

Hierarchical Clustering and Consensus in Trust Networks

Santiago Segarra Alejandro Ribeiro
 Department of Electrical and Systems Engineering
 University of Pennsylvania

Abstract—We apply recent developments in clustering theory of asymmetric networks to study the equilibrium configurations of consensus dynamics in trust networks. We show that reciprocal clustering characterizes the equilibrium opinions of mutual trust dynamics. That is, clusters in the reciprocal dendrogram correspond to different equilibrium opinions of mutual trust consensus for varying trust thresholds. Moreover, for unidirectional trust dynamics, we show that aggregating nonreciprocal clusters into single nodes does not modify reachability of global consensus, thus, simplifying the consensus analysis of large networks.

I. INTRODUCTION

Consensus or agreement problems in networked systems have been extensively studied in the last decade with a wide range of applications [1]. Formation control [2] and flocking [3] among other problems have been approached through a consensus perspective. Furthermore, consensus dynamics have been used to model opinion propagation in social networks [4], [5]. In this context, agents update their own opinion by considering the opinion of their neighbors, i.e. a subset of the community which they trust.

In society, trust relations arise between individuals and can be modeled through a trust network. We all have an idea of how much we trust other members of our society like relatives, friends, or acquaintances. However, it is unclear if we should trust more the friend of a relative or a direct acquaintance. More generally, it is unclear who should we trust within a network or, equivalently, who belongs to our circle of trust. We use hierarchical clustering [6, Ch. 4] to model these concerns. In particular, we apply recent developments [7], [8] to hierarchically cluster asymmetric networks.

When sharing opinions in society, it is reasonable to filter whose opinions to take into account depending on the issue being discussed. E.g., for a simple and public issue, we might trust the opinion of a large set of agents in our community whereas when it comes to intimate and private matters we only rely on close friends. In this context, for every trust threshold we have different consensus dynamics and, thus, different equilibrium configurations. Our main contribution is the relation between the hierarchical clustering of a trust network and the consensus equilibria for different trust thresholds of this network. In particular, reciprocal clustering determines the equilibria when mutual trust is required for propagation and nonreciprocal clustering informs about the equilibria when unidirectional trust is enough for propagation.

II. PRELIMINARIES

Define a network $N = (X, A_X)$ as a set of n nodes X endowed with a real valued dissimilarity function $A_X : X \times X \rightarrow \mathbb{R}_+$ defined for all pairs of nodes $x, x' \in X$. Dissimilarities $A_X(x, x')$ are nonnegative for all $x, x' \in X$, and null if and only if $x = x'$, but need not satisfy the triangle inequality and may be asymmetric, i.e. $A_X(x, x') \neq A_X(x', x)$ for some $x, x' \in X$. When we hierarchically




Fig. 1. Reciprocal clustering. Nodes x, x' cluster at resolution δ if they can be joined with a bidirectional chain of maximum dissimilarity δ [cf. (2)].

cluster a network $N = (X, A_X)$, we obtain a dendrogram D_X , i.e. a nested set of partitions $D_X(\delta)$ indexed by the resolution parameter $\delta \geq 0$; see e.g. Fig. 6 and Fig. 7-(a). Partitions in a dendrogram D_X must satisfy two boundary conditions: for the resolution parameter $\delta = 0$ each node $x \in X$ forms a singleton cluster, i.e., $D_X(0) = \{\{x\}, x \in X\}$, and for some sufficiently large resolution δ_0 all nodes must belong to the same cluster, i.e., $D_X(\delta_0) = \{X\}$. Partitions being nested implies that if any two nodes $x, x' \in X$ are in the same cluster at a given resolution δ' then they stay co-clustered for all larger resolutions $\delta > \delta'$. If two nodes x and x' belong to the same cluster at resolution δ in dendrogram D_X then we write $x \sim_{D_X(\delta)} x'$. For a given dendrogram D_X , we denote by $u_X(x, x')$ the minimum resolution at which nodes x and x' are co-clustered, i.e.

$$u_X(x, x') := \min \{ \delta \geq 0, x \sim_{D_X(\delta)} x' \}. \quad (1)$$

A hierarchical clustering method is then a map $\mathcal{H} : \mathcal{N} \rightarrow \mathcal{D}$ mapping every network in \mathcal{N} to a dendrogram in \mathcal{D} . Two clustering methods of interest are reciprocal and nonreciprocal clustering [7]. The reciprocal clustering method \mathcal{H}^R with output dendrogram $D_X^R = \mathcal{H}^R(X, A_X)$ merges nodes x and x' at resolution $u_X^R(x, x')$ given by

$$u_X^R(x, x') := \min_{C(x, x')} \max_{i | x_i \in C(x, x')} \bar{A}_X(x_i, x_{i+1}), \quad (2)$$

where $\bar{A}_X(x, x') := \max(A_X(x, x'), A_X(x', x))$ for all $x, x' \in X$. Intuitively, in (2) we search for chains $C(x, x')$ linking nodes x and x' . Then, for a given chain, walk from x to x' and determine the maximum dissimilarity, in either the forward or backward direction, across all links in the chain. Then $u_X^R(x, x')$ is the minimum of this value across all possible chains; see Fig. 1.

Reciprocal clustering joins x to x' by going back and forth at maximum cost δ through the same chain. Nonreciprocal clustering \mathcal{H}^{NR} permits different chains. Define the minimum directed cost as

$$\tilde{u}_X^{NR}(x, x') := \min_{C(x, x')} \max_{i | x_i \in C(x, x')} A_X(x_i, x_{i+1}), \quad (3)$$

and the nonreciprocal merging resolution as the maximum of the two minimum directed costs from x to x' and x' to x

$$u_X^{NR}(x, x') := \max \left(\tilde{u}_X^{NR}(x, x'), \tilde{u}_X^{NR}(x', x) \right). \quad (4)$$

In (4) we implicitly consider forward chains $C(x, x')$ going from x to x' and backward chains $C(x', x)$ from x' to x . We then determine the respective maximum dissimilarities and search independently for the

Thanks to AFOSR MURI FA9550-10-1-0567 and NSF CCF-1217963.




Fig. 2. Nonreciprocal clustering. Nodes x, x' cluster at resolution δ if they can be joined in both directions with possibly different chains of maximum dissimilarity δ [cf. (4)].

forward and backward chains that minimize the respective maximum dissimilarities. The nonreciprocal merging resolution $u_X^{\text{NR}}(x, x')$ is the maximum of these two minimum values; see Fig. 2.

In the present paper we also consider directed, unweighted graphs $G = (X, E)$ on the node set X , where $E \subseteq X \times X$ is the edge set. By definition, E does not contain self-loops. For every node x , we define its neighborhood E_x as the subset of nodes where the edges starting at x end, i.e. $E_x := \{x' \in X | (x, x') \in E\}$. We also define the adjacency matrix $\mathcal{A}_G = [a_{ij}]$ of graph G as a binary matrix where $a_{ij} = 1$ if $(x_i, x_j) \in E$ and $a_{ij} = 0$ otherwise. The degree matrix $\Delta_G = [d_{ij}]$ of graph G is a diagonal matrix with $d_{ii} = |E_{x_i}|$ and $d_{ij} = 0$ for $i \neq j$. The Laplacian matrix L_G associated with graph G is then given by

$$L_G = \Delta_G - \mathcal{A}_G. \quad (5)$$

In this paper we focus on the continuous time consensus dynamics given by

$$\dot{p}(t) = -L_G p(t), \quad p(0) = p_0, \quad (6)$$

where $p(t) \in \mathbb{R}^n$ for all times t , with $p_i(t)$ describing the state – or opinion in our context – of node x_i at time t . In the dynamics given by (6), the change in a given node's opinion is determined by the average disagreement with its neighbors. I.e., if my opinion is equal to the average of my neighbors' opinion, then my opinion will remain unchanged for the next time instant. We say that global consensus is reached when every node converges to the same opinion.

III. TRUST NETWORKS AND CONSENSUS DYNAMICS

Define a trust network $N = (X, A_X)$ as one where nodes $x \in X$ represent agents, e.g. people in society, and $A_X(x, x')$ represents a measure of how much x distrusts x' , i.e. $A_X(x, x') < A_X(x, x'')$ implies that x trusts more in x' than in x'' . Notice that the function A_X is inherently asymmetric since trust relations between people need not be bidirectional. Indeed, it is usually the case in social networks that some influential agents, say celebrities or politicians, are heard by a big portion of the network but they do not take into account the opinions of all of their followers.

If we want to model opinion propagation through consensus in a trust network, one possibility is extending (5) and (6) for weighted graphs as considered in, e.g., [9]. In this case, an agent weighs the importance of others' opinions depending on how much he trusts them. However, we consider different dynamics where the neighborhood of an agent is a function of the issue being discussed in the network. E.g., if we need advice on where to have dinner we trust the opinion of a larger set of people than when we need advice on how to approach a relationship problem. Thus, a given trust network N has associated an infinite number of consensus problems indexed by a trust threshold parameter δ . In this way, for low values of δ we only listen to the opinion of our intimate circle of trust whereas for large values of δ we admit the opinion of more distant acquaintances. In this




Fig. 3. Graphs induced by the network N for different resolutions δ under the assumption of mutual trust.

paper, we consider two different ways of determining neighborhoods given a trust parameter δ : mutual and unidirectional trust.

A. Mutual trust consensus

Given a trust network $N = (X, A_X)$, in mutual trust consensus we require that for two agents x and x' to share their opinions they should distrust each other less than a given threshold δ . In other words, communication between agents only occurs when there is a minimum of trust in both directions. Hence, given a trust network N and a threshold δ we define the mutual trust communication graph $G_N(\delta) = (X, E)$ with adjacency matrix $\mathcal{A}_{G_N}(\delta)$ given by

$$[\mathcal{A}_{G_N}(\delta)]_{ij} = \begin{cases} 1 & \text{if } A_X(x_i, x_j) \leq \delta \text{ and } A_X(x_j, x_i) \leq \delta, \\ 0 & \text{otherwise,} \end{cases} \quad (7)$$

for all $i \neq j$ and $[\mathcal{A}_{G_N}(\delta)]_{ii} = 0$ for all i . To facilitate understanding refer to Fig. 3. Note that definition (7) is symmetric implying that the graph $G_N(\delta)$ is undirected for all $\delta \geq 0$.

From (5), we compute the Laplacian matrix $L_{G_N}(\delta)$ for every resolution δ and obtain the mutual trust consensus dynamics

$$\dot{p}(t) = -L_{G_N}(\delta) p(t), \quad p(0) = p_0, \quad (8)$$

which is just a specialization of (6) for a particular Laplacian matrix. Note that in (8) we actually have an infinite number of consensus problems indexed by the resolution parameter δ . We are interested in the equilibrium configuration $\lim_{t \rightarrow \infty} p(t)$ of the consensus, i.e. the opinions of the agents after a long time has elapsed. When $\delta = 0$, we obtain $\mathcal{A}_{G_N}(0) = L_{G_N}(0) = \mathbf{0}$ implying that there is no communication at all. In this situation, every node in the network preserves its original opinion through time. On the other hand, for a sufficiently large resolution $\delta = \delta_0$, the adjacency network is that of a complete graph and global consensus is achieved. For resolutions in between, we want to characterize local equilibrium configurations.

B. Unidirectional trust consensus

In unidirectional trust consensus, for a given agent to be influenced by the opinion of another, the first agent must trust the second one independently of the trust relation in the inverse direction. Given a trust network $N = (X, A_X)$ and a trust threshold δ we

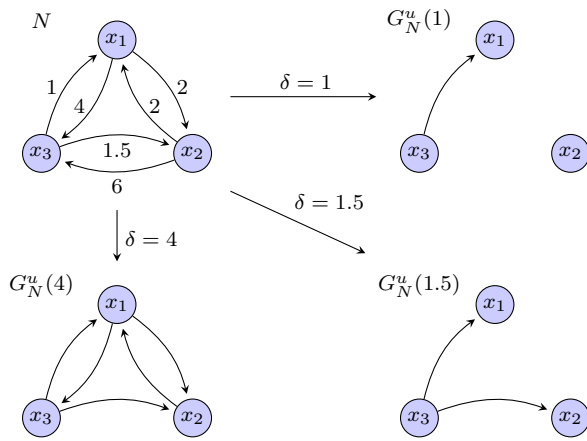


Fig. 4. Graphs induced by the network N for different resolutions δ under the assumption of unidirectional trust.

define the unidirectional communication graph $G_N^u(\delta) = (X, E)$ with adjacency matrix $\mathcal{A}_{G_N^u}(\delta)$ defined as

$$[\mathcal{A}_{G_N^u}(\delta)]_{ij} = \begin{cases} 1 & \text{if } A_X(x_i, x_j) \leq \delta, \\ 0 & \text{otherwise,} \end{cases} \quad (9)$$

for all $i \neq j$ and $[\mathcal{A}_{G_N^u}(\delta)]_{ii} = 0$ for all i ; see Fig. 4. Note that definition (9) is asymmetric entailing a directed graph $G_N^u(\delta)$. Also, edges in the graphs $G_N^u(\delta)$ denote trust relations and, hence, information flows in the opposite direction of the edges. E.g., in graph $G_N^u(1)$ in Fig. 4, x_3 trusts in x_1 implying that the opinion of x_1 influences that of x_3 . Similar to (8), we can define

$$\dot{p}(t) = -L_{G_N^u}(\delta)p(t), \quad p(0) = p_0, \quad (10)$$

which contains an infinite number of consensus problems indexed by the resolution parameter δ . The extremal results for $\delta = 0$ and $\delta = \delta_0$ sufficiently large guaranteeing $G_N^u(\delta_0)$ to be a complete graph, coincide with the mutual trust case in (8). As for mutual trust, we are interested in characterizing the equilibrium for intermediate resolutions.

IV. CIRCLES OF TRUST AND CONSENSUS EQUILIBRIA

We are all part of trust networks in our social lives, thus motivating the question: who should we trust? or equivalently, which nodes form our circle of trust? This question is in fact ill-posed as discussed in Section III since the issue being discussed would determine the extent of our circle of trust, i.e. an intimate matter determines a close circle whereas a trivial matter admits an extended circle of trust. Moreover, it is reasonable for the circles of trust to be nested in the sense that if you trust someone with a very intimate matter then you would trust that same person with a more trivial issue. Hence, dendrograms are natural representations for circles of trust in networks, where the resolution parameter δ denotes the level of intimacy of the issue in discussion, with lower δ denoting more intimate matters. Consequently, we can reinterpret hierarchical clustering methods \mathcal{H} as maps that assign a nested collection of circles of trust $\mathcal{H}(N)$ to every trust network N .

It is reasonable to expect agents in the same circle of trust to converge to the same opinion through consensus dynamics. Indeed, the reciprocal clustering method \mathcal{H}^R solves the mutual trust problem (8) as the following proposition asserts.




Fig. 5. A trust network N with five agents. Undrawn edges correspond to trust dissimilarities greater than 5.




Fig. 6. Reciprocal dendrogram of the network in Fig. 5. At resolution $\delta = 2.5$ clusters $\{x_2, x_5\}$, $\{x_3, x_4\}$ and $\{x_1\}$ are formed whereas at resolution $\delta = 3.5$ nodes are clustered into $\{x_2, x_5, x_3, x_4\}$ and $\{x_1\}$. According to Proposition 1, these clusters correspond to identical equilibrium opinions in the consensus problem (8).

Proposition 1 *In the mutual trust consensus dynamics (8) with parameter δ , for every initial condition p_0 ,*

$$\lim_{t \rightarrow \infty} p_i(t) = \lim_{t \rightarrow \infty} p_j(t) \iff u_X^R(x_i, x_j) \leq \delta, \quad (11)$$

where u_X^R is defined as in (2).

Proof: See [10]. ■

Proposition 1 implies that the reciprocal dendrogram of a given trust network contains information about opinion convergence for the infinite family of consensus problems in (8) indexed by δ . To explain this assertion, consider the five-node trust network N in Fig. 5. Using the algorithms developed in [8] we compute the reciprocal dendrogram $\mathcal{H}^R(N)$ and depict it in Fig. 6. If we want to obtain the equilibrium opinion profile for the consensus problem (8) for a given δ , we just perform a cut in the dendrogram at the desired resolution. The clusters at this resolution coincide with opinion clusters in equilibrium. In this way, from Fig. 6 we can observe that three opinion profiles arise when $\delta = 2.5$ while two profiles arise when $\delta = 3.5$.

For the case of the unidirectional trust consensus problem in (10), an equivalence result as the one found in Proposition 1 is impossible since the clusters of nodes converging to the same opinion are not nested. To see this, consider the graphs $G_N^u(1)$ and $G_N^u(1.5)$ induced by network N in Fig. 4. For $G_N^u(1)$, we have that x_3 listens to x_1 and will eventually converge to his opinion while x_2 preserves his original opinion through time. Thus, there are two opinions in equilibrium.

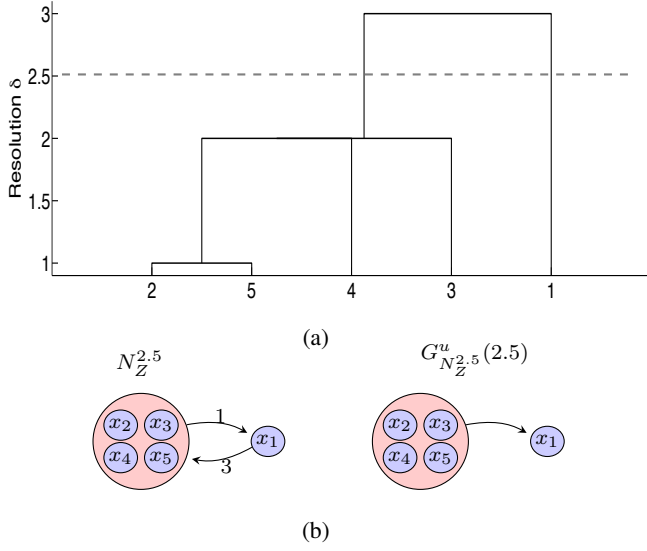


Fig. 7. (a) Nonreciprocal dendrogram of the network in Fig. 5. When cutting the dendrogram at resolution $\delta = 2.5$ we obtain two clusters. (b) The network of equivalence classes $N_Z^{2.5}$ and the corresponding unidirectional trust graph $G_{N_Z^{2.5}}^u(2.5)$. Global consensus in this graph implies global consensus in the network in Fig. 5 for resolution $\delta = 2.5$.

For the larger resolution $\delta = 1.5$, we have that in $G_N^u(1.5)$, x_3 listens to both x_1 and x_2 and will reach their average opinion in equilibrium while x_1 and x_2 maintain their original opinions. This outputs three different equilibrium opinions. Thus, opinion clusters are not nested as we modify the resolution parameter and cannot be represented by a dendrogram.

Nonetheless, nonreciprocal clustering \mathcal{H}^{NR} as defined in (4) *does* provide insight to further understand the consensus dynamics in (10). Indeed, nonreciprocal clustering is the correct way to aggregate data while maintaining the global consensus behavior of the original network. In order to explain this precisely, we need to define the network of equivalence classes N_Z^δ at a given resolution δ . Recall that nodes x and x' belong to the same nonreciprocal cluster at resolution δ , i.e. $x \sim_{D_X^{\text{NR}}(\delta)} x'$, if and only if $u_X^{\text{NR}}(x, x') \leq \delta$. Consider the space $Z^\delta := X \text{ mod } \sim_{D_X^{\text{NR}}(\delta)}$ of equivalence classes and the map $\phi_\delta : X \rightarrow Z^\delta$ that maps each point of X to its equivalence class. Notice that x and x' are mapped to the same point z if they belong to the same cluster at resolution δ , that is

$$\phi_\delta(x) = \phi_\delta(x') \iff u_X^{\text{NR}}(x, x') \leq \delta. \quad (12)$$

We define the network $N_Z^\delta = (Z^\delta, A_Z^\delta)$ by endowing Z^δ with the dissimilarity function A_Z^δ derived from the dissimilarity A_X as

$$A_Z^\delta(z, z') = \min_{x \in \phi_\delta^{-1}(z), x' \in \phi_\delta^{-1}(z')} A_X(x, x'). \quad (13)$$

The dissimilarity $A_Z^\delta(z, z')$ compares all the dissimilarities $A_X(x, x')$ between a member of the equivalence class z and a member of the equivalence class z' and sets $A_Z^\delta(z, z')$ to the value corresponding to the least dissimilar pair. Global consensus reachability of networks N and N_Z^δ is equivalent for every resolution δ as the following proposition asserts.

Proposition 2 *Given a trust network $N = (X, A_X)$, for the unidirectional trust consensus dynamics (10) with parameter δ , the*

graph $G_N^u(\delta)$ reaches global consensus if and only if $G_{N_Z^\delta}^u(\delta)$ reaches global consensus where the network of equivalence classes $N_Z^\delta = (Z^\delta, A_Z^\delta)$ is defined in (12) and (13).

Proof: See [10]. ■

In general, clustering in networks seeks to aggregate data while preserving relevant features of the original network. Proposition 2 shows that nonreciprocal clustering aggregates the data in N into the equivalence class network N_Z^δ while preserving reachability of global consensus in (10) for every resolution δ . To exemplify this, in Fig. 7-(a) we depict the nonreciprocal dendrogram of the network in Fig. 5 computed with the algorithm in [8]. At resolution $\delta = 2.5$, there are two equivalence classes given by $Z^{2.5} = \{\{x_2, x_3, x_4, x_5\}, \{x_1\}\}$. From (13) we compute $A_Z^{2.5}$ and we depict the network $N_Z^{2.5}$ at the leftmost part of Fig. 7-(b). From this network we compute the corresponding directed graph $G_{N_Z^{2.5}}^u(2.5)$ using (9) and illustrate it in the rightmost part of Fig. 7-(b). This two-node graph trivially reaches global consensus. Hence, by Proposition 2 we can assert that the five-node graph $G_N^u(2.5)$ derived from the network in Fig. 5 also reaches global consensus. Furthermore, in Fig. 7-(b) we see that the node containing x_2 through x_5 adopts the opinion of x_1 . Thus, the global consensus opinion of the five-node graph $G_N^u(2.5)$ coincides with the original opinion of agent x_1 .

V. CONCLUSION

We applied a theory for hierarchical clustering of asymmetric networks to study equilibrium configurations of consensus problems. Reciprocal clustering was shown to describe opinion profiles in mutual trust consensus problems whereas nonreciprocal clustering was shown to be the right way to aggregate data while maintaining global consensus reachability in unidirectional trust consensus problems.

REFERENCES

- [1] W. Ren, R. W. Beard, and E. M. Atkins, "A survey of consensus problems in multi-agent coordination," *American Control Conference*, 2005.
- [2] W. Ren, "Consensus based formation control strategies for multi-vehicle systems," *American Control Conference*, 2006.
- [3] H. G. Tanner, A. Jadbabaie, and G. J. Pappas, "Stable flocking of mobile agents, part i: Fixed topology," *Conference on Decision and Control*, 2003.
- [4] M. H. DeGroot, "Reaching a consensus," *Journal of the American Statistical Association*, 1974.
- [5] J. C. Dittmer, "Consensus formation under bounded confidence," *Non-linear Analysis*, vol. 47, 2001.
- [6] A. K. Jain and R. C. Dubes, *Algorithms for clustering data*, ser. Prentice Hall Advanced Reference Series. Englewood Cliffs, NJ: Prentice Hall Inc., 1988.
- [7] G. Carlsson, F. Mémoli, A. Ribeiro, and S. Segarra, "Axiomatic construction of hierarchical clustering in asymmetric networks," *Proc. Int. Conf. Acoustics Speech Signal Process*, 2013.
- [8] —, "Hierarchical clustering methods and algorithms for asymmetric networks," *Asilomar Conf. on Signals, Systems, and Computers*, vol. (accepted), 2013.
- [9] R. Olfati-Saber and R. M. Murray, "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Transactions on Automatic Control*, vol. 49, no. 9, pp. 1520–1533, 2004.
- [10] S. Segarra and A. Ribeiro, "Appendix to hierarchical clustering and consensus in trust networks," <https://fling.seas.upenn.edu/~ssegarra/wiki/index.php?n=Research.Publications>, 2013.