

Detecting the Structure of Social Networks Using (α, β) -Communities^{*,**}

Jing He², John Hopcroft¹, Hongyu Liang²,
Supasorn Suwajanakorn¹, and Liaoruo Wang¹

¹ Department of Computer Science, Cornell University, Ithaca, NY 14853
{jeh17,ss932,lw335}@cornell.edu

² Institute for Theoretical Computer Science, Tsinghua University, Beijing, China
{hejing2929,hongyuliang86}@gmail.com

Abstract. An (α, β) -community is a subset of vertices C with each vertex in C connected to at least β vertices of C (self-loops counted) and each vertex outside of C connected to at most α vertices of C ($\alpha < \beta$) [9]. In this paper, we present a heuristic (α, β) -COMMUNITY algorithm, which in practice successfully finds (α, β) -communities of a given size. The structure of (α, β) -communities in several large-scale social graphs is explored, and a surprising core structure is discovered by taking the intersection of a group of massively overlapping (α, β) -communities. For large community size k , the (α, β) -communities are well clustered into a small number of disjoint cores, and there are no isolated (α, β) -communities scattered between these densely-clustered cores. The (α, β) -communities from the same group have significant overlap among them, and those from distinct groups have extremely small pairwise resemblance. The number of cores decreases as k increases, and there are no bridges of intermediate (α, β) -communities connecting one core to another. The cores obtained for a smaller k either disappear or merge into the cores obtained for a larger k . Further, similar experiments on random graph models demonstrate that the core structure displayed in various social graphs is due to the underlying social structure of these real-world networks, rather than due to high-degree vertices or a particular degree distribution.

1 Introduction

Much of the early work on finding communities in social networks focused on partitioning the corresponding graph into disjoint communities [3, 5, 8, 10–14]. Algorithms often required dense graphs and conductance was taken as the measure of the goodness of a community [4, 7, 8, 15]. To identify well-defined communities in social networks, one needs to realize that an individual may belong to

* Authors are listed alphabetically.

** This research was partially supported by the U.S. Air Force Office of Scientific Research under Grant FA9550-09-1-0675, the National Natural Science Foundation of China under Grant 60553001, and the National Basic Research Program of China under Grant 2007CB807900 and 2007CB807901.

multiple communities at the same time and is likely to have more connections to individuals outside of his/her community than inside. For example, a person in the theoretical computer science community is likely to have many connections to individuals outside of the theoretical computer science community, who may be his/her family members, or enroll in his/her institution, or attend his/her religious group. One approach to finding such overlapping communities is that of Mishra et al. [9], where the concept of an (α, β) -community was introduced and several algorithms were given for finding an (α, β) -community in a dense graph provided there is an advocate for the community. An advocate for a community is an individual who is connected to a large fraction of the members of that community.

In this paper, we discuss the concept of (α, β) -community, and develop a heuristic (α, β) -COMMUNITY algorithm that in practice efficiently finds (α, β) -communities of a given size. Further, we thoroughly explore the structure of (α, β) -communities in several large-scale social networks. Surprisingly, in a Twitter friendship graph with 112,957 vertices and 481,591 edges, there are 6,912 distinct (α, β) -communities of size 200 among the 45,361 (α, β) -communities returned by the algorithm. Moreover, these (α, β) -communities are neatly categorized into a small number of massively overlapping clusters. Specifically, the (α, β) -communities from the same cluster have significant overlap ($> 90\%$) among them, while the (α, β) -communities from distinct clusters have extremely small ($< 5\%$) pairwise resemblance. This leads to the notion of a core which is the intersection of a group of massively overlapping (α, β) -communities, where the core also shares a significant overlap ($> 75\%$) with every member (α, β) -community in that group.

The total number and average size of cores in the Twitter graph as functions of the community size k are given in Table 1. Interestingly, as the size k increases, some cores merge into larger ones while others simply disappear. Moreover, cores may fracture when they merge into larger ones, with a fraction of vertices disappearing from larger cores and reappearing later. Among the interesting questions we explore in this paper are why (α, β) -communities correspond to well-defined clusters and why there is no bridge of (α, β) -communities connecting one cluster to another. A bridge is a sequence of intermediate (α, β) -communities where adjacent subsets of the sequence have substantial overlap but the first and last subsets have little overlap. Other intriguing questions include whether different types of social networks incorporate fundamentally different social structures, and what it is about the structure of social networks that leads to the structure of cores as in the Twitter graph and why some networks do not display this structure as in random graph models.

By taking the intersection of a group of massively overlapping (α, β) -Communities obtained from repeated experiments, we can eliminate the random factors and extract the underlying structure with multiple runs of the (α, β) -COMMUNITY algorithm. In social graphs, for large community size k , the (α, β) -communities are well clustered into a small number of disjoint cores, and there are no isolated (α, β) -communities scattered between these densely-clustered

cores. The number of cores decreases as k increases and becomes relatively small for large k . The cores obtained for a smaller k either disappear or merge into the cores obtained for a larger k . Moreover, the cores correspond to dense regions of the graph, and there are no bridges of intermediate (α, β) -communities connecting one core to another. In contrast, the cores found in several random graph models usually have significant overlap among them, and the number of cores does not necessarily decrease as k increases. Extensive experiments demonstrate that the core structure displayed in various large-scale social graphs is indeed due to the underlying social structure of the networks, rather than due to high-degree vertices or a particular degree distribution.

The rest of this paper is organized as follows. First, we introduce the definition of an (α, β) -community in Section 2 and show their frequent existence. Then, we prove the NP-hardness of finding an (α, β) -community and present the heuristic (α, β) -COMMUNITY algorithm. In Section 3, we apply the algorithm to various large-scale social graphs and random graphs to explore, analyze, and demonstrate the core structure in social networks. We conclude in Section 4 with comments on the problems considered and future work.

2 Preliminaries

The concept of (α, β) -community was proposed by Mishra et al. [9] as a powerful tool for graph clustering and community discovery. In [9], an (α, β) -community refers to a cluster of vertices with each vertex in the cluster adjacent to at least a β -fraction of the vertices in the cluster and each vertex outside of the cluster adjacent to at most an α -fraction of the vertices in the cluster. In this paper, we adopt a slightly different definition. Given a graph $G = (V, E)$ with self-loops added to all vertices, a subset $C \subseteq V$ is called an (α, β) -community when each vertex in C is connected to at least β vertices of C (self-loop counted) and each vertex outside of C is connected to at most α vertices of C ($\alpha < \beta$). Similarly to that of [9], this definition acknowledges the importance of self-loops: although a maximal clique should intuitively be a community, this cannot be guaranteed without self-loops. An (α, β) -community in a graph G is called *proper* if it corresponds to a non-empty proper subgraph of G .

Given a subset $S \subseteq V$, for a vertex $v \notin S$, $\alpha(v)$ is defined as the number of edges between v and vertices of S . Similarly, for a vertex $w \in S$, $\beta(w)$ is defined as the number of edges between w and vertices of S (self-loop counted). Then, $\alpha(S) = \max\{\alpha(v) | v \notin S\}$ and $\beta(S) = \min\{\beta(w) | w \in S\}$.

A maximal clique is guaranteed to be an (α, β) -community since self-loops are counted by the definition. Every graph that is not a clique must contain an (α, β) -community (or, a maximal clique) as a proper subgraph. Starting with any vertex, it is either a proper (α, β) -community (i.e. a maximal clique) or there must be another vertex connected to it (i.e. not a maximal clique). Then, a pair of two vertices connected by an edge is either a proper (α, β) -community (i.e. a maximal clique) or there must be a third vertex connected to both (i.e. not a maximal clique). Continue this argument until a proper (α, β) -community is

found or all vertices are included in a clique, contradicting the assumption that the graph is not a clique. Thus, we have the following theorem:

Theorem 1. *Every graph other than a clique contains a proper (α, β) -community.*

Proof. Omitted due to space limitations (see [6]).

Algorithm 1. (α, β) -COMMUNITY($G = (V, E), k$)

```

1:  $S \leftarrow$  a random subset of  $V$  of  $k$  vertices
2: while  $\beta(S) \leq \alpha(S)$  do
3:    $S \leftarrow$  SWAPPING( $G, S$ )
4:    $A \leftarrow \{v \notin S \mid \alpha(v) = \alpha(S)\}$ 
5:    $B \leftarrow \{v \in S \mid \beta(v) = \beta(S)\}$ 
6:   if  $\{(a_i, b_j) \notin E \mid a_i \in A, b_j \in B\} \neq \emptyset$  then
7:     pick such a pair of vertices  $(a_i, b_j)$ 
8:      $S \leftarrow (S - \{b_j\}) \cup \{a_i\}$ 
9:   else if  $\{a_i \in A \mid (a_i, a_k) \notin E, \forall a_k \in A, k \neq i\} \neq \emptyset$  then
10:    pick such a vertex  $a_i$ 
11:     $S \leftarrow S \cup \{a_i\}$ 
12:   else if  $\{b_j \in B \mid (b_j, b_k) \notin E, \forall b_k \in B, k \neq j\} \neq \emptyset$  then
13:    pick such a vertex  $b_j$ 
14:     $S \leftarrow S - \{b_j\}$ 
15:   else
16:      $S \leftarrow S \cup A$ 
17:   end if
18: end while
19: return  $S$ 

```

Algorithm 2. SWAPPING($G = (V, E), S$)

```

1: while  $\beta(S) < \alpha(S)$  do
2:    $A \leftarrow \{v \notin S \mid \alpha(v) = \alpha(S)\}$ 
3:    $B \leftarrow \{v \in S \mid \beta(v) = \beta(S)\}$ 
4:   pick a vertex  $a \in A$  and a vertex  $b \in B$ 
5:    $S \leftarrow (S - \{b\}) \cup \{a\}$ 
6: end while
7: return  $S$ 

```

Given a graph G with a self-loop added to each vertex and an integer k , define COMMUNITY as the problem of finding an (α, β) -community of size k in G . Given a graph G and an integer k , define CLIQUE as the problem of determining whether there exists a clique of size k in G .

Theorem 2. *The COMMUNITY problem is NP-hard.*

Proof. Omitted due to space limitations (see [6]).

Next, we give a heuristic algorithm for finding an (α, β) -community of size k in a graph $G = (V, E)$. A mathematical description of this (α, β) -COMMUNITY algorithm, along with a subroutine called SWAPPING, is given above. See [6] for more details.

3 Experimental Results

3.1 Social Graphs

Twitter. The Twitter dataset [1, 2] corresponds to a directed friendship graph among a subset of Twitter user accounts. Each vertex represents an individual Twitter user account, and each edge represents a following relation from one user to another. For simplicity, we convert this directed graph into an undirected graph, ignoring the direction of the edges and combining multiple edges with the same pair of endpoints. Further, we iteratively remove from the graph the isolated and degree-one vertices in order to get rid of the insignificant outliers. This effectively reduces the number of vertices and edges, resulting in a smaller graph with 112,957 vertices and 481,591 edges. Then, the average degree of the Twitter graph is 8.52. Finally, self-loops are added to this graph in accordance with the definition of (α, β) -community.

For a given size k , the heuristic (α, β) -COMMUNITY algorithm is applied to the Twitter graph for finding (α, β) -communities starting with a number of (e.g. 500) random subsets of size k . Theoretically, the algorithm is not guaranteed to terminate within a reasonable period of running time, thus we specify an upper bound (e.g. 1,000) on the number of iterations the algorithm can execute. However, from what we have observed in the experiments, the case of not finding any (α, β) -community within 1,000 iterations is extremely rare. In other words, 500 (α, β) -communities are obtained most of the time with 500 runs of the algorithm.




Fig. 1. The overlapping structure

To shed a light on how many (α, β) -communities there are in the Twitter graph, 45,361 runs of the algorithm are performed for $k = 200$ and 6,912 distinct (α, β) -communities are obtained. Surprisingly, many (α, β) -communities are observed to massively overlap with each other and differ only by a few vertices. Moreover, such a great number of (α, β) -communities are all clustered into

Table 1. Cores of the Twitter graph

k	25	50	100	150	200	250	300	350	400	450	500
number of cores	221	94	19	9	4	4	4	3	3	3	3
average core size	23	45	73	112	151	216	276	332	364	402	440

a small number of disjoint groups. Specifically, every pair of (α, β) -communities from the same group shares a resemblance higher than 0.9, while every pair of (α, β) -communities from distinct groups shares a resemblance lower than 0.06. Here, the pairwise resemblance $r(A, B)$ between two sets A and B is defined as:

$$r(A, B) = \frac{|A \cap B|}{|A \cup B|}.$$

The overlapping (α, β) -communities form a “core” structure rather than a “chain” structure, as illustrated in Fig. 1. The intersection of all (α, β) -communities in each group bears an over 75% resemblance with every single (α, β) -community in that group. For $k = 200$, all 6,912 (α, β) -communities found by the 45,361 runs of the algorithm cluster into four “cores”. The “cores” correspond to dense regions of the graph while being exclusively disjoint, and in contrast to what we would have expected, there are no isolated (α, β) -communities scattered between these densely-clustered “cores”.

For a group of pairwise similar (α, β) -communities, we formally define the *core* to be the intersection of those (α, β) -communities. The number of cores can be determined by computing the resemblance matrix of all obtained (α, β) -communities. Intuitively, (α, β) -communities can be categorized according to the resemblance matrix in a way that every pair of (α, β) -communities in the same category is similar to each other, i.e. the pairwise resemblance is large. A pairwise resemblance is considered to be sufficiently large if it is greater than 0.6, while in practice we frequently observe resemblance greater than 0.9. Based on each category, a core is formed by taking the intersection of all member (α, β) -communities. Therefore, the number of cores is equal to that of such intersections, i.e. the number of blocks along the diagonal of the resemblance matrix. The number and average size of cores in the Twitter graph as functions of the community size k are given in Table 1.

The number of cores decreases as the size k increases. This number is relatively small when k becomes large and will eventually decrease to one as k further increases, indicating that (α, β) -communities are well clustered into a small number of cores before gradually merging into one large core. For example, the (α, β) -communities are clustered into 9 cores for $k = 150$ and 4 cores for $k = 200$, where in both cases the cores are disjoint from each other. As the size k increases, the cores obtained for a smaller k either disappear or merge into the cores obtained for a larger k . A layered tree diagram is constructed to illustrate this phenomenon in Fig. 2(a).

Each level of the diagram, indexed by the size k , consists of cores extracted from collections of pairwise similar (α, β) -communities by taking their respective

intersections. For each pair of cores in adjacent levels, a directed edge is added from the lower level to the upper level if the fraction of overlap is significant, i.e. a substantial fraction (e.g. 60%) of vertices in the core of the lower level is contained in the core of the upper level. If the fraction of overlap is smaller than one, a dotted arrow with this fraction labeled is added to represent a partial merge. Otherwise, a solid arrow with the label “1” omitted is added to represent a full merge. As shown in Fig. 2(a), the fraction of overlap is close to one as we move up the levels, that is, a core of some lower level is (almost) entirely merged into a core of the next higher level.

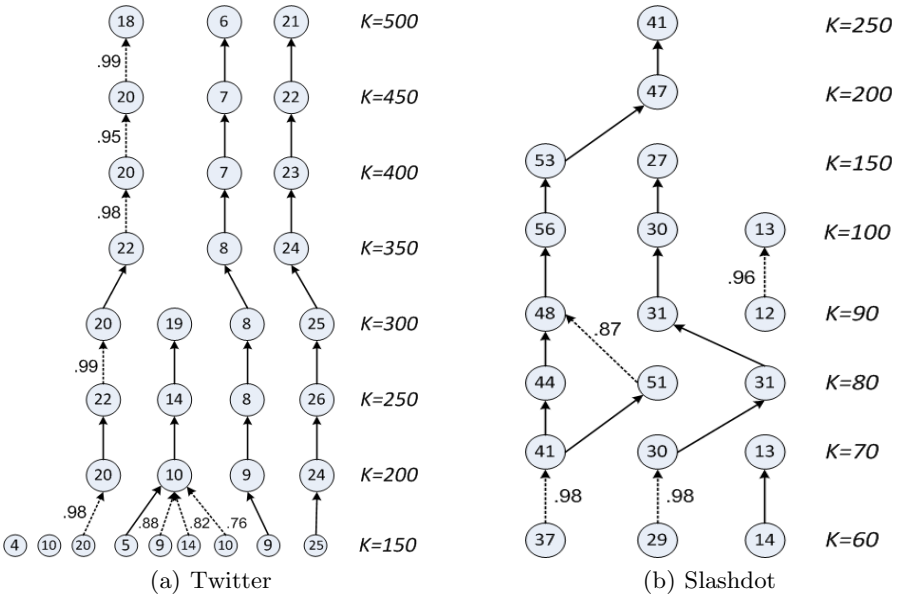


Fig. 2. Tree diagrams indexed by the size k . (Each circle represents a core obtained for a given size, in which the integer denotes the β -value of the core. Each dotted arrow represents a partial merge with the fraction of overlap labeled, and each solid arrow represents a full merge).

The definition of (α, β) -community does not prevent a community from having more edges connecting it to the rest of the graph than those connecting within the community itself. Empirically, there are many more vertices outside of an (α, β) -community, and the edges connecting the community to the rest of the graph are almost always more than those connecting within itself. This definition gives an intuitive criterion as to whether to classify a subset of vertices as a community, i.e. the number of edges connecting each vertex in the community to vertices of the community should be strictly greater than that connecting any vertex outside of the community to vertices of the community. Moreover, by

taking the intersection of a number of massively overlapping (α, β) -communities, the set of (α, β) -communities which differ only by a few vertices is reduced to an underlying core. Thus, each (α, β) -community consists of one of a small number of cores and a few random peripheral vertices, and these peripheral vertices are what gives rise to such a large number of (α, β) -communities.

Before proceeding to our experiments on other social networks, we provide a detailed discussion on the core structure. There might be a generic bias in the (α, β) -COMMUNITY algorithm which is attracted to dense regions of the graph, and thus it is possible that (α, β) -communities located in sparse regions of the graph are never found by the algorithm.

A natural question is what causes the Twitter graph to display this core structure, and further, why the graph shows only a small number of disjoint cores for a large size k . As we will show later, this is due to the fact that a definite social structure, as opposed to randomness, exists in the Twitter network. To take a closer look into this, we simplify the Twitter graph by removing low-degree vertices, i.e. vertices of degree lower than 19, and then obtain a smaller graph with 4,144 vertices and 99,345 edges. The smallest β -value for most (α, β) -communities is given by 19, thus removing vertices of degree lower than 19 will get rid of insignificant low-degree vertices without destroying the fundamental structure of the graph. Again, the (α, β) -COMMUNITY algorithm is applied to this graph with minimum degree 19 for $k = 200, 250, 300, 350, 400$, and exactly two disjoint cores are obtained in each case. Between any two adjacent levels in the corresponding tree diagram, the two cores of the lower level are completely contained in those of the upper level. One possible reason for such a small number of cores could be that the vertices of the cores are more “powerful” in pulling other vertices toward them. If we remove the two cores from the graph and repeat the experiment for $k = 200$, the returned (α, β) -communities are no longer clustered and form a large number of scattered communities.

Another question is why there are exactly two distinct cores in the simplified Twitter graph. For instance, define C_1 and C_2 as the two cores obtained for $k = 200$. C_1 corresponds to a fairly dense subgraph with 156 vertices and 3,029 edges, where the minimum degree is 23 and the average degree is 38.8. C_2 has 159 vertices and 2,577 edges, where the minimum degree is 19 and the average degree is 32.4. Consider the bipartite graph with the two sets of vertices being the vertices of C_1 and the vertices of C_2 . Surprisingly, there are only 105 cross edges between C_1 and C_2 , where 110 (70%) vertices of C_1 and 100 (63%) vertices of C_2 are not associated with any cross edges. Thus, the cores C_1 and C_2 correspond to two subsets of vertices that are densely connected internally but sparsely connected with each other. As a result, they are returned by the algorithm as the cores of two groups of massively overlapping (α, β) -communities.

It is observed that, in addition to the merging of cores, some cores existing for a smaller k simply disappear from the tree diagram as k increases. In other words, few vertices of these disappearing cores are contained in the cores of the next higher level. The cores take on more vertices as the community size k increases, and there may be two cores taking on the same set of vertices. Thus, the

SWAPPING algorithm should run into one of the following two situations: either 1) most vertices of the two cores merge into a new core with some peripheral vertices discarded, or 2) most vertices of one core plus a small fraction of the other form a new core with the latter one disappearing. To verify that one of the above two cases happens, consider the two cores obtained for $k = 150$ that later disappear for $k = 200$, as shown in Fig. 2(a). Let C be one of the two disappearing cores, and recursively perform the following process: enlarge C by adding a random vertex $v \notin C$, run the (α, β) -COMMUNITY algorithm on this enlarged C to find an (α, β) -community of one size larger, and update C to be this obtained (α, β) -community. This process is repeated a number of times until the size of C is increased to 200. Empirically, any obtained (α, β) -community of size 200 contains only a small fraction of vertices of the initial core, while the initial core was completely contained in the (α, β) -communities of size up to about 170. A core may fracture when merging into some other core of a larger size. What happens is that, as vertices are added to one core A , they are also well connected to another core B . As k further increases, these vertices of A will be included in a larger core C that completely contains B , leading to the disappearance of A . If we continue to increase k , the vertices that have disappeared may reappear in a larger core that completely contains C , since they are well connected to the rest of that core.

A *bridge* between two (α, β) -communities or two cores S_1 and S_m is a sequence of intermediate (α, β) -communities S_2, \dots, S_{m-1} , where the pairwise resemblance is large between adjacent subsets but small between the first and last subsets, i.e. $r(S_1, S_m) < 0.3$ and $r(S_i, S_{i+1}) > 0.6$ for all $i \in \{1, 2, \dots, m-1\}$. The *length* of the bridge is thus given by $m-1$. Recall that for $k = 200$, (α, β) -communities are all clustered into four disjoint cores, and there is little overlap between any two (α, β) -communities from distinct cores. It is possible that there exists a bridge in the Twitter graph, but the bias of our algorithm may prevent it from being found. Thus, although no bridge is detected in this experiment, a subsequent question is whether the graph contains any bridge between two cores at all.

Next, the following experiment is designed to determine whether there exists a bridge between two cores. Pick any two cores obtained for $k = 200$ and recursively perform the following steps: randomly choose r vertices from one core and $200 - r$ vertices from the other to form an initial subset of size 200, and apply the (α, β) -COMMUNITY algorithm to this subset. If every iteration returns an (α, β) -community that substantially overlaps with one core but is disjoint from the other, then it implies that there does not exist any bridge between the two cores. During 100 runs of the algorithm, 99 of them return such an (α, β) -community that significantly overlaps with one core but is disjoint from the other. Only one trial returns an (α, β) -community C that contains 95.54% of one core A and 26.22% of the other core B . However, no other intermediate (α, β) -communities can be found between B and C using the above method, which demonstrates the non-existence of a bridge.

Table 2. Cores of the Slashdot graph

k	30	40	50	60	70	80	90	100	150	200	250
number of cores	29	10	3	3	3	3	3	3	2	1	1
average core size	25	33	41	53	62	72	85	97	148	197	244

Another approach to finding a bridge is to search for (α, β) -communities that fall between cores. Generate random subsets of size 200 and run the (α, β) -COMMUNITY algorithm recursively. As we have seen before, four disjoint cores are obtained with 500 runs of the algorithm, and for another 45,361 runs, the (α, β) -community obtained at the end of each iteration is compared with the four cores to check whether it is an intermediate (α, β) -community. This approach is also useful for estimating the total number of (α, β) -communities of a given size. Among the 45,361 runs, no intermediate (α, β) -communities are found, however, only 6,912 distinct (α, β) -communities are returned, which indicates a relatively small number of (α, β) -communities of size 200 and/or a generic bias of our algorithm towards some particular communities over others.

Overall, the above experiments have suggested that there is no bridge between cores, that is, there is not likely to exist a sequence of intermediate (α, β) -communities that connects two cores with substantial overlap between adjacent pairs. The non-existence of such a bridge demonstrates the underlying social structure of the Twitter network with (α, β) -communities neatly categorized into a few densely-clustered disjoint cores.

Slashdot. Slashdot is a technology-related news website known for its professional user community. The website features contemporary technology-oriented news submitted by users and evaluated by editors. In 2002, Slashdot introduced the Slashdot Zoo feature, which allows users to tag each other as friends or foes. The social network based on the common interest shared among Slashdot users was obtained and released by Leskovec et al. [8] in February 2009.

The Slashdot graph contains 82,168 vertices and 504,230 edges, with an average degree of 12.3. Similarly, the (α, β) -COMMUNITY algorithm is applied to this dataset and the statistics are given in Table 2. The number of cores decreases as the community size k increases and becomes relatively small for large k , behaving the same as it did in the Twitter graph. The cores returned by the algorithm are almost disjoint from each other and correspond to dense regions of the graph, with few edges connecting the bipartite graph induced by the vertices of each pair of cores. This indicates that (α, β) -communities are well clustered into a small number of cores for large k , which correspond to dense regions of the graph and share little overlap among them. For example, the (α, β) -communities are clustered into three nearly disjoint cores for $k = 100$, where only 171 edges connect the two cores of size 93 and 100 that have 2,142 and 1,105 internal edges, respectively. Before eventually merging into one large core as k further increases, these densely-clustered cores emerge as the underlying social structure displayed

by the Slashdot network. It is also observed that, as k increases, the cores obtained for a smaller k either disappear or merge into the cores obtained for a larger k . A layered tree diagram is constructed to illustrate this phenomenon in the Slashdot graph, as shown in Fig. 2(b).

arXiv hep-ph Coauthor and Citation. Similar patterns of the core structure can also be observed in the arXiv hep-ph Coauthor and Citation graphs. See [6] for detailed experimental results.

3.2 Random Graphs

To demonstrate that the structure we have found in social graphs is not merely a random artifact, a similar set of experiments is carried out for random graphs. Random graph models do not produce clusters as social graphs do. The cores obtained by the (α, β) -COMMUNITY algorithm usually have significant overlap among them, and correspond to dense regions due to the way the graph was generated. This contrast between social graphs and random graphs again verifies the existence of community structure in various large-scale social networks. See [6] for more details.

4 Conclusion

In social networks, the (α, β) -communities returned by the (α, β) -COMMUNITY algorithm for a given size k are well clustered into a small number of disjoint cores, each of which is the intersection of a group of massively overlapping (α, β) -communities. Two (α, β) -communities from the same group share a significant overlap and differ by only a few vertices, while the pairwise resemblance of two (α, β) -communities from different groups is extremely small. The number of cores decreases as k increases and becomes relatively small for large k . The cores obtained for a smaller k either disappear or merge into the cores obtained for a larger k . Further, the cores correspond to dense regions of the graph. There are no isolated (α, β) -communities scattered between these densely-clustered cores, nor bridges of (α, β) -communities connecting one core to another. Various large-scale social graphs have been explored thoroughly, all of which display the core structure rather than the chain structure.

By constructing random graphs with a power-law degree distribution or the same degree distribution as the social graphs, it is demonstrated that neither high-degree vertices nor a particular degree distribution can result in the core structure displayed in large-scale social graphs. The cores found by the (α, β) -COMMUNITY algorithm in random graphs usually have significant overlap among them and are increasingly scattered across the graph as the size k increases, which implies the non-existence of well-defined clusters in random graphs and verifies the existence of underlying structure in various social networks.

References

1. Choudhury, M.D., Lin, Y.-R., Sundaram, H., Candan, K., Xie, L., Kelliher, A.: How does the sampling strategy impact the discovery of information diffusion in social media? In: Proc. 4th Int'l AAAI Conf. Weblogs and Social Media, ICWSM (2010)
2. Choudhury, M.D., Sundaram, H., John, A., Seligmann, D.D., Kelliher, A.: Birds of a feather: does attribute homophily impact information diffusion on social media? (under review)
3. Clauset, A., Newman, M.E.J., Moore, C.: Finding community structure in very large networks. *Phys. Rev. E* 70, 06111 (2004)
4. Gaertler, M.: Clustering. In: Brandes, U., Erlebach, T. (eds.) *Network Analysis*. LNCS, vol. 3418, pp. 178–215. Springer, Heidelberg (2005)
5. Girvan, M., Newman, M.E.J.: Community structure in social and biological networks. *Proc. Natl. Acad. Sci. USA* 99(12), 7821–7826 (2002)
6. He, J., Hopcroft, J.E., Liang, H., Supasorn, S., Wang, L.: Detecting the structure of social networks using (α, β) -communities. Tech. rep., Cornell University (2011), <http://hdl.handle.net/1813/22415>
7. Lang, K., Rao, S.: A flow-based method for improving the expansion or conductance of graph cuts. In: Bienstock, D., Nemhauser, G.L. (eds.) *IPCO 2004*. LNCS, vol. 3064, pp. 325–337. Springer, Heidelberg (2004)
8. Leskovec, J., Lang, K., Dasgupta, A., Mahoney, M.: Statistical properties of community structure in large social and information networks. In: Proc. 18th Int'l World Wide Web Conf. WWW (2008)
9. Mishra, N., Schreiber, R., Stanton, I., Tarjan, R.E.: Finding strongly-knit clusters in social networks. *Internet Mathematics* 5(1-2), 155–174 (2009)
10. Newman, M.E.J.: Detecting community structure in networks. *The European Physical J. B* 38, 321–330 (2004)
11. Newman, M.E.J.: Fast algorithm for detecting community structure in networks. *Phys. Rev. E* 69, 066133 (2004)
12. Newman, M.E.J.: Finding community structure in networks using the eigenvectors of matrices. *Phys. Rev. E* 74, 036104 (2006)
13. Newman, M.E.J.: Modularity and community structure in networks. *Proc. Natl. Acad. Sci. USA* 103(23), 8577–8582 (2006)
14. Newman, M.E.J., Girvan, M.: Finding and evaluating community structure in networks. *Phys. Rev. E* 69, 026113 (2004)
15. Schaeffer, S.E.: Graph clustering. *Computer Science Review* 1(1), 27–64 (2007)