

Rearranging trees for robust consensus

George Forrest Young, Luca Scardovi and Naomi Ehrich Leonard

Abstract—In this paper, we use the \mathcal{H}_2 norm associated with a communication graph to characterize the robustness of consensus to noise. In particular, we restrict our attention to trees, and by systematic attention to the effect of local changes in topology, we derive a partial ordering for undirected trees according to the \mathcal{H}_2 norm. Our approach for undirected trees provides a constructive method for deriving an ordering for directed trees. Further, our approach suggests a decentralized manner in which trees can be rearranged in order to improve their robustness.

I. INTRODUCTION

The study of linear consensus problems has gained much attention in recent years [1]–[4]. This attention has arisen, in part, from the wide range of applications of linear consensus, including collective decision-making [3], formation control [5], sensor fusion [6], distributed computing [7] and analysis of biological groups [8]. In most of these applications, information passed between agents can be corrupted by noise. It is therefore necessary to understand the robustness of consensus when noise is present, with the goal of designing systems that can efficiently filter noise and remain close to consensus. For a linear system with additive white noise, a natural measure of robustness is the \mathcal{H}_2 norm [9].

For the study of consensus, most of the important details are described by the communication graph. In fact, the properties of the Laplacian matrix of the graph are deeply related to the performance of the consensus protocol. In this way, the study of consensus often reduces to studying the underlying graph, and relating graph properties to the performance of the original system [1], [3], [10], [11].

Communication in a multi-agent system is likely to be of a directed nature, simply because each agent may treat the information it receives differently than its neighbors. Additionally, directed communication can arise when information is transferred through sensing or when agents have limited capabilities and may choose to only receive information from a subset of possible neighbors. In many real systems, the graph between agents may change over time depending on the behavior and decisions of individual agents [1]. Therefore, when searching for ways in which to impose effective graphs on multi-agent systems, it is highly advantageous to consider whether such graphs could be formed in a decentralized manner.

This research was supported in part by AFOSR grant FA9550-07-1-0-0528 and ONR grant N00014-09-1-1074.

G. F. Young and N. E. Leonard are with the Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ 08544, USA. L. Scardovi is with the Department of Electrical Engineering and Information Technology, Technical University of Munich, 80333 Munich, Germany. gfyong@princeton.edu, scardovi@tum.de, naomi@princeton.edu

In this paper, we study the robustness of a particular family of graphs, namely trees, according to their \mathcal{H}_2 norms. We develop a partial ordering among trees that allows us to find a tree with minimal \mathcal{H}_2 norm, given certain constraints. Although most of this partial ordering has already been developed in the literature on Wiener indices [12]–[15], our methods of proof are new. In particular, we rely only on local changes in which one or more leaf nodes are moved from a single location in the tree to a new location. This approach provides insight into ways in which trees can be rearranged in a decentralized manner in order to improve their robustness. Additionally, our methods can be used to derive a similar ordering for directed trees that could not be done using the Wiener index literature.

This paper is organized as follows. In Section II we summarize notation. In Section III we discuss the \mathcal{H}_2 norm in more detail. In Sections IV and V we discuss the relationship between the \mathcal{H}_2 norm and other graph indices. In Section VI we introduce a system of terminology to describe tree graphs, and in Section VII we derive our partial ordering. Finally, in Section VIII, we discuss the potential for a decentralized algorithm to improve the \mathcal{H}_2 norm of a tree.

II. PRELIMINARIES AND NOTATION

The state of the system is $x = [x_1, x_2, \dots, x_N] \in \mathbb{R}^N$, where x_i is the state of agent i . We call the state *consensus* when $x = \gamma 1_N$, where $1_N = [1, 1, \dots, 1]^T \in \mathbb{R}^N$ and $\gamma \in \mathbb{R}$. For each agent i we define the set of neighbors, \mathcal{N}_i , to be the set of agents that supply information to agent i .

We associate to the system a *communication graph* $\mathcal{G} = (\mathcal{V}, \mathcal{E}, A)$, where $\mathcal{V} = \{1, 2, \dots, N\}$ is the set of nodes, $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the set of edges and $A \in \mathbb{R}^{N \times N}$ is a weighted adjacency matrix with nonnegative entries $a_{i,j}$ corresponding to the weight on edge (i, j) . Every node in the graph corresponds to an agent, while the graph contains edge (i, j) when $j \in \mathcal{N}_i$. Then $a_{i,j}$ is the weight given by agent i to the information from agent j .

An edge $(i, j) \in \mathcal{E}$ is said to be *undirected* if (j, i) is also in \mathcal{E} and $a_{i,j} = a_{j,i}$. A graph is undirected if every edge is undirected, that is, if A is symmetric. Two nodes are said to be *adjacent* when there is an edge between them.

The *out-degree* of node k is defined as $d_k^{\text{out}} = \sum_{j=1}^N a_{k,j}$. \mathcal{G} has an associated *Laplacian* matrix L , defined by $L = D - A$, where $D = \text{diag}(d_1^{\text{out}}, d_2^{\text{out}}, \dots, d_N^{\text{out}})$ is the diagonal matrix of out-degrees. The row sums of the Laplacian matrix are zero, that is $L 1_N = 0$. Thus 0 is always an eigenvalue of L with corresponding eigenvector 1_N . Furthermore, all eigenvalues of L have non-negative real part (by Geršgorin's

Theorem). For an undirected graph, L is symmetric, so in addition $1_N^T L = 0$ and all the eigenvalues of L are real.

A *path* in \mathcal{G} is a (finite) sequence of nodes containing no repetitions and such that each node is a neighbor of the previous one. The length of a path is given by the sum of the weights on all edges traversed by the path.

The graph \mathcal{G} is *connected* if it contains a globally reachable node k ; i.e. there is a path in \mathcal{G} from i to k for every node i . It can be shown that 0 will be a simple eigenvalue of L if and only if \mathcal{G} is connected [16]. If an undirected graph is connected, there will be a path in \mathcal{G} between every pair of nodes. We use λ_i to refer to the i^{th} eigenvalue of the Laplacian matrix, when arranged in ascending order by real part. Thus $\lambda_1 = 0$ for any Laplacian matrix, and $\text{Re}\{\lambda_2\} > 0$ if and only if \mathcal{G} is connected.

The *distance*, $d_{i,j}$, between nodes i and j in a graph is the shortest length of any path from i to j . If no such path exists, $d_{i,j}$ is infinite. The *diameter*, d , of a graph is the maximum distance between all pairs of nodes in the graph.

A *tree* on N nodes is a connected undirected graph in which every pair of nodes is connected by a unique path. This implies that a tree contains exactly $N - 1$ undirected edges and that it contains no cycles (paths with positive length connecting a node to itself). A *rooted tree* is a tree in which one particular node has been identified as the root (note that other than being called the root, there is nothing “special” about this node). A *directed tree* is a connected graph containing exactly $N - 1$ directed edges. In a directed tree, the globally reachable node is identified as the root.

III. ROBUST NOISY CONSENSUS AND THE \mathcal{H}_2 NORM

In this paper we assume that every agent is independently affected by white noise of the same intensity. The resulting consensus dynamics are (as in [5], [9], [17])

$$\dot{x}(t) = -Lx(t) + \xi(t) \quad (1)$$

with $x \in \mathbb{R}^N$ and where $\xi(t) \in \mathbb{R}^N$ is a zero-mean mutually independent white stochastic process.

Since (1) is only marginally stable in the noise-free case (corresponding to the fact that there is no “preferred” or “correct” value for the agents to agree upon), we only consider the dynamics on the subspace of \mathbb{R}^N orthogonal to the subspace spanned by 1_N . We let $Q \in \mathbb{R}^{(N-1) \times N}$ be a matrix with rows that form an orthonormal basis of this subspace. This is equivalent to requiring that

$$Q1_N = 0,$$

$$QQ^T = I_{N-1} \text{ and } Q^T Q = I_N - \frac{1}{N} 1_N 1_N^T. \quad (2)$$

Next, we define $y := Qx$. Then $y = 0$ if and only if $x = \gamma 1_N, \gamma \in \mathbb{R}$. A measure of the distance from consensus is the *dispersion* of the system $\|y(t)\|_2$.

Differentiating $y(t)$, we obtain

$$\dot{y}(t) = -\bar{L}y(t) + Q\xi(t) \quad (3)$$

where $\bar{L} = QLQ^T$ is the *reduced Laplacian* matrix.

\bar{L} has the same eigenvalues as L except the zero eigenvalue, which implies that $-\bar{L}$ is Hurwitz precisely when the

graph is connected [9]. Thus, for a connected graph in the absence of noise, system (3) will converge exponentially to zero. In the presence of noise, (3) will no longer converge, but will remain in motion about zero. We define the robustness of consensus to noise as the expected dispersion of the system in steady state. Note that this definition is analogous to the steady-state mean-square deviation used in [7]. Our measure of robustness corresponds to the \mathcal{H}_2 norm of system (3), with output equation $z(t) = I_{N-1}y(t)$. \mathcal{H}_2 norms of consensus systems have also been studied in [5], [17].

In [9] we proved that for a system with an undirected communication graph¹, the \mathcal{H}_2 norm is given by

$$H = \left(\sum_{i=2}^N \frac{1}{2\lambda_i} \right)^{\frac{1}{2}}. \quad (4)$$

In general, the \mathcal{H}_2 norm of a directed graph can be computed as $H = [\text{tr}(\Sigma)]^{\frac{1}{2}}$, where Σ is the solution to the Lyapunov equation [9]

$$\bar{L}\Sigma + \Sigma\bar{L}^T = I. \quad (5)$$

Since this \mathcal{H}_2 norm can be computed using only the communication graph, in the rest of this paper we associate the \mathcal{H}_2 norm with the graph. Thus when we refer to the \mathcal{H}_2 norm of a graph, we mean the \mathcal{H}_2 norm of system (3) (with output $z = y$) with \bar{L} computed from the given graph.

IV. EFFECTIVE RESISTANCE AS A MEASURE OF THE \mathcal{H}_2 NORM

Although equation (4) allows us to compute the \mathcal{H}_2 norm for any undirected graph, it does not readily allow us to infer relationships between structural features of the graph and the \mathcal{H}_2 norm. However, the concept of the *effective resistance*, or *Kirchhoff index*, of a graph can help us in this respect. The effective resistance results from considering a given graph as an electrical network, where every edge corresponds to a resistor with resistance given by the inverse of the edge weight. The resistance between two nodes in the graph is given by the resistance between those two points in the electrical network, and the effective resistance of the graph is the sum of the resistances between all pairs of nodes [18].

The effective resistance of a graph is related to the eigenvalues of the graph Laplacian [18] by the formula

$$K_f = N \sum_{i=2}^N \frac{1}{\lambda_i}, \text{ leading to the relationship } H = \left(\frac{K_f}{2N} \right)^{\frac{1}{2}}. \quad (6)$$

We therefore see that for graphs with equal numbers of nodes, any ordering induced by the effective resistance is the same as that induced by the \mathcal{H}_2 norm.

Although computing the effective resistance can be difficult for most graphs, it is very straightforward for trees. In a tree with unit weights on every edge, the resistance between two nodes is given by the distance between them [19]. Hence, the effective resistance of a tree with unit edge weights is

$$K_f = \sum_{i < j} r_{i,j} = \sum_{i < j} d_{i,j}. \quad (7)$$

¹In fact, the result in [9] is more general and extends to directed graphs with a normality condition on their Laplacian matrix.

Although the concept of effective resistance does not apply to directed graphs, we can define an extension so that equation (6) applies to all graphs. The resistance between two nodes of an undirected graph can be computed as [18]

$$r_{i,j} = (L^\dagger)_{i,i} + (L^\dagger)_{j,j} - 2(L^\dagger)_{i,j} \quad (8)$$

where L^\dagger is the Moore-Penrose pseudoinverse of L . For an undirected graph, we can explicitly write $L^\dagger = 2Q^T \Sigma Q$, where Σ is the solution to the Lyapunov equation (5). Thus, if we let $X = 2Q^T \Sigma Q$, we can compute for any graph

$$r_{i,j} = (X)_{i,i} + (X)_{j,j} - 2(X)_{i,j}. \quad (9)$$

Using equation (9), we can compute “directed resistances” (and hence Kirchhoff indices) for directed graphs. Through this construction, we can show that equation (6) will hold for directed graphs as well. Then, in a directed tree (as in an undirected tree) the resistance between two nodes only depends on the paths between them. The proofs and results for directed trees will appear in a future publication.

In the following sections we determine a partial ordering of undirected trees with unit edge weights. The same ordering will apply to the set of trees with a given constant edge weight, as all resistances will be proportional to those in the corresponding tree with unit weights.

V. THE \mathcal{H}_2 NORM AND OTHER GRAPH INDICES

In addition to the Kirchhoff index, many other “topological” indices of graphs have arisen out of the mathematical chemistry literature [20]. One of the earliest to arise was the Wiener index, W [20]. The Wiener index for any (undirected) graph is defined as

$$W = \sum_{i < j} d_{i,j}. \quad (10)$$

Thus, for trees with unit edge weights, the Kirchhoff and Wiener indices are identical. However, the two indices differ for any other graph. Hence, while the results in Section VII apply equally to Wiener indices, we choose to interpret them only in terms of the Kirchhoff index and \mathcal{H}_2 norm.

Much work has already been done on comparing trees based on their Wiener indices. It is already well-known that the Wiener index of a tree will fall between that of the star and that of the path [12], [21]. For trees with a fixed number of nodes, the 15 trees with smallest Wiener index and the 17 trees with largest Wiener index have been identified [13], [14]. Further, for trees with a fixed number of nodes and a fixed diameter, the tree with smallest Wiener index has been found [15]. Therefore, most of the main results in Section VII have already been derived. Our contribution includes new methods of proof that rely on local changes of topology and provide constructive means to order directed trees and derive decentralized strategies for improving robustness.

A different graph index, developed in the mathematical literature, is the maximum eigenvalue of the adjacency matrix A [22]. Simić and Zhou developed a partial ordering of trees with fixed diameter according to this index in [22]. Their work, in particular the families of trees they considered and the order in which they proved their results, has motivated the approach taken in this paper.

VI. A SYSTEM OF TERMINOLOGY FOR TREES

We first introduce a system of terminology relating to trees. Much of our terminology corresponds to that in [22] and earlier papers. $\mathcal{T}_{N,d}$ is the set of all trees containing N nodes and with diameter d . For $N \geq 3$, a tree must have $d \geq 2$, and $\mathcal{T}_{N,2}$ contains only one tree. This tree is called a star, and is denoted $K_{1,N-1}$. For all positive N , the maximum diameter of a tree is $N - 1$, and $\mathcal{T}_{N,N-1}$ contains only one tree. This tree is called a path, and is denoted P_N .

A *leaf* (or pendant) is a node with degree 1. A *bouquet* is a non-empty set of leaf nodes, all adjacent to the same node. A node which is not a leaf is called an *internal* node.

A *caterpillar* is a tree for which the removal of all leaf nodes would leave a path. The set of all caterpillars with N nodes and diameter d is denoted by $\mathcal{C}_{N,d}$ (see Figure 1). Any caterpillar in $\mathcal{C}_{N,d}$ contains a path of length d , with all other nodes adjacent to internal nodes of this path. In particular, we refer to the caterpillar that contains a single bouquet attached to the i^{th} internal node along this path $P_{N,d,i}$ (see Figure 2). To avoid ambiguity, we require $1 \leq i \leq \lfloor \frac{d}{2} \rfloor$. The tree formed from $P_{N-1,d,\lfloor \frac{d}{2} \rfloor}$ by attaching an additional node to one of the leaves in the central bouquet is denoted by $N_{N,d}$ (see Figure 3).




Fig. 1. General form of a caterpillar in $\mathcal{C}_{N,d}$, with $n_j \geq 0$ additional leaf nodes attached to each internal node j in the path of length d .




Fig. 2. The caterpillar $P_{N,d,i}$, a path of length d with a bouquet containing $N - d - 1$ leaf nodes attached to the i^{th} internal node on the path.




Fig. 3. The tree $N_{N,d}$, formed from $P_{N-1,d,\lfloor \frac{d}{2} \rfloor}$ by attaching an additional node to one of the leaves in the central bouquet. Note that $N - d - 3$ could be 0.

The *double palm tree* (also referred to as a dumbbell in [12]) is a caterpillar with two bouquets, one at each end of the path (see Figure 4). We use $D_{N,p,q}$ to denote the double palm tree on N nodes, with bouquets of sizes p and q . If we take a rooted tree T (with root r) and attach two separate paths containing l and k nodes to the root, we call the resulting tree a *vine* and denote it by $T_{l,k}^r$ (see Figure 5).

VII. MANIPULATIONS TO REDUCE THE EFFECTIVE RESISTANCE OF TREES

We can now start to describe a partial ordering on trees based on their \mathcal{H}_2 norms. Every tree is assumed to have a




Fig. 4. Double palm tree $D_{N,p,q}$, with bouquets of sizes p and q at each end of a path.




Fig. 5. The vine $T_{l,k}^r$, formed from a rooted tree T by separately connecting paths containing l and k nodes to the root.

unit weight on every edge. First, we determine the effect of moving a leaf from one end of a double palm tree to the other, and use this to derive a complete ordering of all trees in $\mathcal{T}_{N,3}$ (Theorem 1). Second, we consider moving a leaf from one end of a vine to the other, and use this to prove that the path has the largest \mathcal{H}_2 norm of any tree with N nodes (Theorem 2), and to derive a complete ordering of $\mathcal{T}_{N,N-2}$ (Theorem 3). Finally, by moving all (or almost all) nodes in a bouquet to an adjacent node, we show that $P_{N,d,\lfloor \frac{d}{2} \rfloor}$ has the smallest \mathcal{H}_2 norm of any tree with diameter d (Theorem 4) and that for any tree that is not a star, we can find a tree of smaller diameter with a smaller \mathcal{H}_2 norm (Theorem 5). From Theorem 5 we also conclude that the star has the smallest \mathcal{H}_2 norm of any tree with N nodes.

A. Double Palm Trees

We begin our partial ordering by showing that the \mathcal{H}_2 norm of a double palm tree is reduced when we move a single node from the smaller bouquet to the larger one.

Lemma 1: Let $1 < p \leq q$ and $p + q \leq N - 2$. Then $\mathcal{H}_2(D_{N,p,q}) > \mathcal{H}_2(D_{N,p-1,q+1})$.

Proof: In $D_{N,p,q}$, let one of the nodes in the bouquet of size p be node 1. The remaining nodes are labelled 2 through N . To form $D_{N,p-1,q+1}$, we take node 1 and move it to the other bouquet. Since all other nodes remain unchanged, we can use equation (7) to write

$$K_f(D_{N,p,q}) - K_f(D_{N,p-1,q+1}) = \left(\sum_{j=2}^N d_{1,j} \right)_{D_{N,p,q}} - \left(\sum_{j=2}^N d_{1,j} \right)_{D_{N,p-1,q+1}}.$$

Now, in $D_{N,p,q}$, the path length between node 1 and any of the remaining $p-1$ nodes in the bouquet of size p is 2. Similarly, the path length between node 1 and any node in the bouquet of size q is $N-p-q+1$. Finally, the path lengths between node 1 and the internal nodes take on each integer value from 1 to $N-p-q$.

Conversely, in $D_{N,p-1,q+1}$, the path length between node 1 and any of the nodes in the bouquet of size $p-1$ is $N-p-q+1$. The path length between node 1 and any of the remaining q nodes in the bouquet of size $q+1$ is 2. Again, the path lengths between node 1 and the internal nodes take on all integer values from 1 to $N-p-q$.

Thus, $D_{N,p-1,q+1}$ (compared to $D_{N,p,q}$) has more nodes at a distance 2 from node 1 and fewer nodes at a dis-

tance $N-p-q+1$, while the sum of distances to all internal nodes remains the same. Therefore $K_f(D_{N,p,q}) > K_f(D_{N,p-1,q+1})$. Hence, by equation (6), the result holds. ■

Although Lemma 1 applies to double palm trees with any diameter, we can apply it to trees with $d = 3$ in order to prove our first main result.

Theorem 1: For $N \geq 4$, we have a complete ordering of $\mathcal{T}_{N,3}$, namely $\mathcal{H}_2(D_{N,1,N-3}) < \mathcal{H}_2(D_{N,2,N-4}) < \dots < \mathcal{H}_2(D_{N,\lfloor \frac{N-2}{2} \rfloor, \lceil \frac{N-2}{2} \rceil})$.

Proof: Any tree with $d = 3$ must have a longest path of length 3. Any additional nodes in the tree must be connected through some path to one of the two internal nodes on this longest path. In addition, any node adjacent to one of the internal nodes of the longest path forms a path of length 3 with the node at the far end of the path. Hence all such nodes must be leaves and so every tree with $d = 3$ is a double palm tree. The ordering follows from Lemma 1. ■

B. Vines

Our next task is to find an ordering of trees with the largest possible diameter. Lemma 2 applies to trees of any diameter, but again we can specialize it to give the results we need.

Lemma 2: Let T be a tree containing more than one node and with a root r , and let l, k be any positive integers such that $1 \leq l \leq k$. Then $\mathcal{H}_2(T_{l,k}^r) < \mathcal{H}_2(T_{l-1,k+1}^r)$.

Proof: Let the total number of nodes in $T_{l,k}^r$ be N (so $N - k - l > 1$), and let the leaf at the end of the path containing l nodes be node 1. Let the remaining nodes in the two paths be nodes 2 through $l+k$, and let the root of T be node $l+k+1$. The remaining nodes in T are labelled $l+k+2$ through N . To form $T_{l-1,k+1}^r$, we take node 1 and move it to the end of the other path. Since all other nodes remain unchanged, we can use equation (7) to write

$$K_f(T_{l-1,k+1}^r) - K_f(T_{l,k}^r) = \left(\sum_{j=2}^N d_{1,j} \right)_{T_{l-1,k+1}^r} - \left(\sum_{j=2}^N d_{1,j} \right)_{T_{l,k}^r}. \quad (11)$$

Now, in both $T_{l,k}^r$ and $T_{l-1,k+1}^r$, the path lengths between node 1 and all nodes along the paths (including the root of T) take on each integer value between 1 and $l+k$. Hence the sum of these path lengths does not change between the two trees. Furthermore, since the root of T lies on every path between node 1 and any other node in T , we can write

$$d_{1,j} = d_{1,l+k+1} + d_{l+k+1,j}, \quad j \geq l+k+2.$$

Therefore, for $T_{l,k}^r$, the sum of the distances from node 1 to all the nodes in T is $(N-l-k+1)l$ plus the sum of the distances from node r to each node in T . However, in $T_{l-1,k+1}^r$, the sum of the distances from node 1 to all the nodes in T is $(N-l-k+1)(k+1)$ plus the sum of the distances from node r to each node in T . Thus $K_f(T_{l-1,k+1}^r) > K_f(T_{l,k}^r)$ and so by equation (6), the result holds. ■

The first consequence of Lemma 2 is that the tree with largest d (i.e. $d = N-1$) also has the largest \mathcal{H}_2 norm.

Theorem 2: The path P_N has the largest \mathcal{H}_2 norm of any tree with N nodes.

Proof: Any tree T_1 which is not a path will contain a node with degree greater than 2. We can locate one such node that has two paths (each with fewer than N nodes) attached. Let T be the tree formed by removing these two paths from T_1 , and let this node be the root of T . Then $T_1 = T_{l,k}^r$, and by Lemma 2 we can find a tree with larger \mathcal{H}_2 norm. ■

We can also use Lemma 2 to derive an ordering of those trees with d one less than its maximum value (i.e. $d = N-2$).

Theorem 3: For $N \geq 4$, we have a complete ordering of $\mathcal{T}_{N,N-2}$, namely $\mathcal{H}_2(P_{N,N-2,\lfloor \frac{N-2}{2} \rfloor}) < \mathcal{H}_2(P_{N,N-2,\lfloor \frac{N-2}{2} \rfloor - 1}) < \dots < \mathcal{H}_2(P_{N,N-2,1})$.

Proof: Every tree in $\mathcal{T}_{N,N-2}$ must contain a path of length $N-2$ (which contains $N-1$ nodes), and one additional node. This node must be adjacent to an internal node of the path, since otherwise we would have a path of length $N-1$. Thus every tree in $\mathcal{T}_{N,N-2}$ is of the form $P_{N,N-2,i}$, for some $1 \leq i \leq \lfloor \frac{N-2}{2} \rfloor$. Now, $P_{N,N-2,i} = T_{i,N-i-2}^r$, with T a path containing 2 nodes (and one identified as the root). Suppose that $i < \lfloor \frac{N-2}{2} \rfloor$. Then $i < N-i-2$, and so by Lemma 2, $\mathcal{H}_2(P_{N,N-2,i}) = \mathcal{H}_2(T_{i,N-i-2}^r) > \mathcal{H}_2(T_{i+1,N-3-i}^r) = \mathcal{H}_2(P_{N,N-2,i+1})$. ■

Each tree in $\mathcal{T}_{N,N-2}$ consists of a path of length $N-2$ with one leaf attached to an internal node. Theorem 3 ensures that the \mathcal{H}_2 norm is smallest when this internal node is at the center of the path.

C. Caterpillars

We now have complete orderings for $\mathcal{T}_{N,2}$ (trivial, since $\mathcal{T}_{N,2}$ contains only the star), $\mathcal{T}_{N,3}$ (by Theorem 1), $\mathcal{T}_{N,N-2}$ (by Theorem 3) and $\mathcal{T}_{N,N-1}$ (trivial, since $\mathcal{T}_{N,N-1}$ contains only the path). We next consider the remaining families of trees with $4 \leq d \leq N-3$ (and hence, $N \geq 7$).

Rather than deriving complete orderings, the main goal of the next two lemmas is to find the tree in $\mathcal{T}_{N,d}$ with lowest \mathcal{H}_2 norm. However, we use two steps to attain our result as this provides greater insight into the ordering amongst the remaining trees. Lemma 5 then allows us to combine the results to prove (Theorem 4) that among trees of diameter d , the one with lowest \mathcal{H}_2 norm is $P_{N,d,\lfloor \frac{d}{2} \rfloor}$. Theorem 5 provides a comparison of trees with different diameter.

Lemma 3: Suppose $N \geq 7$ and $4 \leq d \leq N-3$. If $T \in \mathcal{C}_{N,d}$, then $\mathcal{H}_2(T) \geq \mathcal{H}_2(P_{N,d,\lfloor \frac{d}{2} \rfloor})$, with equality if and only if $T = P_{N,d,\lfloor \frac{d}{2} \rfloor}$.

Proof: Since $d \leq N-3$ and $T \in \mathcal{C}_{N,d}$, a longest path in T contains $N-d-1 \geq 2$ leaves attached to internal nodes (other than the two leaves in the longest path). Suppose that P_T is a longest path in T . For the rest of this proof, when we refer to leaf nodes and bouquets, we mean leaves not part of P_T , and bouquets made up of these leaves.

Suppose T contains a single bouquet. Thus $T = P_{N,d,i}$ for some $1 \leq i \leq \lfloor \frac{d}{2} \rfloor$. If $i \neq \lfloor \frac{d}{2} \rfloor$, then by Lemma 2, $\mathcal{H}_2(P_{N,d,i}) > \mathcal{H}_2(P_{N,d,i+1})$.

Suppose T contains multiple bouquets. Locate a bouquet furthest from the center of P_T , and move every leaf in this

bouquet one node further from the closest end of P_T . Call this new tree T' , and label the nodes that were moved 1 through n . Then between T and T' , the path lengths between each of these leaves and any other leaf decrease by 1. The path lengths between each of these leaves and $\leq \lfloor \frac{d+1}{2} \rfloor$ nodes on P_T increase by 1, and the path lengths between each of these leaves and $\geq \lfloor \frac{d+1}{2} \rfloor$ nodes on P_T decrease by 1. Thus the sum of the path lengths in T' is less than the sum in T , and so by equations (7) and (6), $\mathcal{H}_2(T') < \mathcal{H}_2(T)$.

Thus, if T is not $P_{N,d,\lfloor \frac{d}{2} \rfloor}$, there is a tree in $\mathcal{C}_{N,d}$ with strictly smaller \mathcal{H}_2 norm. ■

Lemma 4: Suppose that $N \geq 7$ and $4 \leq d \leq N-3$. Let T be a tree in $\mathcal{T}_{N,d} \setminus \mathcal{C}_{N,d}$. Then $\mathcal{H}_2(T) \geq \mathcal{H}_2(N_{N,d})$, with equality if and only if $T = N_{N,d}$.

Proof: Let P_T be a longest path in T (of length d), and let m be the number of nodes with distances to P_T greater than 1 (the distance between a node and P_T is the shortest distance between that node and any node on the path).

If $m > 1$, locate a bouquet with the greatest distance from P_T , label the leaves in this bouquet 1 through n , and label the adjacent node $n+1$. Suppose that either the distance between this bouquet and P_T is greater than 2, or the distance is 2 and another bouquet exists the same distance from P_T . Let T' be the tree formed by moving all leaves in this bouquet one node closer to P_T . By our assumptions, $T' \in \mathcal{T}_{N,d} \setminus \mathcal{C}_{N,d}$. Then $d_{i,n+1}$ increases by 1 for $i = 1, \dots, n$. Conversely, $d_{i,j}$ decreases by 1 for $i = 1, \dots, n$ and $j > n+1$. Since there must be at least $d+2 \geq 6$ of these other nodes (with labels above $n+1$), the sum of all distances in T' is smaller than the sum of all distances in T . Thus $\mathcal{H}_2(T') < \mathcal{H}_2(T)$.

If the bouquet we found has a distance of 2 to P_T , and is the only such bouquet, form T' by moving leaves 1 through $n-1$ one node closer to P_T . Then $T' \in \mathcal{T}_{N,d} \setminus \mathcal{C}_{N,d}$. Now, $d_{i,n}$ and $d_{i,n+1}$ both increase by 1 for $i = 1, \dots, n-1$. However, $d_{i,j}$ decreases by 1 for $i = 1, \dots, n-1$ and j one of the remaining $\geq d+1 \geq 5$ nodes. Thus $\mathcal{H}_2(T') < \mathcal{H}_2(T)$.

If $m = 1$, then T contains a single node at a distance 2 from P_T , and all other nodes in T are either on P_T or adjacent to nodes on P_T . Locate a node on P_T with additional nodes attached that is furthest from the center of the path. Label all nodes attached to this node 1 through n (including the node at distance 2 from P_T if it is connected to P_T through this node), and label this node $n+1$. If $n+1$ is not the $\lfloor \frac{d}{2} \rfloor^{\text{th}}$ internal node on the path (i.e. if $T \neq N_{N,d}$), then form T' by moving all nodes not on P_T that are adjacent to $n+1$ (including the node at distance 2, if present) one node further from the closest end of P_T . Then $T' \in \mathcal{T}_{N,d} \setminus \mathcal{C}_{N,d}$. Furthermore, for $i = 1, \dots, n$, $d_{i,j}$ decreases by 1 for any j not on P_T , increases by 1 for $\leq \lfloor \frac{d}{2} \rfloor$ nodes on P_T and decreases by 1 for $\geq \lceil \frac{d}{2} \rceil + 1$ nodes on P_T . Thus the sum of the distances in T' is less than the sum in T , and so $\mathcal{H}_2(T') < \mathcal{H}_2(T)$.

Hence for every tree in $\mathcal{T}_{N,d} \setminus \mathcal{C}_{N,d}$ other than $N_{N,d}$, there exists another tree in $\mathcal{T}_{N,d} \setminus \mathcal{C}_{N,d}$ with smaller \mathcal{H}_2 norm. ■

Lemma 5: Suppose that $N \geq 7$ and $4 \leq d \leq N-3$. Then $\mathcal{H}_2(P_{N,d,\lfloor \frac{d}{2} \rfloor}) < \mathcal{H}_2(N_{N,d})$.

Proof: Label the node in $N_{N,d}$ that is a distance 2 from the longest path as node 1, and label the node it is adjacent to as node 2. Then we can form $P_{N,d,\lfloor \frac{d}{2} \rfloor}$ from $N_{N,d}$ by moving node 1 one node closer to the longest path. Then $d_{1,j}$ decreases by 1 for $j = 3, \dots, N$, and $d_{1,2}$ increases by 1. Since $N \geq 7$, the sum of all path lengths in $P_{N,d,\lfloor \frac{d}{2} \rfloor}$ is less than in $N_{N,d}$. Thus, by equations (7) and (6), $\mathcal{H}_2(P_{N,d,\lfloor \frac{d}{2} \rfloor}) < \mathcal{H}_2(N_{N,d})$. ■

Now, we have enough to determine the tree in $\mathcal{T}_{N,d}$ with smallest \mathcal{H}_2 norm.

Theorem 4: Let $N \geq 4$ and $2 \leq d \leq N - 2$. The tree in $\mathcal{T}_{N,d}$ with smallest \mathcal{H}_2 norm is $P_{N,d,\lfloor \frac{d}{2} \rfloor}$.

Proof: For $d = 2$, $\mathcal{T}_{N,d}$ only contains $K_{1,N-1}$, which is the same as $P_{N,2,1}$. For $d = 3$, the result follows from Theorem 1 since $D_{N,1,N-3} = P_{N,3,1}$. For $4 \leq d \leq N - 3$, this is a simple consequence of Lemmas 3, 4 and 5. For $d = N - 2$, the result follows from Theorem 3. ■

Finally, we can combine several of our earlier results to obtain a basic comparison between trees of different diameters.

Theorem 5: Let $3 \leq d \leq N - 1$. For any tree in $\mathcal{T}_{N,d}$, there is a tree in $\mathcal{T}_{N,d-1}$ with a smaller \mathcal{H}_2 norm. Hence, the star $K_{1,N-1}$ has the smallest \mathcal{H}_2 norm of any tree with N nodes.

Proof: By Lemma 2, $\mathcal{H}_2(P_{N,N-2,i}) < \mathcal{H}_2(P_N)$ (for any $1 \leq i \leq \lfloor \frac{N-2}{2} \rfloor$).

Let $4 \leq d \leq N - 2$. Suppose $T \in \mathcal{T}_{N,d}$. Then by Theorem 4, $\mathcal{H}_2(T) \geq \mathcal{H}_2(P_{N,d,\lfloor \frac{d}{2} \rfloor})$. But by Lemma 2, $\mathcal{H}_2(P_{N,d,\lfloor \frac{d}{2} \rfloor}) > \mathcal{H}_2(P_{N,d-1,\lfloor \frac{d-1}{2} \rfloor})$. Thus $\mathcal{H}_2(P_{N,d-1,\lfloor \frac{d-1}{2} \rfloor}) < \mathcal{H}_2(T)$.

Let $T \in \mathcal{T}_{N,3}$. Then by Theorem 1, $\mathcal{H}_2(T) \geq \mathcal{H}_2(D_{N,1,N-3})$. But by Lemma 2, $\mathcal{H}_2(K_{1,N-1}) < \mathcal{H}_2(D_{N,1,N-3})$. Thus $\mathcal{H}_2(K_{1,N-1}) < \mathcal{H}_2(T)$. ■

VIII. DISCUSSION

We were able to derive the results in Section VII by determining the effect of moving leaves on effective resistances. With our well-defined definition of “directed resistance” for directed graphs, the same calculations can be made for directed trees as well. Thus our approach in this paper provides a constructive method for deriving a partial ordering of directed trees according to their \mathcal{H}_2 norms. Further details will appear in a future publication. Previous known results on the Wiener index of trees do not provide the same opportunity for the examination of directed trees.

Additionally, the manipulations we used to prove our results suggest how trees can be rearranged to improve their \mathcal{H}_2 norms in a decentralized fashion. In particular, we showed that for a non-star tree, the \mathcal{H}_2 norm can always be reduced either by moving a single node to somewhere else in the tree, or by moving a bouquet of nodes to an adjacent node. These manipulations are “local” in the sense that nodes are moved only from a single location in the tree at a time, and the rest of the nodes in the tree are not required to take any additional action. Thus, we can propose the following

decentralized method to improve the robustness of a tree. Consider a large tree connecting many nodes, each of which has only local information. For example, suppose each node knows the graph between it and a fixed number of other nodes, as well as the degrees of each of these nodes. If the local neighborhood of node i connects to the rest of the tree through a single other node j , and node i is a leaf furthest from node j within its local neighborhood, then i is a candidate to be moved one node closer to j . Once such nodes identify themselves, they could move, form a new tree, and repeat the process.

REFERENCES

- [1] R. Olfati-Saber and R. Murray, “Consensus problems in networks of agents with switching topology and time-delays,” *IEEE Trans. Autom. Control*, vol. 49, no. 9, pp. 1520–1533, 2004.
- [2] L. Moreau, “Stability of multiagent systems with time-dependent communication links,” *IEEE Trans. Autom. Control*, vol. 50, no. 2, pp. 169–182, 2005.
- [3] W. Ren, R. Beard, and E. Atkins, “A survey of consensus problems in multi-agent coordination,” in *Proc. ACC*, 2005, pp. 1859–1864.
- [4] V. Blondel, J. Hendrickx, A. Olshevsky, and J. Tsitsiklis, “Convergence in multiagent coordination, consensus, and flocking,” in *Proc. CDC-ECC*, Seville, Spain, 2005, pp. 2996–3000.
- [5] B. Bamieh, M. Jovanovic, P. Mitra, and S. Patterson, “Effect of topological dimension on rigidity of vehicle formations: Fundamental limitations of local feedback,” in *Proc. CDC*, Cancun, Mexico, 2008, pp. 369–74.
- [6] R. Olfati-Saber and J. Shamma, “Consensus filters for sensor networks and distributed sensor fusion,” in *Proc. CDC-ECC*, 2005, pp. 6698–6703.
- [7] L. Xiao, S. Boyd, and S. Kim, “Distributed average consensus with least-mean-square deviation,” *J. Parallel and Distributed Computing*, vol. 67, no. 1, pp. 33–46, 2007.
- [8] D. Sumpter, J. Krause, R. James, I. Couzin, and A. Ward, “Consensus decision making by fish,” *Curr. Bio.*, vol. 18, pp. 1773–1777, 2008.
- [9] G. Young, L. Scardovi, and N. Leonard, “Robustness of noisy consensus dynamics with directed communication,” in *Proc. ACC*, Baltimore, MD, 2010, pp. 6312–6317.
- [10] Z. Wu, Z. Guan, and X. Wu, “Consensus problem in multi-agent systems with physical position neighbourhood evolving network,” *Physica A Stat. Mech. Appl.*, vol. 379, no. 2, pp. 681–690, 2007.
- [11] L. Scardovi and N. Leonard, “Robustness of aggregation in networked dynamical systems,” in *Proc. ROBOCOMM*, Odense, Denmark, 2009, pp. 1–6.
- [12] A. Dobrynin, R. Entringer, and I. Gutman, “Wiener index of trees: Theory and applications,” *Acta Applicandae Mathematicae*, vol. 66, no. 3, pp. 211–249, 2001.
- [13] H. Dong and X. Guo, “Ordering trees by their Wiener indices,” *MATCH Commun. Math. Comput. Chem.*, vol. 56, no. 3, pp. 527–540, 2006.
- [14] H.-Y. Deng, “The trees on $n \geq 9$ vertices with the first to seventeenth greatest Wiener indices are chemical trees,” *MATCH Commun. Math. Comput. Chem.*, vol. 57, no. 2, pp. 393–402, 2007.
- [15] S. Wang and X. Guo, “Trees with extremal Wiener indices,” *MATCH Commun. Math. Comput. Chem.*, vol. 60, no. 2, pp. 609–622, 2008.
- [16] B. Mohar, “The Laplacian spectrum of graphs,” *Graph theory, Combinatorics and Applications*, vol. 2, pp. 871–898, 1991.
- [17] D. Zelazo and M. Mesbahi, “ \mathcal{H}_2 performance of agreement protocol with noise: An edge based approach,” in *Proc. CDC*, Shanghai, China, 2009, pp. 4747–4752.
- [18] W. Xiao and I. Gutman, “Resistance distance and Laplacian spectrum,” *Theor. Chem. Acc.*, vol. 110, no. 4, pp. 284–289, 2003.
- [19] D. Klein and M. Randić, “Resistance distance,” *J. Math. Chem.*, vol. 12, pp. 81–95, 1993.
- [20] D. Rouvray, “The modeling of chemical phenomena using topological indices,” *J. Comp. Chem.*, vol. 8, no. 4, pp. 470–480, 1987.
- [21] R. Entringer, A. Meir, J. Moon, and L. Székely, “On the Wiener index of trees from certain families,” *Australas. J. Combin.*, vol. 10, pp. 211–224, 1994.
- [22] S. Simić and B. Zhou, “Indices of trees with a prescribed diameter,” *Appl. Anal. and Discrete Math.*, vol. 1, pp. 446–454, 2007.