Controllability Characterizations of Leader-Based Swarm Interactions

Jean-Pierre de la Croix and **Magnus Egerstedt**
School of Electrical and Computer Engineering
Georgia Institute of Technology, Atlanta, GA 30308
{jdelacroix,magnus}@gatech.edu

Abstract

In this paper, we investigate what role the network topology plays when controlling a network of mobile robots. This is a question of key importance in the emerging area of human-swarm interaction and we approach this question by letting a human user inject control signals at a single leader-node, which are then propagated throughout the network. Based on a user study, it is found that some topologies are more amenable to human control than others, which can be interpreted in terms of the rank of the controllability matrix of the underlying network dynamics, as well as, measures of node centrality on the leader of the network.

Introduction

When it comes to controlling certain classes of systems, humans have a clear intuition about how input signals translate into motions, e.g., what the effect is when turning the steering wheel of a car. For other types of systems, this intuition becomes less clear and significant training is required to be able to operate the system successfully, as is for example the case when remotely piloting helicopters. In this paper, we investigate an untrained operator's ability to control a large collection of mobile robots.

The trend in human-controlled robotic systems is to move away from the current many-to-one paradigm, where multiple operators are required to control a single system (e.g., (Mekdeci and Cummings 2009)) to a *one-to-many* mode of operation in which a single operator can control a large number of autonomous vehicles. For this transition to be successful, it is critical that we gain a clear understanding of how humans should interact with a large number of robots, which is the human-swarm-interaction problem, e.g., (McLurkin et al. 2006), (Cummings 2004), (Kira and Potter 2009).

The particular approach pursued in this paper is to let the team of mobile robots be structured in a static interaction network over which information is flowing between the robots, as is the case in (Olfati-Saber and Murray 2004), (Ren, Beard, and Atkins 2005). We let the robots execute a coordinated control protocol based on the relative displacements between the robots in the network that ensures that the robots are appropriately spaced. The way human inputs are injected is by selecting a particular robot as the leader in the network and then letting the operator control the movements of that robot. The human inputs are then propagated through the network to the remaining, so-called followers.

This partitioning of the network into leader and follower robots is quite standard and much work has been done focusing on how such networks should be controlled (e.g., (Fax and Murray 2004)). While alternative strategies may exist for controlling mobile-robot networks, this paper focuses explicitly on the leader-follower paradigm due to its wide-spread use. The importance of this work lies in the connection we will make between the system theoretic properties and the user experience, rather than drawing conclusions about whether this is the best control paradigm or not.

The *network topology* itself plays a key role in how effectively networks of robots can be controlled and a significant body of work has emerged trying to connect *controllability* properties of multi-robot networks to the underlying network topology. (See for example (Meshabi and Egerstedt 2010), (Lozano et al. 2009)). However, controllability studies only establish what can theoretically be achieved in the multi-robot networks. But, these properties do not tell us much in terms of how easy or hard it is for human operators to *actually* control certain types of networks. And this is exactly the question under consideration in this paper. *What role does the network topology play when human operators are to control multi-robot systems?*

To answer this question, we conduct a user study where people are to control teams of simulated mobile robots. In particular, the participants were asked to rate the difficulty of the control task across different network topologies. Our findings in this paper indicate that the already established controllability properties do tell part of the story, but they do not tell the whole story.

The User Study

Network Topologies

As the ultimate aim with this paper is to understand what role the network topology plays when human users control multi-robot networks, we first need to discuss which of the many possible network topologies should be considered. We restrict our discussion to a small set of common network topologies in order to generate results that are both repre-

sentative and practically useful.

If the robot network is abstracted to a graph, the *topology* refers to the way in which the nodes in the graph are connected. We only consider connected network topologies in this paper, where it is possible to follow edges from any node in the network to any other node. The least connected (in terms of algebraic connectivity) of the connected graphs over n nodes is the so-called line graph, L_n, where each node, with the exception of the terminating nodes, is connected to two other nodes. This is a very natural organization, found for example in single-file military columns. We consider three leader locations–the head node, a node offset from the head, or the center node of a line graph.

The line graph is a degenerate example of an acyclic (or tree) graph. Another typical example of an acyclic graph is the star graph S_n, consisting of a central node connected to all other nodes in the network. These peripheral nodes do not share edges with any additional nodes, and this topology is found, for example, in certain communication networks where a central hub shares and receives information from the additional nodes. We will let the central as well as a peripheral node be the leader in the user study.

The next canonical graph is the cycle graph, C_n, consisting of a "closed" line graph, where each node shares an edge with two other nodes in the graph. Cycle graphs are found naturally in certain social contexts (such as games).

Finally, the complete graph K_n is a graph where each nodes is connected to all other nodes. This type of structure is common in communication networks (broadcast-based) or when a small number of mobile robots are coordinating their behaviors. As all nodes are connected to all other nodes, it does not matter which node becomes the leader under this topology. Table 1 summarizes the network topologies used in the user study and defines a new notation we use to encode the network topology with a leader location, e.g., $S_{7,p}$ for a star graph with a peripheral leader node.

Experimental Setup

The purpose of the user study is to measure the perceived difficulty of controlling a particular network topology with a single leader and for this, 18 participants are presented with 14 tasks. The order in which the tasks are given is randomized each time, such that any learning and order effects in the data are minimized. Each task consists of moving the network via the leader from its initial configuration to a target formation. For example, the participant may be asked to move a L_7 network from its initial configuration into an ellipse. Table 1 provides a detailed list of the 14 tasks. The tasks are selected such that all networks are paired with each target formation, and a target formation is sufficiently different from a network's initial configuration. These two conditions ensure that none of the tasks are trivial for the participant to complete (e.g., form circle from a C_7 network).

The experiment is structured such that the participant is shown the network topology only prior to the start of the task. Communication links are not visible to the eye and it is up to the participant to infer the behavior of the network from the interactions of the robots. The participant is only able to directly control the movement of the leader during the experiment using a joystick. The participant receives no feedback (e.g., a scoring meter) during the task, such that the focus is entirely on matching the network to the target formation. A score is calculated from a least square fit of the network's current configuration to its target formation. Since we are only concerned about the participant matching the formation, the least square fit is translation, rotation, and assignment invariant, meaning that neither the location of the formation in the workspace, nor the assignment of robots to specific positions in the formation matter, following the developments in (Ji, Azuma, and Egerstedt 2006).

Table 1: Network configuration, leader location, and target configuration for each task.

Tasks	Network	Leader	Notation	Targets
1, 8	L_7	Head	$L_{7,h}$	Ellipse, Wedge
2, 9	L_7	Offset	$L_{7,o}$	Ellipse, Wedge
3, 10	L_7	Center	$L_{7,c}$	Ellipse, Wedge
4, 11	C_7	Any	C_7	Ellipse, Wedge
5, 12	K_7	Any	K_7	Ellipse, Wedge
6, 13	S_7	Center	$S_{7,c}$	Ellipse, Wedge
7, 14	S_7	Periphery	$S_{7,p}$	Ellipse, Wedge

After each task, we measure the participant's experience. The participant rates the difficulty of the task on a scale from very easy (0.0) to very hard (20.0), and completes a NASA Task Load Index (TLX) workload survey (Hart and Staveland 1988). The workload survey consists of six questions that cover physical, mental, and temporal demands, as well as levels of performance, effort, and frustration.

Experimental Results

Mean LSQ, rating, and workload scores are calculated for each task from the ratings provided by the participants, the least squares fit errors, and the total raw score of the TLX survey, respectively. We also record the time and the distance the network traveled for each task. A repeated measures one-way ANOVA statistical test (Girden 1991) reveals that the LSQ score ($p < 0.0000001$), rating score ($p = 0.00138$), and workload score ($p = 0.0256$) are statistically significant at a 0.05 (or 95%) confidence level. We use these measures to draw comparisons between the different tasks, and moreover, between the different network topologies. While the time data ($p = 0.012$) is also statistically significant, the distance data ($p = 0.262$) is not statistically significant enough to use as a measure to distinguish between the tasks. Since we do not ask participants to minimize time or distance, we omit both measures from our analysis.

Figure 1 is a histogram of the mean LSQ scores for each task with error bars denoting the standard error. Standard error is computed by normalizing the standard deviation, σ, with the square root of the number of samples, $\sqrt{n}$. The standard error expresses the region in which we can be confident that the true population mean, μ, lies (see (Cumming, Fidler, and Vaux 2007)). If we want to claim that one task received a lower score than another task, we have to check for a statistically significant difference in the regions denoted by

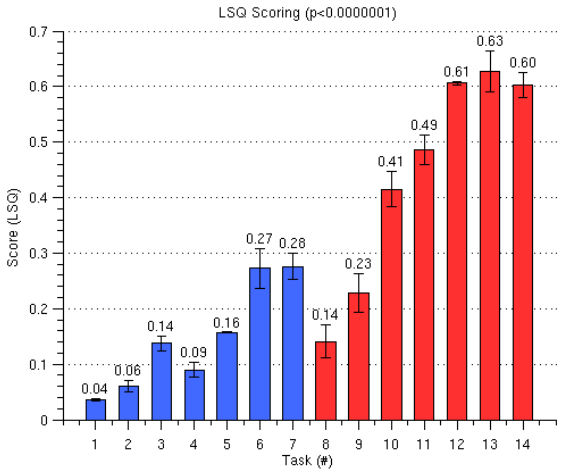

Figure 1: Mean LSQ score for each task.

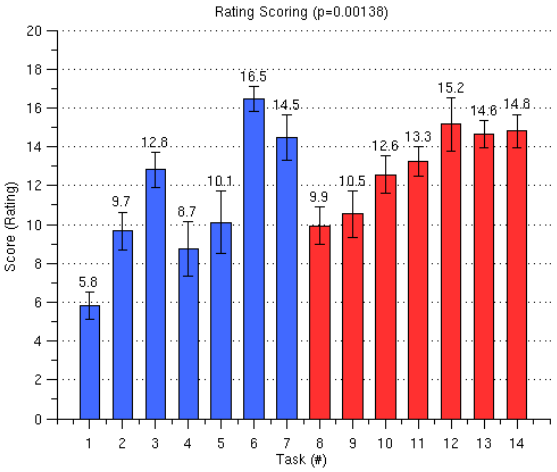

Figure 2: Mean rating score for each task.

the error bars of the respective tasks. The repeated measures ANOVA (analysis of variance) test performs these pairwise comparisons and reports the significance level of any difference. If a difference is statistically significant, we are justified in claiming that one task received a lower (or higher) score than the other task.

Figure 1 shows that the task of moving a network topology to an ellipse formation is generally easier than moving the same network topology to a wedge. The difference in scores between two network topologies are mostly independent of the target formation. We have to be careful and use the modifier *mostly* here, because not all pairwise comparisons yield statistically significant differences. However, almost without exception L_7 networks have a statistically significant lower (better) score than C_7, K_7, and S_7 networks regardless of target formation. Similarly, S_7 networks have in almost all cases a statistically significant higher (worse) score than all other networks.

Figure 2 is a histogram of the mean rating scores for each task. We observe a similar trend as before, where independent of the target formation, line topologies are mostly rated easier than all other topologies. Star topologies are mostly rated as the hardest topologies to move into a particular formation. The $p = 0.0138$ value is larger than the p-value of the LSQ scores, so we see less statistically significant differences between the tasks. For example, the $L_{7,c}$ network does not have a statistically significant advantage over C_7, K_7, or S_7 networks with respect to the rating scores.

Figure 3 is a histogram of the mean workload scores for each task. Each bar is divided into six parts encoding (starting from the bottom) the mental, physical, and time demands, as well as, the levels of performance, effort, and frustration reported by the participant. The size of each part corresponds to the magnitude that each measured response contributes to the total workload score. We observe a similar patter in the workload scores compared to the rating scores. However, the $p = 0.0256$ value is larger for the workload score than for the rating (and LSQ) scores, so we again see less statistically significant differences between the tasks.

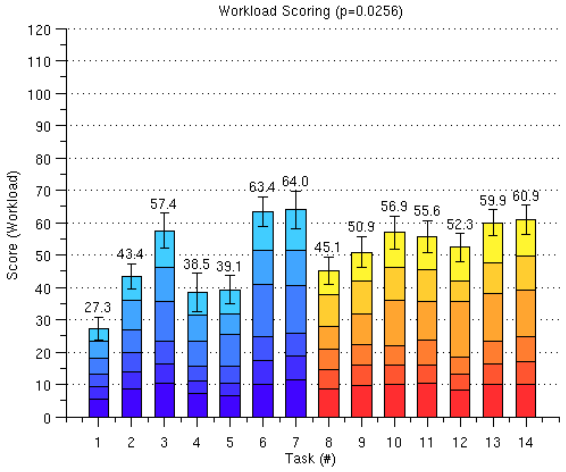

Figure 3: Mean workload score for each task.

From Figures 1, 2, and 3, it is clear that, as expected, some network topologies were significantly more difficult to control than others. For example, we directly see that line topologies are easier to control than star graphs. However, to make these types of observations stand on a more firm mathematical footing, we first need to discuss the actual network dynamics used in the experiments and their corresponding controllability properties. We will return to the results of the user study once this has been done.

System Theoretic Network Properties

To tie the results of our user study to the fundamental concepts of network control theory, we first must fully describe the network dynamics that define our system. Next, we define five measures, the rank of the controllability matrix and the node centrality measures, which we will later correlate to the results of the user study.

Network Dynamics

We begin our description of the system dynamics by defining the low level dynamics that each robot executes locally. Let $p_i(t)$ be the planar position of robot i at time t. We chose to use the standard interaction law

$$\dot{p}_i(t) = -\sum_{j \in N(i)} w(p_i(t), p_j(t))(p_i(t) - p_j(t)), \quad (1)$$

for all robots but the leader. Here $N(i)$ is the set of neighbors to robot i in the network, and where the inter-robot interaction weights are given by

$$w(p_i(t), p_j(t)) = \frac{\|p_i(t) - p_j(t)\| - \delta}{\|p_i(t) - p_j(t)\|}, \quad (2)$$

and δ is the desired distance between robot i and j.

Let us encode a network topology with n robots as an undirected static graph $\mathcal{G} = (V, E)$, where $V = \{v_1, ..., v_n\}$ is the vertex set and the edge set is an unordered set $E \subset V \times V$ such that $(v_i, v_j) \in E$ if and only if information flows between robot i and j. (See, for example (Meshabi and Egerstedt 2010) for this construction.) The neighbor set of robot i, $N(i) = \{j \in \{1, \ldots, n\} \mid (v_i, v_j) \in E\}$, is the set of all robots that share an edge with robot i in E.

Although we only consider undirected networks in this paper, it is useful to associate an orientation with the underlying graph. The way this construction works is by defining a mapping $\sigma : E \to \{-1, 1\}$ with each edge, thus assigning it an orientation and we say that v_i is the tail of edge $(v_i, v_j) \in E$ if $\sigma((v_i, v_j)) = 1$ while it is the head if $\sigma((v_i, v_j)) = -1$, with the interpretation that $\sigma((v_i, v_j)) = -\sigma((v_j, v_i))$.

The corresponding, directed graph becomes $\mathcal{G}_\sigma = (V, E_\sigma)$ and if we number the edges in this graph from 1 to m, the $n \times m$ incidence matrix, $D(\mathcal{G}_\sigma)$, is given by

$$[D(\mathcal{G}_\sigma)]_{ij} = \begin{cases} 1 & \text{if } v_i \text{ is the head to edge } j \\ -1 & \text{if } v_i \text{ is the tail to edge } j \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

Moreover, if we associate a weight with each edge, we can let W be the $m \times m$ diagonal weight matrix, where m is the number of edges, and each entry along the diagonal corresponds to the corresponding edge weight. The weighted graph Laplacian $L_w(\mathcal{G})$ then takes the following form:

$$L_w(\mathcal{G}) = D(\mathcal{G}_\sigma) W D(\mathcal{G}_\sigma)^T. \quad (4)$$

Note here that L_w does not depend on σ even though the incidence matrix does, meaning that the graph Laplacian is orientation independent.

Now, the reason for the introduction of the Laplacian is that it allows us to write down the system dynamics in ensemble form. If we let $p_i = (p_{i,1}, p_{i,2})^T$ we can pull out the positions of the robots along each dimension as $x_j = (p_{1,j}, \ldots, p_{n,j})^T$, $j = 1, 2$. Moreover, if we let the leader be given by the n-th robot in the network, we can let its position be controlled directly as $p_n(t) = u(t)$, or (equivalently from a controllability point-of-view) $\dot{p}_n = u$. We moreover partition the weighted graph Laplacian as

$$L_w(\mathcal{G})(x) = \begin{bmatrix} L_f(x) & l(x) \\ l(x)^T & \lambda(x) \end{bmatrix}, \quad (5)$$

where $L_f(x)$ is a $(n-1) \times (n-1)$ matrix, $l(x)$ is a $(n-1)$ vector, and $\lambda(x)$ is a scalar. It should be noted that since the weights are dependent on the positions of the robot, the weighted Laplacian is state-dependent as well. And, we have used x without any subscript to denote the fact that it contains the state values of all robots along both dimensions.

The final observation (see for example (Meshabi and Egerstedt 2010)) is that under the follower (and leader) dynamics in Equation 1, the ensemble dynamics becomes

$$\dot{x}_j = -L_f(x)x_j - l(x)u_j, \; j = 1, 2. \quad (6)$$

Given the single-leader network dynamics we have derived, we want to answer two questions about a given network: *Is the network controllable?* and *How controllable is the network?*. We will use the rank of the controllability matrix and the node centrality measured applied to the leader to answer these questions.

Controllability

A linear, time-invariant system

$$\dot{x} = Ax + Bu \quad (7)$$

is completely controllable if and only if an input can drive the state from any initial state to any final state. The controllability of such a system can be analyzed by constructing its controllability matrix

$$\Gamma = \begin{bmatrix} B & AB & \ldots & A^{(q-1)}B \end{bmatrix}, \quad (8)$$

where $x \in \Re^q$. Moreover, the rank of this matrix tells us how controllable the system is, in that it is equal to the dimension of the so-called *controllable subspace* which essentially is the subspace in which the control input can drive the system between arbitrary states. As such, the rank of the controllability matrix seems like a promising candidate for understanding what networks are easily controlled.

In fact, if we (for now) let the edge weights be identically equal to one, we get the ensemble dynamics

$$\dot{x}_j = -L_f x_j - l u_j, \; j = 1, 2, \quad (9)$$

which is indeed a linear, time-invariant system. As such, we can construct the network controllability matrix

$$\Gamma_L = \begin{bmatrix} -l & (-L_f)(-l) & \ldots & (-L_f)^{(n-2)}(-l) \end{bmatrix}, \quad (10)$$

and we are going to use the rank of this matrix, $\rho(\Gamma_L)$ as one of the candidate measures of how easy the network is to control. The reason why this is a valid notion is that even though our system has none-unity weights on the edges, the linearized dynamics around the desired inter-robot distances δ is given (almost) by Equation 9, as we will see below.

Since the nonlinear weights achieve a coupling between dimensions that cannot be undone, we consider instead the situation where the weights can indeed be decoupled along the different dimensions. In other words, considering the system dynamics

$$\dot{x}_j = -L_f(x_j)x_j - l(x_j)u_j, \; j = 1, 2, \quad (11)$$

allows us to achieve this decoupling. Note that this is different from Equation 6 in that L_f and l depend on x_j instead of,

as before, on the full state vector x. The decoupled weights are now assumed to be given by

$$w(p_{i,j}(t), p_{k,j}(t)) = \frac{|p_{i,j}(t) - p_{k,j}(t)| - \delta}{|p_{i,j}(t) - p_{k,j}(t)|},\ j = 1, 2. \tag{12}$$

Linearizing this system along the state-input pair $(\hat{x}_j, 0)$, where $\hat{x}_j$ is such that the edge distances are exactly equal to δ gives that the linearized dynamics is equal to

$$\begin{aligned}\dot{\tilde{x}}_j = &- \left.\frac{\partial(L_f(x_j)x_j + l(x_j)u_j)}{\partial x_j}\right|_{(\hat{x}_j,0)} \tilde{x}_j \\ &- \left.\frac{\partial(L_w(x_j)x_j + l(x_j)u_j)}{\partial u_j}\right|_{(\hat{x}_j,0)} u_j,\ j = 1, 2.\end{aligned} \tag{13}$$

After a fair amount of algebra, we get that this is equal to the linear system in Equation 9, and therefore we are (almost) justified in considering the rank of the controllability matrix associated with the linearized system. The reason why we do not consider the controllability properties of the nonlinear system is that a number of connections have already been proven between the network topology and the linear dynamics in Equation 9 when it comes to establishing the controllability properties of the network dynamics. And, as for the "almost" modifier, as the purpose of this paper is ultimately about what human operators *considers* to be easy to control, the question whether or not $\rho(\Gamma_L)$ is an appropriate measure is ultimately an empirical question. And, as will be seen in the next section, $\rho(\Gamma_L)$ will indeed turn out to be a very strong measure of how easy it is for human users to control the underlying network.

Measures of Node Centrality

We will use four classic centrality measure for simple graphs–degree centrality, closeness, betweenness, and eigenvector centrality–to quantify the important of the leader in each of the networks. Degree centrality is equal to the node degree (i.e., number of edges shared with other nodes),

$$C_D(v) = \deg(v), \text{where } v \in V. \tag{14}$$

$C_D(v)$ is simple to calculate, but it only measures the importance of the leader with respect to its immediate neighbors.

Closeness is defined in terms of the shortest paths from a node to all other nodes on the network.

$$C_C(v) = \sum_{u \in V \setminus v} 2^{-d_{\mathcal{G}}(v,u)}, \text{where } u, v \in V. \tag{15}$$

Closeness penalizes a node with long paths to other nodes.

Betweenness is a measure of the fraction of shortest paths between any two nodes that passes through a particular node.

$$C_B(v) = \sum_{u \neq w \in V \setminus v} \frac{\sigma_{u,w}(v)}{\sigma_{u,w}}, \tag{16}$$

where $\sigma_{u,w}(v)$ is the total number of shortest paths between u and w that intersect v and $\sigma_{u,v}$ is the total number of shortest paths between u and w.

Eigenvector centrality measures the influence of a node on the network. This centrality measure is computed by first solving the eigenvalue problem, $A\mathbf{x} = \lambda_{max}\mathbf{x}$, where A is the adjacency matrix and λ_{max} its largest eigenvalue. The i-th entry of the vector $\mathbf{x}$ is the centrality score given to the i-th node in the network.

$$C_E(v) = x_i, \text{where } x_i \text{ is the } i\text{-th entry of } \mathbf{x}. \tag{17}$$

Given the importance of the leader in the leader-follower structure, we can expect that the measures of node centrality for the leader is another indicator of how difficulty it is for a human operator to control a network.

Analysis and Conclusions

We finally connect the network characterization – the rank of the controllability matrix and the node centrality measures – to the user study. The average rating, average LSQ error, total workload score, and the five candidate measures, Γ_L, C_D, C_C, C_B, and C_E are summarized in Table 2. The line graph, L_7, with the leader node located at the head of the network, is completely controllable and, intuitively, it should be easy to move the followers into position by pulling the leader around. This observation is supported by the user study data. It is important to note that controlling this particular L_7 graph can be easily accomplished by the human operator independent of the target formation. In fact, if the leader in the L_7 graph is offset from the head of the network, the score, ratings, and workload measures slightly increase in comparison, even though the controllability remains constant. In this case, examining the measures of node centrality helps us explain for the difference. The leader in the $L_{7,h}$ network has a lower node centrality score than the leader in the $L_{7,o}$ network. The results indicate that a less important (or influential) leader in the network is beneficial for controlling networks in tasks require robots to be moved into a specific formation (as opposed to driving the network from point A to point B collectively).

Selecting a leader in the center of the L_7 graph cuts the rank of the controllability matrix in half, while again increasing the reported measures in comparison to the $L_{7,o}$ network. We can conclude from this observation that a decrease in rank results in an increase in the reported measures. The rank of its controllability matrix is the same as that of the C_7 network; however, its leader's centrality score is larger or equal than the centrality score of the C_7 network's leader. Therefore, we can expect that the C_7 network is easier for a human operator to control than the $L_{7,c}$ network. This conclusion is validated by the user study data in those cases where the difference is statistically significant.

The complete graph K_7 is rank deficient due to its high degree of symmetry. In fact, the rank of the controllability matrix is 1; meaning that the only the network's center of mass can be controlled in this configuration. As such, it is impossible for the participant to move K_7 into a wedge formation. The results from the user study confirm this fact. In contrast, from Table 2 demonstrates that the reported measures are low for tasks 4 and 5, where the human operator has to move a C_7 and K_7 network into an elliptical formation. We can conclude that the perceived difficulty is only low for such rank deficient networks, if the target formation is analogous to the natural formation of the network (e.g.,

Table 2: Mean LSQ, rating, and workload scores with controllability matrix rank, ρ, and node centrality measures for each task.

Task	Network	Target	ρ	C_D	C_C	C_B	C_E	LSQ	Rating	Workload
1	$L_{7,h}$	Ellipse	6	1	0.984	0	0.191	0.035	5.83	27.33
2	$L_{7,o}$	Ellipse	6	2	1.469	10	0.354	0.061	9.65	43.37
3	$L_{7,c}$	Ellipse	3	2	1.750	18	0.500	0.137	12.82	57.40
4	C_7	Ellipse	3	2	1.750	6	0.378	0.090	8.72	38.46
5	K_7	Ellipse	1	6	3.000	0	0.378	0.157	10.11	39.14
6	$S_{7,c}$	Ellipse	1	6	3.000	30	0.707	0.273	16.47	63.42
7	$S_{7,p}$	Ellipse	2	1	1.750	0	0.289	0.276	14.46	63.98
8	$L_{7,h}$	Wedge	6	1	0.984	0	0.191	0.141	9.93	45.14
9	$L_{7,o}$	Wedge	6	2	1.469	10	0.354	0.229	10.54	50.88
10	$L_{7,c}$	Wedge	3	2	1.750	18	0.500	0.415	12.57	56.94
11	C_7	Wedge	3	2	1.750	6	0.378	0.486	13.26	55.59
12	K_7	Wedge	1	6	3.000	0	0.378	0.606	15.16	52.32
13	$S_{7,c}$	Wedge	1	6	3.000	30	0.707	0.627	14.64	59.90
14	$S_{7,p}$	Wedge	2	1	1.750	0	0.289	0.602	14.81	60.86

C_7, a cycle, to an ellipse) and the user can avoid driving the system into an uncontrollable subspace.

In contrast, all tasks involving a $Star_7$ network using either a peripheral or center leader are perceived as being very difficult. Not only are these networks rank deficient, but any input quickly drives the system into an uncontrollable subspace. Interestingly, the extra rank of the $S_{7,p}$ network has no advantage over the fully rank deficient $S_{7,c}$ network.

In fact, the rank of the controllability matrix is negatively correlated ($r^2_{LSQ} = -0.60$, $r^2_{\text{Rating}} = -0.73$, $r^2_{\text{Workload}} = -0.54$) to the scores, which support the claim that *a configuration with a higher rank was almost without exceptions given a better score than a configuration with a lower rank.* We can conclude that the rank of the controllability matrix is a strong predictor of how easy it is to control a team of mobile robots. As such, it is the first thing one should consider when choosing an easily user-controlled network. As a corollary, symmetric configurations (e.g., star graphs and complete graphs) are not particularly well-suited for human control. The node centrality measures of the leader are positively correlated (e.g., for C_E, $r^2_{\text{Rating}} = 0.58$, $r^2_{\text{Workload}} = 0.54$) to the scores. In other words, a small leader-node centrality is a another good indicator that a particular network of mobile robots is easier to control. In fact, *given two configurations with the same ranks, the C_D, C_B, C_C, and C_E serve as reasonable tie breakers for which network is easiest to control.* It is important to note, however, that rank and node centrality are by no means *absolute* measures of the difficulty of controlling a given network, but good predictors of the perceived difficulty.

References

Cumming, G.; Fidler, F.; and Vaux, D. L. 2007. Error bars in experimental biology. *Journal of Cell Biology* 177:7–11.

Cummings, M. 2004. Human Supervisory Control of Swarming Networks. *2nd Annual Swarming: Autonomous Intelligent Networked Systems Conference*.

Fax, J., and Murray, R. 2004. Information flow and cooperative control of vehicle formations. *IEEE Transactions on Automatic Control* 49(9):1465–1476.

Girden, E. R. 1991. *ANOVA: Repeated Measures*. Sage Publications, Inc.

Hart, S., and Staveland, L. 1988. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. In Hancock, A., and Meshkati, N., eds., *Human Mental Workload*. North Holland Press.

Ji, M.; Azuma, S.; and Egerstedt, M. 2006. Role Assignment in Multi-Agent Coordination. *International Journal of Assistive Robotics and Mechatronics* 7(1):32–40.

Kira, Z., and Potter, M. 2009. Exerting human control over decentralized robot swarms. In *4th International Conference on Autonomous Robots and Agents (ICARA), 2009.*, 566–571.

Lozano, R.; Spong, M.; Guerrero, J.; and Chopra, N. 2009. Controllability and observability of leader-based multi-agent systems. In *47th IEEE Conference on Decision and Control, 2008. CDC 2008.*, 3713–3718.

McLurkin, J.; Smith, J.; Frankel, J.; Sotkowitz, D.; Blau, D.; and Schmidt, B. 2006. Speaking swarmish: Human-robot interface design for large swarms of autonomous mobile robots. *AAAI Spring Symposium*.

Mekdeci, B., and Cummings, M. L. 2009. Modeling Multiple Human Operators in the Supervisory Control of Heterogeneous Unmanned Vehicles. *9th Conference on Performance Metrics for Intelligent Systems*.

Meshabi, M., and Egerstedt, M. 2010. *Graph Theoretic Methods in Multiagent Networks*. Princeton University Press.

Olfati-Saber, R., and Murray, R. 2004. Consensus problems in networks of agents with switching topology and time-delays. *IEEE Transactions on Automatic Control* 49(9):1520–1533.

Ren, W.; Beard, R.; and Atkins, E. 2005. A survey of consensus problems in multi-agent coordination. In *Proceedings of the 2005 American Control Conference*, 1859–1864.