

Investigating the properties of a social bookmarking and tagging network

Ralitsa Angelova¹, Marek Lipczak², Evangelos Milios², Paweł Prałat³

¹Max Planck Institut für Informatik, Saarbrücken, Germany

²Faculty of Computer Science, Dalhousie University
Halifax, Canada

³Department of Mathematics and Statistics, Dalhousie University,
Halifax, Canada *†‡

October 2, 2008

Abstract

Social networks and collaborative tagging systems are rapidly gaining popularity as a primary means for storing and sharing data among friends, family, colleagues, or perfect strangers as long as they have common interests. `del.icio.us`¹ is a social network where people store and share their personal bookmarks. Most importantly, users tag their bookmarks for ease of information dissemination and later look up. However, it is the friendship links, that make `del.icio.us` a social network. They exist independently of the set of bookmarks that belong to the users and have no relation to the tags typically assigned to the bookmarks. To study the interaction among users, the strength of the existing links and their hidden meaning, we introduce implicit links in the network. These links connect only highly “similar” users. Here, similarity can reflect different aspects of the user’s profile that makes her similar to any other user, such as number of shared bookmarks, or similarity of their tags clouds. We investigate the question whether friends have common interests, we gain additional insights on the strategies that users use to assign tags to their bookmarks, and we demonstrate that the graphs formed by implicit links have unique properties differing from binomial random graphs or random graphs with an expected power-law degree distribution.

1 INTRODUCTION

Social bookmarking and collaborative tagging services lead to the formation of a new type of organically grown network structure. In such networks, users are linked to other users through social connections (e.g. directed friendship links) and to network specific online resources (e.g. bookmarks, photos, books, etc.) by either explicitly linking to them, tagging them with appropriate terms, or commenting on them. Clustering users in this context is a challenging problem, as it involves accounting for multiple types of social linkage among users and diversity of content ranging from personal photos (`flickr.com`) or bookmarks (`del.icio.us`) to whole libraries of read books throughout the user’s lifetime (The Personal Library, `librarything.com`). The complexity of the clustering problem raises

* Authors listed in alphabetical order

† This article is an expanded version of (Angelova, Lipczak, Milios, and Prałat 2008)

‡ Submitted to the International Journal of Data Warehousing and Mining (<http://www.igi-pub.com/journals/details.asp?id=4291>) on 2008/10/01.

¹<http://del.icio.us>

dramatically if we look at the overall electronic fingerprint of these users after connecting all their profiles from the various social networks they actively contribute to (Moser, Ge, and Ester 2007). Not only is clustering itself challenging but evaluation of the clustering solution is also very hard as reference class assignments are typically missing or very expensive to manually gather. These class assignments (also known as ground truth) are ignored in the clustering process. They are used exclusively in the evaluation phase to compare the groups produced by the clustering technique to the known classes it comprises. Modelling social bookmarking and tagging services is a way to generate synthetic data sets that mimic the behaviour of such social networks. Moreover, the synthetic data generative model also provides the corresponding ground truth for performance evaluation and comparison. A requirement for the design of useful models is an in-depth understanding of the properties of real-life data sets obtained from on-line social networks such as `del.icio.us`.

A collaborative tagging system like `del.icio.us` can be visualized as a tripartite structure (Halpin, Robu, and Shepherd 2007), where links (edges) are established between users, tags and bookmarks. Additionally, the social dimension introduces “friendship” links between users. Several research questions about the structure of the social network and its implications arise:

- What is the role of friendship in relation to interest sharing as reflected in the bookmarks and tags of users. Do friends appear to have more common interests than non-friends?
- Do “highly social” users share more topics of interest with others than the “less social” users?
- What are the structural properties of the friendship graph and the graphs induced by the implicit similarity-based links among users? Is their degree distribution indicative of power-law graphs? What are their connectivity and local density properties, measured by the clustering coefficient, as a function of k in their k -core analysis (Healy, Janssen, Milios, and Aiello 2008)? How do they compare with binomial random graphs (Janson, Łuczak, and Ruciński 2000) and random graphs with an expected power-law degree distribution (Chung and Lu 2004)?
- What are the common properties of the friendship, bookmark-based and tag-based links? In particular, how do the three types of links correlate for individual users?

We are not the first to analyze social collaboration on the Web. Evolution models of two online social networks - Flickr and Yahoo!360 are examined in (Kumar, Novak, and Tomkins 2006). In the experiments performed in (Ne-goescu 2007) on the photo sharing network Flickr, after taking a random subset of photos and their owners or users, it is demonstrated that Flickr exhibits the characteristics of small-world and scale-free networks described earlier by (Barabasi 2002; Cohen, Avraham, and Havlin 2003). Search and ranking techniques applied to social networks are discussed in (Hotho, Jaeschke, Schmitz, and Stumme 2006).

A detailed analysis of three other online social networks is presented in (Ahn, Han, Kwak, Moon, and Jeong 2007). Tagging distributions in `del.icio.us` are shown to stabilize into power law distribution with a limited number of stable tags and a much larger “long-tail” of more idiosyncratic tags. Similar results are noted in (Golder and Huberman 2005). Both results give strong evidence that collaborative tagging systems (like `del.icio.us`) can be exploited for reliable automatic creation of taxonomies. More recently, a study of a tag co-occurrence network was carried out (Schmitz, Grahl, Hotho, Stumme, Cattuto, Baldassarri, Loreto, and Servedio 2007). The nodes of this network are tags, and tags are linked when the two tags occur together in the set of tags assigned to a specific bookmark by a user (the post). A weight is given to a link that depends on the number of posts in which the two tags co-occur. The tag co-occurrence network was shown to reveal spamming behaviour.

2 Definitions of graphs and associated metrics

To answer the questions above, we study various graphs on the data provided by the social bookmarking network `del.icio.us`. For all graphs, vertices correspond to users. The edges are defined differently for each graph as

follows.

The friendship graph. Edges correspond to directed friendship links between users. In this paper, as discussed in Sec. 3, we ignore the direction of friendship links and obtain an undirected friendship graph. Bidirectional edges (i.e. representing mutual friendship) are included only once.

Common entity graphs. In this type of implicit graph two vertices are connected by an undirected weighted edge, if the corresponding users have sufficient entities (tags or bookmarks) in common. The edge weight reflects the number of entities that the connected users have in common. A drawback of the common entity method is that it does not take into consideration the number of terms in both sets - having three common terms has the same meaning, no matter if the total number of user’s terms is ten, or fifty. This similarity metric is symmetric, and its range is from 0 to the maximum number of entities. The two subtypes of graph under the common entity type are the following:

- **The common bookmark graph**, where the set of entities of a user is the set of bookmarks of that user.
- **The common tag graph**, where the set of entities of a user is the set of tags the user assigned to all her bookmarks.

Similarity graphs. In this type of implicit graph, two vertices are connected by an undirected weighted edge. The edge weight reflects the cosine similarity between the entity vectors of the connected users. In these similarity graphs, entities can be either bookmarks or tags. The user entity vector belongs to a vector space defined on the entire data set, but captures only the individual entities belonging to the specific user. The definition of a user entity vector is analogous to that of a document vector in information retrieval, where users are the equivalent of documents and entities (bookmarks or tags) the equivalent of terms.

Based on the discussion in Section 3, neither bookmarks nor tags contain the equivalent of stop words. Therefore, we include the most frequent bookmarks and tags in computing user similarity metrics. To reduce noise, we remove low frequency entities while defining the dimensions in our vector space model. That is, bookmarks used by less than 2 users and tags used by less than 3 users are removed.

The two subtypes of similarity graph that we study are derived as follows:

- **The bookmark similarity graph** is derived by considering bookmarks as entities. The weight of a bookmark in any user vector is binary.
- **The tag similarity graph** The tag similarity graph is derived by considering tags as entities. The weight of a tag in any user vector is the tf-idf score for this tag. The term frequency factor (tf) is the number of times a user used a given tag. The inverted document frequency factor (idf) is the inverse of the natural logarithm of the number of users that used a given tag plus one.

The cosine similarity between entity vectors, Eq. 1 is used to define the distance between users, where U_i is the vector of user i , and $w_{i,j}$ is its j -th coordinate. The cosine distance has the advantage that it is normalized with respect to the length of vectors. Cosine similarity is symmetric, and its range is from 0 to 1.

$$cos_sim(U_1, U_2) = \frac{\sum_{j=1}^n w_{1,j} * w_{2,j}}{\sqrt{\left(\sum_{j=1}^n w_{1,j}^2\right) * \left(\sum_{j=1}^n w_{2,j}^2\right)}} \quad (1)$$

Unweighted graph versions. Many interesting graph properties are defined for unweighted graphs, such as clustering coefficient and diameter. To measure these properties, we transform the weighted common entity and similarity

	0.1%	0.05%	0.02%	0.0125%	0.01%	0.005%
k-common bookmarks	99994	47929	18602	11668	9206	4608
bookmarks cosine similarity	92179	45991	18466	11661	9158	4663
k-common tags	91992	45699	18360	11505	9199	4586
tags cosine similarity	91427	45726	18274	11424	9142	4567

Table 1: Number of edges in thresholded graph, for each threshold value.

	0.1%	0.05%	0.02%	0.0125%	0.01%	0.005%
k-common bookmarks	22	30	44	53	58	75
bookmarks cosine similarity	0.08	0.10	0.14	0.16	0.17	0.21
k-common tags	471	551	668	736	769	880
tags cosine similarity	0.55	0.60	0.65	0.68	0.69	0.72

Table 2: The set of threshold values that produces binary graph with a specific percentage of strongest connections.

graphs into unweighted versions. A set of thresholds is selected for each graph such that, after removing connections with weight lower than the threshold, a specified fraction of edges remains. The number of edges for each thresholded graph is presented in Table 1. Table 2 presents real threshold values that select the expected fraction of strongest connection edges. This approach allows us to compare different graphs at similar density levels.

Clustering coefficient. A quantity of interest in measuring the local density properties of the various graphs is the *clustering coefficient* (Ahn, Han, Kwak, Moon, and Jeong 2007). Given an undirected graph $G = (V, E)$, the clustering coefficient of vertex $i \in V$ is defined as:

$$C_i = |\{(v, w) | (i, v), (i, w), (v, w) \in E\}| / \binom{k_i}{2}$$

where k_i is the degree of vertex i . The clustering coefficient is not defined for vertices of degree at most 1. The clustering coefficient of a graph G is the average over all vertices (of degree at least 2) in the graph, that is, $C(G) = \sum_{i \in V, \deg(v) \geq 2} C_i / |\{v \in V : \deg(v) \geq 2\}|$. Its values always lie between 0 and 1. The clustering coefficient is asymptotically almost surely $(1 + o(1))|E| / \binom{|V|}{2}$ for a binomial random graph. For real-world networks it is usually much larger.

K-cores. A k -core of a graph is a maximal induced subgraph of minimum degree at least k . If no subgraph has this property, then we say that the k -core is empty. It is possible to show (Batagelj and Zaveršnik 2002) that the degree core is unique for a given graph and a given k , and can be obtained by recursively removing all vertices with degree less than k . The k -cores of a graph can consist of multiple components. The difference between the k -core and simply filtering out all vertices with degree less than k is best illustrated by comparing their effects on a simple tree. In the case of a tree, the filtering of all degree-one vertices results in the pruning of all of a tree’s leaves, whereas the degree core with $k = 2$ would prune back the leaves of a tree at each recursion, thus destroying the tree completely. Cores were first introduced in studying social networks by Seidman (1983) and popularized by Wasserman and Faust (1994). Batagelj and Zaveršnik (2002) generalize Seidman’s work beyond simple degree to include any monotone function p . The same authors go on to define the core number of a vertex to be the highest order of a core that contains this vertex.

In this paper, we perform degree-based k -core analysis of the friendship, common entity and similarity graphs, by repeatedly increasing k by one until the k -core is reduced to empty. We then plot various properties of the k -core sequence of graphs, including the diameter, size, average distance and average clustering coefficient of the vertices between vertices of the largest component, and the number of components.

Note that the k -core analysis can become prohibitively expensive as the input graphs get dense. We control density by the use of thresholds on similarity measures.

3 Exploring the data set

We perform our experiments on a data set obtained by an automated Web crawl on the social community platform `del.icio.us`. The subset of `del.icio.us` we have at hand consists of 13,514 users, 4,574,587 bookmarks, and 47,807 friendship connections, 6,876 of which are mutual (bidirectional) connections. Most users (13,238) have at least one bookmark. Most users (13,439) also point to at least one friend. The total number of tags used in our subset is 643,889.

The crawl on `del.icio.us` is a Breadth First Search (BFS) of the friendship graph, starting with the user with the highest number of friends as a seed node. This resembles the so called *snowball sampling* technique which is argued to be the only feasible sampling method for crawling such networks (Ahn, Han, Kwak, Moon, and Jeong 2007).

The friendship graph. It is a straightforward representation of the directional connections between users. We observe that this graph is quite sparse, with 40,931 uni-directional edges, and 6,876 bi-directional edges corresponding to mutual friendships. Only 14% of friendships are mutual, due to the fact that by design friendships do not have to be confirmed in in the online system supporting `del.icio.us`.

Bookmark distributions and properties. Intuitively two users are connected when they point to similar sets of bookmarks. We investigate the hypothesis that bookmarks can be treated as terms in the standard Information Retrieval sense, where each user is viewed as a document, and the user’s bookmarks are viewed as terms. The proportion of bookmark urls used only once over the total number of unique bookmark urls is 78%, where the proportion of unique words in Wikipedia used only once over the total number of unique words is 52% (second and fifth column of Table 4).

We further examine whether the distribution of bookmark urls follows Zipf’s law. The log-log plot of the url frequency against rank of urls sorted by decreasing order of frequency, has the following characteristics (a) a slowly declining part for the top 1000 urls, (b) a fairly straight part up to rank of about 10^6 , and (c) a horizontal part at frequency equal to 1 up to rank of about 0.25×10^7 corresponding to a large number of bookmarks appearing only once (Fig. 1). Inspecting the frequencies of the ten most frequently bookmarked urls shows that their frequency differences are small, which confirms the shape of Fig. 1. Surprisingly, bookmarks to general use sites (e.g., `google.com`) are not among the ten most frequent bookmarks. We would expect such sites to be analogous to stop words in standard text repositories. One possible explanation is that users avoid adding these stop-word-like urls to their bookmarks to keep the bookmark list smaller. The addresses of general purpose websites are easy to remember and do not have to be stored in `del.icio.us`. (Table 3). If we prune urls keeping only their domain, the frequency-against-rank plot shows that the domains of bookmarks do match Zipf’s law (Fig. 2). The proportion of bookmark url domains used only once over the total number of unique bookmark url domains is 56%, a very close match to the proportion of unique words in Wikipedia used only once over the total number of unique words, which is 52% (third column of Table 4). Furthermore, the list of ten most frequently bookmarked domain addresses contains indeed the most popular websites. Despite a better fit to Zipf’s law, we do not think that the list of pruned bookmarks is useful for finding real connections between users, as domains are usually too general to capture users’ interests.

Tag distributions and properties. Shared tags between two users may be interpreted as showing their overlapping interests. Although tags are mostly nouns and adjectives, and they do not form grammatically correct sentences, we conjecture that they match the characteristics of terms in text repositories. The proportion of tags used only once with respect to the total number of unique tags is 53%, matching very closely the proportion of unique words in Wikipedia used only once with respect to the total number of unique words of 52% (fourth column of Table 4). Similarly to

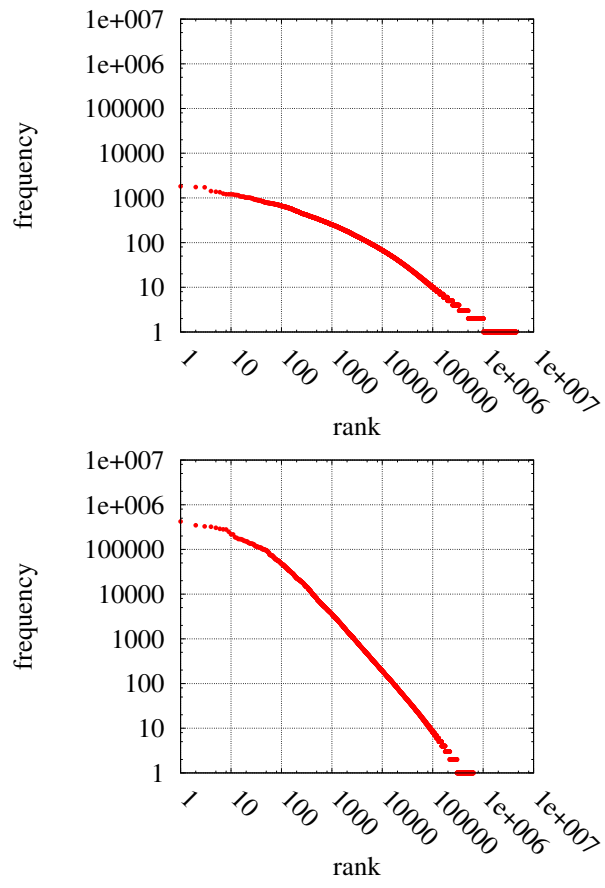


Figure 1: Frequency distribution of bookmark urls (top), and tags (bottom).

Rank	Frequency	Real frequency	Url
1	1815	27359	script.aculo.us/
2	1745	10157	pchere.blogspot.com/2005/02/ absolutely-delicious-complete- tool.html
3	1725	37346	www.pandora.com/
4	1421	29433	www.netvibes.com/
5	1371	21592	www.alvit.de/handbook/
6	1342	24830	www.last.fm/
7	1247	17335	typetester.maratz.com/
8	1207	19069	www.instructables.com/
9	1206	40314	www.flickr.com/
10	1202	6542	johnvey.com/features/ deliciousdirector/

Table 3: List of the ten most frequently bookmarked websites. None of the websites is present in the dataset more than 2000 times, and they also look very specific. Comparing the results to the real frequencies in the complete `del.icio.us` dataset, we see that our crawl is only a very rough approximation of the latter, judging by the relative frequencies. Real frequency data was gathered using `www.xinureturns.com` 28-Jan-08.




Figure 2: Frequency distribution of domain portion of bookmarked urls. The domains reveal the importance of general purpose websites. Users' bookmarks link to some specific information contained in them. The frequency against rank plot shows that the domains of bookmarks match Zipf's law.

Counts (in K)	urls	url do- mains	Tags	Wikipedia words
Total number of unique terms	4,575	1,106	644	4,098
Terms used more than once	1,017	483	303	1,978
Terms used only once	3,558	623	341	2,120

Table 4: Counts (in thousands) of bookmark urls and tags, treating bookmarks and tags as “terms”, the user as the “document”, and the set of all users as the corpus. The same counts are shown for the Wikipedia text corpus. We see that the proportion of terms used only once is much higher in the bookmark url corpus than either the tag or the Wikipedia corpus.

bookmarks, tag frequency is not distributed according to Zipf’s law for the highest ranks, indicating that there are no stop-word-like tags (Fig. 1). The most frequently used tags correspond to general categories, therefore they still convey content information. Interestingly, Sigurbjörnsson and Zwol (2008) show that tags in the popular photography sharing social network `flickr.com` follow a power law distribution much closer than what the observed one in the bookmarking social network `del.icio.us`.

4 Relating friendship with bookmark and tag similarity

We now study the question whether friends have more similar bookmarks or tags than non-friends. We calculate the average connection strengths (according to the previously defined user similarity metrics) over pairs of friends and pairs of non-friends, shown in Table 5. We observe that friends have significantly stronger connections than non-friends based on the bookmark similarity metrics, whereas this is not true for the tag based similarity metrics. We conjecture that the reason for this is that the majority of the tags are individualized, dependent on the individual user’s way of organizing their bookmarks and their background in the corresponding areas. The number of tags that capture the generic meaning of a bookmarked page is a small fraction of the total number of tags associated with that page. In other words, friends who share bookmarks choose mostly different tags for them. To further explore this conjecture, we inspect the Zipf distribution of tags for individual users, Fig. 3, and bookmarked web sites, Fig. 4. In the plots for individual users, we notice the most frequently used tags, which are likely to be the ones recommended by `del.icio.us`, followed by the tail of infrequently used tags, which are likely to be individualized tags, specific to the user’s organization of her bookmarks. A significant fraction of the tags of a user is used only once. In the plots for individual bookmarked web sites, we also notice a small number of most frequently used tags, followed by a large number of user-specific tags, used only once.

This can be easily explained with the difference of user backgrounds and their depth. People, who are more knowledgeable in a particular area are highly likely to use very specific tags while describing a bookmark. On the other hand, if the same URL is bookmarked by a less knowledgeable person, he or she will highly likely associate very generic tags to ease search for the same URL later. Take as an example the URL `http://www.birds.cornell.edu/AllAboutBirds/BirdGuide/Red-winged_Blackbird.html`. Suppose a person, whose hobby is birdwatching, visits this URL and wants to store it in his bookmarks collection. He would highly likely use a tag ‘red-wing blackbird’ for describing the page. Contrary to that, if a person who is not knowledgeable in birds visits that URL and wants to save it, he would highly likely just tag it as ‘wild birds’, which is a very generic description of the page. Therefore, though both users bookmarked the same page, their tags differ significantly.

	Friend pair average	Non-friend pair average
k-common bookmarks	1.931	0.372
bookmark cosine sim	0.011	0.004
k-common tags	54.157	41.816
tag cosine sim	0.081	0.085

Table 5: Average strength of the connection between friends (column 2) and non-friends (column 3). We observe that friends have significantly stronger connections than non-friends based on the bookmark similarity metrics, whereas this is not true for the tag-based similarity metrics.

Finally, we investigate the correlation between the similarity metrics of two users based on tags and bookmarks, using a scatter plot constructed as follows. An edge is sampled randomly from a bookmark similarity graph. The bookmark similarity of the edge is the x-coordinate of the point corresponding to the edge in the scatter plot. For the same edge, the tag similarity is computed, serving as the y-coordinate of the point in the scatter plot. The scatter plot shows practically no correlation between bookmark similarity and tag similarity between two users (Fig. 5).

5 Density properties of the friendship graph and the content-based social graphs

We now apply the concept of k-core introduced in Sec. 2 to discover the density properties of the friendship graph and our content-based social graphs (common bookmark/tag, and bookmark/tag similarity graphs), and compare them with the properties of binomial random graphs (Janson, Łuczak, and Ruciński 2000) and random graphs with power-law degree distribution (Chung and Lu 2004). We focus on plots of properties of the graphs as a function of increasing k in the k-core computation. For fairly small k the smaller components disappear and we are left with a single component, which becomes progressively denser with increasing k, and eventually disappears. The properties of interest are: the number of components, and, for the largest component, its size (number of vertices), its average distance between vertex pairs, and its diameter. We furthermore generate scatter plots of the clustering coefficients versus degree for the graphs' vertices, and the average clustering coefficient of the vertices of the largest component.

Basic k-core properties of the content-based social graphs The average distance between pairs of points of the largest component in the k-core analysis of common bookmark graphs is shown in Fig. 6. The linear nature of the curves is also worth noting.

The diameter of the largest component in the k-core analysis of common bookmark graphs is lower than 7.

The number of components in the k-core analysis of common bookmark graphs drops very quickly to one, as k increases in the k-core analysis. Having a single component is similar to binomial random graph behaviour.

It is worth noting that this behaviour is unlike the behaviours observed in (Healy, Janssen, Milios, and Aiello 2008) for the majority of the social graphs examined there, including the .gov web graph, patent citation graph and the DBLP co-authorship graph. It is only similar to that of an organizational email communication graph.

The size of the largest component (number of vertices) in the k-core analysis of common bookmark graphs is shown in Fig. 6. The size of the last non-trivial k-core is very close to k, and the average distance between nodes is close to 1, so these vertices form an induced subgraph that is close to a clique. This behaviour is very different from binomial random graphs, as shown in Fig. 8. It is also different from power-law graphs, because the last non-trivial k-core is not as dense. The average distance between nodes and graph size as a function of k for the friendship graph are shown in Fig. 7. We note that the friendship graph behaves very similarly to the power law graph.





Figure 3: Frequency vs. rank distribution of tags for two individual users. We notice the most frequently used tags, which are likely the ones recommended by `del.icio.us`, and the tail of infrequently used tags, which are likely the individualized tags, specific to the user's organization of her bookmarks. A significant fraction of the tags is used only once.



(a) flickr.com

Figure 4: Zipf distribution of tags associated with the url `flickr.com` by all users in our crawled data set. The flat part of the curve corresponds to the most frequent tags, which are likely the ones proposed by `del.icio.us`. Similar curves were obtained for urls `del.icio.us`, `pandora.com` and `youtube.com`.




Figure 5: Scatter plot of tag similarity versus bookmark similarity for 10,000 randomly sampled pairs of the bookmark similarity graph, showing no correlation between bookmarks and tags.

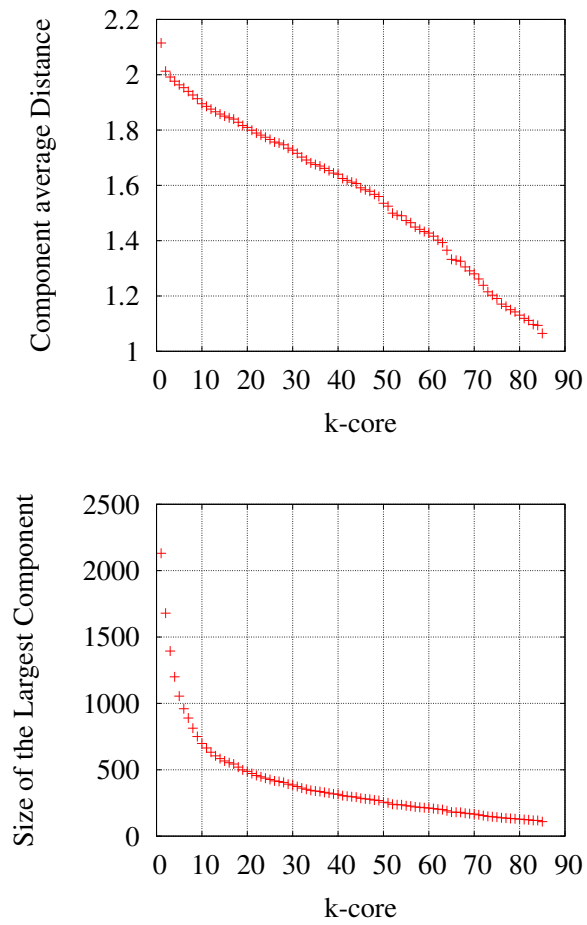


Figure 6: This plot focusses on the largest component of the common bookmark graph, and it depicts its average distance between nodes (top) and its size (bottom) as a function of k in the k -core analysis. This behaviour is typical of the other common entity and entity similarity graphs.

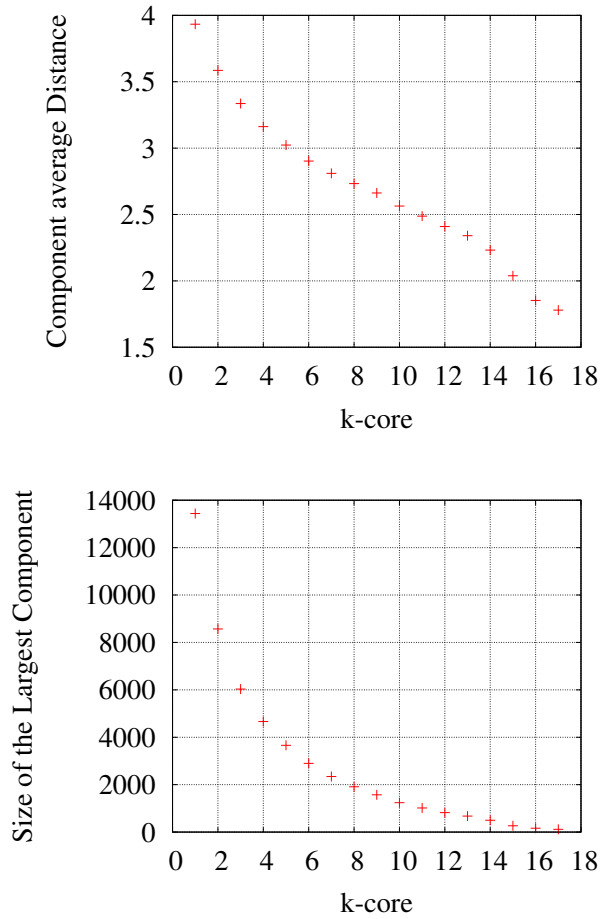


Figure 7: This plot focusses on the largest component of the friendship graph, and it depicts its average distance between nodes (top) and its size (bottom) as a function of k in the k -core analysis.

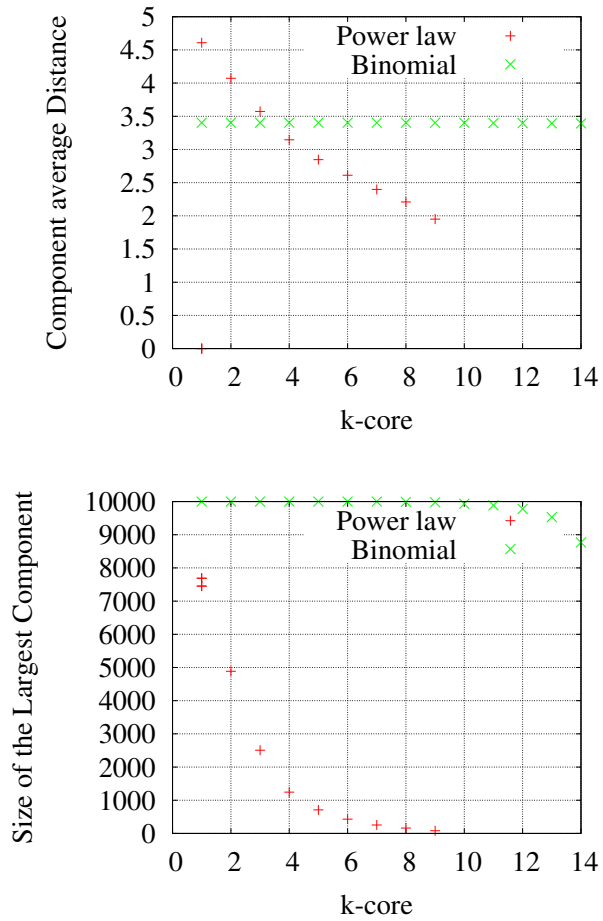


Figure 8: The average distance between nodes and the size of the largest component in the k-core analysis of binomial and power-law random graphs.

Clustering coefficient of k-cores In this subsection we examine the clustering coefficient and its relation to the degree of the graph vertices. A plot of the average clustering coefficient for vertices of a given degree is shown as a function of degree in the left column of Fig. 9 and Fig. 10. For a binomial random graph and random graph with expected power law degree distribution, the corresponding theoretical curve would be horizontal, since the fact that two vertices are friends of a third vertex does not affect the probability of them being linked.

The average clustering coefficient over all vertices as a function of k in the k -core analysis is shown in the right column of Fig. 9 and Fig. 10. These plots are consistent with the plots in the left column, considering that the low-degree vertices are dropped first as k increases in the k -core analysis. For low values of k , low degree vertices are lost. When such vertices have high clustering coefficient, we see a drop in the average clustering coefficient, down to a minimum for a value of k . As k increases past that minimum, the densification process prevails and we see the expected monotonic increase in the average clustering coefficient. In the friendship graph, shown in Fig. 11, we observe the same trend of low degree vertices having high clustering coefficients, with the average monotonically decreasing with increasing degree. This implies that friends of users (represented by vertices) with low degrees, or equivalently, few friends, tend to be friends themselves, while friends of users with large degrees are not necessarily connected.

For comparison, the clustering coefficient as a function of degree, and the average clustering coefficient as a function of k in the k -core analysis are shown in Fig. 12. Here we note that the behaviour of the binomial random graph is very different. The power law graph shows an average clustering coefficient that increases monotonically as a function of k .

6 Discussion

In this paper, we closely examined the properties of a typical social bookmarking and tagging data set, to obtain insights for facilitating creation of models for such data. We summarize our observations as follows:

- Friendship correlates well with common bookmarks or similar bookmark vectors of users, but not well with common tags or similar tag vectors. This implies that the majority of tags that users assign to bookmarks are user-specific tags, and only a small fraction of the tags capture the generic meaning of the bookmarked web page.
- Tags behave more like words in text, while bookmarks less so, in the sense that a much higher proportion of bookmarks are used only once compared to the proportion of tags and words in the Wikipedia corpus that are used only once.
- There are no tags or bookmarks behaving like stop words in text. Even the highest frequency tags or bookmarks do not appear that frequently to deserve the characterization of stop words.
- The clustering coefficient as function of k in the k -core analysis displays a U-shaped curve. This is not consistent with binomial random graphs or random graphs with a power law degree distribution.

Acknowledgements. We are grateful for the financial support of the Natural Sciences and Engineering Research Council of Canada, the MITACS Network of Centres of Excellence, and Genieknows.com. Tom Crecelius, and Mouna Kacimi of the Databases and Information Systems Department, Max Planck Institut für Informatik directed by Prof. Gerhard Weikum, kindly provided the *del.icio.us* data set for this project. Jacek Wołkowicz gave us the Wikipedia statistics. The computations required for the k -core results would not be possible without access to the Atlantic Computational Excellence Network ACEnet and the Shared Hierarchical Academic Research Computing Network SHARCNET computing clusters.

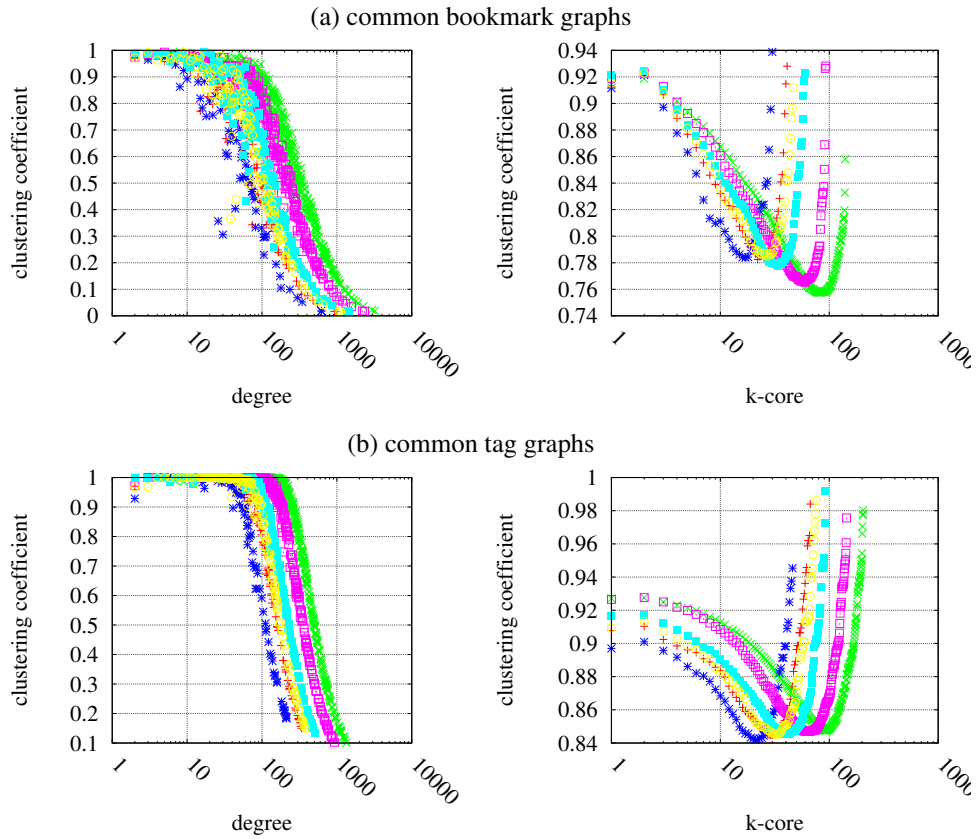


Figure 9: The average clustering coefficient over the vertices of a given degree as a function of degree (left). The average clustering coefficient over all vertices of the largest component in the k-core analysis, as a function of k (right). We observe that less “social” individuals tend to have a closely knit set of neighbours, who are likely to be connected to each other, because of their very specialized interests.

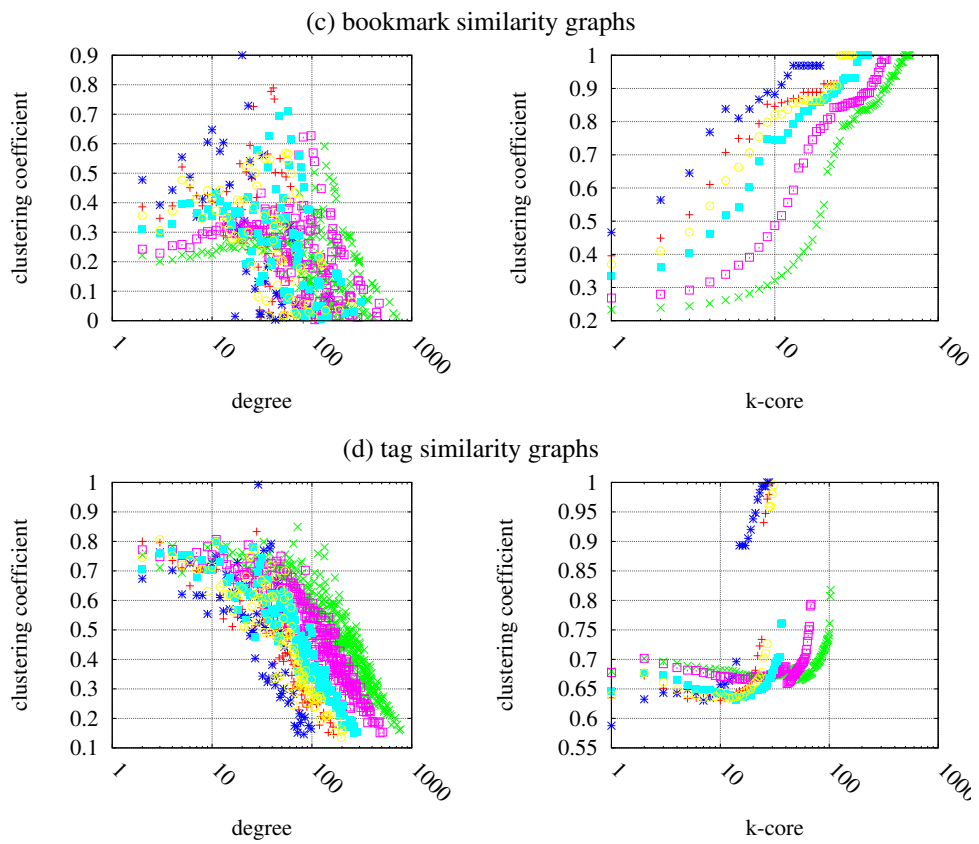


Figure 10: The average clustering coefficient over the vertices of a given degree as a function of degree (left). The average clustering coefficient over all vertices of the largest component in the k-core analysis, as a function of k (right).

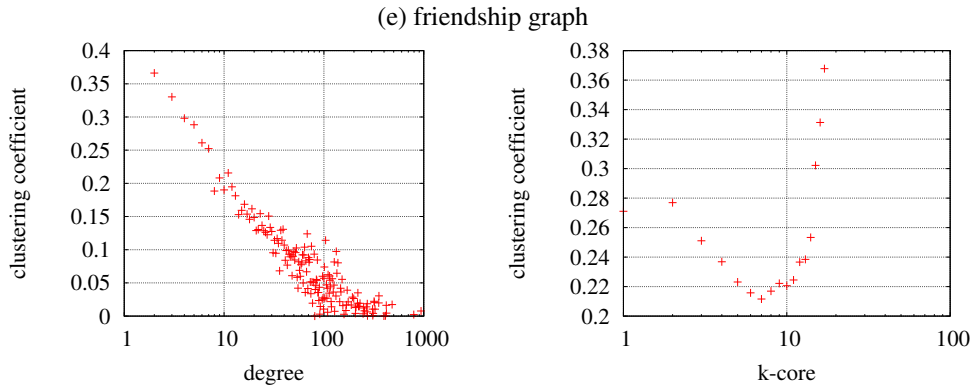


Figure 11: The average clustering coefficient over the vertices of a given degree as a function of degree (left) for the friendship graph. The average clustering coefficient over all vertices of the largest component in the k-core analysis, as a function of k (right), for the friendship graph.

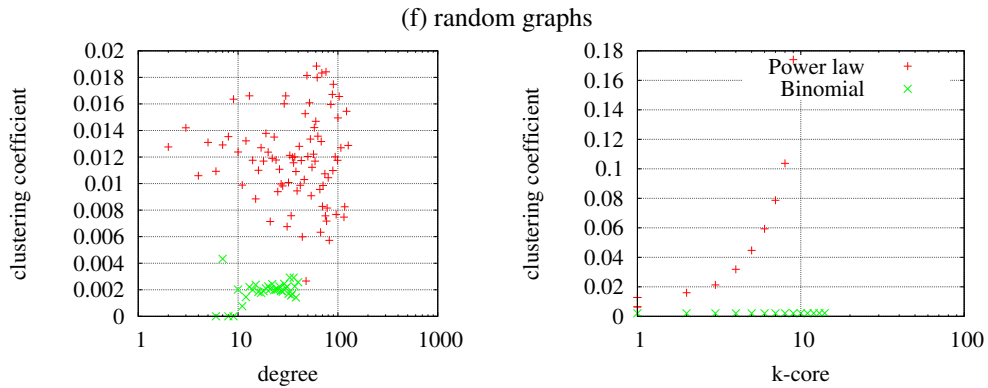


Figure 12: The average clustering coefficient over the vertices of a given degree as a function of degree (left) for binomial and power-law random graphs. The average clustering coefficient over all vertices of the largest component in the k-core analysis, as a function of k (right), for binomial and power-law random graphs.

References

- Ahn, Y.-Y., S. Han, H. Kwak, S. Moon, and H. Jeong (2007). “Analysis of Topological Characteristics of Huge Online Social Networking Services”. In: *World Wide Web (WWW) Conference*. Banff, Alberta. 835–844.
- Angelova, Ralitsa, Marek Lipczak, Evangelos Milios, and Pawel Pralat (2008). “Characterizing a social bookmarking and tagging network”. In: *Mining Social Data (MSoDa), a Workshop of the 18th European Conference on Artificial Intelligence (ECAI 2008)*. <http://mlkd.csd.auth.gr/msoda/program.html>. University of Patras, Greece.
- Barabasi, A.-L. (2002). *Linked: The New Science of Networks*. Cambridge, MA: Perseus.
- Batagelj, C., and M. Zaveršnik (2002). “Generalized Cores”. In: *ArXiv Computer Science e-prints*.
- Chung, F.R.K., and L. Lu (2004). *Complex graphs and networks*. U.S.A.: American Mathematical Society.
- Cohen, R., D. Avraham, and S. Havlin (2003). “Structural properties of scale free networks”. In: *Handbook of graphs and networks*. Wiley-VCH. chap. 4. URL: citeseer.ist.psu.edu/cohen02structural.html.
- Golder, S., and B. Huberman (2005). The Structure of Collaborative Tagging Systems. Tech. rep. Information Dynamics Lab, HP Labs.
- Halpin, H., V. Robu, and H. Shepherd (2007). “The Complex Dynamics of Collaborative Tagging”. In: *World Wide Web (WWW) Conference*. Banff, Alberta. 211–220.
- Healy, John, Jeannette Janssen, Evangelos Milios, and William Aiello (2008). “Characterization of graphs using degree cores”. In: *Algorithms and Models for the Web-Graph: Fourth International Workshop, WAW 2006*. Vol. LNCS-4936. Lecture Notes in Computer Science. Banff, Canada: Springer Verlag.
- Hotho, Andreas, Robert Jaeschke, Christoph Schmitz, and Gerd Stumme (2006). “Information Retrieval in Folksonomies: Search and Ranking”. In: *Proceedings of the 3rd European Semantic Web Conference*. Lecture Notes in Computer Science. Springer.
- Janson, S., T. Łuczak, and A. Ruciński (2000). *Random Graphs*. New York: Wiley.
- Kumar, R., J. Novak, and A. Tomkins (2006). “Structure and Evolution of Online Social Networks”. In: *Knowledge Discovery in Databases (KDD) Conference*. ACM. Philadelphia, PA. 611–617.
- Moser, Flavia, Rong Ge, and Martin Ester (2007). “Joint Cluster Analysis of Attribute and Relationship Data Without A-Priori Specification of the Number of Clusters”. In: *Knowledge Discovery in Databases (KDD) Conference*. San Jose, CA: ACM. 510–519.
- Negoescu, Radu (2007). An Analysis of the Social Network of Flickr. Tech. rep. Switzerland: Laboratory of nonlinear systems - LANOS, School of Computer, Communication Sciences, École Polytechnique Fédérale de Lausanne.
- Schmitz, Christoph, Miranda Grahl, Andreas Hotho, Gerd Stumme, Ciro Cattuto, Andrea Baldassarri, Vittorio Loreto, and Vito D.P. Servedio (2007). “Network Properties of Folksonomies”. In: *World Wide Web Conference (WWW)*. Banff, Canada.
- Seidman, S. B. (1983). “Network structure and minimum degree”. In: *Social Networks*. 269–287.
- Sigurbjörnsson, Börkur, and Roelof van Zwol (2008). “Flickr tag recommendation based on collective knowledge”. In: *WWW '08: Proc. the 17th international conference on World Wide Web*. Beijing, China: ACM. ISBN 978-1-60558-085-2. DOI: <http://doi.acm.org/10.1145/1367497.1367542>. 327–336.
- Wasserman, S., and K. Faust (1994). *Social network analysis: Methods and applications*. Cambridge: Cambridge University Press.