CONGRESSUS NUMERANTIUM 237 (2026) 93–98

Congressus Numerantium
www.combinatorialpress.com/cn

Combinatorial Press

Weakly negative circles and best clustering in signed graphs

Michael J. Gottstein*, Leila Parsaei-Majd and Thomas Zaslavsky

ABSTRACT

Clustering a signed graph means partitioning the vertices into clusters so that every positive edge, and no negative edge, is within a cluster. The obstruction to clustering is circles with exactly one negative edge ("weakly negative circles"). The correlation clustering problem is to cluster with the minimum number Q of edges that violate the clustering rule. A lower bound is w, the maximum number of edge-disjoint weakly negative circles. If every two such circles are edge disjoint, then $Q = w$. We characterize the signed graphs in which no two weakly negative circles share any edges. A corollary is a straightforward recognition algorithm for such signed graphs. An unsolved problem is to characterize the signed graphs with $Q = w$.

Keywords: signed graph clustering, correlation clustering, weakly negative circle

2020 Mathematics Subject Classification: Primary 05C22; Secondary 05C85.

We elucidate the structure of a class of signed graphs that meets a simple lower bound for inability to be fully clustered.

A signed graph $\Sigma = (\Gamma, \sigma)$ consists of a graph $\Gamma = (V, E)$ and a sign function $\sigma : E \to \{+, -\}$. (Our graphs may have multiple edges but not loops.) A standard question about a signed graph is whether it is clusterable, i.e., whether it has a partition of V such that every positive edge is inside a part of the partition and every negative edge extends between parts. Davis [2] introduced clustering and proved that Σ has a clustering if and only if no circle has exactly one negative edge. (A circle is a connected 2-regular graph, also known as a cycle, circuit, or polygon.) Much later, Bansal et al. [1] introduced

* Corresponding author.

Received 09 Feb 2026; Revised 05 Jun 2026; Accepted 12 Jun 2026; Published Online 17 Jun 2026.

DOI: 10.61091/cn237-06

correlation clustering for signed graphs that are not clusterable; this is the problem of finding an optimal clustering of V, which means a partition with the smallest possible number of edges that disagree with the definition of a clustering. That minimum number was introduced by Doreian and Mrvar [3]; we call it the "inclusterability index", $Q(\Sigma)$.

Call a circle C "all positive" if every edge in it is positive, and "weakly negative" if exactly one edge in it is negative. A clear lower bound for $Q(\Sigma)$ is $w(\Sigma)$, the maximum number of "non-overlapping" (i.e., pairwise edge-disjoint) weakly negative circles, since at least every such circle has to be eliminated for clustering to occur—although this bound is rarely attained. Parsaei-Majd [4] noted this bound and proved (her Lemma 2.5) that $Q(\Sigma) = w(\Sigma)$ if each pair of weakly negative circles is non-overlapping. We call a signed graph with that last property "weakly non-overlapping" and we give a complete structural characterization of weakly non-overlapping signed graphs. As a byproduct, both $Q(\Sigma)$ and an optimal clustering are readily computable

The problem of computing $Q(\Sigma)$ is NP-hard [1] and therefore thought to be intrinsically difficult. Most correlation clustering literature focuses on approximation algorithms for arbitrary graphs, whereas we derive an exact structural result: we characterize all weakly non-overlapping signed graphs. Thus, the paper complements approximation-based work by identifying a graph class for which the optimum clustering problem admits a precise combinatorial description rather than merely an approximation guarantee. This is the only such completely structured and perfectly computable case we know of, other than Davis's original characterization of clusterable signed graphs.

For completeness we begin with a short proof and a converse of Parsaei-Majd's result. Let $t(\Sigma) :=$ the total number of weakly negative circles.

Proposition 0.1. *The following properties of a signed graph Σ are equivalent:*

(i) $Q(\Sigma) = t(\Sigma)$.
(ii) $w(\Sigma) = t(\Sigma)$.
(iii) *No two weakly negative circles share a common edge.*

Proof. It is clear that $w \leq Q \leq t$, so (ii) implies (i), and that $w = t$ if and only if no two weakly negative circles overlap, (ii) $\iff$ (iii). Now, suppose (i) is true but (iii) is false; that is, there are weakly negative circles C and C' that have an edge e in common. By removing e and at least one edge from every other weakly negative circle we obtain a signed graph without weakly negative circles; but the number of deleted edges is less than t, contradicting the assumption (i). Thus, (i) implies (iii). □

It is easy to make $w < Q < t$. Consider K_4 with two nonadjacent edges signed negative. The weakly negative circles are the triangles, so we have $w = 1$, $Q = 2$ (delete the negative edges), and $t = 4$.

We digress slightly to note a property of some weakly negative circles in any signed graph. It generalizes an argument in the proof of the Theorem.

Lemma 0.2. *Suppose W is a weakly negative circle in Σ that is edge-disjoint from every other weakly negative circle. Then W has at most one vertex in common with each all-*

positive circle and each other weakly negative circle.

Proof. Let W have negative edge e^- and let C be another circle that is all positive or weakly negative and has two or more common vertices with W. The common vertices divide C into paths that are internally disjoint from W (and, if C is all positive, possibly also edges common to W and C). At least one such path P_{vw} is all positive. The union $W \cup P_{vw}$ is a theta graph consisting of two all-positive paths and one path with a single negative edge; these three paths form two weakly negative circles with the common edge e^-. □

Returning to the main topic, the description of weakly non-overlapping graphs is slightly complicated. Let $\Sigma^+ = (V, E^+)$ be the positive subgraph of Σ, where $E^+ = \sigma^{-1}(+)$ is the set of positive edges. A block of a graph or signed graph is a maximal inseparable subgraph, i.e., it has no cutpoints. A block is "trivial" if it is isomorphic to K_1 or K_2. A nontrivial block is a multiple edge or it is 2-connected, and every two edges are in a circle. An "isthmus" is an edge whose deletion separates one component into two. A "cactus" is a connected graph all of whose blocks are circles or isthmi. A "leaf" is a vertex of degree 1. A "theta graph" consists of three internally disjoint paths with the same two endpoints. The union of all trivial blocks in a graph Γ is a forest, which we call $F(\Gamma)$; this graph is a disjoint union of trees (which may have any positive number of vertices). The subgraph of Σ induced by a vertex set X is denoted by $\Sigma{:}X$, and similarly $\Gamma{:}X$.

Theorem 0.3. *In a signed graph Σ every two weakly negative circles are edge disjoint if and only if Σ has the following form:*

For each non-trivial block B of Σ^+, the induced subgraph $\Sigma{:}V(B)$ is all positive. For each tree component T of $F(\Sigma^+)$, $\Sigma{:}V(T)$ is a cactus. There are no negative edges with both ends in the same component of Σ^+ except for those in a cactus $\Sigma{:}V(T)$. There is no restriction on negative edges between components of Σ^+.

We note that there are no restrictions on negative edges between components of Σ^+. Also, the negative edges in the induced cactus $\Sigma{:}V(T)$ consist of exactly one in each circle block of the cactus, because the circle blocks in the cactus are the fundament circles of the negative edges with respect to the tree T.

Proof. First, we assume Σ has only edge-disjoint weakly negative circles and prove it has the described structure.

Consider a connected component Ψ of Σ^+. We observe that Ψ may have both trivial and nontrivial blocks. Those of its edges that are isthmi of Ψ are the edges in $F(\Sigma^+) \cap \Psi$; that is, $F(\Psi)$ is the disjoint union of the trees T of $F(\Sigma^+)$ that are contained in Ψ. The remaining edges of Ψ are contained in nontrivial blocks of Ψ.

Add to Ψ a negative edge vw. If there is more than one path in Ψ from v to w, then vw creates two weakly negative circles with a common edge. If there is only one such path, then v and w must belong to the same tree T of $F(\Sigma^+)$. Thus, any negative edge of Σ with both endpoints in Ψ has both endpoints in a tree T of $F(\Psi)$.

Suppose there are two negative edges, uv and wx, with endpoints in the same tree T.

The endpoints are joined by unique paths T_{uv} and T_{wx} in the tree, and both $T_{uv} \cup uv$ and $T_{wx} \cup wx$ are weakly negative circles. Therefore, there can be no common edge between T_{uv} and T_{wx}. This implies that T along with all negative edges having both endpoints in T must be a cactus. (This is essentially a direct proof of a special case of the Lemma.) Since a cactus has only edge-disjoint circles, such negative edges may exist, but there can be only one in each circle of the cactus.

We have shown the form of the negative edges with both endpoints in a single component of Σ^+. On the other hand, negative edges between components of Σ^+ can never form a weakly negative circle, so they are unrestricted. That completes the proof in this direction.

For the opposite direction, assume Σ has the described structure. A weakly negative circle must consist of an all-positive path together with one negative edge closing the path; that is, the negative edge has both ends in a component of Σ^+. By assumption, the only such negative edges are those with both ends in the same tree component T of $F(\Gamma)$. Let e be such an edge; then $T \cup e$ contains a single circle, which is a weakly negative circle in the induced subgraph $\Sigma{:}V(T)$. Two such circles in $\Sigma{:}V(T)$ intersect in at most a single vertex, because $\Sigma{:}V(T)$ is a cactus. That implies no two weakly negative circles in Σ have a common edge, completing the proof. □

A weakly non-overlapping signed graph has a disagreeing set that consists of the negative edges inside clusters. (The converse is not true; a signed graph with a minimum disagreeing set that consists of the negative edges within the clusters of a best clustering may have weakly negative circles that are not edge disjoint; simple examples are in [4, Figures 7(a, b)], reproduced in Figure 1.)

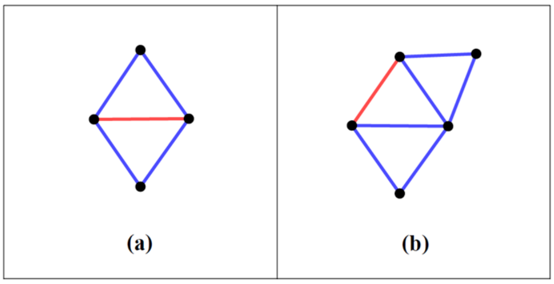

Fig. 1. Two counterexamples to a converse

Corollary 0.4. *In a weakly non-overlapping signed graph Σ, the vertex partition $\pi(\Sigma^+)$ induced by the connected components of Σ^+ is a best clustering. The value of $Q(\Sigma)$ is the number of negative edges inside clusters of $\pi(\Sigma^+)$.*

Proof. Each negative edge within a cluster belongs to a different weakly negative circle. □

The partition $\pi(\Sigma^+)$ may not be the only optimal clustering. For example, if Σ itself is a weakly negative circle, then $Q(\Sigma) = 1$ and $\pi(\Sigma^+)$ has only one cluster, but any partition of $V(\Sigma)$ into two parts, each of which induces an all-positive path, is also optimal.

Algorithm 0.5. Although finding an optimal clustering for general signed graphs is NP-hard, in our special case both $Q(\Sigma)$ and the optimal clustering $\pi(\Sigma^+)$ are readily computable. We estimate the computational complexity of testing Σ for having only edge-disjoint weakly negative circles, by testing for the structure in the Theorem. For simplicity we assume a simple graph Γ, whose order is $n = |V(\Sigma)|$ and with size $m = |E(\Sigma)|$. The overall time requirement is $O(m + n)$.

Step 1. Find the positive subgraph Σ^+. This requires checking the signs of the edges, which takes time $O(m)$.

Step 2. Find the connected components Ψ of the positive subgraph and construct the induced subgraphs $\Sigma{:}V(\Psi)$. Time $O(m + n)$.

Step 3. Find the set E_F of isthmi of each Ψ and identify the trees in each forest $F(\Psi) = (V(\Sigma), E_F)$. Time $O(m)$ [6].

Step 4. For each Ψ, find the negative edges within Ψ. Test each such negative edge to see if its endpoints are contained within a tree of $F(\Psi)$. If any is not, stop. Time $O(m + n)$ using BFS or DFS search.

Step 5. If not stopped, test each tree T of $F(\Sigma^+)$ with its induced negative edges to see if it is a cactus. If any such tree with negative edges is not a cactus, stop. This requires time $O(m)$ [7, Lemma 4].

Step 6. If (and only if) not stopped, success. Σ has no overlapping weakly negative circles and an optimal clustering is the partition induced by the components of Σ^+.

We conclude with a research problem that looks hard, but should be the next step.

Problem 0.6. *Characterize the signed graphs for which $Q(\Sigma) = w(\Sigma)$.*

Two of the authors have drafted a paper about this [5].

References

[1] N. Bansal, A. Blum, and S. Chawla. Correlation clustering. *Machine Learning*, 56(1–3):89–113, 2004. https://doi.org/10.1023/B:MACH.0000033116.57574.95.

[2] J. A. Davis. Clustering and structural balance in graphs. *Human Relations*, 20(2):181–187, 1967. https://doi.org/10.1177/001872676702000206.

[3] P. Doreian and A. Mrvar. A partitioning approach to structural balance. *Social Networks*, 18(2):149–168, 1996. https://doi.org/10.1016/0378-8733(95)00259-6.

[4] L. Parsaei-Majd. Correlation clustering for general graphs. *Discrete Mathematics*, 349(2):114747, 2026. https://doi.org/10.1016/j.disc.2025.114747.

[5] L. Parsaei-Majd and T. Zaslavsky. Correlation clustering and edge-disjoint weakly negative cycles. Submitted, 2026.

[6] R. E. Tarjan. A note on finding the bridges of a graph. *Information Processing Letters*, 2(6):160–161, 1974. https://doi.org/10.1016/0020-0190(74)90003-9.

[7] B. Zmazek and J. Žerovnik. Computing the weighted Wiener and Szeged number on weighted cactus graphs in linear time. *Croatica Chemica Acta*, 76(2):137–143, 2003.

Michael J. Gottstein
Mathematics and Computer Science Department, Marywood University, Scranton, PA 18509, U.S.A.
E-mail mjgottstein@maryu.marywood.edu

Leila Parsaei-Majd
University of Potsdam, Germany
E-mail leila.parsaeimajd@uni-potsdam.de

Thomas Zaslavsky
Department of Mathematics and Statistics, Binghamton University, Binghamton, NY 13902-6000, U.S.A.
E-mail zaslav@math.binghamton.edu