

# A sequential algorithm for fast clique percolation

Jussi M. Kumpula<sup>1,\*</sup>, Mikko Kivelä<sup>1</sup>, Kimmo Kaski<sup>1</sup>, and Jari Saramäki<sup>1</sup>

<sup>1</sup>*Department of Biomedical Engineering and Computational Science,  
Helsinki University of Technology,  
P.O. Box 9203, FIN-02015 HUT, Finland*

(Dated: May 10, 2008)

In complex network research clique percolation, introduced by Palla *et al.*, is a deterministic community detection method, which allows for overlapping communities and is purely based on local topological properties of a network. Here we present a sequential clique percolation algorithm (SCP) to do fast community detection in weighted and unweighted networks, for cliques of a chosen size. This method is based on sequentially inserting the constituent links to the network and simultaneously keeping track of the emerging community structure. Unlike existing algorithms, the SCP method allows for detecting  $k$ -clique communities at multiple weight thresholds in a single run, and can simultaneously produce a dendrogram representation of hierarchical community structure. In sparse weighted networks, the SCP algorithm can also be used for implementing the weighted clique percolation method recently introduced by Farkas *et al.* The computational time of the SCP algorithm scales linearly with the number of  $k$ -cliques in the network. As an example, the method is applied to a product association network, revealing its nested community structure.

PACS numbers: 89.75.Fb,89.75.Hc,89.75.-k

## I. INTRODUCTION

Over the last decade, complex networks have become a standard framework in the study of complex systems [1, 2]. The simplicity of the network representation, where the interactions and interacting elements are mapped to links and nodes, respectively, facilitates its use on a number of systems, ranging from human societies to biological systems. One prominent feature of complex networks is related to their mesoscopic properties. Networks often display *modular* structure, *i.e.*, are structured in terms of modules or communities, which are, in general, sets of densely interconnected nodes. Such communities are often closely related to functional units of the system, for example groups of individuals interacting with each other in society [3, 4, 5, 6], or functional modules in metabolic networks [7, 8, 9].

The problem of detecting communities in complex networks has received a lot of attention during the last years [10]. This problem is twofold: first, there is no unique way to rigorously define what constitutes a community. For any definition, several choices have to be made: whether communities are defined using local or global network properties, whether nodes can participate in several communities, and whether the definition allows for weighted networks and nested hierarchy of communities. Second, any definition is useful in practice only if it can be reformulated as an algorithm which scales well enough to allow processing networks of large enough size. As a result, a large number of community definitions and their algorithmic implementations have been proposed

over the recent years [11, 12, 13, 14, 15, 16].

In this paper we focus on a fast algorithmic implementation of the *clique percolation* (CP) method, originally introduced by Palla *et al.* [9]. The CP method is deterministic and it is based solely on local topological properties, defining a  $k$ -clique community as a set of nodes belonging to adjacent  $k$ -cliques. This allows for overlapping communities, *i.e.*, nodes having multiple community memberships. The CP method has earlier been successfully applied to various community detection problems: detection of protein communities related to cancer metastasis [17], analysis of communities in co-authorship, word-association and protein-interaction networks [9], and time evolution of social groups [6]. In contrary to existing implementations ([www.cfindex.org](http://www.cfindex.org)), which detect  $k$ -clique communities for all values of  $k$  by first finding the maximal cliques by an exponentially scaling algorithm [9], we focus on rapid detection of communities for a chosen value of  $k$ . Our sequential clique percolation (SCP) algorithm is based on sequentially inserting links to the network and keeping track of the emerging community structure. It has specifically been designed for weighted networks containing hierarchical communities which are (mostly) encoded in the link weights. When links are inserted in decreasing order of weight, the algorithm allows for detecting  $k$ -clique communities at chosen threshold levels in a single run and simultaneously produces a dendrogram representation of a possible hierarchical structure. In addition, the algorithm can be used for very fast community detection for unweighted networks.

This paper is structured as follows: first, we present our algorithm for the simplest, unweighted case, and discuss its scaling properties. We then move on to detecting nested communities in weighted networks, applying the algorithm to a product association network gener-

---

\*Electronic address: jkumpula@lce.hut.fi

ated from data on sellers and products on an online auction site. Finally, we discuss a variation of the algorithm which is based on ordering  $k$ -cliques according to their weighted properties, and present our conclusions.

## II. THE SCP ALGORITHM

Let us begin by defining  $k$ -cliques and  $k$ -clique communities [9, 18]:

- A  $k$ -clique is a set of  $k$  nodes which are all connected to each other. A  $k$ -clique community, or  $k$ -community, is a set of nodes which can be reached by a series of overlapping  $k$ -cliques, where overlap means that the  $k$ -cliques share  $k - 1$  nodes.

It should be noted that 2-cliques correspond to pairs of nodes connected by single links and 1-cliques to single nodes. Given a network  $\Gamma$ , the goal is then to find the  $k$ -communities defined as above. In our case, we restrict ourselves to some specific values of  $k$ . If the network in question is not too dense, choosing  $k = 3$  or  $k = 4$  usually yields useful information. Our algorithm is based on detecting and storing  $k$ -communities as they emerge and consolidate when links are sequentially inserted into the network. One can think of the process as first "removing" each link  $l$  from the network  $\Gamma$ , and then inserting them back one by one. For unweighted networks, the links can be inserted in any order, whereas for weighted networks, it may be desirable to sort the links by weight.

Our algorithm for detecting  $k$ -communities consists of two phases: the first phase of the algorithm detects  $k$ -cliques which form when a link is inserted. These are then fed to the second phase, which keeps track of formation and merging of  $k$ -communities by processing the found  $k$ -cliques. The two parts of the algorithm are described in detail below.

### A. Phase I: Detecting the $k$ -cliques

The first part of the algorithm involves detecting  $k$ -cliques which are formed when a link is inserted into the network. Suppose now that the inserted link connects nodes  $v_i$  and  $v_j$  (see Fig. 1). The minimum requirement for a new  $k$ -clique to form is that nodes  $v_i$  and  $v_j$  both have degree of at least  $k - 1$ . If this is the case, the algorithm proceeds by collecting all nodes that are neighbors of both nodes,  $\mathcal{N}_{ij} = \mathcal{N}_i \cap \mathcal{N}_j$ , where  $\mathcal{N}$  denotes neighborhood. Now, when the link  $l_{ij}$  is added, each  $(k - 2)$ -clique contained in the set  $\mathcal{N}_{ij}$  will give rise to a new  $k$ -clique. Therefore, all newly formed  $k$ -cliques are found by detecting all the  $(k - 2)$ -cliques in the  $\mathcal{N}_{ij}$ . For commonly used small clique sizes, this is very fast: for 3-cliques,  $(k - 2)$ -cliques are single nodes, while for  $k = 4$ , all connected pairs of nodes in  $\mathcal{N}_{ij}$  give rise to a new 4-clique.




FIG. 1: Schematic illustration of the process for detecting the  $k$ -cliques a newly inserted link completes. The dashed line depicts the new link, inserted between nodes  $v_i$  and  $v_j$ . The common neighbors of nodes  $v_i$  and  $v_j$  are  $\mathcal{N}_{ij} = \{v_m, v_n, v_p, v_q\}$ . For detecting newly formed 4-cliques, all pairs of nodes in  $\mathcal{N}_{ij}$  are checked to see if they are connected, that is, if they form a 2-clique. Each 2-clique in the set gives rise to a 4-clique, so in total the link  $l_{ij}$  will generate three 4-cliques. In the case  $k = 5$ , only one 3-clique is found, which contains the nodes  $v_m, v_n$  and  $v_p$ . It will give rise to a single 5-clique including these nodes in addition to  $v_i$  and  $v_j$ .

Next the  $k$ -cliques detected as above are fed one by one into the second phase of the algorithm.

### B. Phase II: Detecting the $k$ -communities

The second phase of the algorithm detects and keeps track of  $k$ -communities which form and merge when new  $k$ -cliques are input from the first phase. Because a  $k$ -community is defined as a set of nodes which all can be reached by a series of overlapping  $k$ -cliques, the crucial issue here is the efficient detection of overlap between  $k$ -cliques. A naive approach would be to search for shared sets of  $k - 1$  nodes between the newly input clique and all existing cliques. However, the required computational effort makes this approach unpractical. Instead, we take advantage of the sequential nature of the process by "locally" detecting possible overlap of each new  $k$ -clique with existing  $k$ -communities and by updating the community structure accordingly.

Let us begin by noting that the  $k$ -community structure of a network can be represented by a bipartite network, where the two types of nodes represent  $k$ -cliques and  $(k - 1)$ -cliques. In this network, a link exists between two nodes of different type if the  $k$ -clique contains the  $(k - 1)$ -clique as a sub-clique. This is illustrated in Fig. 2. The usefulness of this representation becomes apparent in the following: each connected component in this bipartite network corresponds to a  $k$ -clique community, because by definition  $k$ -cliques belonging to the same community are connected through shared  $(k - 1)$ -cliques. Furthermore, connected components of the unipartite projections of the bipartite network similarly correspond to  $k$ -clique com-

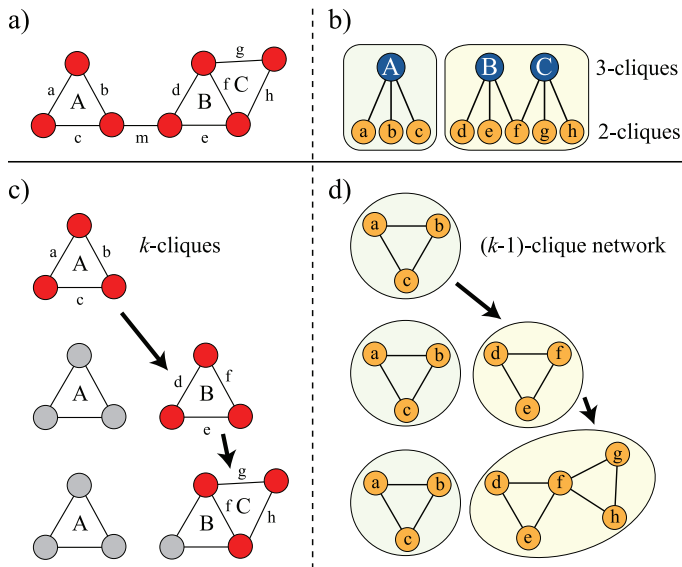


FIG. 2: Illustration of the algorithm for detecting  $k$ -clique communities in a simple example network. Here,  $k = 3$ . a) The original network consists of three 3-cliques labeled  $A$ ,  $B$ , and  $C$ . 2-cliques, *i.e.*, nodes connected by single links, are labeled with lower case letters. b) Bipartite network presentation of the clique structure. Note that in the bipartite network, the 3-cliques  $B$  and  $C$ , which form a 3-clique community, are connected by the shared 2-clique  $f$ . Clique  $A$  forms another 3-clique community. c) 3-cliques detected by the first part of the algorithm as links are sequentially inserted into the network. Each new  $k$ -clique is denoted by red nodes whereas nodes associated with existing  $k$ -cliques appear gray. d) Corresponding updates to the  $(k - 1)$ -clique network as a result of the second part of the algorithm.  $k$ -clique communities correspond to connected components of this network (shaded areas).

munities. In the following, we focus on the  $(k - 1)$ -clique projection of this bipartite network. We denote the network resulting from this projection by  $\Gamma^*$ . In this unipartite network, nodes  $v^*$  represent the  $(k - 1)$ -cliques of  $\Gamma$ , and links  $l^*$  exist between nodes which are sub-cliques of the same  $k$ -clique.

For the sake of clarity, we will first present a "physical" interpretation of Phase II of the algorithm, and then discuss the algorithmic implementation where certain shortcuts can be made. Similarly to Phase I, where the original network  $\Gamma$  is reconstructed link by link, Phase II of the SCP algorithm sequentially builds up  $\Gamma^*$  from the  $k$ -cliques brought forward from Phase I. At the same time, it keeps track of the connected components of  $\Gamma^*$  (see Fig. 2, panels c and d). These correspond to  $k$ -clique communities. When a new  $k$ -clique is input from Phase I, its constituent  $(k - 1)$ -cliques are first extracted; obviously there are always  $k$  of such sub-cliques. Each of these  $(k - 1)$  cliques corresponds to a node in  $\Gamma^*$ . Some of these nodes may already be present, if the corresponding  $(k - 1)$ -cliques have been handled earlier as part of

another  $k$ -clique; if not, they are created at this stage. Finally, links are created between members of this set of  $k$  nodes, and resulting changes in the connected component structure of  $\Gamma^*$  are recorded.

In the algorithmic implementation, things can be done somewhat more efficiently, resembling techniques used in link percolation. The actual network  $\Gamma^*$  does not need to be constructed, as it is enough to keep track of its connected components, *i.e.*, the component indices of its nodes  $v^*$ . This is equal to link percolation in  $\Gamma^*$ , which can be implemented for example with disjoint-set forests [19]. At this stage it is enough to ensure that all  $(k - 1)$ -clique-nodes corresponding to the new  $k$  clique are marked to belong to the same component (the new  $(k - 1)$ -cliques and their links may either form a new connected component, merge with an existing component, or join together at most  $k$  existing components).

The above process is then repeated for each  $k$ -clique input from Phase I. Finally, once all links have been inserted (Phase I) and the subsequently formed  $k$ -cliques handled (Phase II), the  $k$ -communities of the original network  $\Gamma$  can be read from the component indices of  $v^*$ , assigning nodes of  $\Gamma$  to their corresponding communities.

In theory, it would also be possible to keep track of the connected components of the whole bipartite network or alternatively project the bipartite network to  $k$ -cliques instead of  $(k - 1)$ -cliques. Both representations contain the same connected components and would thus yield the same  $k$ -clique communities. However, the former alternative is unnecessarily complicated as it involves nodes of two types. The latter implementation is not as computationally effective as the current choice in cases where a newly inserted  $k$ -clique overlaps with a large number of existing  $k$ -cliques.

### C. Scaling of the algorithm

Let us next discuss the performance of the SCP algorithm, before moving on to its application to weighted network analysis. Obviously, the computational time required to process a network depends on its properties; here, we wish to investigate the performance as a function of network size and the number of  $k$ -cliques contained in the network. To do this, we have applied the SCP algorithm on three types of networks with adjustable sizes. The first test case (GN), introduced by Girvan and Newman [20], contains built-in communities and has often been used for similar purposes. The GN networks used here consist of groups of 32 nodes, where each node has on the average 12 links to nodes of the same group and 4 links to other groups. The network size  $N$  is varied by changing the number of such groups. The second type of networks (WSN) is generated using a recently published model of weighted social networks with communities [21], using parameter values similar to the original reference. As the third type, we have used co-authorship networks based on the *cond-mat* archive (CM), constructed sim-

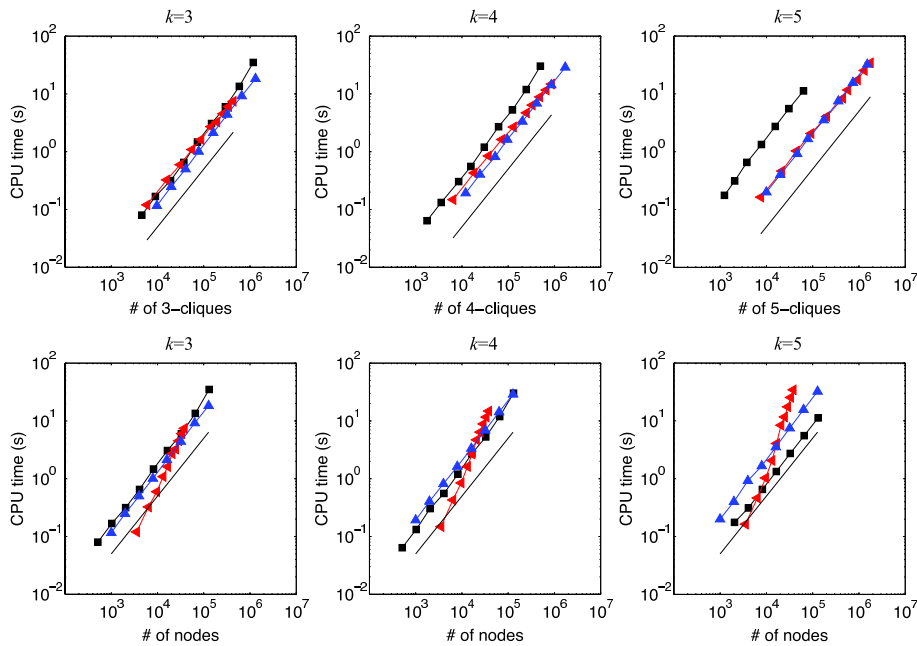


FIG. 3: Computation time of the algorithm for three values of  $k$ , as a function of the number of  $k$ -cliques (upper row) and network size (lower row). Symbols denote different test networks: GN (■), WSN (▲), and CM (◄), see text for details. The solid line is a linear reference.

ilarly to e.g. [22]. However, in order to vary the network size, we have used time windows of varying length, such that two authors are connected if they have published a joint paper during the time window. It should be noted that although the WSN networks are inherently weighted, and the CM networks can also be considered such, here we consider binary versions of both types for the performance analysis.

Results in Fig. 3 show that the computational time grows practically linearly as a function of the number of  $k$ -cliques for all networks. This is as expected, because the algorithm detects and processes the cliques sequentially. This is also reflected in the network size dependence of the required computational time for both types of model networks (GN, WSN). However, for the CM networks, the computational time grows faster than linearly as a function of network size. This is because the CM network is in fact a projection of a bipartite author-publication network, and therefore contains large cliques. These grow in size when the network size increases, and thus the number of  $k$ -cliques grows faster than linearly (see Section III B).

However, when very large cliques are not abundant in the network under investigation, the algorithm is very fast even for networks of large size. For example, detecting 4-clique communities in a mobile phone call network having approximately 4 million nodes and 6 million links [23] takes approximately one minute on a standard desktop computer. The main practical limitations of this al-

gorithm seem to be related to the memory consumption as it requires keeping all  $(k-1)$ -cliques of the network in memory.

### III. SCP FOR WEIGHTED NETWORKS

#### A. Thresholding and nested communities

Let us move on to weighted networks, where the concept of community structure becomes somewhat more complicated. Perhaps only for the very simplest cases, where the networks are sparse, weights can be disregarded, such that communities are associated with the pure topology of the network. However, this is usually not feasible, as weighted networks can be rather dense, even to such an extent that the topology no longer matters, as any modular structure is encoded in the link weights only. This is the case for example in stock interaction networks [24], whose natural representation is a weight matrix with only nonzero elements.

For such networks, one is essentially left with two choices: the first is to threshold the network, such that links whose weights are considered insignificantly small are removed and communities in the resulting sparse network are detected. It is evident that choosing the right threshold is a non-trivial task; in fact, for many cases it may be better to take a multi-resolution approach, by investigating the resulting community structure for

a range of thresholds. Another option is to consider the weights directly when defining what constitutes a community, and to apply a method which is based on this definition [24, 25].

In the original formulation of the clique percolation algorithm, Palla *et al.* suggested a rule for choosing a weight threshold  $w^*$  for the network, such that the resulting  $k$ -community structure would be as diverse as possible [9]. More specifically,  $w^*$  is chosen such that the largest community is twice the size of the second largest one, *i.e.*, below the percolation threshold where a giant  $k$ -clique community appears. For the original implementation, the algorithm had to be run from the beginning for each threshold level. One of the benefits of our approach is that it allows for obtaining the  $k$ -communities at any point of the process of adding links, which is just thresholding done in reverse: If the links of the original network  $\Gamma$  are sorted and processed in descending order of weight, the algorithm yields for each link the  $k$ -community structure of  $\Gamma$  thresholded by the weight of the link. This is very useful for selecting the threshold, as all threshold values can be processed in a single run. Note that for dense networks, sweeping through the entire range of weights is not needed: the algorithm can be stopped before (or immediately after) communities are entirely "smeared out" by a giant community. Stopping the algorithm in time can greatly reduce the workload in dense networks as usually only a small fraction of all links need to be added before the percolating component is found, after which adding more links does not increase the number of nodes in the communities, but only makes the community denser in cliques.

However, by focusing on a single threshold weight, valuable information of the community structure contained in the correlations between weights can be lost. Often, the modular structure of networks is inherently *hierarchical* – denser and stronger communities are nested inside weaker ones, which may further be embedded inside even weaker ones [16, 26, 27, 28]. It is then natural to investigate this nestedness by considering the development of the community structure when the weight threshold is swept through the range of interest. Evidently, this requires book-keeping of the emergence and merger of communities as the threshold is progressively lowered. For the SCP algorithm, this book-keeping is in-built: all necessary information can directly be recorded in Phase II of the algorithm. In particular, it is easy to store when a  $k$ -community appears, which nodes belong to it, how its size grows as new  $k$ -cliques join it, and when it merges with other  $k$ -communities. This information on the nested community structure is best visualized with a *dendrogram*, which is a common presentation format in agglomerative community detection (see, e.g., [28]). In a dendrogram, horizontal lines correspond to communities, and a branching of the lines denotes communities merging. Choosing a single weight threshold would correspond to taking a vertical slice of the dendrogram.

To illustrate the above, Fig. 4 shows two examples of

the nested community structure within a product category network, for  $k = 3$  and  $k = 4$ . This network is constructed from online trading data, downloaded from the Finnish auction website Huuto.net. In this network, nodes correspond to product categories ( $N = 345$ ), and the weights of links connecting two categories to the number of individuals who have been trading in both of them. This network is very dense, the number of links is 52536, corresponding to a link density  $\rho = 0.89$ , and thus the network can be considered as a suitable test case for the evolution of community structure while sweeping the threshold weight. In Fig. 4 the labels associated with each community describe their dominant product categories. Although the dendrograms formed by using  $k = 3$  and  $k = 4$  are not identical, several similar communities appear for both values. From the commonsensical point of view, these appear natural: electronic devices and computer components merge to a single community, as do music and movies, and children's and women's clothing.

In practice we have observed that it is not always possible nor meaningful to include all  $k$ -communities in such a visualization, because it would be dominated by single  $k$ -cliques joining larger  $k$ -communities. Therefore, for practical purposes, it is useful to threshold the dendrogram such that only  $k$ -communities which are larger than a threshold size  $N_{th}$  appear in the plot.

## B. Weighted $k$ -clique percolation

As pointed out above, considering the weights in the definition of what constitutes a community is an alternative to simply discarding low-weight links. Such an extension for clique percolation has recently been introduced by Farkas *et al.* in [25]. In this method, each  $k$ -clique is assigned a "weight", which equals the *intensity* [29] of its edge weights. The intensity is defined as the geometric mean of the link weights in the  $k$ -clique. The community structure is then obtained by choosing an intensity threshold  $I^*$  and taking into account only those  $k$ -cliques whose intensity is above  $I^*$ .

For our SCP algorithm, a simple modification allows for weighted clique percolation according to the above scheme. To achieve this, instead of building the  $k$ -communities simultaneously as the  $k$ -cliques emerge, all links are first inserted to the network and the resulting  $k$ -cliques are stored. Then, the intensity of each of these  $k$ -cliques is calculated, and the cliques are sorted with respect to the intensity. Finally, the sorted  $k$ -cliques are processed one by one by the second part of the algorithm, until the intensity threshold is reached. Multiple thresholding levels are obtained as before, but now with respect to  $k$ -clique intensities, and a dendrogram can be constructed similarly. Note that in addition to intensity, any other measure describing the "weight" of the cliques can be used, *e.g.*, if homogeneous cliques are sought for, one could also take the clique *coherence* [29] into account. Sorting cliques according to their intensities was briefly

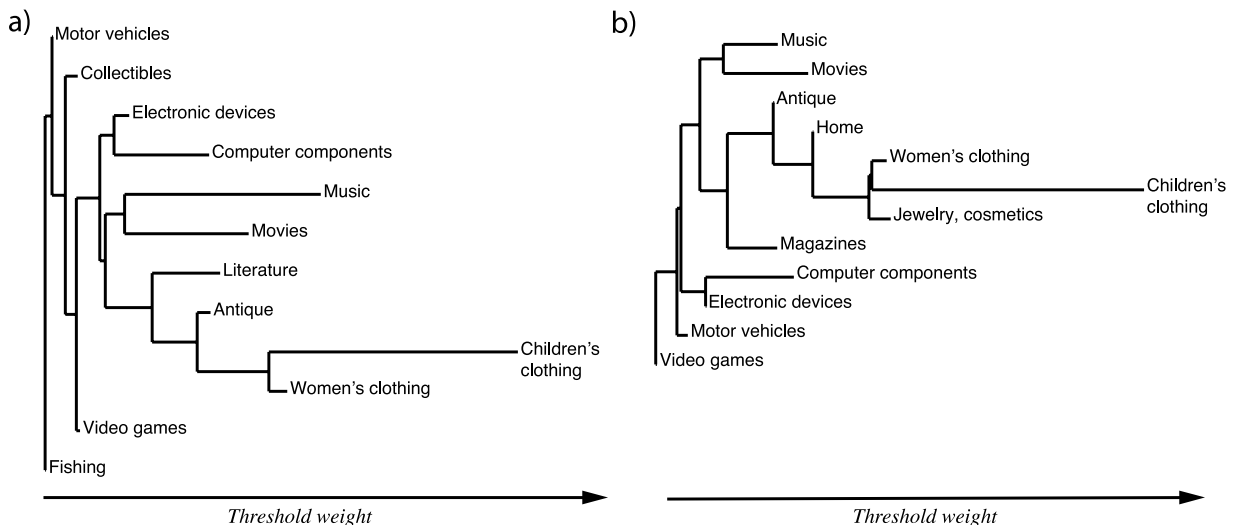


FIG. 4: Dendrogram visualization of the nested  $k$ -community structure of the trading categories of the Finnish online auction site Huuto.net for  $k = 3$  (a) and  $k = 4$  (b).

described by Farkas et al. in [25]; their construction appears somewhat similar to ours as the intensity-sorted cliques are handled in succession, and the method for obtaining overlapping  $k$ -communities seems to correspond to building the whole bipartite network between  $k$ - and  $(k - 1)$ -cliques.

The above procedure requires keeping all  $k$ -cliques in the memory in addition to the  $(k - 1)$ -cliques. In most cases the loss of speed is minimal, as the additional computational load is related to the memory consumption and sorting of cliques, which can be done in log-linear time. However, a possible problem related to the SCP algorithm – and the weighted clique percolation method in general – is that all  $k$ -cliques have to be processed individually, and their number can be very large in dense networks. When the link weight thresholding procedure (Section III A) is applied, this problem can be somewhat circumvented by simply stopping the algorithm as soon as enough links have been inserted for obtaining the community structure at the desired “resolution.” However, for intensity-based clique percolation this cannot be done, as all  $k$ -cliques have to be detected and sorted first, rendering the performance of the method poor for dense weighted networks. As pointed out by Palla *et al.* [9], the problem is that the number of sub-cliques of size  $k$  within a clique of size  $s$  is  $\binom{s}{k}$ . In the limit  $s \gg k$  this leads to

$$\binom{s}{k} \approx \frac{s^k}{k!}. \quad (1)$$

Hence for constant  $k$  and large  $s$ , the number of  $k$ -cliques grows as  $k^{\text{th}}$  power of  $s$ .

#### IV. CONCLUSIONS

We have introduced a sequential clique percolation algorithm for detecting  $k$ -clique communities in a network by sequentially inserting its edges and keeping track of the emerging community structure. This algorithm has specifically been designed for (dense) weighted networks, where weight-based thresholding of either the links or the cliques formed by them is necessary for obtaining meaningful information on the structure. By applying the algorithm on test networks, we have shown that the computational time required to process a network scales linearly with the number of  $k$ -cliques in the network. The sequential nature of the algorithm allows run-time construction of a dendrogram presentation of the nested hierarchical  $k$ -community structure, which we have illustrated using a product category network.

The main tradeoff for our algorithm is that in a single run, it detects  $k$ -communities for a chosen value of  $k$ , instead of starting from the maximal cliques and proceeding towards smaller ones. Hence the SCP algorithm can be considered complementary to earlier presented solutions [9]. Neither of these algorithms can be argued to be strictly better or faster than the other as their performance depends heavily on the network topology and other aspects of the problem they are solving. The SCP algorithm is particularly useful when a small clique size  $k$  is used and when multiple weight threshold levels need to be studied, or no prior knowledge of the proper threshold level of a dense weighted network is at hand. The algorithm can also be considered as a reasonable choice for very large sparse networks as suggested by the short computation times of the community structure of a mobile telephony network having millions of nodes and links.

**Acknowledgements:** We thank J. Hyvönen and

J. Kertész for useful discussions, and acknowledge programming assistance by J. Hyvönen. We acknowledge support by the Academy of Finland, the Finnish Center of Excellence program 2006-2011, project no. 213470.

J.K. is partly supported by the GETA graduate school. J.S. and M.K. acknowledge support by the European Commission NEST Pathfinder initiative on Complexity through project EDEN (Contract 043251).

- 
- [1] G. Caldarelli, *Scale-Free Networks: Complex Webs in Nature and Technology* (Oxford University Press, New York, 2007).
- [2] M. Newman, A.-L. Barabási, and D. Watts, *The Structure and Dynamics of Networks* (Princeton University Press, 2006).
- [3] M. Girvan and M. E. J. Newman, Proc. Natl. Acad. Sci. **99**, 7821 (2002).
- [4] D. Lusseau and M. E. J. Newman, Proc. R. Soc. B **271**, 477 (2004).
- [5] A. Arenas, L. Danon, A. Díaz-Guilera, P. M. Gleiser, and R. Guimerá, Eur. Phys. J. B **38**, 373 (2004).
- [6] G. Palla, A.-L. Barabasi, and T. Vicsek, Nature **446**, 664 (2007).
- [7] P. Holme, M. Huss, and H. Jeong, Bioinformatics **19**, 532 (2003).
- [8] R. Guimerá and L. A. N. Amaral, Nature **433**, 895 (2005).
- [9] G. Palla, I. Derenyi, I. Farkas, and T. Vicsek, Nature **435**, 814 (2005).
- [10] S. Fortunato and C. Castellano, arXiv:0712.2716 (2007).
- [11] M. E. J. Newman and M. Girvan, Phys. Rev. E. **69**, 026113 (2004).
- [12] M. E. J. Newman, Eur. Phys. J. B **38**, 321 (2004).
- [13] F. Radicchi, C. Castellano, F. Cecconi, V. Loreto, and D. Parisi, Proc. Natl. Acad. Sci. **101**, 2658 (2004).
- [14] M. Rosvall and C. T. Bergstrom, Proc. Natl. Acad. Sci. **104**, 7327 (2007).
- [15] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre, arXiv:0803.0476 (2008).
- [16] A. Lancichinetti, S. Fortunato, and J. Kertész, arXiv:0802.1218 (2008).
- [17] P. F. Jonsson, T. Cavanna, D. Zicha, and P. A. Bates, BMC Bioinformatics **7** (2006).
- [18] I. Derenyi, G. Palla, and T. Vicsek, Phys. Rev. Lett. **94**, 160202 (2005).
- [19] R. R. T.H. Cormen, C.E. Leiserson, *Introduction to Algorithms* (The MIT Press, 1990).
- [20] M. Girvan and M. E. J. Newman, Proc. Natl. Acad. Sci. (USA) **99**, 7821 (2002).
- [21] J. M. Kumpula, J. P. Onnela, J. Saramäki, K. Kaski, and J. Kertész, Phys. Rev. Lett. **99**, 228701 (2007).
- [22] M. E. J. Newman, Phys. Rev. E **64**, 016131 (2001).
- [23] J. P. Onnela, J. Saramäki, J. Hyvönen, G. Szabó, D. Lazer, K. Kaski, J. Kertész, and A. L. Barabási, Proc. Natl. Acad. Sci. **104**, 7332 (2007).
- [24] T. Heimo, J. M. Kumpula, K. Kaski, and J. Saramäki, arXiv:0804.3457 (2008).
- [25] I. Farkas, D. Ábel, G. Palla, and T. Vicsek, New Journal of Physics **9**, 180 (2007).
- [26] A. Clauset, C. Moore, and M. E. J. Newman, Lect. Notes Comput. Sc. **4503**, 1 (2007).
- [27] M. Sales-Pardo, R. Guimerá, and L. A. N. Amaral, Proc. Natl. Acad. Sci. (USA) **104**, 15224 (2007).
- [28] A. Clauset, C. Moore, and M. E. J. Newman, Nature **453**, 98 (2008).
- [29] J.-P. Onnela, J. Saramäki, J. Kertész, and K. Kaski, Phys. Rev. E **71**, 065103(R) (2005).