This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Ordozgoiti, Bruno; Matakos, Antonis; Gionis, Aristides
Finding large balanced subgraphs in signed networks

Published in:
The Web Conference 2020 - Proceedings of the World Wide Web Conference, WWW 2020

DOI:
[10.1145/3366423.3380212](https://doi.org/10.1145/3366423.3380212)

Published: 20/04/2020

Document Version
Publisher's PDF, also known as Version of record

Published under the following license:
CC BY

Please cite the original version:
Ordozgoiti, B., Matakos, A., & Gionis, A. (2020). Finding large balanced subgraphs in signed networks. In *The Web Conference 2020 - Proceedings of the World Wide Web Conference, WWW 2020* (pp. 1378-1388). ACM.
<https://doi.org/10.1145/3366423.3380212>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

Finding large balanced subgraphs in signed networks

Bruno Ordozgoiti
Aalto University
bruno.ordozgoiti@aalto.fi

Antonis Matakos
Aalto University
antonis.matakos@aalto.fi

Aristides Gionis*
KTH Royal Institute of Technology
argioni@kth.se

ABSTRACT

Signed networks are graphs whose edges are labelled with either a positive or a negative sign, and can be used to capture nuances in interactions that are missed by their unsigned counterparts. The concept of balance in signed graph theory determines whether a network can be partitioned into two perfectly opposing subsets, and is therefore useful for modelling phenomena such as the existence of polarized communities in social networks. While determining whether a graph is balanced is easy, finding a large balanced subgraph is hard. The few heuristics available in the literature for this purpose are either ineffective or non-scalable. In this paper we propose an efficient algorithm for finding large balanced subgraphs in signed networks. The algorithm relies on signed spectral theory and a novel bound for perturbations of the graph Laplacian. In a wide variety of experiments on real-world data we show that our algorithm can find balanced subgraphs much larger than those detected by existing methods, and in addition, it is faster. We test its scalability on graphs of up to 34 million edges.

CCS CONCEPTS

• Theory of computation → Graph algorithms analysis.

KEYWORDS

graph mining, signed graphs, dense subgraph, community detection

ACM Reference Format:

Bruno Ordozgoiti, Antonis Matakos, and Aristides Gionis. 2020. Finding large balanced subgraphs in signed networks. In *Proceedings of The Web Conference 2020 (WWW '20)*, April 20–24, 2020, Taipei, Taiwan. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3366423.3380212>

1 INTRODUCTION

Social-media platforms have taken hold as one of the main forms of communication in today's society. Despite having served to facilitate connections between individuals, in recent years we have observed an array of negative phenomena associated to these technologies. Among other, these platforms seem to contribute to the polarization of political deliberation, which can be detrimental to the health of democracy. Thus, the study of methods to detect and mitigate polarization in online debates is becoming an increasingly compelling topic [14, 15, 26, 28, 29, 33].

*This work was done while the author was with Aalto University.

This paper is published under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW '20, April 20–24, 2020, Taipei, Taiwan

© 2020 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC-BY 4.0 License.

ACM ISBN 978-1-4503-7023-3/20/04.

<https://doi.org/10.1145/3366423.3380212>

Figure 1: The four possible signed triangles. The two on the left are balanced, while the two on the right are not.

Many social-media platforms can be represented by graphs. Thus, graph theory has found a variety of applications in this domain over the last few decades, such as community detection [13], partitioning [4], and recommendation [30]. One limitation of the graph representations usually employed in the literature is that they can capture the existence, or even the strength, of connections between vertices, but not their disposition. For instance, in a social network, vertices may represent people and edges interactions between them. By relying just on this information we cannot know whether each interaction is friendly or hostile.

Signed graphs can be used to overcome this limitation. In signed graphs, each edge is labeled with either a positive or negative sign. If a graph represents social interactions, signs can be employed to determine whether these interactions are friendly or not. Thus, signed graphs constitute a good representation for detecting polarized groups in online debates. Signed graphs were first introduced by Harary to study the concept of *balance* [17]. A signed graph is said to be balanced if its vertices can be partitioned into two sets in perfect agreement with the edge signs; that is, every edge within each set is positive and every edge between the two sets is negative. Equivalently, a signed graph is balanced when the product of the signs of every cycle is positive. This is analogous to the commonplace notion “the friend of a friend is a friend,” “the enemy of a friend is an enemy,” etc., as illustrated in Fig. 1.

A substantial body of work has been devoted to studying the spectral properties of signed graphs, which have strong connections to the concept of balance. In particular, the spectrum of the Laplacian matrix of a signed graph reveals whether it is balanced [23]. Graphs found in real applications are often not balanced, and therefore the question of finding a balanced subgraph arises naturally. The problem of finding a *maximum balanced subgraph* be formulated in terms of vertex cardinality (MBS) or edge cardinality (MBS-EDGE). Both formulations lead to NP-hard problems, thus, the development of efficient heuristics to approximately solve this problem is well motivated.

In this paper we present an algorithm to find large balanced subgraphs in signed networks. The algorithm works in two stages. First, we rely on spectral theory, as well as on a novel bound for perturbations of the Laplacian, to develop a greedy method to remove vertices and uncover a balanced subgraph. Then, any removed vertices that do not violate the balance of the located structure are restored. We derive analytical properties that allow us to efficiently

implement the algorithm. Finally, we devise a random sampling strategy to significantly enhance the scalability of the method.

In a variety of experiments on real-world and synthetic data, we show that the proposed algorithm finds larger balanced subgraphs than alternative heuristics from the literature, both in terms of vertex and edge count. Furthermore, the proposed algorithm has the desirable properties that (i) it runs faster than any other competing method, (ii) it can be tuned to trade off running time and solution quality, and (iii) produces a vertex-removal sequence, which can be used to trade off balance and graph size. We validate the scalability of the method by testing it on graphs of up to 34 million edges. Our contributions are summarized as follows:

- We propose an algorithm for finding large balanced subgraphs in signed graphs, based on spectral theory and its connections to balance.
- We give an upper bound for the smallest Laplacian eigenvalue after removing a set of vertices, which allows us to efficiently trim the graph to find a balanced subgraph.
- We experimentally show that our algorithm finds subgraphs much larger than those found by state-of-the-art methods.
- We devise a random sampling strategy to significantly enhance the scalability of the method, and show empirically that the quality of the output is not affected.

The rest of this paper is structured as follows. We discuss related work in Section 2. In Section 3 we introduce our notation and relevant notions. In Section 4 we describe our algorithm in detail and discuss relevant considerations, and in Section 5 we show our experimental results. Finally, Section 6 is a short conclusion.

2 RELATED WORK

Signed graphs and balance theory. Signed graphs were first studied by Harary, who was interested in the notion of *balance* [17]. In 1956, Cartwright and Harary generalized Heider’s psychological theory of balance in triangles of sentiments to the theory of balance in signed graphs [7]. Early work on signed graphs focused mainly on properties related to balance theory. For example, Harary and Kabell develop a simple linear-time algorithm to test whether a given signed graph satisfies the *balance property* [18]; while Akiyama et al. [1] study properties of the minimum number of sign changes required so that a signed graph satisfies the balance property.

A more recent line of work develops spectral properties of signed graphs, still related to the balance theory. Hou et al. [19] prove that a connected signed graph is balanced if and only if the smallest eigenvalue of the Laplacian is 0. Hou [20] also investigates the relationship between the smallest eigenvalue of the Laplacian and the unbalancedness of a signed graph.

Maximum balanced subgraphs. The problem studied in this paper is to find a maximum balanced subgraph (MBS) in a given signed graph. Poljak and Turzík [31] show that every connected signed graph with n vertices and m edges has a balanced subgraph with at least $\frac{m}{2} + \frac{n-1}{4}$ edges, and this bound is tight. They give an algorithm to find such a subgraph that requires at least $O(n^3)$ computations. Notice that this algorithm gives a 2-approximation for the MBS-EDGE problem, but it is not practical for large graphs. The MBS-EDGE problem can be formulated as a SIGNEDMAXCUT problem, which is

a generalization of the standard MAXCUT problem, and thus, NP-hard. To obtain an exact solution, the problem has been studied in the context of *fixed-parameter tractability* (FPT). Hüffner et al. [21] propose an FPT algorithm for deciding whether the maximum balanced subgraph has size at least $m - k$, where k is the parameter. Motivated by the lower bound of Poljak and Turzík, Crowston et al. [9] give an FPT algorithm for deciding whether the maximum balanced subgraph has at least $\frac{m}{2} + \frac{n-1}{4} + \frac{k}{4}$ edges, where k is the parameter. These algorithms are not practical, as their running time is exponential in k and the degree of the polynomial in n is large.

The MBS-EDGE problem has also been considered in application-driven studies, and different heuristics have been proposed. Das-Gupta et al. [10] consider an edge-deletion formulation of the MBS-EDGE problem in the context of biological networks. Motivated by the MAXCUT connection, the authors develop an algorithm based on semidefinite programming relaxation (SDP); the approach, however, is not scalable and tested only on very small networks. Figueiredo and Frota [12] ask to find a balanced subgraph that maximizes the number of vertices. They propose and experiment with a branch-and-cut exact approach, a heuristic based on minimum-spanning tree computation, and a heuristic combining a greedy algorithm and local-search. We experimentally compare the proposed method with these heuristics in our empirical evaluation.

Community detection in signed graphs. Different approaches have been proposed for community detection in signed graphs, some of which try to incorporate balance theory. Anchuri et al. [2] propose a spectral approach to partition a signed graph into a number of non-overlapping balanced communities. Yang et al. [34] propose a random-walk-based approach for partitioning a signed graph into communities, where in addition to positive edges within clusters and negative edges across clusters, they also want to maximize cluster densities. Doreian and Mrvar [11] propose an algorithm for partitioning a signed *directed* graph so as to minimize a measure of *imbalance*. The approach is evaluated only on very small networks. Signed directed graphs are also considered by Lo et al. [27], who search for strongly-connected positive subgraphs that are negative bi-cliques. Chu et al. [8] propose a constrained-programming objective to find k *warring factions*, as well as an efficient algorithm to find local optima. Bonchi et al. formulate the problem of finding subgraphs in signed networks that are dense but allow for imperfect balance [6]. Several other methods have been proposed for identifying communities in signed graphs, some of which incorporating notions related to balance. A detailed survey on those methods is provided by Tang et al. [32]. The main difference of our work in comparison with all these approaches is that they are mainly interested in communities or graph partitioning, which are different than the MBS problem.

3 PRELIMINARIES

Before describing the proposed algorithm, we introduce our notation and review some basic results from the literature.

We consider an undirected simple signed graph $G = (V, E^+, E^-)$ where $V = \{1, \dots, n\}$ is the set of vertices and E^+ (respectively, E^-) is the set of positive (respectively, negative) edges. We sometimes simplify this notation and write $G = (V, E)$, where $E = E^+ \cup E^-$. Throughout this paper we denote vectors with boldface letters ($\mathbf{v}$)

and matrices with uppercase letters (A). We use $\mathbf{v}_i$ to denote the i -th entry of a vector $\mathbf{v}$, and A_{ij} to denote the element in the i -th row and j -th column of matrix A . Given a signed graph, we define its adjacency matrix A as follows: $A_{ij} = 1$ if $\{i, j\} \in E^+$, $A_{ij} = -1$ if $\{i, j\} \in E^-$ and 0 otherwise. Further, we define the diagonal degree matrix D as $D_{ii} = d(i)$, where $d(i)$ is the degree of vertex i , i.e., the number of edges (either positive or negative) adjacent to i . The *signed Laplacian* of G is defined as $L(G) = D - A$. We also refer to this matrix simply as *Laplacian*, and will denote $L = L(G)$ when there is no ambiguity. Given a set of vertices S such that $S \subseteq V$, $G \setminus S$ denotes the graph that results from removing from G the vertices in S , as well as all adjacent edges.

We now define the concept of *balance in signed networks*, which is central to our paper.

Definition 3.1 (Balanced graph). Given a connected signed graph $G = (V, E^+, E^-)$, G is balanced if there exists a partition $V = V_1 \cup V_2$, $V_1 \cap V_2 = \emptyset$ such that every edge with both endpoints in V_1 is positive, every edge with both endpoints in V_2 is positive, and every edge with one endpoint in V_1 and the other in V_2 is negative.

In other words, a graph is balanced if we can divide it into two sets in a way that every edge sign agrees with the partition. For instance, if the vertices in graph G represent users in a social network and the edges interactions between them (friendly or hostile, depending on the sign), a dense, balanced graph suggests that there are two polarized communities.

It is easy to decide whether a given signed graph is balanced. In addition to a simple combinatorial algorithm, there is also an interesting characterization of balanced graphs based on the spectrum of the signed Laplacian. This is shown by the following result, which is key in the derivation of our algorithm.

THEOREM 3.2 ([19]). Consider a connected signed graph $G = (V, E^+, E^-)$, with signed Laplacian L . Let $\lambda_1(L) \leq \dots \leq \lambda_n(L)$ be the eigenvalues of L . Then G is balanced if and only if $\lambda_1(L) = 0$.

The smallest eigenvalue of the Laplacian reveals not only whether a graph is balanced, but also how far it is from being balanced. This is established by the following result of Li and Li [25].

THEOREM 3.3 ([25]). Consider a connected signed graph $G = (V, E^+, E^-)$, with signed Laplacian L . Let $\lambda_1(L) \leq \dots \leq \lambda_n(L)$ be the eigenvalues of L . Then

$$\lambda_1(L) \leq \min_{G'} \{\lambda_n(L(G')) : V_{G'} \subseteq V, G \setminus V_{G'} \text{ is balanced}\}.$$

Here, $V_{G'}$ denotes the set of vertices of graph G' . Intuitively, Theorem 3.3 says that if we can make G balanced with just a small modification, then $\lambda_1(L)$ is small. Note that $\lambda_n(L(G')) \leq 2\Delta(G')$, where $\Delta(G')$ denotes the maximum degree of the vertices of G' .

Graphs found in practical applications are usually not balanced. The question that arises naturally is thus whether we can find the maximum balanced subgraph (MBS) of a given signed graph efficiently. We formalize this objective next.

PROBLEM 1 (MBS). Given a signed graph $G = (V, E)$, find the graph G' induced by $V' \subseteq V$ such that G' is balanced and the cardinality of V' is maximized.

A solution to Problem 1 would reveal the frustration number, that is, the minimum number of vertices to remove to make the graph

Algorithm 1 TIMBAL

Input: signed graph G

```

1:  $R \leftarrow \emptyset$ 
2: Optionally:  $R \leftarrow \text{SUBSAMPLE}(G)$ ,  $G \leftarrow G \setminus R$ 
3: while  $G$  is not balanced do
4:   Compute  $L$ , the signed Laplacian of  $G$ .
5:   Compute bound vector  $\mathbf{r}$ .
6:   Compute the set of vertices to remove  $S$ , based on  $\mathbf{r}$ .
7:    $G \leftarrow G \setminus S$ ;  $R \leftarrow R \cup S$ .
8:    $G \leftarrow$  largest connected component in  $G$ .
9: end while
10: for  $v \in R$  do
11:   if  $G \cup \{v\}$  is balanced then
12:      $G \leftarrow G \cup \{v\}$ 
13:   end if
14: end for
15: Output  $G$ 

```

balanced, and is thus NP-hard [36]. In this paper we approach this problem based on Theorems 3.2 and 3.3. In particular, we address the following question: what vertices should we remove from G so that the minimum eigenvalue of the resulting graph is as small as possible? This question inspires the algorithm to find balanced subgraphs described in the next section.

4 ALGORITHM

Our algorithm works in two stages. First, it greedily removes vertices from the graph to improve balance as much as possible, until it obtains a balanced subgraph. Then, it does a single pass over the removed vertices and restores the ones that do not violate balance. In this section we describe these two stages in detail, as well as several optimizations and a procedure to enable the processing of huge graphs. Throughout this section, we assume the input graph to be connected.

The procedure, which we dub TIMBAL (Trimming Iteratively to Maximize Balance), is summarized in Algorithm 1. This section explains each of its steps in detail.

4.1 First stage: removing vertices

In the first stage of the algorithm, we iteratively remove vertices from the graph. The challenge is to determine which vertices to remove at each step. Our criterion for selecting vertices to remove is based on Theorem 3.2, and more precisely on Theorem 3.3. In particular, the smallest eigenvalue of the signed Laplacian measures how far the graph is from being balanced.

Given a graph G with signed Laplacian L , define $L^{(i)}$ to be the signed Laplacian of $G \setminus \{i\}$, that is, of the graph resulting from removing vertex i . We want to find the vertex that minimizes the smallest eigenvalue of the resulting Laplacian, that is, we want to find the vertex j such that

$$j = \arg \min_i \lambda_1(L^{(i)}). \quad (1)$$

Naturally, computing $\lambda_1(L^{(i)})$ for every vertex i using the spectral decomposition of the modified Laplacian $L^{(i)}$ is costly. To overcome this challenge, we present the following result, which gives an upper bound on the smallest eigenvalue of the perturbed Laplacian.

LEMMA 4.1. *Given a graph G with signed Laplacian L , let $\lambda_1(L)$ be the smallest eigenvalue of L and $\mathbf{v}$ an eigenvector of L satisfying $L\mathbf{v} = \lambda_1(L)\mathbf{v}$. Then*

$$\lambda_1(L^{(i)}) \leq \frac{\lambda_1(L)(1 - 2\mathbf{v}_i^2) - \sum_{j \in N(i)} \mathbf{v}_j^2 + \mathbf{v}_i^2 d(i)}{1 - \mathbf{v}_i^2}, \quad (2)$$

where $N(i)$ denotes the set of neighbours of i in G .

PROOF. We can obtain the matrix $L^{(i)}$ by applying the following operations to L : (1) for every j in $N(i)$, subtract 1 from L_{jj} ; (2) remove row and column i . Thus, if we define the vector $\hat{\mathbf{v}}$ to be equal to $\mathbf{v}$ after removing the i -th entry, it is

$$\hat{\mathbf{v}}^T L^{(i)} \hat{\mathbf{v}} = \mathbf{v}^T L \mathbf{v} - d(i) \mathbf{v}_i^2 - \sum_{j \in N(i)} \mathbf{v}_j^2 - 2\mathbf{v}_i \sum_{j \in N(i)} \mathbf{v}_j L_{ij}.$$

Now, observe that $\mathbf{v}^T L \mathbf{v} = \lambda_1(L)$ and $\sum_{j \in N(i)} \mathbf{v}_j L_{ij} = \lambda_1(L) \mathbf{v}_i - d(i) \mathbf{v}_i$. Therefore,

$$\hat{\mathbf{v}}^T L^{(i)} \hat{\mathbf{v}} = \lambda_1(L) - \left(\sum_{j \in N(i)} \mathbf{v}_j^2 + \mathbf{v}_i^2 (2\lambda_1(L) - d(i)) \right).$$

Since $\lambda_1(L^{(i)}) = \min_{\mathbf{x}} \frac{\mathbf{x}^T L^{(i)} \mathbf{x}}{\mathbf{x}^T \mathbf{x}}$, we have

$$\begin{aligned} \lambda_1(L^{(i)}) &\leq \frac{\hat{\mathbf{v}}^T L^{(i)} \hat{\mathbf{v}}}{\hat{\mathbf{v}}^T \hat{\mathbf{v}}} \\ &= \frac{\lambda_1(L) - \left(\sum_{j \in N(i)} \mathbf{v}_j^2 + \mathbf{v}_i^2 (2\lambda_1(L) - d(i)) \right)}{1 - \mathbf{v}_i^2}. \end{aligned}$$

Elementary computations yield the desired result. $\square$

We define the vector $\mathbf{r}$ whose entries are the values of the right-hand side of Inequality (2), for each i , i.e.,

$$\mathbf{r}_i = \frac{\lambda_1(L)(1 - 2\mathbf{v}_i^2) - \sum_{j \in N(i)} \mathbf{v}_j^2 + \mathbf{v}_i^2 d(i)}{1 - \mathbf{v}_i^2}. \quad (3)$$

In order to choose which vertex to remove from the graph, we can therefore take the one minimizing $\mathbf{r}_i$. The first stage of our algorithm removes vertices according to this criterion — as shown in Section 4.4, we can remove several vertices at once — until a balanced subgraph is found.

4.2 Second stage: restoring vertices

Once we have found a balanced induced subgraph $G' = (V', E')$, it is straightforward to obtain the corresponding partition $V_1 \cup V_2 = V'$ that agrees perfectly with the edge signs, e.g., by performing a breadth-first search. At this point, we can inspect the vertices removed in the first stage to see if some of them agree with the obtained partition and they can thus be reinserted into the graph. That is, if our input graph is $G = (V, E)$, we consider the set of removed vertices $R = V \setminus V'$. For every $v \in R$, if adding v to either V_1 or V_2 results in a balanced graph, we add it back to G' — restoring as well its edges with endpoints in G' — and proceed. We inspect this set of vertices in the order they were removed from the graph in the first stage of the algorithm.

In the remainder of this section we discuss the different considerations that must be taken into account to implement our algorithm.

4.3 Computing the bound efficiently

A key advantage of the bound from Lemma 4.1 is that it can be computed efficiently. Given a graph G , we define $\bar{L}$ to be the matrix whose entries are the absolute values of the entries of $L(G)$.

Define the matrix $W = \bar{L} + 2\lambda_1(L)I$ and $\mathbf{w} = \mathbf{v} \circ \mathbf{v}$, where $\circ$ denotes the element-wise product of two vectors. Then

$$\mathbf{r}_i = \frac{(\lambda_1(L)\mathbf{1} - W\mathbf{w})_i}{1 - \mathbf{v}_i^2}. \quad (4)$$

That is, the computation of the vector $\mathbf{r}$ reduces essentially to a matrix-vector multiplication operation.

4.4 Removing several vertices at once

Every time we remove a vertex, we need to compute the smallest eigenvalue and corresponding eigenvector of the updated Laplacian. Even though this can be done efficiently (see Section 4.5), when dealing with large graphs the overall computation time may be too high. Therefore, it might be desirable to remove several vertices at once, before updating the eigenpair, in order to find a balanced subgraph more quickly.

The most straightforward way to accomplish this batch operation is to simply consider the k smallest entries of $\mathbf{r}$ and remove the corresponding k vertices. However, we argue that this might have undesirable consequences. Consider the graph on the left of Figure 2. Removing either vertex 1 or 2 will make the graph balanced. Thus, both $\mathbf{r}_1$ and $\mathbf{r}_2$ — where $\mathbf{r}$ is the ranking vector defined in Equation (3) — are bound to be equally small. Removing these two vertices at the same time will result in a subgraph of size 2, but we could have obtained a balanced subgraph of size 3 by removing only one vertex.

To partially alleviate this shortcoming, we propose considering independent — i.e., non-neighbouring — vertices for simultaneous removal only. Formally, consider we want to remove k vertices, and assume we have so far chosen $q < k$ of these, to form the set R . Then, the next chosen vertex is defined as

$$\arg \min_{i \notin \bigcup_{j \in R} N(j)} \mathbf{r}_i.$$

Choosing independent vertices has the additional advantage that the bound given in Lemma 4.1 is additive in the following sense.

LEMMA 4.2. *Given a graph G with signed Laplacian L , let $\lambda_1(L)$ be the smallest eigenvalue of L and $\mathbf{v}$ an eigenvector of L satisfying $L\mathbf{v} = \lambda_1(L)\mathbf{v}$. Let R be a set of independent vertices in G and let L^R be the signed Laplacian of $G \setminus R$. Then*

$$\lambda_1(L^R) \leq \frac{\sum_{i \in R} \lambda_1(L)(1 - 2\mathbf{v}_i^2) - \sum_{j \in N(i)} \mathbf{v}_j^2 + \mathbf{v}_i^2 d(i)}{1 - \sum_{i \in R} \mathbf{v}_i^2}. \quad (5)$$

The lemma is easily verified by similar techniques as employed in the proof of Lemma 4.1.

Furthermore, the value of this upper bound can be tracked as we add vertices to the set R to decide how many of them to remove. Note that as the set R grows, the upper bound becomes less reliable (as the perturbation of L is more significant). At some point, the magnitude of the denominator will become too small and the bound will decrease very slowly, or even increase. This can be used as a criterion to choose a cut-off point for the removal.

Figure 2: Illustration of why we must only consider independent vertices for simultaneous removal. Solid edges are positive, while dashed ones are negative. Removing either 1 and 2 will make the graph balanced, so both will get an equal value in the ranking. However, we only need to remove one.

Note that this does not completely solve the problem of adequately choosing a vertex set for simultaneous removal. Consider, for instance, a cycle graph with arbitrary signs. Removing any vertex results in a balanced subgraph, since every tree is balanced [35]. However, simply discarding neighbouring vertices is not sufficient to limit the number of removed vertices to one in this case. In general, determining this set to optimality might constitute a hard problem in and of itself, and is therefore left for future work. Nevertheless, in our experiments we show that discarding neighbours provides good results in practice.

4.5 Updating the eigenpair

A remaining concern is the computation of the eigenvalue $\lambda_1(L^{(i)})$ and the corresponding eigenvector each time we remove a set of vertices. We propose two alternatives for this purpose.

Locally-optimal block preconditioned conjugate gradient method: Since the smallest eigenvalue of the Laplacian is in the “flat” part of the spectrum, that is, where consecutive eigenvalues are close to each other, the methods usually employed to compute eigenvalue decompositions can be slow to converge when dealing with large matrices. To speed up the process, we use the method described by Knyazev [22] to estimate the desired eigenpair. Since our goal is to rank the vertices according to the corresponding upper bound from Lemma 4.1, an approximation of $\lambda_1(L^{(i)})$ and the corresponding eigenvector is enough. Our experimental results support this claim.

Low-rank updates of the eigenvalue decomposition: An alternative to the use of the conjugate gradient method is to rely on known results concerning low-rank updates of the spectral decomposition. In particular, consider a matrix L with eigenvalue decomposition $L = Q\Lambda Q^T$. Now consider a matrix $\tilde{L} = L + W$, where W is a rank- ρ matrix satisfying $W = VV^T$, i.e., positive semi-definite. It is well known [3] that an eigenvalue λ of $\tilde{L}$ — not in the spectrum of L — makes the following expression equal to zero:

$$\det(I - V^T(\lambda - L)^{-1}V), \quad (6)$$

where $\det$ denotes the determinant of a matrix. Elementary operations yield the following, equivalent expression:

$$\det(I - U^T(\lambda - \Lambda)^{-1}U), \quad (7)$$

where $U = Q^T V$. Since $(\lambda - \Lambda)^{-1}$ is diagonal, its inversion is cheap. Moreover, the number of eigenvalues of $\tilde{L}$ below any real number

can be inferred exactly. That is, if one can afford to compute the eigenvalue decomposition of L and the rank ρ of the perturbation W is small, the eigenvalue decomposition of $\tilde{L}$ can be computed efficiently by means of a bisection algorithm on Expression (7).

It only remains to show how to compute the matrix V , so that we can construct U . We now show that if the vertices to remove are chosen carefully, the entries of V depend only on the degree of the removed vertices, and thus can be permanently stored and queried on execution, instead of being computed at each iteration.

First, consider a graph G with signed Laplacian L . We remove a single vertex and obtain a new graph with Laplacian $L^{(i)} = L - W$. It is easily verified that W — if we exclude zero rows and columns — is the Laplacian of a star graph, that is, a connected k -tree with $k - 1$ leaves. Now, consider we remove a set R of vertices, to obtain $L^R = L - W$, satisfying the following condition:

$$\text{for all } i \neq j \in R, N(i) \cap N(j) = \emptyset. \quad (8)$$

Then the matrix W can be permuted so that it is block-diagonal, with each block corresponding to the Laplacian of a star graph. The following result establishes that its eigenvalues are easily inferred from the size of each block.

LEMMA 4.3. *Let G be a signed star graph comprised of k vertices, with signed Laplacian L . Then the eigenvalues of L are*

- 0 with multiplicity 1,
- k with multiplicity 1 and
- 1 with multiplicity $k - 2$.

PROOF. First, since every tree is balanced [35], from Theorem 3.2 we know that 0 is an eigenvalue of L .

To infer the rest of the eigenvalues, consider the structure of L : it is $L_{11} = k - 1$, and $L_{ii} = 1, L_{1i} = L_{i1} = \pm 1$ for all $i > 1$. Thus, the vector $\mathbf{v} = (x, \gamma_2, \dots, \gamma_k)$ where x is an arbitrary real number and $\gamma_i = x/(k - 1) \times \text{sign}(L_{i1})$ is an eigenvector of eigenvalue k . That 1 is an eigenvalue with multiplicity $k - 2$ is easily verified from the fact that the rank of $L - I$ is 2. $\square$

Given the structure of L , once the eigenvalues are known, the eigenvectors can be efficiently computed. Thus, the eigenvalue decomposition of W need not be computed explicitly. The eigenvectors only need to be computed once for each value of k , and can then be reused whenever a vertex of degree k is removed.

To summarize, if we know the spectral decomposition of L and choose to remove vertices satisfying Condition (8), we can easily make use of Expression (7) to update the eigenpair.

4.6 Handling various connected components

During the execution of the first stage — see Section 4.1 — after removing the chosen vertices the graph might become disconnected. In this case, we can consider various alternatives. If among the resulting connected components only one is large enough, we can discard the rest. If various connected components are large enough to warrant further analysis, we can simply apply the algorithm recursively on each of them.

Nevertheless, in our experiments we observed that the connected components resulting from the removal, except from the largest one, are small, comprised of a handful of vertices in almost every case. This is consistent with the principle behind the algorithm —

i.e., the reduction of the smallest eigenvalue of the Laplacian. Note, for instance, that graphs comprised of two vertices are always balanced. Therefore, if removing a vertex results in a two-vertex structure becoming disconnected, the obtained graph will have a smallest eigenvalue equal to zero, and this vertex will thus be a good choice according to the criterion defined by (1). Notice that this is not undesirable behaviour: if a small subgraph becomes disconnected by removing one (or a few) vertices, then it is not densely connected to the rest of the graph and would therefore not contribute significantly to the density of the balanced subgraph found by the algorithm.

4.7 Scaling to big graphs

The proposed algorithm provides an efficient method to detect vertices to remove from a graph in order to improve its balance. The method, however, requires the estimation of an eigenvalue-eigenvector pair from the “flat” part of the spectrum, where algorithms for this purpose take longer to converge. This operation is linear in the number of entries of the adjacency matrix, that is, potentially quadratic in the number of vertices. Thus, the analysis of a graph comprised of millions of nodes can remain impractical.

To alleviate this shortcoming we propose a randomized preprocessing algorithm. Given the quadratic complexity of our algorithm’s iterations, processing a large number of subgraphs independently and then combining the results can result in significant time savings. To justify the approach, we rely on the following, easily verified, statement:

PROPOSITION 4.4. *Let G be a signed graph. G is balanced if and only if all of its subgraphs are balanced.*

The statement guarantees, on one hand, that by balancing random subgraphs we will never remove vertices from a balanced substructure of G , and therefore the process is safe in this regard. On the other hand, it implies that we cannot obtain a balanced subgraph of G before all of its subgraphs are balanced. Thus, any unbalanced subgraph we sample needs to become balanced in G . Based on these facts, we propose the preprocessing algorithm summarized as Algorithm 2. The algorithm randomly samples a large number of connected subgraphs and then runs Algorithm 1 on them. The set of vertices that have been removed from at least one of these subgraphs is then removed from the main graph, which is then processed normally.

To sample subgraphs, we take the following approach. We sample a vertex uniformly at random, and then perform a randomized breadth-first search (RBFS). RBFS works as follows: we first take all neighbours of the sampled vertex. Then, at each step of the search, we only take a random fraction of the corresponding vertex’s neighbours. The size of the fraction and the depth of the search are set by the user to achieve subgraphs of a certain size. This sampling strategy can produce dense subgraphs very efficiently.

4.8 Complexity analysis

The running time of our algorithm is dominated by the following operations: the computation of the smallest eigenvalue and corresponding eigenvector of the Laplacian, done essentially in $O(|E|)$ operations; the computation of the bounds in Eq. (2), which can be done in time $O(|V|^2)$, as shown in Eq. (4); the test of balance

Algorithm 2 SUBSAMPLE

Input: signed graph G

- 1: Sample s subgraphs $G_1, \dots, G_s$ of size k from G .
 - 2: **for** $i = 1 \dots s$ **do**
 - 3: Run first stage of TIMBAL on G_i , resulting in set of vertices to remove R_i .
 - 4: **end for**
 - 5: Output $\bigcup_i R_i$.
-

and location of connected components, a breadth-first search done in $O(|V| + |E|)$ time operations; the vertex restoration phase can be implemented using two inner product operations per vertex to check compliance with the balanced subgraph, which results $O(|V|\Delta)$ operations for the whole stage, where Δ is the maximum degree over all graph vertices. All these steps can be implemented to exploit the efficiency of sparse matrix operations.

The optional SUBSAMPLE procedure computes the first stage of the algorithm for each sampled subgraph, whose size can be controlled by limiting the number of sampled neighbours in the RBFS step. Our experimental results shows that this approach remains effective even when the size of the sampled subgraphs is small, and their number moderate.

5 EXPERIMENTS

This section presents the evaluation of the proposed algorithm. Our main purpose is to determine the effectiveness of our method in finding large balanced subgraphs. In particular, we evaluate the following aspects:

- We measure the size of the subgraphs found by our method, in both vertex and edge cardinality.
- We assess whether the vertex removal sequence produced by our method can be exploited to trade off solution quality and size, where by quality we mean that we allow our method to return not perfectly balanced subgraphs.
- We measure the running time of our implementation.
- We visualize some of the obtained results.

For our experimental evaluation, we use a variety of real-world data, which we briefly describe below.

Datasets. We select publicly-available real-world signed networks, whose main characteristics are summarized in Table 1. HighlandTribes¹ represents the alliance structure of the Gahuku–Gama tribes of New Guinea. Cloister¹ contains the esteem/disesteem relations of monks living in a cloister in New England (USA). Congress¹ reports (un)/favorable mentions of politicians speaking in the US Congress. Bitcoin² and Epinions² are who-trusts-whom networks of the users of Bitcoin OTC and Epinions, respectively. WikiElec-tions¹ includes the votes about admin elections of the users of the English Wikipedia. Referendum [24] records Twitter data about the 2016 Italian Referendum: an interaction is negative if two users are classified with different stances, and is positive otherwise. Slash-dot² contains friend/foe links between the users of Slashdot. The edges of WikiConflict² represent positive and negative edit conflicts between the users of the English Wikipedia. WikiPolitics¹

¹konekt.cc

²snap.stanford.edu

Table 1: Signed networks used in the experiments: number of vertices and edges; ratio of negative edges ($\rho_- = \frac{|E_-|}{|E_+ \cup E_-|}$); and ratio of non-zero elements of A ($\delta = \frac{2|E_+ \cup E_-|}{|V|(|V|-1)}$).

Real-world datasets	$ V $	$ E_+ \cup E_- $	ρ_-	δ
HIGHLANDTRIBES	16	58	0.50	0.48
CLOISTER	18	125	0.55	0.81
CONGRESS	219	521	0.20	0.02
BITCOIN	5 k	21 k	0.15	$1.2e-03$
WIKIELECTIONS	7 k	100 k	0.22	$3.9e-03$
TWITTERREFERENDUM	10 k	251 k	0.05	$4.2e-03$
SLASHDOT	82 k	500 k	0.23	$1.4e-04$
WIKICONFLICT	116 k	2 M	0.62	$2.9e-04$
EPINIONS	131 k	711 k	0.17	$8.2e-05$
WIKIPOLITICS	138 k	715 k	0.12	$7.4e-05$
WIKICONFLICT-4	1.1 M	34.7 M	0.62	$3.4e-05$
EPINIONS-4	1.1 M	12.2 M	0.17	$9.5e-06$

represents interpreted interactions between the users of the English Wikipedia that have edited pages about politics.

5.1 Proposed baselines

We compare the results of our method to heuristics proposed in the literature for the MBS, as well as a non-trivial spectral baseline. We now describe these methods.

EIGEN: The spectral approach from [6]. We take the dominant eigenvector $\mathbf{v}$ of the adjacency matrix A of the input graph. For a given threshold $\tau \in \mathbb{R}_+$, we construct a vector $\mathbf{x}$ as follows: $\mathbf{x}_i = \text{sign}(\mathbf{v}_i)$ if $|\mathbf{v}_i| \geq \tau$, and $\mathbf{x}_i = 0$ otherwise. We try all possible values of τ , that is, all values in $\{|\mathbf{v}_i| : i = 1, \dots, n\}$, where n is the number of nodes in the input graph. Note that $\mathbf{x}$ defines a partition of the graph into V_1, V_2 , such that a vertex i is in V_1 if and only if $\mathbf{v}_i = 1$, and $i \in V_2$ if and only if $\mathbf{v}_i = -1$. If the graph is balanced, $\mathbf{v}$ reveals the perfect partition.

GRASP: The GRASP heuristic proposed by Figueiredo and Frota [12]. The method consists of a construction phase, which greedily builds a balanced partition inspecting the vertices one by one in random order, and a local-search phase. The local-search phase is computationally costly and yielded negligible improvements in our experiments if kept within reasonable running times. Therefore, we only report the results of the construction phase.

GGMZ: The heuristic proposed by Gülpinar et al. [16], which functions as follows. First, a minimum spanning tree of the input graph is computed. This tree is then switched — i.e., a subset of vertices is chosen and all cut edges change sign — so that all edges become positive. The same switch is applied to the entire graph and a set of vertices that are independent in the negative edge set is returned.

All baselines, as well as our algorithm, are implemented in Python3, using sparse matrix operations when possible.³

Other proposals in the literature, such as the work of Chu et al. [8], tackle similar problems, but they tend to find small or imbalanced subgraphs and therefore are not comparable to our methods.

5.2 TIMBAL implementation details

As explained in section 4.4, our algorithm can be tuned to decide how many vertices are removed at each iteration. Even though better results might be obtained by fine-tuning this value for each problem instance, in our experiments we take a simple approach to balance quality and efficiency. In the small datasets — HIGHLANDTRIBES, CLOISTER and CONGRESS — we remove one vertex at a time. In the rest of the datasets, we remove at most 100. Note that the set of removed vertices can be smaller if we cannot find 100 independent vertices.

In the first stage of the algorithm, whenever the removal of a vertex subset results in various connected components, we discard all but the largest. In all our experiments, all other components are very small, comprised of a handful of vertices at most, and are not worth analysing further.

In all experiments with large data sets ($\geq 82k$ vertices) we randomly sample 1000 subgraphs of approximately 200 vertices and process them as explained in section 4.7.

5.3 Finding balanced subgraphs

The main purpose of the algorithm described in this paper is to find large, preferably dense, balanced subgraphs in signed networks. Therefore, we first evaluate the ability of our method and the proposed baselines to accomplish this goal. We run all algorithms on all datasets, and measure the size of the obtained subgraph in both vertices ($|V|$) and edges ($|E|$). We report the results in Table 2. Since the algorithms are sensitive to the choices made in their initial stages, we run them ten times and report the maximum values achieved in both vertex and edge count. TIMBAL obtains larger results in all datasets, significantly so in most of them.

Although our algorithm finds balanced subgraphs of large size, we cannot know how far off we are from an optimal solution. To evaluate this, we do a further experiment where we create alternative versions of our datasets with planted balanced subgraphs of selected size. In order to do this, we make a randomly selected part of the graph balanced, by switching the edge signs accordingly, while randomizing the signs of all the other edges. As an additional performance test, we also generated some power-law graphs, using the Barabási-Albert model [5] and planted balanced subgraphs in them using the same procedure. The first power-law graph has 20 000 nodes and 59 991 edges ($m = 3$),⁴ while the second one has 20 000 and 79 984 edges ($m = 4$). We experiment with planted balanced subgraphs of approximately half the size of total graph. The results can be seen in Table 3. $|V_p|$ denotes the size of the planted balanced subgraph. POWERLAW-3 and POWERLAW-4 are Barabási-Albert graphs with $m = 3$ and $m = 4$ respectively. We also report the percentage of the planted subgraph that was recovered by the algorithm (notice that it may be bigger than the planted subgraph, due to the added random noise).

We observe that in all cases, our algorithm manages to return a balanced subgraph of size at least the size of the planted balanced subgraph, while the baselines fail to do so, for all datasets.

³Source code: <https://github.com/justbruno/finding-balanced-subgraphs>

⁴The parameter m here refers to the one used in the Barabási-Albert model.

Table 2: Largest balanced subgraph found by each method for each dataset

	HIGHLANDTRIBES		CLOISTER		CONGRESS		BITCOIN		TWITTERREFERENDUM	
method	V	E	V	E	V	E	V	E	V	E
TIMBAL	13	35	10	33	208	452	4 208	10 158	8 944	166 243
GRASP	10	18	6	11	115	145	2 167	3 686	5 425	49 105
GGMZ	10	21	5	7	153	238	1 388	1 683	2 501	2 821
EIGEN	12	37	8	27	11	16	7	17	132	6 140

	WIKIELECTIONS		SLASHDOT		WIKICONFLICT		WIKIPOLITICS		EPINIONS	
TIMBAL	3 786	18 550	42 205	96 460	48 136	356 204	63 252	218 360	62 010	169 894
GRASP	1 752	4 416	23 289	40 511	18 576	82 726	31 561	81 557	28 189	63 250
GGMZ	713	771	16 389	17 867	6 137	9 145	23 342	37 098	21 009	25 013
EIGEN	11	41	35	491	11	28	10	45	6	14

Table 3: Solution found by each method for each dataset, compared to the size of the planted balanced subgraph in the graph

method	WIKIELECTIONS		SLASHDOT		WIKICONFLICT		WIKIPOLITICS		EPINIONS		POWERLAW-3		POWERLAW-4	
$ V_p $	3500	%	41 000	%	58 000	%	69 000	%	65 500	%	10 000	%	10 000	%
TIMBAL	4 097	117%	65 963	160%	88 529	152%	123 367	178%	103 136	156%	11 491	114%	11 346	113%
GRASP	1 072	30%	9 640	23%	15 268	26%	11 253	16%	9 313	14%	4 858	48%	5 834	58%
GGMZ	1 952	55%	14 320	34%	12 671	21%	25 202	36%	16 944	25%	9 341	93%	9 344	93%

Figure 3: Edge agreement ratio of all subgraphs visited by TIMBAL and EIGEN.**Table 4: Mean running times in seconds for each algorithm on the larger datasets, with the corresponding variance reported in the brackets**

	SLASHDOT	WIKICONFLICT	WIKIPOLITICS	EPINIONS
TIMBAL	117 (5.85)	159 (36.35)	210 (33.94)	244 (31.11)
GRASP	59 (0.87)	105 (3.07)	154 (13.27)	116 (0.69)
GGMZ	318 (0.20)	461 (0.56)	528 (1.36)	670 (0.32)
EIGEN	14	82	25	48

5.4 Trading off balance and graph size

Even though our method is intended to find balanced subgraphs, in some applications a small number of mistakes might be acceptable if that means we can find a larger, denser subgraph. An advantage

of our algorithm is that the first stage produces a sequence of vertices to remove so that the graph becomes increasingly balanced. Therefore, we can inspect the subgraph obtained at every step of the process and keep the one that best suits our purposes.

To evaluate the extent to which we can benefit from this, we inspected all the graphs generated during the first stage of TIMBAL and measured their quality in terms of balance. We compared its performance to EIGEN, which is the only baseline that can produce a removal sequence.⁵ Figure 3 shows the results. For each visited graph, we indicate its size on the x axis and plot its corresponding *edge agreement ratio*, which we define as follows. Given a graph G

⁵GRASP inspects the vertices at a random order, and GGMZ finds an independent set, the complement of which does not have an intrinsic order.

Figure 4: Average degree of all subgraphs visited by TIMBAL and EIGEN.

Table 5: Running times in seconds for our algorithm on the artificially augmented datasets. We report the size of each augmented dataset ($|V|$, $|E|$), as well as the size of the solutions ($|V_S|$, $|E_S|$).

Dataset	$ V $	$ E $	Time (s)	$ V_S $	$ E_S $
WIKICONFLICT-1	233 434	6.1 M	260	67 962	718 455
WIKICONFLICT-2	350 151	10.1 M	431	75 024	759 280
WIKICONFLICT-3	583 585	18.3 M	798	99 506	808 804
WIKICONFLICT-4	1.05 M	34.7 M	2 059	152 789	964 446
EPINIONS-1	263 160	2.1 M	306	111 973	348 502
EPINIONS-2	394 740	3.6 M	590	153 419	481 378
EPINIONS-3	657 900	6.5 M	1 776	231 709	695 614
EPINIONS-4	1.1 M	12 M	5 628	385 478	1 081 607

Figure 5: Visualization of the result for CONGRESS. Solid blue edges are positive, while dashed, red ones are negative.

with adjacency matrix A , let $\mathbf{x}$ be the indicator vector obtained by taking the signs of the dominant eigenvector of A . Then the edge agreement ratio of G is

$$\frac{\mathbf{x}^T A \mathbf{x}}{\|A\|_F^2}, \quad (9)$$

where $\|A\|_F$ denotes the Frobenius norm. Thus, the edge agreement ratio quantifies the proportion of edges in the graph that agree with the eigenvector-based partition. In a balanced graph, this quantity is 1. Note that in most cases, TIMBAL provides significantly better subgraphs than EIGEN throughout the entire process. However,

TIMBAL is able to locate a large balanced subgraph at some step, whereas EIGEN cannot. Recall that all plots discussed in this section correspond to the first stage of TIMBAL.

Additionally, Figure 3 provides some insights about the behaviour of TIMBAL. For instance, in various cases the edge agreement ratio increases sharply at some point. This suggests that once the graph is close to balance, our method can quickly find which vertices to remove to achieve perfect balance.

Figure 4 shows analogous results, but plotting average degree instead of edge agreement ratio. As expected, EIGEN achieves higher density than TIMBAL. This can be explained by the fact that EIGEN can be seen as optimizing the following objective [6]:

$$\max_{\mathbf{x} \in \{-1, 0, 1\}^n} \frac{\mathbf{x}^T A \mathbf{x}}{\mathbf{x}^T \mathbf{x}}.$$

However, this comes at a noticeable cost in edge agreement ratio. TIMBAL, on the contrary, tries to achieve balance, which results in reduced density. Nevertheless, Figure 4 also illustrates how one can trade off balance for density in the different stages of the execution of TIMBAL. If desired, one can take the graph visited at some iteration of the first stage of the algorithm and then execute the second stage, adding the vertices that agree with the current best partition, even if perfect balance cannot be achieved.

5.5 Running times

We report the running times of the algorithms on the larger datasets in Table 4. The experiments are executed on a machine equipped with an Intel Xeon E5-2670 with 24 cores and 256 GB of RAM. We run each algorithm ten times on each dataset and report averages, as well as the variance of the running times. TIMBAL is executed removing 100 vertices at each iteration.

We only report results for the larger subgraphs, as the running times are similar for all algorithms on the rest. As expected, EIGEN is the fastest method, as its running time is dominated by the computation of the dominant eigenvector of the adjacency matrix of the input graph. However, its results in terms of balanced subgraph size are poor. Among the competitive methods, TIMBAL arguably achieves the best combination of quality and running time ratio.

Figure 6: Visualization of the result for BITCOIN. Solid blue edges are positive, while dashed, red ones are negative.

5.6 Scalability

In order to assess the scalability of our algorithm, we augment two of the larger datasets (i.e., WIKICONFLICT and EPINIONS) by artificially injecting vertices with a number of randomly-connected edges equal to the average degree of the original network, while maintaining ρ_- (i.e., the ratio of negative edges). The largest obtained datasets are comprised of about 1.1 million vertices and 34 million edges in the case of WIKICONFLICT, and 1.1 million vertices and 12 million edges in the case of EPINIONS.

We execute our algorithm on these datasets. For each dataset, we make five iterations and report the average running time, as well as the size of the best solution found during the iterations. To improve running times, we dynamically set the number of vertices to remove in the i -th iteration to $n_i/100$, where n_i is the size of the subgraph at iteration i . Notice that the graphs found by the algorithm are large. The results are shown in Table 5.

5.7 Examples

In order to gain insight on the results of our methods, we plot the discovered balanced subgraphs for two of the datasets. The graphs discussed in this section are those found in the first stage of TIMBAL.

CONGRESS: A balanced subgraph comprised of 208 vertices was found. The first subset in the partition had 95 vertices and 173 edges, while the second had 113 vertices and 199 edges. There are 80 edges between the two sets. Note that this dataset contains 218 vertices in total. This result reveals that the individuals represented by the vertices of this graph are very polarized, as they can be perfectly partitioned by removing just nine of them. The result is depicted in Figure 5.

BITCOIN: A balanced subgraph comprised of 3254 vertices was found. The first subset in the partition had 2986 vertices and 6504 edges, while the second had 268 vertices and 173 edges. There

are 807 edges between the two sets. In this case, notice that the second set is more densely connected to the other set than within itself. This reveals that the graph contains a large community of affine users towards which many users feel negatively. The result is depicted in Figure 6.

6 CONCLUSIONS AND FUTURE WORK

In this paper we presented a novel, efficient algorithm for finding large balanced subgraphs in signed networks. By relying on signed spectral theory and a novel bound for perturbations of the graph Laplacian, we derived an efficient implementation. Through a wide variety of experiments on real-world and synthetic data we showed that our method achieves better results, in shorter or comparable time, than state-of-the-art methods. We tested scalability on graphs of up to 34 M edges.

Our work leaves several open avenues of further inquiry. First of all, it would be interesting to study the problem of optimally choosing a constrained subset of vertices to decrease the smallest Laplacian eigenvalue as much as possible. To the best of our knowledge, this problem has not been considered before in the literature. Second, we would like to carry out a thorough analysis of the impact of the number of removed vertices on the quality of the results and running time of the algorithm. A more interesting question to answer is whether we can determine this number optimally at each iteration. Finally, can we further improve the scalability of our algorithm? How efficiently can we find a large balanced subgraph in massive networks?

ACKNOWLEDGMENTS

This work was supported by three Academy of Finland projects (286211, 313927, 317085), the EC H2020RIA project “SoBigData++” (871042), and the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by Knut and Alice Wallenberg Foundation.

REFERENCES

- [1] Jin Akiyama, David Avis, Vasek Chvátal, and Hiroshi Era. 1981. Balancing signed graphs. *Discrete Applied Mathematics* 3, 4 (1981), 227–233.
- [2] Pranay Anchuri and Malik Magdon-Ismael. 2012. Communities and balance in signed networks: A spectral approach. In *Proceedings of the 2012 International Conference on Advances in Social Networks Analysis and Mining (ASONAM 2012)*. IEEE Computer Society, 235–242.
- [3] Peter Arbenz and Gene H Golub. 1988. On the spectral decomposition of Hermitian matrices modified by low rank perturbations with applications. *SIAM J. Matrix Anal. Appl.* 9, 1 (1988), 40–58.
- [4] Sanjeev Arora, Satish Rao, and Umesh Vazirani. 2009. Expander flows, geometric embeddings and graph partitioning. *Journal of the ACM (JACM)* 56, 2 (2009), 5.
- [5] Albert-László Barabási and Réka Albert. 1999. Emergence of scaling in random networks. *Science* 286, 5439 (1999), 509–512.
- [6] Francesco Bonchi, Edoardo Galimberti, Aristides Gionis, Bruno Ordozgoiti, and Giancarlo Ruffo. 2019. Discovering polarized communities in signed networks. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. 961–970.
- [7] Dorwin Cartwright and Frank Harary. 1956. Structural balance: a generalization of Heider’s theory. *Psychological review* 63, 5 (1956), 277.
- [8] Lingyang Chu, Zhefeng Wang, Jian Pei, Jiannan Wang, Zijin Zhao, and Enhong Chen. 2016. Finding gangs in war from signed networks. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 1505–1514.
- [9] Robert Crowston, Gregory Gutin, Mark Jones, and Gabriele Muciaccia. 2013. Maximum balanced subgraph problem parameterized above lower bound. *Theoretical Computer Science* 513 (2013), 53–64.
- [10] Bhaskar DasGupta, German Andres Enciso, Eduardo Sontag, and Yi Zhang. 2007. Algorithmic and complexity results for decompositions of biological networks into monotone subsystems. *Biosystems* 90, 1 (2007), 161–178.
- [11] Patrick Doreian and Andrej Mrvar. 1996. A partitioning approach to structural balance. *Social networks* 18, 2 (1996), 149–168.
- [12] Rosa Figueiredo and Yuri Frota. 2014. The maximum balanced subgraph of a signed graph: Applications and solution approaches. *European Journal of Operational Research* 236, 2 (2014), 473–487.
- [13] Santo Fortunato. 2010. Community detection in graphs. *Physics reports* 486, 3–5 (2010), 75–174.
- [14] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. 2017. Reducing controversy by connecting opposing views. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*. ACM, 81–90.
- [15] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. 2018. Quantifying controversy on social media. *ACM Transactions on Social Computing* 1, 1 (2018), 3.
- [16] Nalán Gülpinar, Gregory Gutin, Gautam Mitra, and Alexey Zverovitch. 2004. Extracting pure network submatrices in linear programs using signed graphs. *Discrete Applied Mathematics* 137, 3 (2004), 359–372.
- [17] Frank Harary. 1953. On the notion of balance of a signed graph. *The Michigan Mathematical Journal* 2, 2 (1953), 143–146.
- [18] Frank Harary and Jerald A Kabell. 1980. A simple algorithm to detect balance in signed graphs. *Mathematical Social Sciences* 1, 1 (1980), 131–136.
- [19] Yaoping Hou, Jiongsheng Li, and Yongliang Pan. 2003. On the Laplacian eigenvalues of signed graphs. *Linear and Multilinear Algebra* 51, 1 (2003), 21–30.
- [20] Yao Ping Hou. 2005. Bounds for the least Laplacian eigenvalue of a signed graph. *Acta Mathematica Sinica* 21, 4 (2005), 955–960.
- [21] Falk Hüffner, Nadja Betzler, and Rolf Niedermeier. 2007. Optimal edge deletions for signed graph balancing. In *International Workshop on Experimental and Efficient Algorithms*. Springer, 297–310.
- [22] Andrew V Knyazev. 2001. Toward the optimal preconditioned eigensolver: Locally optimal block preconditioned conjugate gradient method. *SIAM journal on scientific computing* 23, 2 (2001), 517–541.
- [23] Jérôme Kunegis, Stephan Schmidt, Andreas Lommatzsch, Jürgen Lerner, Ernesto W De Luca, and Sahin Albayrak. 2010. Spectral analysis of signed graphs for clustering, prediction and visualization. In *Proceedings of the 2010 SIAM International Conference on Data Mining*. SIAM, 559–570.
- [24] Mirko Lai, Viviana Patti, Giancarlo Ruffo, and Paolo Rosso. 2018. Stance Evolution and Twitter Interactions in an Italian Political Debate. In *International Conference on Applications of Natural Language to Information Systems*. Springer, 15–27.
- [25] Hui Shu Li and Hong Hai Li. 2016. A note on the least (normalized) laplacian eigenvalue of signed graphs. *Tamkang Journal of Mathematics* 47, 3 (2016), 271–278.
- [26] Q Vera Liao and Wai-Tat Fu. 2014. Can you hear me now?: mitigating the echo chamber effect by source position indicators. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*. ACM, 184–196.
- [27] David Lo, Didi Surian, Kuan Zhang, and Ee-Peng Lim. 2011. Mining direct antagonistic communities in explicit trust networks. In *Proceedings of the 20th ACM international conference on Information and knowledge management*. ACM, 1013–1018.
- [28] Yelena Mejova, Amy X Zhang, Nicholas Diakopoulos, and Carlos Castillo. 2014. Controversy and sentiment in online news. *arXiv preprint arXiv:1409.8152* (2014).
- [29] AJ Morales, Javier Borondo, Juan Carlos Losada, and Rosa M Benito. 2015. Measuring political polarization: Twitter shows the two sides of Venezuela. *Chaos: An Interdisciplinary Journal of Nonlinear Science* 25, 3 (2015), 033114.
- [30] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1999. *The PageRank citation ranking: Bringing order to the web*. Technical Report. Stanford InfoLab.
- [31] Svatopluk Poljak and Daniel Turzík. 1986. A polynomial time heuristic for certain subgraph optimization problems with guaranteed worst case bound. *Discrete Mathematics* 58, 1 (1986), 99–104.
- [32] Jiliang Tang, Yi Chang, Charu Aggarwal, and Huan Liu. 2016. A survey of signed network mining in social media. *ACM Computing Surveys (CSUR)* 49, 3 (2016), 42.
- [33] VG Vinod Vydiswaran, ChengXiang Zhai, Dan Roth, and Peter Pirolli. 2015. Overcoming bias to learn about controversial topics. *Journal of the Association for Information Science and Technology* 66, 8 (2015), 1655–1672.
- [34] Bo Yang, William Cheung, and Jiming Liu. 2007. Community mining from signed social networks. *IEEE Transactions on Knowledge and Data Engineering* 19, 10 (2007), 1333–1348.
- [35] Thomas Zaslavsky. 1982. Signed graphs. *Discrete Applied Mathematics* 4, 1 (1982), 47–74.
- [36] Thomas Zaslavsky. 2012. A mathematical bibliography of signed and gain graphs and allied areas. *The Electronic Journal of Combinatorics* 1000 (2012), 8–6.