2004 Conference on Decision and Control
http://www.cds.caltech.edu/~murray/papers/2004i_sm04-cdc.html

Robust Connectivity of Networked Vehicles

Demetri P. Spanos and Richard M. Murray

Control and Dynamical Systems MC 107-81
California Institute of Technology
1200 E. California Blvd
Pasadena, CA 91125
{demetri,murray}@caltech.edu

***Abstract*— We present a simple geometric analysis of wireless connectivity in vehicle networks. We introduce a localized notion of connectedness, and construct a function that measures the robustness of this local connectedness to variations in position. Under a mild feasibility hypothesis, this function provides a sufficient condition for *global* connectedness of the network. Further, it is *distributed*, in the sense that both the function and its gradients can be calculated using only neighbor-to-neighbor communications. It can thus form the basis for distributed motion-control algorithms which respect connectivity constraints. We conclude with two simple examples of target applications.**

I. Introduction

Networked vehicles have become a prominent research theme in the control literature. They have opened several challenging problems in distributed motion control, and illustrated the need for algorithms that can be implemented across networks with little message passing. However, while much work has been done on "networkable" motion-control algorithms, relatively little attention has been paid to the motion-control problems posed by the network itself.

In particular, it is common in the literature to assume continuous availability of an underlying connected network. Clearly, motion of the vehicles comprising the network will have a significant impact on whether the network remains connected. However, we do not as yet have useful tools for understanding the relationship between motion control and preservation of network connectivity.

This paper takes a first step toward understanding the connectedness problem in a formalism amenable to standard (nonlinear) control techniques. Our main contribution is the construction of a quantity we call the *geometric connectivity robustness*. We believe that this function provides a relatively intuitive framework in which to consider motion constraints imposed by connectivity requirements.

We wish to make clear that we present this work as a *proxy* for a realistic (and necessarily more complicated) network model; it is intended to be complementary to mobile networking research, but does not aim to answer the same kinds of questions. In particular, we will say nothing about networking algorithms, routing, channel characteristics, or any of the other usual issues that arise in the wireless networking community. Instead, we hope that our simple model captures enough of the *connectivity* characteristics of most networks to be useful in control design.

II. Outline of Paper

The next section provides some background discussion on related work in networking and control of vehicle formations.

Sections IV and V discuss the three central aspects of our analysis: the communication network, the information flow, and the geometric connectivity robustness.

Section VI demonstrates some properties of the robustness function, and its relationship to overall connectedness of the communication network.

Section VII shows how to construct sparse information flow graphs with a cheap on-line distributed algorithm, a trick that will enable us to create *adaptable* information topologies in response to vehicle reconfiguration.

Finally, Section VIII discusses two target applications with connectivity constraints: broadcast-range optimization, and maximal sensor coverage.

III. Background

The recent literature on multi-vehicle systems has grown rapidly and so has general interest in development of control systems implemented across networks.

The main application of current interest in multi-vehicle systems is formation control, and several papers have appeared in recent years treating this problem from a distributed (networked) perspective. While they do not explicitly address networking concerns, the approaches taken by Leonard and Fiorelli [10] and Olfati-Saber and Murray [13] are distributed, and would require message passing on some wireless network in order to be implementable.

The motion control work most relevant to our own is Fax [6] and Fax and Murray [7], which address the dependence of control performance on network features, particularly delay time. It is this work that motivated the sparse information flow algorithm presented in Section VII, as well as our overall concern for wireless networking issues in multi-vehicle systems.

Wireless networking itself is a major research effort, and we only mention a few recent developments that have informed our work. Power-control is a blossoming field due to the energy limitations inherent to wireless networking applications, especially sensor networks. Recent developments in power-control include the work of Rodoplu and Meng [15] and that of Li and Halpern [9], which use the notion of a *relay region*. This approach is similar in motivation to our notion of *connectivity robustness*. Both constitute an attempt at geometrizing the routing issue, which is particularly relevant to multi-vehicle systems because of the intrinsically geometrical nature of motion control.

Another research effort which is very closely related to our work, both in spirit and in method, is that of Ramanathan and Rosales-Hain [14]. This work presents a constrained optimization problem for the construction of a connected network. Our work differs in two important ways. First, our constraint is the local guarantee of feasibility of a given information flow, rather than (bi-)connectivity. Second, we use a continuous-variable approach (rather than a combinatorial one), which lends itself to different types of optimization problems (see Section VIII for examples).

IV. Basic Setup and Notation

We begin with a set V of n vehicles in the plane, and associate with each vehicle a position vector $q_i \in \mathbf{R}^2$. We suppose that each vehicle can communicate in some circular domain with radius $r_i \in \mathbf{R}$. We suppose that the ranges are subject to a constraint $r_i \leq r_{\max}$, which is a simple model for bounded transmission power. We will occasionally use the stacked vector of positions $\mathbf{q} \in \mathbf{R}^{2n}$ and the vector of radii $\mathbf{r} \in \mathbf{R}^n$. We will also use the quantities $d_{ij} = \|q_i - q_j\|$, the Euclidean distances between the vehicles.

The entire discussion will center around two (undirected) graphs. The first graph is the communication network, C, in which an edge exists between i and j if and only if *bidirectional* communication between the vehicles is possible. Under our assumption above, this means that

$$ij \in C \text{ iff } \min\{r_i, r_j\} - d_{ij} \geq 0. \qquad (1)$$

Under this definition of the graph, the graph-theoretic neighborhood of a vehicle i is given by

$$N_C(i) = \{j \in V \mid \min\{r_i, r_j\} - d_{ij} \geq 0\} \qquad (2)$$

which we will call the *communication neighborhood* of i. We employ this terminology in distinction to the (to be defined) *information neighborhood*.

The second graph we will use is the *information flow* graph, I, which indicates which vehicles need to share information for their control calculations. We assume that this relationship is symmetric, i.e. that if i requires data from j, then j requires data from i. We denote the graph-theoretic neighborhood of node i in the information flow graph by $N_I(i)$, and will refer to it as the *information neighborhood*.

For the moment we suppose the information flow is given, which amounts to a choice of a motion control algorithm. We will later discuss how to construct sparse connected

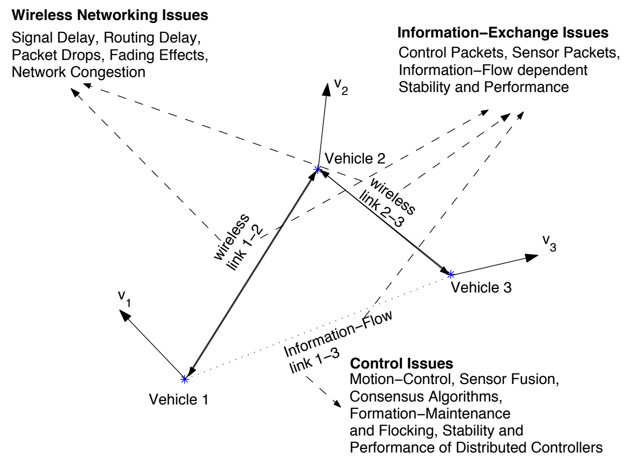

Fig. 1. A schematic showing the communication network (solid lines) and an information flow edge (dashed line) which can be implemented by the communication network in two hops.

information flows (and why one would want to do such a thing).

However we obtain it, we will view the information flow as a design requirement on the communication network. In particular, a connected information flow will be a *de facto* requirement that the communication network be connected.

The schematic in Figure 1 illustrates the setup we have described.

V. Geometric Connectivity Robustness: Definition

We will now study the robustness of network connectivity under perturbations in position. We will ultimately obtain an inequality relationship reminiscent of (1) that will characterize connectedness of the network in terms of relative positions and broadcast ranges.

We begin with a trivial observation about (1): if node i is displaced by a distance δ, then node i and j will remain connected if $\delta < \min\{r_i, r_j\} - d_{ij}$. We call this the *robustness* of the edge ij. Applying this idea to a two-edge path from i to j through k, then we see that the *path robustness*[1] is given by:

$$P(i,j,k) = \min\{\min\{r_i, r_j\} - d_{ij}, \min\{r_j, r_k\} - d_{jk}\}$$

This quantity can be computed locally at node j, and hence i and k can know it using communication to node j. If this quantity is positive, then both i and k will be connected to j, and so this communication will be possible.

Now we will discuss notions of robustness for an *information flow*, I. Recall that a vehicle i requires access to data from all members of $N_I(i)$. This access will be

[1]Clearly, similar definitions can be made for longer length paths, but full generality with respect to path length makes the notation quite cumbersome. Thus, we will restrict ourselves to two-edge paths, and trust that the reader will be able to duplicate the upcoming analysis for paths of other lengths.

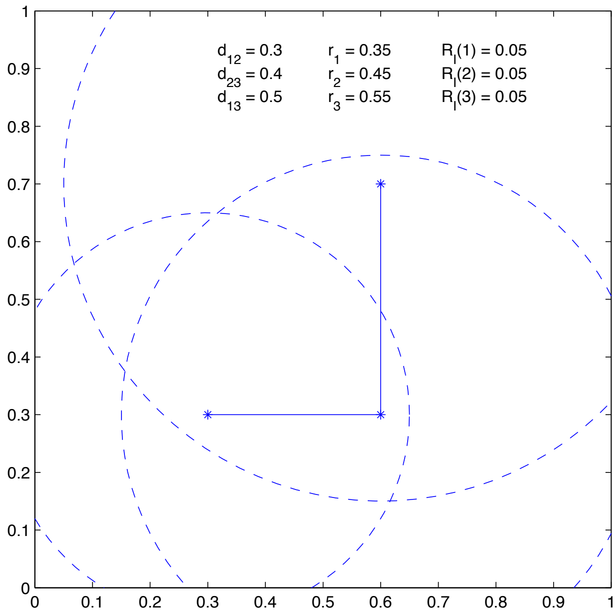

Fig. 2. Calculating the robustness for a simple network. Solid lines indicate edges in the communication network. The information flow I (not shown) is a complete graph (all vehicles connect to each other).

possible if a path exists between i and j *in the graph C*. However, this will depend on global properties of C, and this does not suit our desire for distributed characterizations of connectedness.

We will say that a vehicle i is *locally connected* to j if there exists a path in C between i and j of length at most two. Again, we could easily have defined local connectedness in terms of any other number of edges, but allowing for full generality creates an additional notational burden. Furthermore, we can calculate all two-edge quantities using only information from *communication neighbors*.

We now define the *geometric connectivity robustness* of node i (relative to the information flow I):

$$R_I(i) = \min_{j \in N_I(i) \cup N_C(i)} \left[\frac{1}{2} \max_{k \in N_C(i)} P(i,j,k) \right]. \quad (3)$$

The maximum over the communication neighbors represents the fact that vehicle i remains locally connected to vehicle j as long as *at least one* two-edge path through a mutual (communication) neighbor is available. The minimum over the communication and information neighbors reflects the fact that the motion control algorithm modeled by the graph I can be implemented *using local connectivity* as long as vehicle i is locally connected to *all* of its information neighbors. We again remark that this quantity can be computed at node i using only communication to its neighbors in C.

Figure 2 illustrates the above quantities in a simple network.

VI. Geometric Connectivity Robustness: Properties

We start this section with a simple proposition that illustrates the utility of the robustness function as a localized measure of connectedness.

Proposition 1: If $R_I(i) > 0$ for all i, C is connected.

Proof: Since the graph I is connected by hypothesis, there is a path in I from each vehicle i to any other vehicle j. Each edge in this path, say kl, links two information neighbors, and since the connectivity robustness of each vehicle is positive, there exists a path of at most two edges in C that links k to l. Through the concatenation of such paths, we can obtain a path in C which links i to j. ■

We thus have a *locally computable* function which gives a sufficient condition for connectedness of the network. The reader may feel that we have skirted the issue, in that we have pushed the difficulty of assessing connectedness to having a connected information flow graph. This does not pose a practical problem, as the information flow is frequently specified in advance by the motion control algorithm. For cases where an information flow is not pre-specified, we will later show a distributed algorithm for creating a sparse information flow which is "as connected as physically possible".

The next proposition addresses the question of robustness of the communication network to perturbations in vehicle positions.

Proposition 2: If $R_I(i) > 0$ for all i, then each vehicle can be displaced arbitrarily by a distance $R_I(i)$ while maintaining the connectedness of C.

Proof: Each vehicle's robustness is determined by a single least-robust edge in the information flow graph. This edge corresponds to a path in the communication graph with $P(i,j,k) = 2R_I(i)$. No vehicle on the path (i,j,k) can have robustness greater than $R_I(i)$, and so if no vehicle in this path moves more than this distance, the pair-wise distances d_{ij}, d_{jk} change at most by $2R_I(i)$. But, the path robustness was precisely $2R_I(i)$. Thus, the path robustness will remain positive, and all information flow edges will remain implementable as two-edge paths in C. From above, this implies that C is connected. ■

Thus, the robustness function quantifies (perhaps conservatively) the freedom of each vehicle to move relative to other vehicles. We will see a simple application of this result in Section VIII.

We now characterize the continuity and differential properties of the robustness function. In order to do so, we will introduce some additional terminology. Here, and in the remainder of the paper, we will make the assumption that all distance and robustness quantities are distinct. This amounts to ignoring a set of zero measure, and greatly simplifies the upcoming discussion [2].

[2]All results presented can be recovered without this assumption, but require some tedious bookkeeping.

We observe that the connectivity robustness of a particular vehicle is determined by a maximally robust path to a minimally robust neighbor. The robustness of this path, in turn, is determined by a minimally robust edge. Thus, for each vehicle i, there exist j and k such that

$$R_I(i) = \min\{r_j, r_k\} - d_{jk}. \quad (4)$$

We will use the notation L_i for the set $\{j,k\}$ (the *limiting set*), and the notation l_i for the member of L_i with smallest broadcast range (the *limiting transmitter*).

Now, note from the definition (3) that the robustness function is piece-wise continuous (it is obtained through minimum and maximum operations over families of continuous functions). In fact, it is a discontinuous function, due to the fact that the sets $N_C(i)$ depend on the positions and ranges. This may seem to detract from its utility as a connectivity measure, but the following proposition shows that it is in fact continuous whenever the network is locally connected.

Proposition 3: $R_I(i)$ is continuous (as a function of $\mathbf{q}$ and $\mathbf{r}$) wherever it is positive.

Proof: From Proposition 2, we see that the robustness function is Lipschitz, and hence continuous, whenever it is positive. ■

Proposition 4: Wherever they exists, the partial derivatives of the robustness function are given by

$$\frac{\partial R_I(i)}{\partial r_j} = \begin{cases} 1 \text{ iff } j = l_i \\ 0 \text{ else} \end{cases} \quad (5)$$

$$\nabla_{q_j} R_I(i) = \begin{cases} \nabla_{q_j} d_{jk} \text{ iff } j,k \in L_i \\ 0 \text{ else.} \end{cases} \quad (6)$$

Proof: The function is differentiable whenever the distances and robustness quantities are distinct (as we have assumed) [3]. The formulae above follow directly from (4). ■

The advantage of these formulae is that they can be computed locally at each vehicle, using only data from *communication neighbors*.

VII. Constructing Sparse Connected Information Flows

Many of the emerging control algorithms proposed for vehicle networks do not rely on any *particular* information flow, but rather on its connectedness. It is frequently not essential to specify which vehicles share variables, provided that the resulting information flow is connected. Interestingly, the *number of connections* in the information flow has been shown to have significant stability and performance implications for some control algorithms, and one would like to have an on-line distributed procedure for constructing an information flow which is "sparse".

We will exhibit one such scheme here, which provably provides a connected information flow with bounded degree. The degree bound provides a bound on the largest eigenvalue of the *Laplacian* matrix, which frequently appears in distributed motion control algorithms. This matrix is defined in terms of the adjacency matrix A, and the diagonal degree matrix D, as follows:

$$\begin{aligned} A_{ij}(G) &= 1 \text{ iff } ij \in G, 0 \text{ else}, \\ D_{ii}(G) &= \sum_j A_{ij}, \\ L(G) &= D - A. \end{aligned}$$

The work of Fax and Murray in [7] explores the relevance of the eigenvalues of the Laplacian to stability and performance issues in multi-vehicle systems.

The information flow we propose is a subgraph of the maximally connected graph, C_M, in which all nodes broadcast at their maximum broadcast range r_{max}. The information flow we propose, I_S, is as follows: an edge ij is in I_S if and only if it is the maximally robust path (in C_M) from i to j.

Proposition 5: Let I_S be as described above. Then, the maximal degree in I_S is at most five.

Proof: Let i be any node in V. The broadcast radius r_{max} defines a circle centered at q_i. Consider any sixty-degree sector of this circle. We will first show that there is at most one neighbor of i in this sector.

Let j be the node in the aforementioned sector which is closest to i. Let k be any other node in the sector. By construction, the angle between $q_j - q_i$ and $q_k - q_i$ is bounded by $\frac{\pi}{3}$. Applying the law of cosines, we find that $d_{ik} > d_{jk}$. Thus, the path (i,j,k) is more robust than the edge ik. Hence, i is not connected to k.

We can cover the entire circle with six such sectors, and hence bound the degree of i by six. However, one can readily see that the only way to achieve this bound is for all the neighbors to lie exactly sixty degrees apart and lie on vertices of equilateral triangles. This is impossible under our assumption that all the distances are distinct. Thus, the degree of i is at most 5. ■

We can now obtain a simple bound on the eigenvalues of the Laplacian matrix.

Proposition 6: The maximum eigenvalue of the Laplacian matrix associated with the information flow I_S is at most 10.

Proof: This follows immediately from Gershgorin's theorem, Proposition 5 (which guarantees $L_{ii} \leq 5$), and the construction of L, which implies $\sum_{j \neq i} |L_{ij}| = L_{ii}$. ■

Finally, we show that this algorithm will construct a connected information flow, assuming that the communication network at maximum broadcast range is connected (i.e. if connectedness is feasible).

Proposition 7: If C_M is connected, then so is I_S.

Proof: Consider a path between any two nodes i and j in C_M, and let kl be an edge in this path which is not in I_S. By construction of I_S, there is a path in C_{max} from

[3] This quantity is a *generalized gradient*, in the sense of [4], when the function is not smooth.

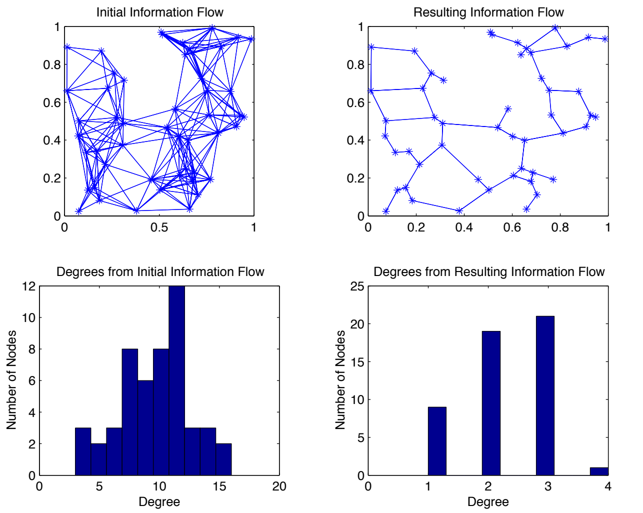

Fig. 3. An application of the distributed sparse information flow algorithm. Note that the degree bound of five is not achieved in this network.

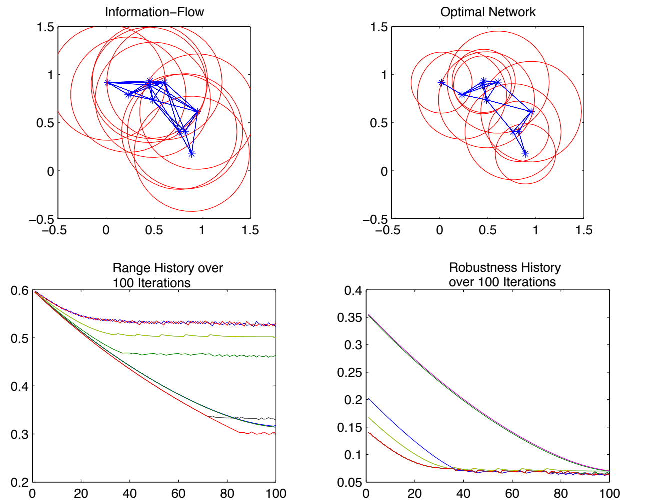

Fig. 4. A sample run of the power-optimization algorithm. The cost functions were $c_i = r_i^2$, the required robustness was $b = 0.05$, the barrier parameter was 0.01 and the stepsize was 0.001. The "chattering" behavior is a consequence of non-smoothness. This topology has been shown, by enumeration of all possible topologies, to be the global optimum.

k to l beginning with an edge which is also in I_S, say kh. Reapplying the previous argument on the new path (from h to l), we can construct a path in C_M from i to j beginning with *two* edges in I_S, and so on. Since the graphs are finite, this process will yield a path in I_S from k to l. ∎

VIII. Some Preliminary Applications with Connectivity Constraints

Here we examine two simple problems with connectivity constraints that are amenable to the analysis we have presented.

A. Broadcast-Power Optimization

We suppose that each vehicle has some convex cost function $c_i(r_i)$ associated with its broadcast range, and we wish to minimize the aggregate (additive) cost in the network subject to an information flow constraint. In particular, we would like to solve the following problem for a given information flow I and vehicle configuration $\mathbf{q}$:

$$\text{minimize } \sum c_i(r_i)$$
$$\text{subject to } R_I(i)(\mathbf{q}, \mathbf{r}) \geq b \text{ for all } i.$$

Here b is some positive scalar (known to all vehicles) quantifying how much robustness is required of the network (for example, depending on how much inter-vehicle distances vary during typical operation). It can be shown that the feasible set is compact and connected, and so the problem has a global optimum which is accessible from any initial point. Unfortunately, the set is not convex, and so in general local optima are not global optima [4].

We solve this problem using a logarithmic barrier and a gradient algorithm (see [2] and [3] for details). In particular, each vehicle implements the following (synchronous) update algorithm for its broadcast range:

$$r_i \leftarrow r_i - \gamma \left(\frac{dc_i}{dr_i} - \mu \sum_{j \in N_C(i)} \frac{\partial R_I(j)}{\partial r_i} \right). \tag{7}$$

Here μ parametrizes the strength of the logarithmic barrier (a measure of how conservative the solution will be), and γ is a stepsize parameter. Under suitable conditions [2], this algorithm will converge to a neighborhood of a local optimum (the size of the neighborhood depends on the barrier parameter).

Note that, as per our previous comments, the above algorithm can be implemented using only local information (i.e. data from communication neighbors).

B. Maximal Sensor Coverage

This is a simplified version of coverage problems studied in [12]. We assume that the vehicles are each capable of sensing a square[5] domain centered on themselves, $S_i = \{p \in \mathbf{R}^2 \mid \|q_i - p\|_\infty < w\}$. We suppose that there is a function $f(x, y)$ which quantifies the "significance" of a particular point in the plane, and that the vehicles wish to maximize the integral of this function over the *union* of their sensing domains.

[4] The gradient algorithm we use is only guaranteed to converge to a local optimum. However, our numerical simulations suggest that convergence to a global optimum can be achieved if the network is initialized at C_M. Unfortunately, we have found no way to make a formal statement about global convergence.

[5] We make this choice only for calculational convenience. In principle, other domains can be handled equally well.

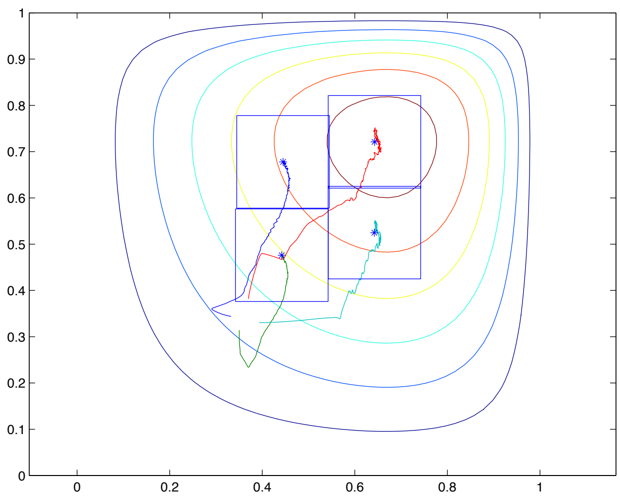

Fig. 5. A sample run of the optimal sensor coverage algorithm. The sensing width was 0.05, and the communication range was 0.15. The information flow is the complete information flow. The contours of the function $f(x, y)$ are shown. Note that although the gradients initially pull the vehicles apart, the barrier terms force them to stay together. Again, we observe chattering due to the non-smooth constraints.

$$\text{maximize} \quad g(\mathbf{q}) = \int_{\cup S_i} f(x, y) \tag{8}$$

The motion of the vehicles will be driven by a gradient algorithm. The gradient of the objective function can be shown to be:

$$\nabla_{q_i} g = \int_{\partial S_i \cap \cup_{j \neq i} S_j} f(x, y) n_i(x, y) \tag{9}$$

Here, $n_i(x, y)$ is the outward normal to the set S_i.

In addition to the coverage objective, we will require that the vehicles preserve an information flow, and we will do so by controlling positions rather than broadcast ranges. This will be accomplished with barrier terms as in the previous section. The resulting algorithm is:

$$\dot{q}_i = \nabla_{q_i} g - \mu \sum_{j \in N_C(i)} \nabla_{q_i} R_I(j). \tag{10}$$

Again, this algorithm is implementable using only local communication. Figure 5 shows a sample run of this algorithm.

IX. Conclusion and Future Work

We have presented a simple geometric analysis of connectivity in vehicle networks under a distance-based radio communication model. Using our notion of connectivity robustness, we were able to obtain a locally computable function which provided a sufficient condition for connectedness of the network. We have also used robustness as a tool for constructing sparse information flow graphs with a cheap distributed algorithm. Two simple applications were discussed, and demonstrated that a meaningful and *mathematically tractable* connectivity constraint could be appended to standard problems arising in vehicle and sensor networks. In principle, this connectivity constraint can be appended to any nonlinear motion control problem.

Our plans for future work center on answering a fundamental motion control question: what classes of motion control algorithms will accomplish their objective in the presence of the network connectivity constraint we have formulated in this paper? Our sensor coverage application is a simple instance in which the connectivity constraint only impeded the transient behavior of the algorithm, but not the ultimate objective of optimal placement. However, it is easy to construct motion control algorithms (for example, group-splitting maneuvers) that will not accomplish their objective under this constraint.

Finally, we hope to find useful extensions of our robustness notion that will include more realistic connectivity phenomena. This may include intermittent link failures, as well as dynamic behaviors such as link acquisition and termination in mobile networks.

References

[1] S. Agarwal, R. Katz, S. Krishnamurthy, and S. Dao, *Distributed Power Control in Ad Hoc Wireless Networks*, Proc. of IEEE International Symposium on Personal, Indoor and Mobile Radio Communications, 2001

[2] D. Bertsekas, *Nonlinear Programming*, Athena Scientific, 1995

[3] D. Bertsekas and J. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*, Prentice-Hall, 1989

[4] F. Clarke, *Optimization and Nonsmooth Analysis*, Wiley Interscience, 1983

[5] T. ElBatt and A. Ephremides, *Joint Scheduling and Power Control for Wireless Ad hoc Networks*, Proceedings of IEEE Infocom 2002

[6] J.A. Fax, *Optimal and Cooperative Control of Vehicle Formations*, PhD Thesis, California Institute of Technology, 2002

[7] J.A. Fax and R. M. Murray, *Graph Laplacians an Stabilization of Vehicle Formations*, Proceedings of IFAC World Congress, 2002

[8] P. Gupta and P. R. Kumar, *Critical Power for Asymptotic Connectivity*, Proceedings of 37th IEEE Conference on Decision and Control

[9] L. Li and J. Y. Halpern, *Minimum-Energy Mobile Wireless Networks Revisited*, Proceedings of IEEE International Conference on Communications (ICC), 2001

[10] N.E. Leonard and E. Fiorelli, *Virtual Leaders, Artifical Potentials, and Coordinated Control of Groups*, Proceedings of 40th IEEE Conference on Decision and Control, 2000

[11] N.A. Lynch, *Distributed Algorithms*, Morgan Kauffmann Publishers, 1996

[12] P. Ogren, E. Fiorelli, and N.E. Leonard *Cooperative Control of Mobile Sensor Networks: Adaptive Gradient Climbing in a Distributed Environment*, 2003

[13] R. Olfati-Saber and R. M. Murray, *Flocking with Obstacle Avoidance: Cooperation with Limited Communication in Mobile Networks*, Submitted to IEEE Conference on Decision and Control, 2003

[14] R. Ramanathan and R. Rosales-Hain, *Topology Control of Multihop Wireless Networks using Transmit Power Adjustment*, Proceedings of the IEEE Conference on Computer Communications (INFOCOM), 2000

[15] V. Rodoplu and T. Meng, *Minimum Energy Mobile Wireless Networks*, Proceedings of the 1998 IEEE International Conference on Communications, 1998

[16] R. Wattenhofer, L. Li, P. Bahl, and Y. Wang, *Distributed Topology Control for Power Efficient Operation in Multihop Wireless Ad Hoc Networks*, Proceedings of IEEE Infocom 2001