Quantifying trust dynamics in signed graphs, the S-Cores approach

Christos Giatsidis, Bogdan Cautis, Silviu Maniu, Michalis Vazirgiannis,
Dimitrios M. Thilikos

► To cite this version:

Christos Giatsidis, Bogdan Cautis, Silviu Maniu, Michalis Vazirgiannis, Dimitrios M. Thilikos. Quantifying trust dynamics in signed graphs, the S-Cores approach. SDM: SIAM Data Mining, Aug 2014, Philadelphia, United States. pp.668-676, 10.1137/1.9781611973440.77 . lirmm-01083529

HAL Id: lirmm-01083529

<https://hal-lirmm.ccsd.cnrs.fr/lirmm-01083529>

Submitted on 17 Nov 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Quantifying trust dynamics in signed graphs, the S-Cores approach

Christos Giatsidis^{*} Bogdan Cautis[†] Silviu Maniu[‡] Dimitrios M. Thilikos^{§¶}

Michalis Vazirgiannis^{||}

Abstract

Lately, there has been an increased interest in *signed networks* with applications in trust, security, or social computing. This paper focuses on the issue of defining models and metrics for reciprocity in signed graphs. In unsigned directed networks, reciprocity quantifies the predisposition of network members in creating mutual connections. On the other hand, this concept has not yet been investigated in the case of signed graphs. We capitalize on the graph degeneracy concept to identify subgraphs of the signed network in which reciprocity is more likely to occur. This enables us to assess reciprocity at a global level, rather than at an exclusively local one as in existing approaches. The large scale experiments we perform on real world data sets of trust networks lead to both interesting and intuitive results. We believe these reciprocity measures can be used in various social applications such as trust management, community detection and evaluation of individual nodes. The global reciprocity we define in this paper is closely correlated to the clustering structure of the graph, more than the local reciprocity as it is indicated by the experimental evaluation we conducted.

Keywords: Graph mining, trust networks, graph degeneracy, signed networks.

1 Introduction

Online social networks have become increasingly popular in today's Web landscape. Compared to the rest of the Web, their main difference is that they focus on the user instead of the content. Relationships (links or edges in a

social network) between the users generate a rich context that has formed a basis for extensive research in the last years, driven by applications such as search, recommendations or access control. All these applications can benefit from the knowledge of the social network associated to the data content as well.

Going beyond a paradigm of simple links between users, an emerging research direction is the study of networks in which links convey richer semantics. In particular, certain social networks, usually called *signed*, can contain both links with a positive interpretation (trust, friendship, similarity) and links with a negative interpretation (distrust, opposition, antagonism) [12]. Indeed, several social applications support and publish such links (e.g., *Epinions* and *Slashdot*). Furthermore, recent research tackled the problem of extracting *implicit* signed networks, through the analysis of user interactions in the application [14, 13, 3, 18].

This paper deals with the issue of reciprocity in signed graphs and introduces reciprocity in signed networks both as a local and a global property illustrating the value of the latter. The basic definition of reciprocity is a local property based on mutuality in pairs of nodes in directed graphs [21, 16, 19]: $r \equiv \frac{L^{\leftrightarrow}}{L}$, where $L^{\leftrightarrow}$ is the number of links pointing in both directions and L is the total number of links. Thus, the highest value of r is 1, when the network is fully bi-directional, and the lowest is 0 when the network is completely unidirectional.

We argue that the concept of reciprocity as defined in existing works does not offer an adequate descriptive capability (especially for signed networks) for measuring reciprocity at the graph level. Reciprocity, in unsigned directed networks, quantifies the predisposition that the members of a network display in creating mutual connections. In signed trust networks, reciprocity would have different interpretations based on the pairs of signs we examine. By in^+/out^+ pairs, we would get an indication on the level of trust. In contrast, in^-/out^- pairs would indicate distrust or vindictiveness. Moreover, the in^+/out^- and in^-/out^+ pairs may reveal interesting aspects as well. The reciprocity of the former would describe impartiality under positive votes (trust), while the latter would describe impartiality under negative votes (distrust). A more strict version of reciprocity could be viewed by the account of only the number of bidirectional

^{*}LIX – École Polytechnique, France. xristosakamad@gmail.com

[†]Université Paris-Sud – INRIA. bogdan.cautis@u-psud.fr

[‡]University of Hong Kong. smaniu@cs.hku.hk

[§]AIGCo project-team, CNRS, LIRMM, Montpellier, France and National and Kapodistrian University of Athens, Athens, Greece. sedthilk@thilikos.info

[¶]co-financed by the E.U. (European Social Fund - ESF) and Greek national funds through the Operational Program "Education and Lifelong Learning" of the National Strategic Reference Framework (NSRF) - Research Funding Program: "Thales. Investing in knowledge society through the European Social Fund".

^{||}Department of Informatics, Athens Univ. of Economics & Business, Greece & LIX – École Polytechnique, Palaiseau Cedex, France. mvazirg@lix.polytechnique.fr

Figure 1: Example digraph of signed networks.

links bearing the same sign in both ends. This would describe only the most basic nature of reciprocity without taking into account any further context.

Disregarding for now what would be the best definition of reciprocity for signed networks, in Figure 1 we see a sample toy signed directed graph (to keep it simple we assume only positive signs on the links) representing trust relations. As we can observe there is no pairwise mutuality in terms of mutual plus links among pairs of nodes – therefore the local reciprocity as defined in directed and signed network is zero. On the other hand it is clear from the graph that there is a global reciprocity as for each node we observe a balanced in / out positive trust. For example, in this case each node offers two outgoing positive (+) trust links to the community and at the same time receives two positive (+) incoming trust links – although they do not emanate from the same node as those it gives trust to. Thus it is evident that there is a challenge in representing and dealing with reciprocity at a more global level. This issue becomes even greater when we consider the added complexity of signs; since possible combinations of local reciprocities would lessen their importance when looking at the graph at a node level.

Local reciprocity is a property that represents the mutuality among relationships. Global reciprocity would be defined as the average reciprocity at graph level in terms of the portion of positive (+) to negative (-) trust within the community rather than at the node level. Hence, it is our objective to measure graph level reciprocity for dense communities.

Following, we list the contributions of the paper:

- We introduce the notion of degeneracy in signed graphs – to the best of our knowledge this is the first attempt towards this direction.
- We introduce the notion of reciprocity for signed graphs.
- We capitalize on the former to define novel metrics to measure reciprocity in signed graphs in the context of trust at a global graph level as opposed to local one.
- We conduct large scale experimental evaluations and we interpret the results – an interesting issue in the absence of previous similar work.

The reader may also consider viewing the supplementary material for a complete picture over the extend of the conducted experiments and the concepts introduced here.

2 Related Work

Many algorithms for mining social networks have been proposed in the literature, but none of them takes into account the concept of degeneracy in the context of community evaluation in signed networks. In the area of signed graphs, a machine learning-based approach for inferring negative/positive links in Epinions was published in [13], whose techniques rely on an existing signed network complemented by user interactions. In [14], a signed network over the editors of the *Wikipedia*, denoted *Wiki-Signed*, is inferred exclusively from interactions; it is evaluated, at both local and global level, in relation with social theories and existing signed networks on the Web. We rely in this paper on networks built as in [14]. Another approach for detecting positive and negative interactions in *Wikipedia* was presented in [3], showing the emergence of polarization in *Wikipedia* articles.

Several papers have also studied the prediction of links and link signs, when only the signed network is known, a problem also known as *trust propagation*. The first rigorous treatment of this problem is given in [10], where the authors define four atomic operators to predict link signs (direct propagation, co-citation, transpose trust and trust coupling). This approach was extended in [12, 11], where trust propagation was studied through the lens of social theories such as balance and status, and a prediction model based on the number of triangles involving each candidate link was proposed.

For undirected signed graphs, the theory of Social Balance [2] is a model for the dynamics friendship and enmity through time. The weakness of this model is that it assumes that all relationships are reciprocal. A more sophisticated model (for directed signed edges) called Status Model is introduced in [10] and elaborated in [12]. The main point of this model is that a directed signed edge signifies someone of either higher or lower status and thus predicts that the flipping of a direction should flip the sign as well. But this would not account for the relationships of trust (that we attempt to study). In principle (and shown by our results), it is counter intuitive to assume that showing trust or distrust to others would lead to the opposite assumptions of others to us.

Reciprocity is used to examine directed networks of various kinds [16, 19, 6], an extension for weighted networks is in the recent work of [1]. In [6], reciprocity is extended in order to take into account the density of the network. The model of reciprocity presented in that work is described by two equations that are supposed to be equivalent but, on deeper examination, do not match. In our work, we build upon the original definition of reciprocity as it is a clear one with intuitive interpretation and is widely accepted.

Additional work has been done in the area of extending k -cores, a fundamental concept in graph theory whose study goes all the way back to the 60s [4, 20, 15]. The existence of k -cores of large size in sufficiently dense graphs has been formally studied in [17], for random graphs generated by the Erdős-Rényi model [5]. Extensions of the k -cores structure have been studied for weighted undirected graphs in [8], and for directed graphs in [7]. Both extensions are targeted in the evaluation of community graphs and their members.

3 Preliminaries

In this section we define the fundamental concept of degeneracy for signed graphs that will be exploited towards the definition of reciprocity in such graphs. First we define the notion of S-core – an extension of D-core [7] – that represents degeneracy in signed graphs. Moreover we define additional concepts that quantify the robustness of the graph under degree based degeneracy.

3.1 S-cores: Degeneracy in signed graphs A *signed digraph* is a triple $G = (V, E, \mathbf{w})$ such that (V, E) is a simple directed graph and $\mathbf{w} : E \rightarrow \{+, -\}$ is a labeling of E , assigning either a positive or a negative sign on the edges of G . The existence of a positive signed edge $e = (x, y)$ from a vertex x to a vertex y represents the fact that “ x trusts y ” or “ x likes y ”, while the existence of the same edge with negative sign means that “ x distrusts y ” or “ x dislikes y ”.

Given a vertex v of G , we denote by $\deg_{\text{in}}^+(v, G)$ (resp. $\deg_{\text{out}}^+(v, G)$) the *positive in-degree* (resp. *positive out-degree*) of v in G , i.e., the number of positive-signed edges tailing (resp. heading) on v . The *negative in-* and *out-degrees* of the vertices of G are defined analogously and are denoted by $\deg_{\text{in}}^-(v, G)$ and $\deg_{\text{out}}^-(v, G)$.

Let $G = (V, E, \mathbf{w})$ be a signed graph. Let also $s, t \in \{+, -\}$ and $k, l \in \mathbb{N}$. We define the (l^t, k^s) -dicore of G as the maximum size subgraph H of G where, for each vertex v of H , it holds that $\deg_{\text{in}}^s(v, H) \geq k$ and $\deg_{\text{out}}^t(v, H) \geq l$. Throughout this paper, we use the generic term S-core when we do not need to make explicit the values of the pair (l^t, k^s) .

Notice that the (l^t, k^s) -dicore of G can be computed by the following greedy procedure, similar to those described in [8, 7]: *remove from G a vertex v where $\deg_{\text{in}}^s(v, H) < k$ or $\deg_{\text{out}}^t(v, H) < l$, until this is not possible anymore*. It is straightforward that the resulting sub-graph is well-defined – i.e., it is the same regardless of the order of elimination of vertices – and it is indeed the (k^s, l^t) -dicore of G .

(s, t) -degeneracy: Given a pair $(s, t) \in \{+, -\}^2$, we define the (s, t) -degeneracy of G as follows.

$$\delta^{s,t}(G) = \max \left\{ \frac{k+l}{2} \mid G \text{ contains a non-empty } (l^t, k^s)\text{-dicore} \right\}$$

For convenience, we define the sign function $\mathbf{s} : \mathbb{Z} \rightarrow$

$\{+, -\}$ that, given an integer i , outputs $-$ or $+$ depending whether i is negative or not.

Thus the (s, t) -degeneracy of the graph represents its robustness under degeneracy in the four different combinations of edge-direction and sign. In the case of trust networks the (s, t) -degeneracy represents the degeneracy of the graph for each of the combinations of incoming/outgoing and positive/negative trust. For instance refer to Figure 2 where the four cases of degeneracy are depicted as $\delta_{\text{max}}^{++}(G)$ etc.

Signed dicore diagram: For every signed digraph G , we define its *signed dicore diagram* (or S-core diagram) as a matrix $A = (\alpha_{i,j})_{(i,j) \in \mathbb{Z}^2}$ where for each $(i, j) \in \mathbb{Z}^2$, $\alpha_{i,j}$ is the size (i.e., the number of vertices) of the $(i^{\mathbf{s}(i)}, j^{\mathbf{s}(j)})$ -dicore of G .

Signed Graph Extension: As the above definition of $A = (\alpha_{i,j})_{(i,j) \in \mathbb{Z}^2}$ produces an infinite matrix, it is sufficient to consider its finite portion, which contains all its non empty dicores. For this, we restrict i and j to belong in the *frame of G* that is the set $F_G = \{-(b+1), \dots, 0, \dots, b+1\}$ where

$$b = \max\{l, k \mid G \text{ has a non-empty } (l^t, k^s)\text{-dicore for some } (s, t) \in \{+, -\}^2\}.$$

We call the value b the *extension* of the signed graph G and we denote it by $\mathbf{b}(G)$.

S-core region: Given a signed graph G , we define its *core region* as the set $R_G = \{(i, j) \in F_G^2 \mid \alpha_{i,j} > 0\}$, that is all the pairs that correspond to a non-empty S-core.

S-core frontier: The *core-frontier* of G , denoted by B_G , is the set of all entries $(i, j) \in F_G^2$ with the property that $\alpha_{i,j} > 0$ and $\alpha_{i+(s(i)1), j+(s(j)1)} = 0$. These are the extreme non-empty S-cores, in the sense that any further shift of their coordinate results in an empty S-core.

Here we attempt an intuitive presentation of the above definitions based on Figure 2. There the reader may see the extension of the degeneracy of the four cases as the areas (i.e. $R^{++}(G)$ for the positive in and out trusts edges degeneracy) enclosed by the respective frontiers. In Figure 3 we see the aforementioned metrics and concepts depicted for two real works graphs. Apparently the size of the graph shrinks as the thresholds for in/out for positive/negative signs increase but we notice that graphs are much more robust under degeneracy for the positive/positive trust case. We notice that the bottom diagram indicates that the *Wikipedia-politics* graph is more robust as its extension and therefore its region R_G is larger that the *Slashdot* graph. Evidently the S-core frontier for the second graph is well more extended than the other one. Many of the above definitions, (and some that follow in latter sections) can be seen in a visual representation in Figure 3.

3.2 Reciprocity in signed graphs Here, we define different notions of reciprocity in signed graphs. Initially we build up on existing definitions of reciprocity for directed graphs

based on local criteria and extend them towards signed ones. Moreover we define novel notions of reciprocity that do not depend only on local binary reciprocity but represent this concept in an aggregate manner at the graph level.

3.2.1 Signed graph reciprocity – local definition First, we need to adapt the existing definitions [21, 16, 19] to signed digraphs. We consider trust networks as a prominent example of signed digraphs where a node can either trust or distrust another. Additionally, since self-trust is trivial, self loops are excluded.

The intuition of reciprocity in signed networks must also be examined. We assume two different options: i. *Contextual local reciprocity* where we examine all four possible sign permutations between two reciprocal edges where each sign permutation defines a context of trust and ii. *Simple local reciprocity* where we consider only the mutuality under trust and distrust – i.e. we consider only the pairs of nodes with the same sign that represents the coarse level of trust reciprocity.

Contextual local reciprocity: Following we define the reciprocity emanating from all the possible signs permutations on reciprocal links on a pair on nodes.

$$\begin{aligned} r^{++} &\equiv \frac{L^{+\leftrightarrow+}}{L^+} \text{ for the } in^+/out^+ \\ r^{+-} &\equiv \frac{L^{+\leftrightarrow-}}{L^+} \text{ for the } in^+/out^- \\ r^{-+} &\equiv \frac{L^{-\leftrightarrow+}}{L^-} \text{ for the } in^-/out^+ \\ r^{--} &\equiv \frac{L^{-\leftrightarrow-}}{L^-} \text{ for the } in^-/out^- \end{aligned}$$

where $L^{+/-}$ is the count of positive/negative edges and the signs on the double arrow of $L^{\leftrightarrow}$ (links pointing both ways i.e. reciprocations) indicate the sign of in and out edges respectively. Notice that the denominator is not the same for all the definitions. We could have used L instead of $L^{+/-}$ but, in our trust model, the second and fourth reciprocity would have had identical values. The identical values would not be an issue for a more relaxed model that would allow two edges of different sign to have the same source and target.

The rationale for this definition is that, since reciprocation quantifies mutuality, we only care about the type of actions that are being mutual. For each type of reciprocity above, we select a different set of actions that we are interested to see if they are being reciprocated. For example, when we study the reciprocation of trust by trust (in^+/out^+) it is more intuitive (and more expressive) to compute reciprocity as only the portion of the positive edges and not the total number of them. More over this way, in our assumptions about the network, we can have distinguishable values between the in^+/out^- and in^-/out^+ cases of reciprocity.

Simple local reciprocity: Moving on to the second definition, we strictly consider only the same sign reciprocations and thus we have the following :

$$r^s \equiv \frac{L^{+\leftrightarrow+} + L^{-\leftrightarrow-}}{L}$$

For the rest of the document we will refer to r^s as *simple reciprocity* and the former set of four signed reciprocities as *contextual reciprocity*. We also consider the average of the local reciprocities over the individual nodes (e.g. $r_a^{+\leftrightarrow+}$ is the average of ratios of reciprocal positive edges in individuals over all vertices). These *average local reciprocities* are utilized only for comparison (see Section 5.3). All references to local reciprocities in the text correspond to the original five definitions unless specified otherwise.

OBSERVATION 1. Invariance under sign flipping: A crucial difference between the two definitions is that the first one is not invariant to sign flipping while the second one is. With contextual reciprocity we try to quantify different behaviors under a particular context. For example, trust and distrust are two opposite concepts and their measurement should change if we have a different count of reciprocal signs. On the other hand, simple reciprocity remains the same since we count one type of behavior. Therefore, flipping the signs should not (and does not) change the measurement of simple reciprocity.

3.2.2 Signed graph reciprocity – global definition The concept of reciprocity as defined in existing works capitalizes on the local property of mutuality among pairs of nodes and does not offer an adequate descriptive capability for measuring reciprocity at the graph level.

We proceed here to define metrics that represent signed graph reciprocity at graph level. Figure 3 is a visual aid to those definitions (the S-core frontier here is the irregular shape outlined with the thick line). In this diagram the trust axes (in, out) and signs $(+, -)$ define respective quadrants Q_{out_sign, in_sign} , where $out_sign, in_sign \in \{+, -\}$. Each of the quadrants bears specific semantics regarding the in/out trust. For instance $Q_{+,+}$ represents degeneracy in graphs where the criterion is the mutual trust incoming and outgoing. On the other hand $Q_{+,-}$ represents degeneracy under outgoing trust but incoming distrust. The graphs in the S-Core frontier (in $Q_{+,-}$) represent situations where users maximally trust others in the graph but they receive distrust from others. The interpretations are analogous for the remaining two quadrants $Q_{-,-}$, $Q_{-,+}$.

Maximum degeneracy on the trust axes: We now discuss the extreme degeneracy on each of the four trust axes – representing the robustness of the graph for each type of trust. For instance the $\delta_{max}^{(out,+)}(G)$ represents the extreme graph with regards to outgoing positive trust degeneracy, i.e.

Figure 2: The S-core diagram, displaying the core sizes, for the *Slashdot* and *Wikipedia-politics* digraphs. Each point, of integer coordinates (i, j) , corresponds to the size of the $(i^{s(i)}, j^{s(j)})$ -dicore. The color variation reflects the sizes of the dicores.

the last non empty graph when we increase the threshold for the outgoing positive trust. Similarly we define the rest of extreme degeneracies $\delta_{max}^{(in,+)}(G)$, $\delta_{max}^{(out,-)}(G)$, $\delta_{max}^{(in,-)}(G)$ on the other trust axes.

Quadrant bounding box: For each of the four aforementioned quadrant there is the previously defined frontier that has a respective bounding box which is defined by the maximal degeneracies on the relevant axes. For instance the bounding box for $Q_{+,+}$ is the dotted rectangle defined by the points: $0, 0$, $(\delta_{max}^{(out,+)}(G), \delta_{max}^{(in,+)}(G))$. This bounding box would be the S-core frontier of G if all its vertices (or at least a subset of the vertices in G) had degrees of at least $\delta_{max}^{(out,+)}(G)$ and $\delta_{max}^{(in,+)}(G)$ and moreover their out/in edges connected them with vertices having the same property.

Quadrant maximal degeneracy: By utilizing the bounding box, we then define the *max degeneracy* of each quadrant as the *intersection point between the diagonal of the bounding box*. For instance for the $Q_{+,+}$ quadrant the maximal degeneracy $\delta_{max}^{(+,+)}(G)$ is defined by the intersection of the diagonal $(0, 0, (\delta_{max}^{(out,+)}(G), \delta_{max}^{(in,+)}(G)))$ and the respective frontier $F^{++}(G)$. The max degeneracy of each quadrant corresponds to the most extreme core in relevance to the

Figure 3: Metrics on the S-core graph.

natural ratio of the maximum degeneracies that characterize the quadrant. This metric signifies the overall activity of the signed network (i.e. evaluation of how much the users interact for each type of relationship) while, at the same time, taking into account the over all outward or inward tendencies (e.g. are the users more prone to giving or receiving trust).

Contextual reciprocity: In accordance to the local reciprocities, defined in previous sections, we design the definition of reciprocity at a graph level. As we can observe in Figure 3 the S-core frontier covers the bounding boxes reaching different levels of degeneracy in each. We utilize this to measure graph level reciprocity. Thus *contextual reciprocity* at the graph level is defined (per quadrant) as the ratio of the area under the respective quadrant frontier over the corresponding bounding box surface. For instance the contextual reciprocity $GR^{+,+}$ of the quadrant $Q_{+,+}$ is defined as:

$$R^{++}(G) / (\delta_{max}^{(out,+)}(G) * \delta_{max}^{(in,+)}(G)).$$

These *Global Reciprocities* have the same contextual meaning as their node level equivalents. For example, GR^{++} measures global trust reciprocity and would reach a maximum value of one if everybody gave as much trust as they received (but not necessarily to the same people).

Global reciprocity: Much like the node level reciprocities we need to define a more strict version of reciprocity for reciprocation of the same types of actions. For this purpose, we consider the quadrants of same sign (i.e. $Q_{+,+}$, $Q_{-,-}$) as those capturing this type of reciprocity, and the quadrants having different in/out signs as ones capturing inverse reciprocity. We can then define *Graph Reciprocity* as such :

$$GR = \frac{GR^{++} + GR^{--}}{GR^{++} + GR^{+-} + GR^{-+} + GR^{--}}.$$

Ideally, the value of this metric would reach a maximum of one only when the reciprocation is in the same sign quadrants (i.e. GR^{+-} and GR^{-+} are equal to zero) thus keeping the same range of values as the rest of reciprocations. By

taking into account the inversely reciprocal quadrants, we differentiate between cases where the signed graph is highly reciprocal only in the same sign quadrants and cases where the same applies and simultaneously there is also high reciprocation in the other quadrants as well.

4 Datasets Description and Methodology

In this paper, we have used two kinds of signed networks: *explicit* ones, from existing Web applications publishing such networks, and *implicit* ones, *inferred* from interactions that can be interpreted as positive or negative.

Explicit signed networks. We have used two explicit signed networks, available on the SNAP website¹, *Epinions* and *Slashdot*. The *Epinions* network is extracted from the epinions.com website in which any user of the site can indicate if they trust or distrust other users. Similarly, in the *Slashdot* signed network, extracted from the slashdot.org website, users declare friends or foes. We present the main properties of these explicit signed networks in Table 1a. For a more in-depth description, we refer the reader to [12].

Implicit (inferred) signed networks. We adopt in this paper the network *Wiki-Signed*, which is a signed network built with the methodology of [14], over the Wikipedia editors, based on the articles of the English Wikipedia and the revision history thereof. In short, this network tracks the various interactions between contributors, either in text editing, in votes for adminship of pages, or in acknowledgments of contributions (so called barnstars). From the global Wiki-Signed network, we selected four subsets of articles, from the following domains History, Politics, Religion and Mathematics. Each of these sets gave a corresponding subgraph of Wiki-Signed. We give in Table 1b the properties of the four signed networks – corresponding to the four domains of articles: number of extracted articles for each domain, number of nodes, number of edges and ratio of negative edges

In Figure 2, we show the distributions of S-cores sizes along the degeneracy coordinates for the *Wikipedia*-Politics (bottom) and *Slashdot* (top) digraphs. Each point in the plot corresponds to the size of the $(i^{s(i)}, j^{s(j)})$ -dicore and the color variation reflects the sizes of the cores. It is clear in both cases that the trust pattern is robust to degeneracy at the positive in and out trust. On the contrary, the negative trust constituents are very minor. This conveys that (i) users tend to engage more in trust actions than distrust ones, (ii) it is unlikely that a user that trusts another to receive a negative trust link in reciprocation (and/or the opposite), and (iii) a very small portion of the graph concerns mutual distrust, as the $(-, -)$ quadrant has a very small relative volume overall.

Network	Nodes	Edges	Negative
<i>Epinions</i>	119,217	841,200	15.0%
<i>Slashdot</i>	82,144	549,202	22.6%

(a)

Domain	Articles	Nodes	Edges	Negative
History	3,331	141,983	534,693	17.5%
Politics	12,921	453,116	2,428,945	13.5%
Religion	6,459	277,482	1,423,279	12.6%
Mathematics	9,610	158,671	651,450	15.9%

(b)

Table 1: a. Properties of the explicit signed networks. b. The signed networks extracted from the four Wikipedia domains.

5 Experimental Evaluation

We present in this section the experiments we performed on the explicit signed graphs (*Epinions* and *Slashdot*) and all the inferred *Wikipedia* networks. The algorithm for computing the S-cores of a signed digraph is linear to the number of the graph edges. As the signed graphs we examine are sparse, the construction of the S-cores is hence very fast. The computation is quite straightforward: for a given pair of in- and out-degree thresholds, we remove iteratively the vertices having degrees that are below the desired threshold and update the degrees of the remaining nodes. We repeat until there are no more nodes in the graph to remove.

The algorithm follows the same logic as in [8, 7]. In particular, we optimize the efficiency of computing all the S-cores by utilizing the following property. A $(i^{s(i)}, j^{s(j)})$ -dicore is a subgraph of every $(i'^{s(i')}, j'^{s(j')})$ -dicore where $i' \leq i$ and $j' \leq j$ and for the signs: $s(i') = s(i)$ or $i'=0$ and $s(j') = s(j)$ or $j'=0$ (i.e. both of the dicores are in the same quadrant). Thus, we can compute e.g. $(2^+, 0)$ -dicore having computed and stored in memory the $(1^+, 0)$ -dicore. Moreover, we can compute the entire S-core diagram by computing firstly the S-cores on the axes. Note that two S-cores upon different axes are not correlated so we need to compute all of them across the axes but we need on each quadrant only one of the axes to fill in the rest.

5.1 Slashdot and Epinion graphs The graphs derived from the *Slashdot* and *Epinions* networks are explicitly defined by the users thus providing ground truth examples for the S-cores and their metrics. Figure 4a displays comparatively the frontiers of the *Slashdot* and *Epinion* graphs. It can be seen that the *Epinions* network has a larger negative area, which is interpreted as a more distrustful community. For more details we look at Table 2, containing the general trends of the two networks. For example we see the higher activity of mutual trust ($Q_{+,+}$) displayed from *Slashdot* (com-

¹<http://snap.stanford.edu/data/index.html#signednets>

Graph	$Q_{+,+}$	$Q_{+,-}$	$Q_{-,-}$	$Q_{-,+}$	r^s/GR
Epinions					
max degeneracy ($\delta_{\max}$)	(19,18)	(1,-4)	(-5,-5)	(-5,1)	-
local reciprocity (r)	0.347	0.003	0.038	0.022	0.302
mean local reciprocity	0.230	0.002	0.046	0.013	0.160
graph reciprocity (GR)	0.976	0.109	0.886	0.197	0.859
Slashdot					
max degeneracy ($\delta_{\max}$)	(37,35)	(2,-2)	(-4,-4)	(-3,1)	-
local reciprocity (r)	0.197	0.004	0.072	0.016	0.169
mean local reciprocity	0.228	0.004	0.091	0.025	0.133
graph reciprocity (GR)	0.978	0.067	0.8	0.108	0.911
History					
max degeneracy ($\delta_{\max}$)	(17,17)	(1,-1)	(-2,-2)	(-1,1)	-
local reciprocity (r)	0.065	0.004	0.010	0.020	0.055
mean local reciprocity	0.050	0.004	0.030	0.029	0.010
graph reciprocity (GR)	0.938	0.158	1.0	0.205	0.842
Politics					
max degeneracy ($\delta_{\max}$)	(64,65)	(1,-2)	(-2,-2)	(-3,1)	-
local reciprocity (r)	0.105	0.006	0.020	0.040	0.094
mean local reciprocity	0.059	0.006	0.067	0.048	0.014
graph reciprocity (GR)	0.955	0.535	1.0	0.564	0.640
Religion					
max degeneracy($\delta_{\max}$)	(42,43)	(1,-2)	(-2,-2)	(-2,1)	-
local reciprocity (r)	0.084	0.006	0.022	0.044	0.076
mean local reciprocity	0.058	0.007	0.064	0.050	0.015
graph reciprocity (GR)	0.952	0.396	1.0	0.540	0.676
Mathematics					
max degeneracy ($\delta_{\max}$)	(46,47)	(1,-1)	(-2,-2)	(-2,1)	-
local reciprocity (r)	0.099	0.006	0.011	0.032	0.085
mean local reciprocity	0.063	0.006	0.037	0.029	0.012
graph reciprocity (GR)	0.963	0.327	1.0	0.356	0.742

Table 2: The calculated metrics for all graphs. The first four columns (of the calculated values) present the four contextual reciprocities (e.g. r in column $Q_{+,+}$ is the contextual local reciprocity in the $+, +$ quadrant r^{++}). The last column is Simple Local Reciprocity r^s or the Global Reciprocity GR (depending on the corresponding row).

pared to *Epinions*) evident by the higher valued coordinates of $\delta_{\max}^{(+,+)}(G)$ in that quadrant.

Additionally, we start to see a pattern concerning the equivalent values of reciprocity at node and graph level. We see that both types of reciprocity agree in the inversely reciprocal cases with low values. The more distinct differences are in the other quadrants and when comparing *simple reciprocity* (r^s) with the graph reciprocity (GR). It is quite evident that, concerning the remaining reciprocities, on a node level there is much less reciprocity. On the other hand, the general reciprocities are quite large which is attributed to users reciprocating more at a community level than a local one.

Figure 4: a. The frontiers for *Slashdot* and *Epinions* networks. b. The frontiers for the *Wikipedia* topics.

5.2 Wikipedia topics We analyze the S-cores structure on the four Wikipedia topics selected above: *politics*, *history*, *mathematics* and *religion*. In Figure 4b we can see the corresponding S-core frontiers, and in Table 2 the values for the defined metrics. In terms of *maximal degeneracy* the topic of *Politics* has by far the largest trust quadrant. This is expected since there is more activity in that topic, in comparison with the others, evident by the larger number of articles, resulting in a larger overall graph. On the other hand, the *History* graph has the smallest value (a direct result from the smaller number of articles). It is quite interesting that, despite the difference in size, all four networks present the same behavior in many of their aspects. The network derived from the topics under *History* seems to display a slightly different behavior. The larger value of GR expresses a general tendency for the users to reciprocate edges of only the same sign back to the community, which in turn can be assigned to a larger bias in the actions of a user.

Again, the node level reciprocities cannot describe the bigger picture of the community's collective actions. For example, *Contextual Global Reciprocity* GR^{++} has a high value – indicating a lot of trustworthiness coming and going from the user to the community– while at the same time

the equivalent node level reciprocities for all four topics are very low. We should also point out that GR^{--} has reached maximum value for all four topics as well. Meaning that, despite the fact we have less unbiased actions (evident by the higher values of GR^{+-} and GR^{-+} and with the exception of *History*), there is a (small) remaining part of the community where distrust is at its maximum – but not directly as the node level *contextual reciprocities* values are very small.

5.3 Local vs. Global reciprocity It is visible from the comparison in Table 2, that node level reciprocity, although it captures somewhat the different trends in trust and distrust, it can not evaluate the wider concept of general (“social”) trust/distrust. This is more visible in quadrant $Q_{+,+}$ as reciprocity r^{++} is low in most of the cases while the respective GR^{++} is close to maximum. The large value of *quadrant maximal degeneracy* $\delta_{max}^{(+,+)}(G)$ indicates that there is a strong community of individuals that reciprocate their trust with their community.

In Table 2 we also compare the *Graph Reciprocities* to the average local reciprocities over all vertices in the presented data graphs. The motivation for that is to establish that *Global Reciprocity* represents the reciprocal behavior better than the local one. The *local reciprocities* are ratios over the edges while the *average local reciprocities* are the mean values over individual behaviors and could be perhaps better at capturing the average behavior over the entire graph. From the comparison in Table 2, it is clear that a) the values of *average local reciprocities* and of *local reciprocities* display more or less the same trends and b) *Graph Reciprocity* can capture collective behaviors the other two metrics cannot.

5.4 S-core reciprocity vs clustering structure In this section we display the correlation between triangles formed by the nodes of signed networks and the properties of reciprocity. The number of triangles a node participates is important to tasks like graph clustering. As graph clustering in signed networks is still quite unexplored, the results presented here could be the seed for further research in graph mining algorithms of signed networks.

In Figure 5 we see the four possible combinations of edges in a triangle formation of a directed graph. From left to right we have: a) the in-triangle 5a, b) the out-triangle 5b, c) the through triangle 5c and d) the cycle triangle 5d. In a signed network the types of triangles (for a given node) is sixteen. In our case though, the types of triangles can be seen as four sets of the four directed types in Figure 5. Each set is the aforementioned four types that is “restricted” by the “properties” imposed from each quadrant. For example, for $Q_{+,+}$ we have the four types displayed in Figure 5 where the incoming edges must be positive and the outgoing as well. While for $Q_{+,-}$ only the outgoing must be positive and

Figure 5: The four possible triangle configurations (for a node – the green one) in a directed network. The dashed line indicates that the direction is not important (as the two possible directions create a “mirror” of one another). From left to right: a) In-triangle, b) Out-triangle, c) through triangle, d) cycle triangle.

the incoming negative.

In Table 3 we see the correlation between a node’s triangle count and the respective reciprocities of the nodes for each quadrant. Much like the frontier for each individual node we can get the global reciprocity (GR) for that node for each of the quadrants. Afterwards we measure the correlation coefficient between a nodes triangle count (four types for each quadrant) and the GR . All the values displayed in Table 3 have $p - value < 0.05$ (from comparison to the no correlation null hypothesis) thus indicating that the correlation is significant.

The global reciprocity can be computed by the cores of a node’s S-core frontier and a connection has been found between the k -core properties of an undirected graph and its clustering coefficient (e.g. see [9]).

Despite that, to our knowledge, this is the first time that the correlation, between core properties in a signed digraph (or even just a digraph) and clustering properties, has been explored. And, even though the local reciprocity of a node is not connected by intuition to clustering properties, we also presented the correlation of the local reciprocity of a node to the number of triangles for comparison and completeness of the exploration.

It is very interesting that, with the exception of the in-triangles, the global reciprocity presents a consistently higher that the local one correlation with the triangles count. In the in-triangle case the reason for this inconsistency (in the $Q_{-,-}$) could be that newcomers into a social network –since they are new to the network and less knowledgeable and/or more eager to participate– are more likely to receive negative votes due to inexperience (which in turn could have a direct reciprocation with a negative vote for the same reason). Never the less the high correlation of GR with triangle count is an indication that GR could be used in signed network clustering (for perhaps a selection phase of seed nodes). Furthermore this indicates the validity and superiority of the GR measure over the local reciprocity.

Reciprocity	in triangle	out triangle	through triangle	cycle triangle
GR^{++}	0,56	0,79	0,96	0,87
r^{++}	0,29	0,09	0,08	0,09
GR^{-+}	0,17	0,26	0,60	0,47
r^{-+}	0,10	0,02	0,01	0,02
GR^{--}	0,05	0,79	0,97	0,87
r^{--}	0,46	0,03	0,04	0,04
GR^{+-}	0,17	0,19	0,42	0,36
r^{+-}	0,10	0,01	0,04	0,03

Table 3: Correlation coefficient values between the four types of triangles and reciprocities for each quadrant.

6 Conclusions

We have investigated in this paper models and metrics for reciprocity in signed graphs, capitalizing on the graph degeneracy concept to identify the most reciprocal parts of signed networks and thus be able to assess reciprocity at a global level, rather than at a local one as existing approaches manifest. Starting from previous work, we have extended the D-core structure – degeneracy in directed graphs – to the concept of S-core, to handle degeneracy in signed graphs. We utilize this tool to better represent reciprocity in the context of trust networks. Our experimental evaluation shows properties of signed networks, both explicit and implicit ones, that are better captured by the new reciprocity measures, suggesting its potential as an objective for optimization algorithms in the context of directed and/or signed graphs, such as graph clustering or link formation models. There are several direction we intend to focus our future work on, such as:

- *Weighted signed networks:* We will extend the notion of degeneracy in signed graphs with weights, i.e., where the link is defined by a real-valued weight representing partially positive or negative endorsement.
- *Link formation models for signed networks:* Based on the existing models for directed graphs we are interested in producing models that fit the structure and behavior of signed graphs and can be used as generators of such realistic graphs. Specifically we are interested in investigating the issue of preferential attachment in signed graphs.
- *Utilizing global reciprocity for clustering:* Graph clustering is achieved by maximizing intra-cluster density, while having as few links as possible between clusters. Intuitively, the global reciprocity measure could be used as an alternative measure for graph clustering, especially in the context of signed networks, where simple

intra-cluster density would not suffice as a measure of realistic user communities.

References

- [1] L. Akoglu, P. O. S. V. de Melo, and C. Faloutsos. Quantifying reciprocity in large weighted communication networks. In *PAKDD* (2), 2012.
- [2] T. Antal, P. Krapivsky, and S. Redner. Dynamics of social balance on networks. *Physical Review E*, 72, 2005.
- [3] U. Brandes, P. Kenis, J. Lerner, and D. van Raaij. Network analysis of collaboration structure in wikipedia. In *International conference on World wide web, (WWW)*, 2009.
- [4] P. Erdős. On the structure of linear graphs. *Israel J. Math.*, 1, 1963.
- [5] P. Erdős and A. Rényi. On the evolution of random graphs. *Magyar Tud. Akad. Mat. Kutató Int. Közl.*, 5:17–61, 1960.
- [6] D. Garlaschelli and M. I. Loffredo. Patterns of link reciprocity in directed networks. *Physical Review Letters*, 93, 2004.
- [7] C. Giatsidis, D. M. Thilikos, and M. Vazirgiannis. D-cores: Measuring collaboration of directed graphs based on degeneracy. In *ICDM*, 2011.
- [8] C. Giatsidis, D. M. Thilikos, and M. Vazirgiannis. Evaluating cooperation in communities with the k-core structure. In *ASONAM*, 2011.
- [9] D. F. Gleich and C. Seshadhri. Vertex neighborhoods, low conductance cuts, and good seeds for local community methods. In *KDD*, 2012.
- [10] R. Guha, R. Kumar, P. Raghavan, and A. Tomkins. Propagation of trust and distrust. In *International conference on World Wide Web, (WWW)*, 2004.
- [11] J. Leskovec, D. Huttenlocher, and J. Kleinberg. Predicting positive and negative links in online social networks. In *International conference on World wide web, (WWW)*, 2010.
- [12] J. Leskovec, D. Huttenlocher, and J. Kleinberg. Signed networks in social media. In *International conference on Human factors in computing systems, (CHI)*, 2010.
- [13] H. Liu, E.-P. Lim, H. W. Lauw, M.-T. Le, A. Sun, J. Srivastava, and Y. A. Kim. Predicting trusts among users of online communities: an epinions case study. In *ACM Conference on Electronic commerce, (EC)*, 2008.
- [14] S. Maniu, B. Cautis, and T. Abdesslem. Building a signed network from interactions in wikipedia. In *Databases and Social Networks, (DBSocial)*, 2011.
- [15] D. W. Matula. A min-max theorem for graphs with application to graph coloring. *SIAM Reviews*, 10, 1968.
- [16] J. B. M.E. Newman, S. Forrest. Email networks and the spread of computer viruses. *Phys Rev E Stat Nonlin Soft Matter Phys*, 66, 2002.
- [17] B. Pittel, J. Spencer, and N. Wormald. Sudden emergence of a giant k-core in a random graph. *J. Combin. Theory Ser. B*, 67, 1996.
- [18] H. Sepehri Rad, A. Makazhanov, D. Rafiei, and D. Barbosa. Leveraging editor collaboration patterns in wikipedia. In *The 23rd ACM conference on Hypertext and social media*, 2012.

- [19] M. A. Serrano and M. Boguñá. Topology of the world trade web. *Phys. Rev. E*, 68, 2003.
- [20] G. Szekeres and H. S. Wilf. An inequality for the chromatic number of a graph. *J. Combinatorial Theory*, 4, 1968.
- [21] S. Wasserman and K. Faust. *Social Networks Analysis: Methods and Applications*. Cambridge: Cambridge University Press., 1994.