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 $A_X : X \times X \to \mathbb{R}_+$ 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 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., $D_X(0) = \big\{\{x\}, x \in X\big\}$, and for some sufficiently large resolution δ_0 all nodes must belong to the same cluster, i.e., $D_X(\delta_0) = \big\{X\big\}$. 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.

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

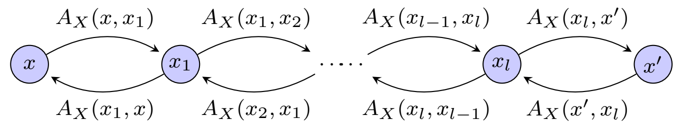

Fig. 1. Reciprocal clustering. Nodes x, x' cluster at resolution δ if they can be joined with a bidirectional chain of maximum dissimilarity δ [cf. (2)].

A hierarchical clustering method is then a map $\mathcal{H} : \mathcal{N} \to \mathcal{D}$ 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 $\mathcal{H}^{\text{R}}$ with output dendrogram $D_X^R = \mathcal{H}^{\text{R}}(X, A_X)$ merges nodes x and x' at resolution $u_X^{\text{R}}(x, x')$ given by

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

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 $u_X^{\text{R}}(x, x')$ 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 $\mathcal{H}^{\text{NR}}$ permits different chains. Define the minimum directed cost as

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

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

$$u_X^{\text{NR}}(x, x') := \max\Big(\tilde{u}_X^{\text{NR}}(x, x'),\ \tilde{u}_X^{\text{NR}}(x', x)\Big). \tag{4}$$

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.

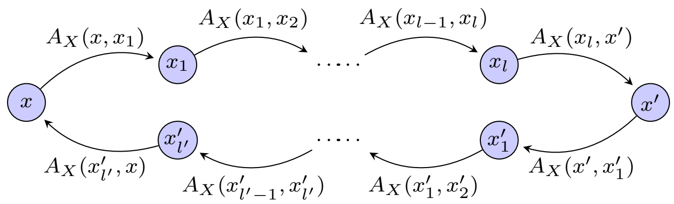

Fig. 2. Nonreciprocal clustering. Nodes x, x' cluster at resolution δ if they can be joined in both directions with possibly different chains of maximum dissimilarity δ [cf. (4)].

forward and backward chains that minimize the respective maximum dissimilarities. The nonreciprocal merging resolution $u_X^{\mathrm{NR}}(x, x')$ 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

$$L_G = \Delta_G - \mathcal{A}_G. \tag{5}$$

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

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

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

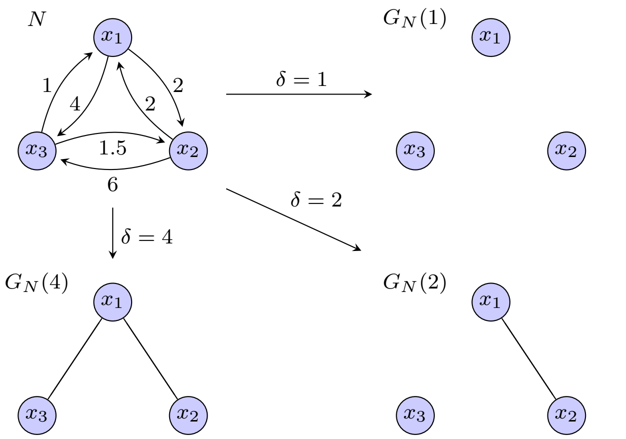

Fig. 3. Graphs induced by the network N for different resolutions δ under the assumption of mutual trust.

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 communication graph $G_N(\delta) = (X, E)$ with adjacency matrix $\mathcal{A}_{G_N}(\delta)$ given by

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

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.

From (5), we compute the Laplacian matrix L_{G_N} and obtain the mutual trust consensus dynamics

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

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 $\lim_{t\to\infty} p(t)$ 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

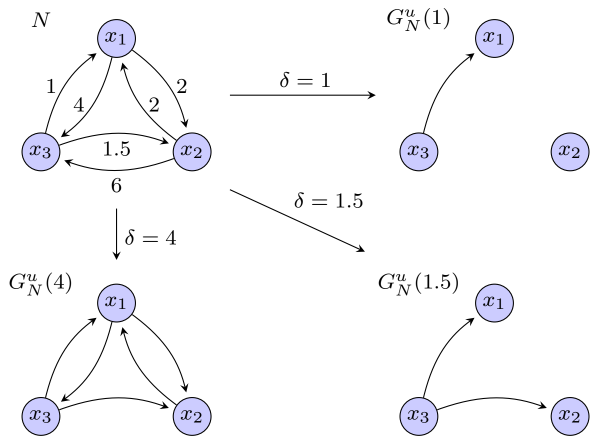

Fig. 4. Graphs induced by the network N for different resolutions δ under the assumption of unidirectional trust.

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

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

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

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

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 $\mathcal{H}^{\mathrm{R}}$ solves the mutual trust problem (8) as the following proposition asserts.

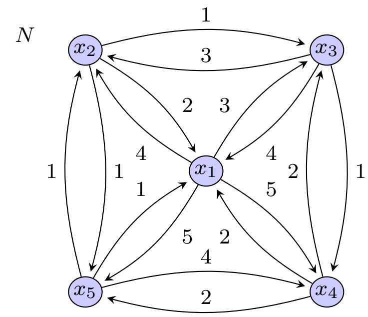

Fig. 5. A trust network N with five agents. Undrawn edges correspond to trust dissimilarities greater than 5.

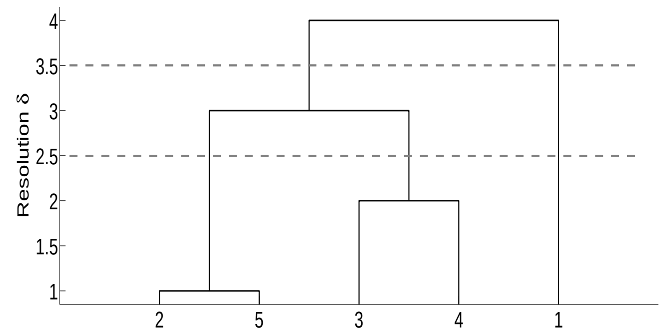

Fig. 6. Reciprocal dendrogram of the network in Fig. 5. At resolution $\delta = 2.5$ clusters $\{x_2, x_5\}$, $\{x_3, x_4\}$ and $\{x_1\}$ are formed whereas at resolution $\delta = 3.5$ nodes are clustered into $\{x_2, x_5, x_3, x_4\}$ and $\{x_1\}$. According to Proposition 1, these clusters correspond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

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

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

Proof: See Appendix A. ■

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.

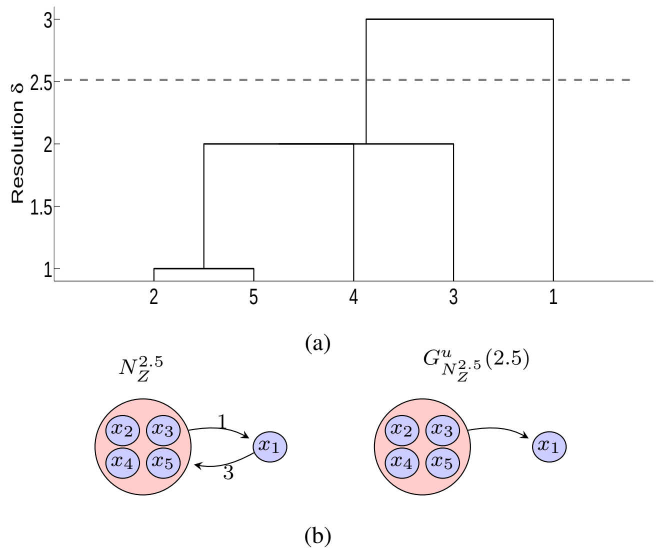

Fig. 7. (a) Nonreciprocal dendrogram of the network in Fig. 5. When cutting the dendrogram at resolution $\delta = 2.5$ we obtain two clusters. (b) The network of equivalence classes $N_Z^{2.5}$ and the corresponding unidirectional trust graph $G^u_{N_Z^{2.5}}(2.5)$. Global consensus in this graph implies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

For the larger resolution $\delta = 1.5$, we have that in $G^u_N(1.5)$, 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 $\mathcal{H}^{\rm NR}$ 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 $x \sim_{D_X^{\rm NR}(\delta)} x'$, if and only if $u_X^{NR}(x,x') \leq \delta$. Consider the space $Z^\delta := X \mod \sim_{D_X^{\rm NR}(\delta)}$ of equivalence classes and the map $\phi_\delta : X \to Z^\delta$ 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

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

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

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

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 $G^u_N(\delta)$ reaches global consensus if and only if $G^u_{N_Z^\delta}(\delta)$ 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 Appendix A. ∎

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 $G^u_{N_X^{2.5}}(2.5)$ 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 $G^u_N(2.5)$ 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 $G^u_N(2.5)$ 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.

Appendix A
Proofs

Proof of Proposition 1: We start by showing that nodes $x, x' \in X$ are connected in $G_N(\delta)$ if and only if $u_X^R(x,x') \leq \delta$. To show sufficiency, assume that $u_X^R(x,x') \leq \delta$. From (2) there must be a minimizing chain $C^*(x,x') = [x = x_0, x_1, x_2, \ldots, x_l = x']$ with reciprocal cost not greater than δ such that we can write that

$$\max_{i | x_i \in C^*(x,x')} \max\left(A_X(x_i, x_{i+1}), A_X(x_{i+1}, x_i)\right) \leq \delta. \tag{14}$$

This implies that $\max(A_X(x_i, x_{i+1}), A_X(x_{i+1}, x_i)) \leq \delta$ for every $x_i \in C^*(x,x')$. Hence, from (7) me must have that $[\mathcal{A}_{G_N}(\delta)]_{i\,i+1} = 1$ and, consequently, every pair of consecutive nodes in $C^*(x,x')$ is connected in $G_N(\delta)$, connecting the extreme nodes x and x', as wanted.

In order to show necessity, assume that two arbitrary nodes $x, x' \in X$ are connected in $G_N(\delta)$ for a given resolution δ. Then, there must exist at least one path joining these two nodes. Pick any of these paths and form a chain $C_*(x,x') = [x = x'_0, x'_1, \ldots, x'_{l'} = x']$ where there are edges in $G_N(\delta)$ between consecutive nodes of $C_*(x,x')$. From (7), this means that

$$\max\left(A_X(x'_i, x'_{i+1}), A_X(x'_{i+1}, x'_i)\right) \leq \delta, \tag{15}$$

for all $i = 0, \ldots, l' - 1$. Since $C_*(x, x')$ is a particular chain joining x and x' and in the definition of u_X^{R} (2) we minimize over all possible chains, it follows that

$$u_X^{\mathrm{R}}(x, x') \leq \max_{i | x'_i \in C_*(x,x')} \max\left(A_X(x'_i, x'_{i+1}), A_X(x'_{i+1}, x'_i)\right) \leq \delta, \tag{16}$$

where we used (15) for the rightmost inequality and bounded the value of $u_X^{\mathrm{R}}(x, x')$ as wanted.

The proposition statement follows immediately since in undirected graphs, like $G_N(\delta)$, the equilibrium opinion of two agents coincides for every initial opinion state if and only if they belong to the same connected component [10]. ■

Proof of Proposition 2: Instrumental to our proof is the definition of a spanning converging tree. We say that a directed graph G contains a spanning converging tree (s.c.t.) if there exists a spanning tree where one node – denominated the root – is reachable through directed paths by every other node in the tree. E.g., graph $G_N^u(1.5)$ in Fig. 4 does not contain a s.c.t. because, although it contains a spanning tree, in this tree node x_1 is not reachable by x_2 and vice versa, and node x_3 cannot be reached either by x_1 or x_2. Having this definition, we divide the proof of the proposition into two parts. First, we show that a directed graph G reaches global consensus if and only if it contains a spanning converging tree. In the second part, we show that for every resolution δ the graph $G_N^u(\delta)$ contains a s.c.t. if and only if $G_{N_X^\delta}^u(\delta)$ contains a s.c.t. Combining these two observations, the statement of the proposition follows. Thus, we begin by showing the first claim.

Claim 1 *A directed, unweighted graph G reaches global consensus in dynamics* (6) *for all initial opinion profiles p_0 if and only if it contains a spanning converging tree.*

Proof: In order to prove this claim we will use two known results about the spectral features of directed graphs. First, from [11] we know that the rank of the Laplacian $\text{rank}(L_G) = n - 1$ where n is the number of nodes in G if and only if G contains a spanning converging tree. Notice that in [11], the authors do the analysis for diverging trees because in their case opinions flow in the direction of the edges whereas in our model they flow in the opposite one. The other result we use is that the eigenvalues of the Laplacian L_G for a directed graph G have all nonnegative real part. This can be shown using Gershgorin theorem as in, e.g., [9].

The solution of $p(t)$ for any time t in dynamics (6) is given by

$$p(t) = e^{-L_G\, t}\, p_0 = V\, e^{-J\, t}\, V^{-1}\, p_0, \tag{17}$$

where J is the Jordan form of L_G and V is a matrix containing as columns the corresponding generalized eigenvectors of L_G. From the definition of the Laplacian (5), notice that $L_G\mathbf{1} = 0$, thus, every nontrivial vector parallel to the vector of all ones is an eigenvector associated with the null eigenvalue. If the graph G contains a spanning converging tree then $\text{rank}(L_G) = n - 1$ implying that the algebraic and geometric multiplicity of the null eigenvalue is 1. Since the other Jordan blocks contain positive real parts, their negative exponential converges to zero when time goes to infinity. Thus, if we denote the i^{th} column of V as v_i and the i^{th} row of V^{-1} as w_i^T then, from (17), we obtain

$$\begin{aligned}\lim_{t\to\infty} p(t) &= \begin{pmatrix} v_1 & v_2 & \cdots & v_n \end{pmatrix} \begin{pmatrix} 1 & 0 & \cdots & 0 \\ 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 \end{pmatrix} \begin{pmatrix} w_1^T \\ w_2^T \\ \vdots \\ w_n^T \end{pmatrix} p_0 \\ &= w_1^T p_0\ v_1, \end{aligned} \tag{18}$$

where v_1 is the eigenvector associated with the null eigenvalue. However, we know that v_1 is parallel to $\mathbf{1}$, hence the limiting opinion in (18) is the same for every node and global consensus is achieved. Now, suppose that the graph G does not contain a spanning converging tree. Then, $\text{rank}(L_G) < n-1$ and the algebraic multiplicity of the null eigenvalue is greater than one, say equal to $d > 1$. If the geometric multiplicity of the null eigenvalue is less than the algebraic multiplicity, then the system is unstable and global consensus is trivially not reached. If the geometric multiplicity is equal to d, meaning that all the eigenvectors associated with zero are linearly independent, then we can state an equivalence similar to (18) but with d ones in the diagonal of the exponential matrix, to obtain

$$\begin{aligned}\lim_{t\to\infty} p(t) &= \begin{pmatrix} v_1 & v_2 & \cdots & v_n \end{pmatrix} \begin{pmatrix} 1 & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 \end{pmatrix} \begin{pmatrix} w_1^T \\ w_2^T \\ \vdots \\ w_n^T \end{pmatrix} p_0 \\ &= \sum_{k=1}^{d} w_k^T p_0\ v_k, \end{aligned} \tag{19}$$

which cannot be a uniform vector for every initial condition p_0. Indeed, assume that (19) is uniform for every initial condition p_0 and consider the n different initial condition $p_0 = e_i$ for $i = 1, \ldots, n$ where e_i contains a 1 in the i^{th} position and 0 in the rest, the equilibrium opinions for these cases are

$$\lim_{t\to\infty} p(t) = \sum_{k=1}^{d} w_k^T e_i\ v_k = \alpha_i \mathbf{1}. \tag{20}$$

Notice that the vectors w_k for $k = 1, \ldots, d$ must be linearly independent since they correspond to the left eigenvectors of L_G associated with the null eigenvalue. Thus, the vectors w_k must span d dimensions and in these dimensions α_i must be different from zero, otherwise, the linear independence of the vectors v_k would be violated since a non trivial sum of them would equal zero. Without loss of generality, assume $\alpha_1 > 0$, $\alpha_2 > 0$ and the vectors w_k linearly independent in these two dimensions, then from (20), and denoting by $[a]_i$ the i^{th} component of a given vector a, we may write

$$\frac{1}{\alpha_1}\left([w_1]_1\, v_1\ + [w_2]_1\, v_2 + \ldots + [w_d]_1\, v_d\right) = \mathbf{1}, \tag{21}$$

$$\frac{1}{\alpha_2}\left([w_1]_2\, v_1\ + [w_2]_2\, v_2 + \ldots + [w_d]_2\, v_d\right) = \mathbf{1}. \tag{22}$$

By replacing (21) into (22) we violate the linear independence of vector v_k. Hence, (20) cannot hold for all i and (19) is not uniform for every initial condition. Consequently, global consensus is not reached in general. ■

Having shown that a spanning converging tree is a necessary and sufficient condition for global consensus, we show our second claim towards proving the proposition.

Claim 2 *For every resolution $\delta \geq 0$ and every asymmetric network N, the graph $G_N^u(\delta)$ contains a spanning converging tree if and only if the graph $G_{N_Z^\delta}^u(\delta)$ contains one.*

Proof: For an arbitrary network $N = (X, A_X)$ and resolution δ, suppose that the graph $G_N^u(\delta)$ contains a s.c.t. with root $x_r \in X$. Consider the node $z_r \in Z$ in the graph $G_{N_Z^\delta}^u(\delta)$ where $z_r = \phi_\delta(x_r)$ with the map ϕ_δ defined in (12). Pick an arbitrary node $z \in Z$ and pick any $x \in X$ such that $x = \phi_\delta^{-1}(z)$. From the existence of a s.c.t. in $G_N^u(\delta)$ there must be a directed path from x to x_r, namely $[x = x_0, x_1, ..., x_l = x_r]$. Moreover, (13) implies that in the directed path $[\phi_\delta(x) = z, \phi_\delta(x_1), ..., \phi_\delta(x_r) = z_r]$ every two consecutive nodes are either coincident or connected in $G_{N_Z}^u(\delta)$. Hence, there is a directed path from z to z_r in $G_{N_Z}^u(\delta)$. Since z was chosen arbitrarily, there must exist a path from every node $z \in Z$ to $z_r \in Z$, implying that $G_{N_Z}^u(\delta)$ contains a s.c.t. rooted at z_r.

Assume now that $G_{N_Z}^u(\delta)$ contains a s.c.t. rooted at a particular node $z_r \in Z$ and denote the edge set of the s.c.t. by $E_Z(\delta)$. We first show that for any $z \in Z$ the subgraph $G_N^u(\delta)|_{\phi_\delta^{-1}(z)}$, i.e. the graph $G_N^u(\delta)$ restricted to the pre image of a particular node z, contains a s.c.t. rooted at any node $x \in \phi_\delta^{-1}(z)$. To see this, pick an arbitrary node $x \in \phi_\delta^{-1}(z)$ and call it x_r^z. From (12) we know that

$$u_X^{\mathrm{NR}}(x, x_r^z) \leq \delta, \tag{23}$$

for all $x \in \phi_\delta^{-1}(z)$. Combining the definitions of nonreciprocal clustering (4) and unidirectional influence (9), inequality (23) implies the existence of directed paths from every x to x_r^z in $\phi_\delta^{-1}(z)$ and, hence, the existence of a s.c.t. rooted at x_r^z. We now construct a s.c.t. for $G_N^u(\delta)$ given a s.c.t. for $G_{N_Z}^u(\delta)$ rooted at $z_r \in Z$. Pick any node $x \in \phi_\delta^{-1}(z_r)$ and denote it by $x_r^{z_r}$. We already know that it is possible to construct a s.c.t. spanning $\phi_\delta^{-1}(z_r)$ rooted at $x_r^{z_r}$. Consider the nodes $z \in Z$ linked to z_r in the s.c.t., i.e. nodes z such that $(z, z_r) \in E_Z(\delta)$. For these nodes z, consider as a root x_r^z for a s.c.t. spanning $\phi_\delta^{-1}(z)$ the node achieving the minimum for $A_Z^\delta(z, z_r)$ in (13). In this way, we are guaranteed the existence of an edge in $G_N^u(\delta)$ from x_r^z to some node $x \in \phi_\delta^{-1}(z_r)$. Repeat this process for every node $z \in Z$ linked to z_r and then continue the process with the nodes z linked to nodes which are linked to z_r and so on. Once this process is terminated, we obtain a spanning converging tree in $G_N^u(\delta)$, as wanted. ■

From combining claims 1 and 2, Proposition 2 follows. ■

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," *Nonlinear 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] J. A. F. R. Olfati-Saber and R. M. Murray, "Consensus and cooperation in networked multi-agent systems," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 215–233, 2007.

[11] R. P. Agaev and P. Y. Chebotarev, "The matrix of maximum out forests of a digraph and its applications," *Automation and remote control*, vol. 60, no. 9, pp. 1424–1450, 2000.