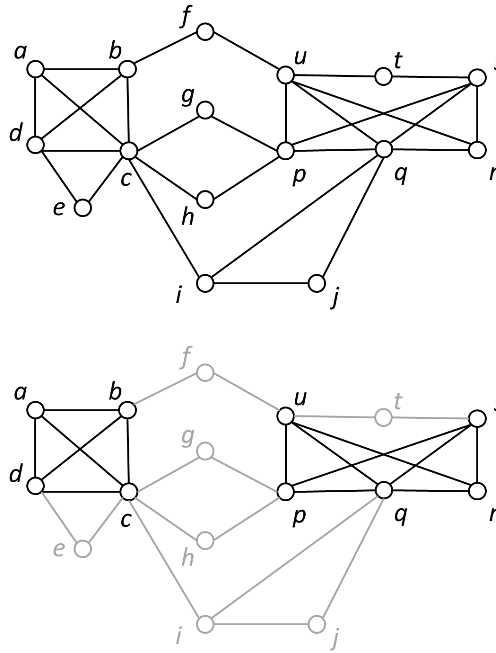

Fig. 1. Example Graph (top). The 3-core of the graph(bottom) [12].

For an example see Figure 1 (from [12]). By definition cores are nested, meaning 3-core also belongs to 2-core and 1-core. Each of the vertices in the graph is in 2-core and no vertices have coreness greater than 4. Even though vertices such as b, c, and q have a degree of 4 or more, their neighbours have degree less than 4, thus they do not belong in 4-core.

2.1 Pseudocode

[12] implements and optimizes an algorithm from [22] on GraphChi. It takes advantage of the vertex-centric programming model proposed in Pregel [19].

This model requires programmers to "think like a vertex" that can send messages to its neighbours by writing on to outgoing edges and receive messages from incoming edges, that is, in a graph, when computing an update function on a vertex, it enables the value on the vertex to be sent to its adjacent vertices. The computations in the update function go on iteratively until there are no more messages sent by vertices.

Computing the k-core decomposition becomes very natural in the vertex-centric model. First all vertices store an estimation of their coreness number in their vertex value, which initially is their degree. Messages can be used to propagate the estimation from the vertex itself to its neighbours using outgoing edges. Algorithm 1 and Algorithm 2 show the flow of the computation in the vertex-centric Giraph model. They are adaptations of corresponding GraphChi algorithms in [12]. Even though both Giraph and GraphChi follow a vertex-centric model, there are differences in the way operations are expressed in each of them.

In Algorithm 1, the first superstep is a special case, where each vertex initializes its vertex value with its degree number which equals the number of its out-going edges. Then it sends this value to all its neighbours. Any vertex that is not being halted will be rescheduled to the next superstep (Lines 2-5).

In the next superstep, a function that computes an upper bound of the coreness of the vertex is assigned to a local estimate called *localEstimate*. The vertex value will be updated to the local estimate if the vertex value is greater than the upper bound of the vertex. In such a case, this new vertex value will also be sent to all the neighbours of the vertex (Lines 7-10).

The vertex will have a chance to lower its core estimate if any of its neighbours has a lower coreness value, then this vertex will be scheduled to the next superstep. Otherwise, this vertex will not be scheduled, and then it can switch to inactive state by voting to halt (Lines 11-20).

Algorithm 2 displays the details of the computation of a tighter upper bound of the vertex. It uses a *count array* which is indexed by the value of the upper bound of its neighbours. The value of each element is the number of the neighbours which have an estimate equal to the index. The largest index of the count array is the value of the vertex's the current coreness estimate. Any neighbour that has an upper bound greater than the current vertex value will be added to the last element of the array (Lines 5-8).

Algorithm 1 Update function running at a vertex

```
1: function Update(Vertex vertex, Iterable messages)
2:     if superstep = 0 then
3:         vertex.value ← vertex.numOutEdges
4:         sendMessageToAllEdges(vertex, vertex.value)
5:     else
6:         localEstimate ←computeUpperBound(vertex, messages)
7:         if localEstimate < vertex.value then
8:             vertex.value ← localEstimate
9:             sendMessageToAllEdges(vertex, vertex.value)
10:        end if
11:        halt ← true
12:        for all message in messages do
13:            if vertex.value ≥ message then
14:                halt ← false
15:                break
16:            end if
17:        end for
18:        if halt = true then
19:            vertex.voteToHalt()
20:        end if
21:    end if
22: end function
```

In order to compute the new upper bound of the vertex, a loop is used to add the array elements beginning from the largest index down to 2 until the summation of the array elements is greater or equal to the corresponding index. This index is the new coreness estimate for the vertex in the superstep (Lines 9-15).

3 Experimental Evaluation

All the experiments are conducted on Amazon Web Services (AWS) using the Amazon Elastic Compute Cloud (EC2) platform. We configured twenty-one virtual machines, with one master machine and twenty slaves. All of the virtual machines have two cores, Intel Xeon Family, 2.4 GHz CPU with 8GB RAM running Ubuntu Linux System.

To explore the relationship between the number of the slaves and the running time, we respectively use two, five, ten, fifteen, and twenty slaves to handle different datasets. The datasets we used in this experiment were chosen from Stanford Large Network Dataset Collection. They are Astro Physics (ca-AstroPh), Gnutella P2P network (p2p-Gnutella31), Amazon product co-purchasing network (amazon0601), California road network (roadNet-CA), and Live-Journal social network (soc-LiveJournal1). The detailed information of these datasets is described in Table 1. From the table, we can observe that the first two datasets are small; they only have few thousands of nodes and a hundred thousands of

Algorithm 2 computeUpperBound function for a vertex

```
1: function COMPUTEUPPERBOUND(Vertex vertex, Iterable messages)
2:     for all i ← 1 to vertex.value do
3:         c[i] ← 0
4:     end for
5:     for all message in messages do
6:         j ←min(message, vertex.value)
7:         c[j] + +
8:     end for
9:     cumul ← 0
10:    for all i ← vertex.value to 2 do
11:        cumul ← cumul + c[i]
12:        if cumul ≥ i then
13:            return i
14:        end if
15:    end for
16: end function
```

edges. The medium sized datasets are amazon0601 and roadNet-CA with around three million edges. The largest dataset is soc-LiveJournal1, which has 4.8 million nodes and approximately 69 million edges.

Table 1. Datasets used for the experiments

Dataset Name	**Numbers of Nodes**	**Numbers of Edges**
ca-AstroPh	18,772	198,110
p2p-Gnutella31	62,586	147,892
amazon0601	403,394	3,387,388
roadNet-CA	1,965,206	2,766,607
soc-LiveJournal1	4,847,571	68,993,773

Results for the Giraph implementation are shown in Table 2. Column "Sent Messages" gives the total numbers of messages that were sent during the whole computation. Column "Update Times" gives the average vertex update times for each dataset. "K-Max" and "K-Ave" are the maximum and average k-core numbers for each dataset. From the table, we can observe that the number of sent messages and vertex update times are not only dependent on the size of the datasets, but also on K-Max and K-Ave. The larger the latter numbers are, the more frequent the message sending and vertex updates will be.

Figure 2 (left) shows the number of iterations executed on Giraph and GraphChi. The reason why the iteration numbers for the same dataset are different is that we cannot control the order of running each vertex in distributed cluster-based Giraph. However, because of the selective scheduling feature, the order is fixed when running GraphChi on a single machine. Except for the largest dataset and

Table 2. Results for the Giraph implementation

Dataset Name	$\lvert V\rvert$	$\lvert E\rvert$	Sent Messages	Update Times(ms)	K-Max	K-Ave
ca-AstroPh	18.7 K	198.1 K	5,104,983	5414	17	2.01
p2p-Gnutella31	62.6 K	147.9 K	322,906	280	50	1.143
amazon0601	0.4 M	2.4 M	12,122,458	2284	10	2.51
roadNet-CA	2.0 M	2.8 M	11,035,492	785	6	1.999
soc-LiveJournal1	4.8 M	43.1 M	888,141,866	3507	434	1.689

the smallest dataset, Giraph needs fewer iterations than GraphChi with its advantage of running on multiple machines. The percentage of updated nodes over several iterations is shown in Figure 2 (right).

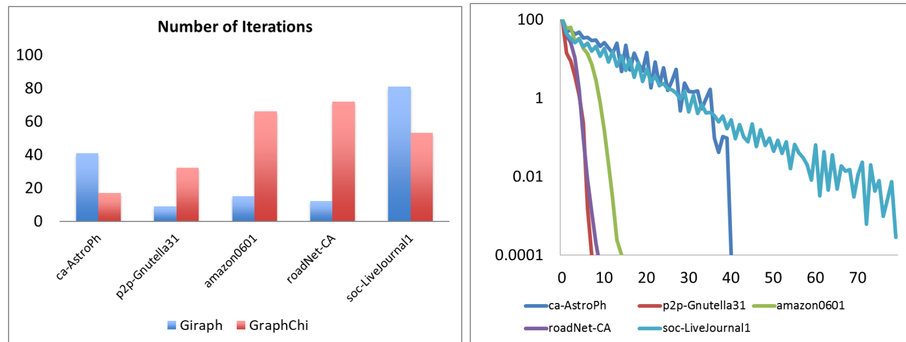

Fig. 2. Number of iterations (left). Percentage of updated vertices in Giraph vs. number of iterations (right).

Figure 3, left and right, shows the running time (in milliseconds) of Giraph versus the number of slave machines used. In the left, we see that with the increase in the number of machines, the running time also increased. The more machines we have, the fewer the number of tasks assigned to each one are. However, the more machines we have, the more the amount of time spent on communication is. That is why the running time does not decrease when we configure more machines for Giraph. For the largest dataset shown in the right, we can notice that Giraph with two slave machines needs the most running time for the computation, which is around 800 seconds. On the contrary, it takes the least running time with ten machines, which can finish the computation within around 700 seconds.

To compare the running time with Giraph and GraphChi, we select the least running time with the proper number of machines for each dataset for Giraph. Figure 4 shows the running time comparisons between Giraph and GraphChi. Giraph spent more time than GraphChi on running the algorithm for all datasets.

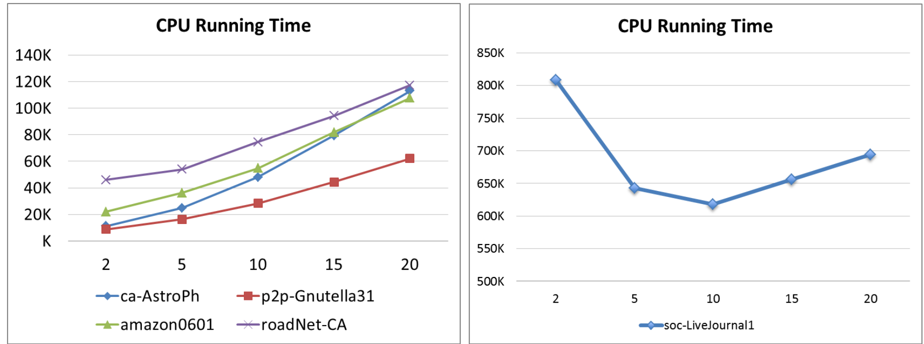

Fig. 3. Running time (ms) in Giraph vs. number of machines.

To be specific, the running time of Giraph and GraphChi are very close when dealing with medium data amazon0601 and roadNet-CA. However, GraphChi is more efficient in computing k-core for the small size and large size datasets on a single machine than Giraph with multiple machines.

However, the conclusion is that the performance of Giraph, with a relatively small number of machines, is quite close to the performance of GraphChi. This is in contrast to what was reported in 2012 in [23], where the situation was quite different. Then, a cluster of 1000 machines was an order of magnitude slower than GraphChi running on a single machine.

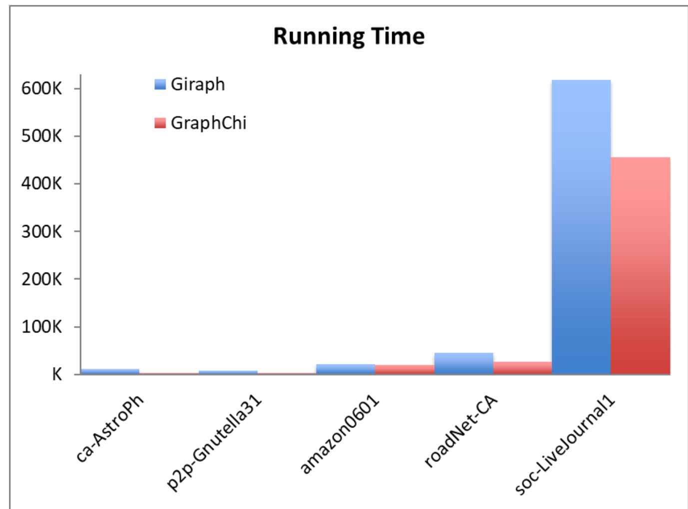

Fig. 4. Running time (ms) in Giraph compared to GraphChi.

4 Related Work

The Pregel distributed graph processing framework was introduced by Malewicz et. al in [19]. Apache Giraph (http://giraph.apache.org) is an open source implementation of Pregel based on Hadoop. An excellent reference on Giraph is the recent book by Martella, Shaposhnik, and Logothetis [20].

GraphChi was created by Aapo et. al [16]. Its excellent speed compared to distributed vertex-centric systems at the time (2012) was commented with awe at MIT Technology Review [23].

Around the same time, a group of Facebook researchers introduced several optimizations to Giraph [1]. These and other optimizations to Giraph are described in a recent paper by Ching et. al in [2].

Thorough analysis of distributed vertex-centric systems have been presented by Han et. al [10] and Lu et. al in [18]. A recent survey of vertex-centric frameworks is by McCune et. al [21].

5 Conclusions

From the experiments for k-core computation on Giraph and GraphChi, we observe that Giraph is suitable for analyzing medium- and large-size data since it can synchronously implement the computation by assigning the tasks to each slave. We observe that the performance of Giraph for computing k-core using a relatively small number of machines is in the same range as the performance of GraphChi for the same vertex-centric algorithm. As such, this is in contrast to the situation described in [23], where a cluster of 1000 machines was slower by an order of magnitude than GraphChi running on a single machine.

As future work, we would like to analyze more specialized graphs with their edges being labeled and/or weighted (c.f. [8, 9]). It will be interesting to see how to devise vertex-centric algorithms for computing k-core on such graphs. Also, adapting the algorithms for environments with many machines failures (c.f. [26, 28]) is another avenue to explore. Finally, we would like to explore the behaviour of Giraph vs. single machine systems for computing other complex graph analytics, such those for trust propagation and probabilistic graph summarization (c.f. [3, 14, 11]), as well as, complex analytics for special kind of graphs, e.g. user-item bipartite graphs for recommendation systems (c.f. [5, 13, 30]).

References

1. Scaling apache giraph to a trillion edges. http://bit.ly/1TomAkh. Accessed: 2016-05-22.
2. A. Ching, S. Edunov, M. Kabiljo, D. Logothetis, and S. Muthukrishnan. One trillion edges: graph processing at facebook-scale. *PVLDB*, 8(12), 2015.
3. M. Chowdhury, A. Thomo, and W. W. Wadge. Trust-based infinitesimals for enhanced collaborative filtering. In *COMAD*, 2009.

4. 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. ACM, 2014.
5. S. Ebrahimi, N. M. Villegas, H. A. Müller, and A. Thomo. Smarterdeals: a context-aware deal recommendation system based on the smartercontext engine. In *Proceedings of the 2012 Conference of the Center for Advanced Studies on Collaborative Research*, pages 116–130. IBM Corp., 2012.
6. S. Fortunato. Community detection in graphs. *Physics reports*, 486(3):75–174, 2010.
7. A. V. Goldberg. *Finding a maximum density subgraph*. University of California Berkeley, CA, 1984.
8. G. Grahne and A. Thomo. Algebraic rewritings for optimizing regular path queries. *Theoretical Computer Science*, 296(3), 2003.
9. G. Grahne, A. Thomo, and W. Wadge. Preferentially annotated regular path queries. In *ICDT*. Springer, 2007.
10. M. Han, K. Daudjee, K. Ammar, M. T. Özsu, X. Wang, and T. Jin. An experimental comparison of pregel-like graph processing systems. *Proceedings of the VLDB Endowment*, 7(12):1047–1058, 2014.
11. N. Hassanlou, M. Shoaran, and A. Thomo. Probabilistic graph summarization. In *International Conference on Web-Age Information Management*, pages 545–556. Springer, 2013.
12. W. Khaouid, M. Barsky, V. Srinivasan, and A. Thomo. K-core decomposition of large networks on a single pc. *Proceedings of the VLDB Endowment*, 9(1):13–23, 2015.
13. M. Khezrzadeh, A. Thomo, and W. W. Wadge. Harnessing the power of favorites lists for recommendation systems. In *Proceedings of the third ACM conference on Recommender systems*, pages 289–292. ACM, 2009.
14. N. Korovaiko and A. Thomo. Trust prediction from user-item ratings. *Social Network Analysis and Mining*, 3(3):749–759, 2013.
15. G. Kortsarz and D. Peleg. Generating sparse 2-spanners. *Journal of Algorithms*, 17(2):222–236, 1994.
16. A. Kyrola, G. Blelloch, and C. Guestrin. Graphchi: Large-scale graph computation on just a pc. In *USENIX OSDI*, pages 31–46, 2012.
17. J. Lu and A. Thomo. An experimental evaluation of giraph and graphchi. In *Advances in Social Networks Analysis and Mining (ASONAM), 2016 IEEE/ACM International Conference on*, pages 993–996. IEEE, 2016.
18. Y. Lu, J. Cheng, D. Yan, and H. Wu. Large-scale distributed graph computing systems: An experimental evaluation. *PVLDB*, 8(3):281–292, 2014.
19. 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 *SIGMOD*. ACM, 2010.
20. C. Martella and R. Shaposhnik. Practical graph analytics with apache giraph. 2015.
21. R. R. McCune, T. Weninger, and G. Madey. Thinking like a vertex: a survey of vertex-centric frameworks for large-scale distributed graph processing. *ACM Computing Surveys (CSUR)*, 48(2):25, 2015.
22. A. Montresor, F. De Pellegrini, and D. Miorandi. Distributed k-core decomposition. *IEEE Transactions on parallel and distributed systems*, 24(2):288–300, 2013.
23. J. Pavlus. Your laptop can now analyze big data. *MIT Technology Review*, July 17, 2014.

24. H. Saito, M. Toyoda, M. Kitsuregawa, and K. Aihara. A large-scale study of link spam detection by graph algorithms. In *Proceedings of the 3rd international workshop on Adversarial information retrieval on the web*, pages 45–48. ACM, 2007.
25. S. B. Seidman. Network structure and minimum degree. *Social networks*, 5(3):269–287, 1983.
26. M. Shoaran and A. Thomo. Fault-tolerant computation of distributed regular path queries. *Theoretical Computer Science*, 410(1):62–77, 2009.
27. Statista. Leading social networks worldwide, ranked by number of active users (in millions), 2017. https://www.statista.com/statistics/272014/global-social-networks-ranked-by-number-of-users/.
28. D. C. Stefanescu, A. Thomo, and L. Thomo. Distributed evaluation of generalized path queries. In *SAC*. ACM, 2005.
29. C. Tsourakakis, F. Bonchi, A. Gionis, F. Gullo, and M. Tsiarli. Denser than the densest subgraph: extracting optimal quasi-cliques with quality guarantees. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 104–112. ACM, 2013.
30. N. Yazdanfar and A. Thomo. Link recommender: Collaborative-filtering for recommending urls to twitter users. *Procedia Computer Science*, 19:412–419, 2013.
31. H. Zhang, H. Zhao, W. Cai, J. Liu, and W. Zhou. Using the k-core decomposition to analyze the static structure of large-scale software systems. *The Journal of Supercomputing*, 53(2):352–369, 2010.