Health Monitoring of Networked Systems

P. Kingston, M. Egerstedt, and E. Verriest

{*pkingston3,magnus,erik.verriest*}*@ece.gatech.edu*

School of Electrical and Computer Engineering

Georgia Institute of Technology

Atlanta, GA 30332, USA

One of the main motivating forces behind the rapidly emerging theory of decentralized, cooperative control is the idea that we can deploy large collections of agents over wide spatial domains to solve otherwise intractable surveillance and monitoring problems. However, it is inevitable that in such networks, the performance will deteriorate over time as resources, such as battery power, are depleted. In this paper, we address this problem by investigating methods to measure the health of a network in a distributed fashion.

1 Introduction

In this paper we consider the problem of establishing certain properties of a networked system associated with its health. Here, the health (or degradation thereof) manifests itself through a reduced response rate of the individual nodes. As such, the health monitoring problem involves solving a partial system identification problem, where the overall interaction topology of the network is known, while the node weights are not.

Other definitions of the networked health monitoring problem are for instance given in [5, 8, 9], while various aspects of the network system identification problem are discussed in [1, 2, 11, 16]. In fact, the contribution of this paper is in part the formulation of the health monitoring problem as a structured system identification problem. Specifically, we will assume that the health of an individual agent is reflected by the gain (e.g. the available power) associated with that agent's linear, nearest neighbor control law.

The outline of this paper is as follows: In Section 2 we introduce the weighted linear consensus protocol that describes the interaction dynamics. Following this, we identify two classes of networked health monitoring problems, in Sections 3 and 4 respectively. In Section 3, the global network health monitoring problem is introduced and solved in closed form from both instantaneous measurements and sampled data. In Section 4, the individual health monitoring problem is introduced. This problem does in general not have a unique solution. Finally, the conclusions are given in Section 5.

2 Linear Consensus Protocols with Unknown Weights

2.1 Linear Consensus Protocols

Consider a system of n networked agents connected in an undirected graph G, where each agent $i \in \{1, 2, \ldots, n\}$ has state $x_i \in \mathbb{R}^p$. Under the standard, linear consensus protocol (e.g. [3, 4, 7, 10, 12, 13, 15, 17]), the i^{th} agent's dynamics may be written as

$$\dot{x}_i = \sum_{j \in N_i} (x_j - x_i) \tag{1}$$

where N_i is the set of all agents in the neighborhood of – or adjacent to – agent i, given a static, undirected network topology defined through the graph G.

We note that (1) is decoupled along each dimension of x_i and, as such, we can simply assume that $x_i \in \mathbb{R}$ without loss of generality. Under this assumption we can rewrite (1) to obtain,

$$\dot{x} = (\mathcal{A} - \mathcal{D})x = -Lx, \tag{2}$$

where $x = [x_1, \ldots, x_n]^T$, $\mathcal{A}$ and $\mathcal{D}$ are G's *adjacency* and *degree* matrices, respectively, and L is the *Graph Laplacian*. It is known from algebraic graph theory (e.g. [6]) that if G is connected, the smallest eigenvalue of L is zero and all others are strictly positive. Additionally, $\mathcal{N}(L) = \text{span}(\mathbf{1})$. Thus, as shown in [7, 13, 15], this system is critically stable and converges to $x_1 = x_2 = \ldots = x_n = c$, where c is the time-invariant centroid, given by

$$c = \frac{1}{n}\mathbf{1}^T x. \tag{3}$$

2.2 Health Monitoring

As an agent's power level decreases or it accumulates damage, we assume that it will be able to exert less control effort and so react more slowly to changes in its environment. To represent this, we introduce a constant factor, γ_i, to the dynamics described in (1), which represents the *health* of agent i. Healthy agents – like robots with full batteries and no mechanical problems – move faster than unhealthy ones and have larger healths. With this modification, (1) is replaced by

$$\dot{x}_i = \gamma_i \sum_{j \in N_i} (x_j - x_i), \tag{4}$$

or equivalently

$$\dot{x} = -\Gamma L x, \tag{5}$$

where $\Gamma \triangleq \text{diag}(\gamma_1, \ldots, \gamma_n)$ is the *health matrix* of the network. It is shown in [7] that this system converges to the weighted centroid

$$c_\Gamma = \frac{1}{\text{tr}(\Gamma^{-1})} \mathbf{1}^T \Gamma^{-1} x \tag{6}$$

which is also invariant in time.

Further suppose that a single agent, which we may think of as "doing the health monitoring," is given access to the state of one other agent, whose state it may monitor relative to its own. Without loss of generality we may say that the last agent monitors the first agent, in which case we have access to the signal

$$y = x_1 - x_n = [\begin{array}{cccccc} 1 & 0 & \ldots & 0 & -1 \end{array}] x. \tag{7}$$

We are now interested in whether it is possible to determine Γ for the system defined by (5) and (7) given,

1. the measurement signal y, and
2. known, static (or piecewise static), connected network topology (i.e., L is given).

2.3 Alternate Interpretations

The determination of unknown consensus gains is equivalent to a number of analogous physical, signal processing, and communications problems:

1. **Sensor networks:** A network of sensors may be deployed to measure a single underlying quantity and arrive at a consensus – e.g., the sensors are thermometers distributed over a region in which the temperature is assumed to be constant, and we wish to know what that temperature is. Especially in a heterogeneous network, we may expect some sensors to be more accurate than others, in which case we would like them to influence the consensus point more than the less accurate nodes. From (6), we can achieve this by reinterpreting the *health* γ_i of agent i as the reciprocal of the *accuracy*, a_i, of that agent; i.e., $a_i = 1/\gamma_i$. In this case, we may think of the problem not as "health monitoring" but instead as "accuracy monitoring."
2. **RC electrical networks:** If we replace each edge of the graph G by a 1Ω resistor, and connect a capacitor between each vertex and a common ground, then we arrive at an electrical system whose dynamics are identical to those of weighted linear consensus, where the capacitances – which are unknown – are analogous to the reciprocals of the consensus gains; i.e., $C_i = 1/\gamma_i$.[1]
3. **Broadcast protocols:** One may view each γ_i most generally as some datum that agent i would like every other agent to receive.

[1] This neglects units. More generally, if capacitors are connected by resistances R, then $\gamma_i = 1/(RC_i)$ and has units of angular frequency.

2.4 Egocentric Coordinates

The n^{th}-order system given by (5) and (7) can always be expressed by a smaller $(n-1)^{th}$-order system which we introduce in this section. We do this by expressing the states of the first $n-1$ agents in the egocentric coordinates of the last agent, which performs the health monitoring. To begin, we block-partition (5) as,

$$\begin{bmatrix} \dot{x}_v \\ \dot{x}_n \end{bmatrix} = -\begin{bmatrix} \Gamma_m & \mathbf{0} \\ \mathbf{0} & \gamma_n \end{bmatrix} \begin{bmatrix} L_m & L_v \\ L_v^T & l_{nn} \end{bmatrix} \begin{bmatrix} x_v \\ x_n \end{bmatrix} \tag{8}$$

where $x_n \in \mathbb{R}$ is the state of the n^{th} agent, $x_v \in \mathbb{R}^{n-1}$ is the vector of the first $n-1$ agents' states, and the other quantities are defined so as to have compatible dimensions, as in [14][2].

Next, we define a state vector $w \in \mathbb{R}^{n-1}$ so that w represents x_v in agent n's egocentric coordinate frame; i.e.

$$w \triangleq x_v - \mathbf{1}x_n \tag{9}$$

and the measurement y that agent n makes of agent 1 can be expressed in terms of w simply as

$$y = [1, 0, \dots, 0]w = cw. \tag{10}$$

From (8) and (9),

$$\begin{aligned} \dot{w}_v &= \dot{x}_v - \mathbf{1}\dot{x}_n \\ &= (-\Gamma_m L_m x_v - \Gamma_m L_v x_n) - \mathbf{1}\left(-\gamma_n L_v^T x_v - \gamma_n l_{nn} x_n\right) \\ &= \left(-\Gamma_m L_m + \gamma_n \mathbf{1} L_v^T\right) x_v + \left(-\Gamma_m L_v + \gamma_n l_{nn}\mathbf{1}\right) x_n \\ &= \left(-\Gamma_m L_m + \gamma_n \mathbf{1} L_v^T\right) w + \left(-\Gamma_m L_m \mathbf{1} + \gamma_n \mathbf{1} L_v^T \mathbf{1} - \Gamma_m L_v + \gamma_n l_{nn}\mathbf{1}\right) x_n \end{aligned} \tag{11}$$

and since $L_m\mathbf{1} = -L_v, L_v^T\mathbf{1} = -\deg(n), l_{nn} = \deg(n)$, then (11) becomes

$$\begin{aligned} &\left(-\Gamma_m L_m + \gamma_n \mathbf{1} L_v^T\right) w + \left(\Gamma_m L_v - \gamma_n \deg(n)\mathbf{1} - \Gamma_m L_v + \gamma_n \deg(n)\mathbf{1}\right) x_n \\ &= \left(-\Gamma_m L_m + \gamma_n \mathbf{1} L_v^T\right) w = A_\Gamma w \end{aligned} \tag{12}$$

where A_Γ is defined by the above. (Note that $A_\Gamma = A_\Gamma^T \prec 0$, as shown in [14].)

We thus obtain a new, autonomous linear system whose dynamics may be written as

$$\dot{w}_v(t) = \left(-\Gamma_m L_m + \gamma_n \mathbf{1} L_v^T\right) w = A_\Gamma w(t) \tag{13}$$

$$y(t) = cw(t) \tag{14}$$

In fact, this is the system with which we will be concerned in this paper.

[2] That is, $L_v \in \mathbb{R}^{n-1}, l_{nn} \in \mathbb{R}$, and $\Gamma_m, L_m \in \mathbb{R}^{(n-1)\times(n-1)}$.

3 Global Health Monitoring

One question we may ask is, "How healthy is the network as a whole?" A natural definition of such a *global network health measure*, J_Γ, is as a linear combination of individual agent healths, weighted by each agent's degree (which we may take to represent the "importance" of an agent in the network):

$$J_\Gamma \triangleq \sum_{i=1}^{n} \gamma_i \deg(i). \tag{15}$$

This definition agrees with the intuition that, if an agent which communicates with few other agents has low health, this will affect the network less adversely than if the health of an agent which communicates with many others is low.

3.1 Instantaneous Global Health Monitoring

Theorem 1. *It is possible to determine J_Γ instantaneously from the measurement signal y as long as (A_Γ, w_0, c) is* minimal.[3]

Proof. Let $y^{(i)}$ be the i^{th} derivative of $y(t)$ at $t = 0$, which we can in principle measure, and let Y and $\vec{Y}$ be the Hankel and shifted Hankel matrices of these derivatives: [4] [5]

$$Y = \begin{bmatrix} y^{(0)} & y^{(1)} & \dots & y^{(n-2)} \\ y^{(1)} & y^{(2)} & \dots & y^{(n-1)} \\ \vdots & & & \vdots \\ y^{(n-2)} & y^{(n-1)} & \dots & y^{(2n-3)} \end{bmatrix} \quad \vec{Y} = \begin{bmatrix} y^{(1)} & y^{(2)} & \dots & y^{(n-1)} \\ y^{(2)} & y^{(3)} & \dots & y^{(n)} \\ \vdots & & & \vdots \\ y^{(n-1)} & y^{(n)} & \dots & y^{(2n-2)} \end{bmatrix} \tag{16}$$

We can factor Y and $\vec{Y}$ as,

$$Y = O[x^{(0)}, \dots, x^{(n-2)}] \tag{17}$$

$$\vec{Y} = OA[x^{(0)}, \dots, x^{(n-2)}] \tag{18}$$

where A is the system matrix for an arbitrary similar realization, $[x^{(0)}, \dots, x^{(n-2)}]$ contains the $0, \dots, (n-2)^{th}$ derivatives of the state in this realization, and O is this realization's observability matrix.

Now, consider the *similar* system in observability canonical form. We know that the observability matrix in this form is the identity (for SISO) systems. So long as (A_Γ, w_0, c) is minimal, Y is invertible and hence, we get

$$A = O^{-1}\vec{Y}Y^{-1}O = \vec{Y}Y^{-1}. \tag{19}$$

[3] Here, w_0, the initial state, takes the place of the usual control matrix b.

[4] The beginning of this proof follows the same development as Section 3 of [18], which establishes a connection between Networked Health Monitoring and Symmetric Symmetrizers.

[5] Recall that our system is $(n-1)^{th}$ order

This matrix is similar to A_Γ and is of the form

$$A = \begin{bmatrix} 0 & 1 & \dots & 0 \\ 0 & 0 & \ddots & 0 \\ \vdots & & & \vdots \\ 0 & \dots & 0 & 1 \\ -a_{n-1} & -a_{n-2} & \dots & -a_1 \end{bmatrix} \tag{20}$$

Thus, for some invertible T, [6]

$$A_\Gamma = TAT^{-1} \tag{21}$$

Since the trace of a matrix is preserved under similarity transformations, this implies that

$$\operatorname{tr}(A_\Gamma) = \operatorname{tr}(A) = -a_1, \tag{22}$$

where

$$\begin{aligned} \operatorname{tr}(A_\Gamma) &= \operatorname{tr}\left(-\Gamma_m L_m + \gamma_n \mathbf{1} L_v^T\right) \\ &= -\sum_{i=1}^{n-1} \gamma_i \deg(i) + \sum_{i=1}^{n-1} \begin{cases} -\gamma_n & \text{if } i \in N_n \\ 0 & \text{otherwise} \end{cases} \\ &= -\sum_{i=1}^{n-1} \gamma_i \deg(i) - \gamma_n \deg(n) \\ &= -J_\Gamma. \end{aligned} \tag{23}$$

Hence,

$$J_\Gamma = -\operatorname{tr}(A_\Gamma) = -\operatorname{tr}(A) = a_1. \tag{24}$$

Since we know A from instantaneous measurement of $y(\cdot)$ at $t = 0$, we also instantaneously know J_Γ, the global health. □

Thus, provided (A_Γ, w_0, c) is minimal, global health monitoring is possible by having a single agent, which need not know its own health, watch just one other agent in the network. (For a discussion of network minimality, see for example [14].)

3.2 Discrete-Time Global Health Monitoring

The method described in the proof of Theorem 1 establishes that in principle global health monitoring is possible, but since it relies on differentiating $y(t)$ many times, it will generally not itself be a practical means to do so. In this section, we describe how to find the global network health instead from sampled data.

The discrete-time system obtained by sampling the output of (13) with a sampling period τ is given by

$$\begin{aligned} w_{k+1} &= (e^{A_\Gamma \tau}) w_k = \bar{\Phi}_\tau w_k \\ y_k &= c w_k, \end{aligned} \tag{25}$$

[6] In fact, $T = O(A_\Gamma, c)^{-1}$.

where $w_{k+1} \triangleq w(k\tau)$, $y_k \triangleq y(k\tau)$, and $k \in \{0, 1, 2, \ldots\}$. As in the continuous-time case, we construct Hankel matrices of these samples,

$$Y_\tau = \begin{bmatrix} y_0 & y_1 & \cdots & y_{n-2} \\ y_1 & y_2 & \cdots & y_{n-1} \\ \vdots & & & \vdots \\ y_{n-2} & y_{n-1} & \cdots & y_{2n-3} \end{bmatrix} \quad \vec{Y}_\tau = \begin{bmatrix} y_1 & y_2 & \cdots & y_{n-1} \\ y_2 & y_3 & \cdots & y_n \\ \vdots & & & \vdots \\ y_{n-1} & y_n & \cdots & y_{2n-2} \end{bmatrix} \tag{26}$$

and find the discrete-time observability-canonical-form transition matrix

$$\Phi_\tau = \vec{Y}_\tau Y_\tau^{-1}, \tag{27}$$

where Φ_τ is of the form,

$$\Phi_\tau = \begin{bmatrix} 0 & 1 & \ldots & 0 \\ 0 & 0 & \ddots & 0 \\ \vdots & & & \vdots \\ 0 & \ldots & 0 & 1 \\ -\alpha_{n-1} & -\alpha_{n-2} & \ldots & -\alpha_1 \end{bmatrix}. \tag{28}$$

Since the matrix exponential has the property that, for some matrix Z, $\det e^Z = e^{\operatorname{tr} Z}$, then

$$\det \Phi_\tau = \det \bar{\Phi}_\tau = \det e^{A_\Gamma \tau} = e^{\operatorname{tr} A_\Gamma \tau} = e^{-J_\Gamma \tau} \tag{29}$$

and

$$J_\Gamma = -\frac{1}{\tau} \ln |\Phi_\tau| = -\frac{1}{\tau} \ln [(-1)^n \alpha_{n-1}]. \tag{30}$$

In this manner, the global network health can be obtained from samples of the output.

3.3 Interpretations of the Global Network Health

State-Space Volumes

The relationship in (29) between J_Γ and $|\Phi_\tau|$ gives an interpretation of the global network health as the reciprocal of the time constant with which volumes in the state-space “shrink.” For instance, suppose it is known that in the worst case the initial condition w_0 will lie on some closed surface S_0 in state-space – say, an $(n-1)$-ball – enclosing volume V_0, and that all points on S_0 evolve by (12) to give a time-varying surface $S(t)$ enclosing volume $V(t)$. Then, from (29),

$$V(t) = V_0 e^{-J_\Gamma t} \qquad \forall t \geq 0. \tag{31}$$

Thus $V(t)$ obeys a first-order model, and we know from the global health alone that the volume of state space which may contain the state decreases at a known rate. Moreover, this rate increases monotonically with increased global health.

System Poles

The global network health can also be understood in terms of the eigenvalues of the system matrix A_Γ. Since A_Γ is trivially similar to its Jacobi form Λ_Γ, we have

$$J_\Gamma = -\operatorname{tr}(A_\Gamma) = -\operatorname{tr}(\Lambda_\Gamma) = -\sum_{i=1}^{n-1} \lambda_i. \tag{32}$$

This observation makes clear why J_Γ is available from the output signal y: It is the sum of the poles of the Laplace transform of the output, $\mathcal{L}\{y\}$.

4 Individual Health Monitoring

The global health measure J_Γ does not provide any information about how well the individual agents are performing. A natural question to ask is thus: *Is knowledge of $y(.)$, L, and γ_n sufficient for agent n to deduce all of the other agent healths $\gamma_1, \ldots, \gamma_{n-1}$?* It turns out that in general the answer is no, and we will show this by counterexample by demonstrating that there is more than one health matrix that can explain the measured output.

In the context of this paper, we will define a *plausible* health matrix $\tilde{\Gamma}$ as any diagonal matrix such that the system $(A_{\tilde{\Gamma}}, c)$ is *similar* to the system (A_Γ, c), which we call the *true* system. The significance of this is that initial conditions exist for the *plausible* system such that it produces the same output y as the *true* system.

Consider the observability canonical form realization of the true system, (13). This realization has system matrix (20), and is related to the system in natural coordinates by (21), i.e. $z = O(A_\Gamma, c)w$. It has dynamics

$$\begin{aligned} \dot{z} &= Az = O(A_\Gamma, c)A_\Gamma O(A_\Gamma, c)^{-1} z \\ y &= [\ 1 \quad 0 \quad \ldots \quad 0\ ]z \end{aligned} \tag{33}$$

Next suppose that we can find a *plausible* health matrix $\tilde{\Gamma}$, such that the system $(A_{\tilde{\Gamma}}, c)$ has the same system matrix A in observability canonical form. The state vector $\tilde{w}$ for this system is likewise related to z by $z = O(A_{\tilde{\Gamma}}, c)\tilde{w}$, and hence

$$z = O(A_\Gamma, c)w = O(A_{\tilde{\Gamma}}, c)\tilde{w},$$

and hence

$$\tilde{w} = O(A_{\tilde{\Gamma}}, c)^{-1} O(A_\Gamma, c)w = Mw, \tag{34}$$

where M is defined by the above. Therefore, if we choose the initial condition $\tilde{w}_0 = Mw_0$ for the plausible system, then its output $\tilde{y}$ will match that of the true system, i.e.

$$\begin{aligned} y &= ce^{A_\Gamma t} w_0 = ce^{M^{-1} A_{\tilde{\Gamma}} M t} w_0 = cM^{-1} e^{A_{\tilde{\Gamma}} t} M w_0 \\ &= cO(A_\Gamma, c)^{-1} O(A_{\tilde{\Gamma}}, c) e^{A_{\tilde{\Gamma}} t} \tilde{w}_0 = ce^{A_{\tilde{\Gamma}} t} \tilde{w}_0 = \tilde{y} \end{aligned} \tag{35}$$

We see that if such a plausible health matrix exists, then there are initial conditions $\tilde{w}_0 = O(A_{\tilde{\Gamma}}, c)^{-1} O(A_\Gamma, c) w_0$ for which the *plausible* system will produce the same output as the *true* system. We discuss how to find plausible health matrices in the next section.

Finding Counterexamples

We consider $\bar{\gamma} \triangleq [\gamma_1 \ldots \gamma_{n-1}]^T$ to represent the true healths of the agents, and will find another vector of healths $\tilde{\gamma}$ which maps to the same observability-canonical-form system matrix.

The observability-canonical form system matrix A (20) is entirely specified by its bottom row, which we will call $\vec{a}^T \in \mathbb{R}^{n-1}$. The observability matrix $O(A_\Gamma, c)$ gives the similarity transformation which relates A to A_Γ, and hence

$$A = O(A_\Gamma, c) A_\Gamma O(A_\Gamma, c)^{-1},$$

i.e.

$$\begin{aligned} \vec{a}^T &= (cA_\Gamma^{n-1}) A_\Gamma O(A_\Gamma, c)^{-1} = cA_\Gamma^n O(A_\Gamma, c)^{-1} \\ &= c\left(-\Gamma_m L_m + \gamma_n \mathbf{1} L_v^T\right)^n O(-\Gamma_m L_m + \gamma_n \mathbf{1} L_v^T, c)^{-1}. \end{aligned} \tag{36}$$

Now, we view the search for counterexamples as the root-finding problem in $\tilde{\gamma}$,

$$f(\tilde{\gamma}) - f(\bar{\gamma}) = 0, \tag{37}$$

where $f : \mathbb{R}^{n-1} \longmapsto \mathbb{R}^{n-1}$ is the function,

$$f(\xi) = c\left(-\operatorname{diag}(\xi) L_m + \gamma_n \mathbf{1} L_v^T\right)^n O(-\operatorname{diag}(\xi) L_m + \gamma_n \mathbf{1} L_v^T, c)^{-1}. \tag{38}$$

This is a coupled system of polynomial equations in $\tilde{\gamma}_1 \ldots \tilde{\gamma}_{n-1}$, which we solve numerically to produce the counterexamples which follow.

Counterexample 1: Line Topology

As a limiting example we begin with four agents connected in a straight line as shown in Figure 1; this is the connected topology with the fewest edges. We represent the states of each of the agents in the egocentric coordinates of the fourth agent, and measure the state of the first agent. The agents behave according to (13), with $\gamma_1, \ldots, \gamma_4$ as indicated in Figure 1, and compute the system matrix in observability canonical form, A, as shown in the same figure.[7]

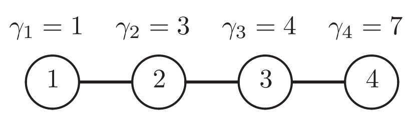

$$A = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ -145 & -124 & -22 \end{bmatrix}$$

Figure 1. *Network 1 - Line Topology*

We can locate other healths $\gamma_1, \ldots, \gamma_{n-1}$ which also map to the same A; e.g.,

$$\gamma_1 \approx 7.370 \quad \gamma_2 \approx -0.780 \quad \gamma_3 \approx 4.600 \tag{39}$$

Thus, we have shown by counterexample that even in the limiting case of straight line topology, A and hence our measurement signal $y(.)$ does not provide enough information to uniquely identify the individual agent healths. Such counterexamples are not unique: For a given topology there exist many $\gamma_1, \ldots, \gamma_{n-1}$ which can explain the measurement $y(\cdot)$.

[7]Note that as expected from (24), $a_1 = 22 = J_\Gamma$.

Counterexample 2: A More Complex Network

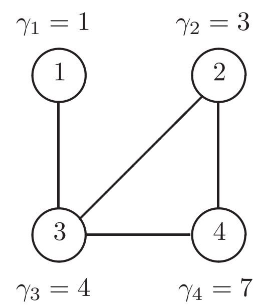

$$A = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ -435 & -291 & -33 \end{bmatrix}$$

Figure 2. *Network 2*

Consider a more complicated network, shown in Figure 2. As before, states are represented in agent 4's egocentric coordinates, while the state of agent 1 is measured, and A is computed. Also as before, we can find another $\tilde{\Gamma}$ which produces the same A and so can explain our measurements:

$$\gamma_1 \approx 10.17 \quad \gamma_2 \approx -0.821 \quad \gamma_3 \approx 3.562. \tag{40}$$

Thus, even under more general topologies, $\gamma_1, ..., \gamma_{n-1}$ are not unique.

5 Concluding Remarks

In this paper we discuss how to establish health properties of linear networks with nearest-neighbor interaction topologies, under the assumption that the health is reflected by the gain in the consensus equation. In particular, we show that a global health measure can be found by only observing a single agent while local (individual) health measurements are not in general available. We also give a method for global health monitoring using sampled data, and present a geometric interpretation of the global network health.

Acknowledgments

The work by Magnus Egerstedt was supported by the U.S. Army Research Office through Grant #99838.

Bibliography

[1] B.R. Chen, G. Peterson, G. Mainland, and M. Welsh. LiveNet: Using Passive Monitoring to Reconstruct Sensor Network Dynamics. In *Proceedings of the 4th IEEE/ACM International Conference on Distributed Computing in Sensor Systems (DCOSS 2008)*, Santorini Island, Greece, June 2008.

[2] A. Das and M. Mesbahi. "Distributed linear parameter estimation in sensor networks," *IEEE Conference on Sensor, Mesh, and Ad Hoc Communications and Networks*, 2006.

[3] J. A. Fax and R. M. Murray. Graph laplacian and stabilization of vehicle formations. In *Proceedings of the 15th IFAC Conf*, pages 283288, 2002.

[4] J.A. Fax and R.M. Murray. Information flow and cooperative control of vehicle formations. *IEEE Trans. Automat. Contr.*, 49:14651476, Sept 2004.

[5] M. Franceschelli, A. Giua, and M. Egerstedt. Motion Probes for Fault Detection and Recovery in Networked Control Systems. *American Control Conference*, Seattle, WA, June 2008.

[6] C. Godsil and G. Royle. *Algebraic graph theory*. Springer, 2001.

[7] M. Ji and M. Egerstedt. Distributed Coordination Control of Multi-Agent Systems While Preserving Connectedness. IEEE Transactions on Robotics, Vol. 23, No. 4, pp. 693-703, Aug. 2007.

[8] J. Kurien, X. Koutsoukos, and F. Zhao, "Distributed Diagnosis of Networked Embedded Systems", Proceedings of the 13th *International Workshop on Principles of Diagnosis (DX-2002)*, pp. 179-188, Semmering, Austria, May 2-4, 2002.

[9] W. Leal, S. Bapat, T. Kwon, P. Wei, and A. Arora. Stabilizing Health Monitoring for Wireless Sensor Networks. In *Book Stabilization, Safety, and Security of Distributed Systems*, Springer Verlag Lecture Notes in Computer Science, pp. 395-410, Berlin 2006.

[10] Z. Lin, M. Broucke, and B. Francis. Local control strategies for groups of mobile autonomous agents. *IEEE Trans. Automat. Contr.*, 49(4):622629, 2004.

[11] Abubakr Muhammad and Ali Jadbabaie, "Decentralized Computation of Homology Groups in Networks by Gossip", *American Control Conference*, 2007.

[12] R. Olfati-Saber and R. M. Murray. Distributed structural stabilization and tracking for formations of dynamic multi-agents. In *Proceedings of the 41st IEEE Conference on Decision and Control*, volume 1, pages 209215, December 2002.

[13] R. Olfati-Saber and R. M. Murray. Consensus problems in networks of agents with switching topology and time-delays. *IEEE Trans. Automat. Contr.*, 49:15201533, Sept 2004.

[14] A. Rahmani, M. Ji, M. Mesbahi, and M. Egerstedt. Controllability of Multi-Agent Systems from a Graph-Theoretic Perspective. *SIAM Journal on Control and Optimization*. To appear 2008.

[15] W. Ren and R.W. Beard. Consensus of information under dynamically changing interaction topologies. In *Proceedings of the American Control Conference*, volume 6, pages 49394944, June 30-July 2 2004.

[16] J. Sandhu, M. Mesbahi, and T. Tsukamaki. "Relative sensing networks: observability, estimation, and the control structure," *IEEE Conference on Decision and Control*, 2005.

[17] H. Tanner, A. Jadbabaie, and G.J. Pappas. Stable flocking of mobile agents, part II : Dynamic topology. In *Proceedings of the 42nd IEEE Conference on Decision and Control*, pages 20162021, 2003.

[18] E. Verriest. "Symmetric Symmetrizers and Some Applications," In *Proceedings of the Eighteenth International symposium on Mathematical Theory of Networks and Systems (MTNS2008)*, Blacksburg, VA, 2008.