A Constrained Optimization Framework for Wireless Networking in Multi-Vehicle Applications

Demetri P. Spanos and Richard M. Murray
Control and Dynamical Systems
California Institute of Technology
{demetri,murray}@caltech.edu

September 21, 2003

Abstract

This paper presents an optimization framework for broadcast power-control, specifically addressed at wireless networking issues arising in implementing information flows for multi-vehicle systems. We formulate an optimization problem for the minimization of an aggregate cost subject to a constraint on a quantity we call the *geometric connection robustness*, which is a locally computable numerical assessment of the robustness of the an information flow to perturbations in position. Our main result is a location-aided distributed power-control algorithm based on a gradient-like optimization scheme. We also use geometric connection robustness to develop a cheap distributed heuristic for the construction of sparse connected information flow.

1 Introduction

The analysis and design of distributed systems has been a prominent theme in much of the recent literature on automation and control systems. Indeed, decentralization is a prominent area of current research in several diverse fields, ranging from computer science to numerical analysis. Computer networking research in particular has had to face this problem since its inception, as the very notion of a distributed system is moot without an underlying mechanism to implement interconnections or otherwise mediate interactions.

The networking problem of present interest is that of a wireless ad hoc network implemented among several vehicles carrying out some kind of distributed motion-control scheme. These kinds of multi-vehicle applications typically require the continuous exchange of position and velocity data according to a graph defining an *information flow*. We deliberately distinguish between the information flow and the communication network. In particular, we are interested in generating ad hoc networks which implement a given information flow while achieving a suitable level of network performance (as measured by an aggregate cost functional and maximum hop-count).

We pursue the construction of ad hoc wireless networks by formulating the problem as a constrained optimization. We take a continuous rather than combinatorial approach by using a geometric formulation based on the exchange of position information. Network connectivity is assessed (locally) in terms of a quantity we call the *geometric connection robustness*. Within the wireless networking community, this approach falls under the category of location-aided protocols. While this is a relatively small portion of the wireless networking work, it is ideally suited to multi-vehicle applications, since position sensors and exchange of position information are inherent to distributed motion-control algorithms.

Wireless Networking Issues

Signal Delay, Routing Delay,
Packet Drops, Fading Effects,
Network Congestion

Information–Exchange Issues

Control Packets, Sensor Packets,
Information–Flow dependent
Stability and Performance

v_2

Vehicle 2

wireless
link 2–3

wireless
link 1–2

v_3

Vehicle 3

v_1

Information–Flow

link 1–3

Control Issues
Motion–Control, Sensor Fusion,
Consensus Algorithms,
Formation–Maintenance
and Flocking, Stability and
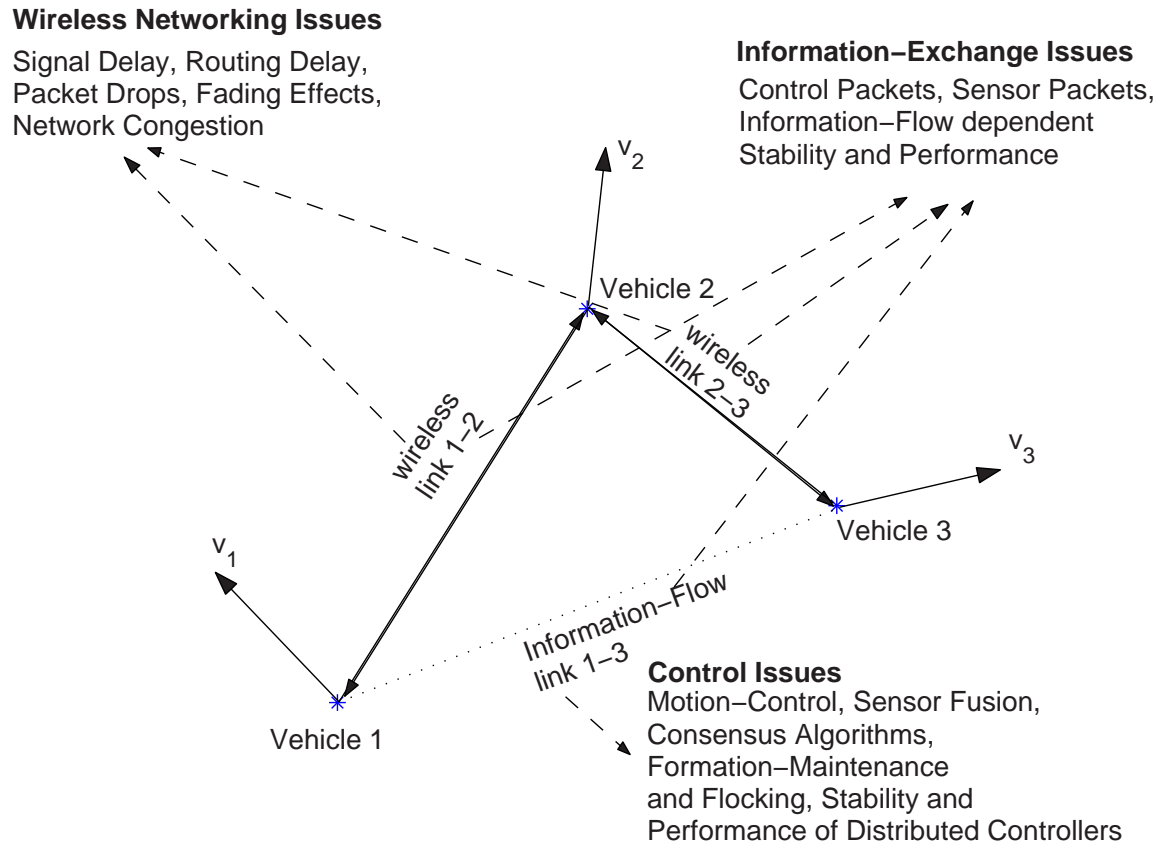Performance of Distributed Controllers

Vehicle 1

Figure 1: A schematic illustration of the interplay between networking, information flow, and control in a multi-vehicle system. Solid arrows indicate wireless communication links, and the dotted line indicates an information flow between two vehicles.

Wireless networks among multi-vehicle systems present unique challenges to communications protocol designers, not the least of which being the problem of determining the network topology. Whereas a wired network has a fixed and known topology, dictated by which stations are physically connected to each other, a wireless network has no intrinsic topology except insofar as each station has only finite transmission power (and hence finite transmission range). This is the primary motivation behind ad hoc networking protocols, in which nodes make local decisions about establishing links with other stations. As a consequence of the unpredictability of wireless channels and the mobility of the nodes, the links comprising the network topology are intrinsically dynamic, switching on and off as deemed necessary by the networking protocol.

The recent literature on multi-vehicle systems has grown rapidly and so has general interest in development of control systems implemented across networks. For discussion of current hardware implementations in the academic community, see the work of Cremean *et al.* [6], D'Andrea and Murray [7, 8], and Stubbs and Dullerud [27]. These references describe the physical devices used to implement networked multi-vehicle systems, and identify several of the networking performance issues we presented in the introduction.

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 perspective. While they do not explicitly address networking concerns, the approaches taken by Leonard and Fiorelli [16] and Olfati-Saber and Murray [21] are distributed, and would require message passing on some wireless network in order to be implementable. Work on flocking such as Olfati-Saber [22] is also particularly relevant, as artificial flocking will definitely necessitate ad hoc networks.

The work most relevant to our own is Fax [10] and Fax and Murray [11], which addresses the dependence of control performance on network performance, particularly delay time. It is this work that motivated the sparse information flow heuristic presented in Section 2.2.1, as well as our overall concern for wireless networking issues in multi-vehicle systems.

Power-control is a blossoming field due to the energy limitations inherent to wireless networking applications, especially sensor webs. Recent developments in power-control include the work of Rodoplu and Meng [24] and Li and Halpern [14], which uses the notion of a *relay region*. This approach is similar in motivation to our notion of *routing robustness* which we present later. 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.

Other related work in power-control and energy-optimization includes Monks, Bharghavan, and Hwu, [18], ElBatt and Ephremides [9], Toh [29] and Agarwal, Katz, Krishnamurthy, and Dao [1]. One paper which is very closely related to our work, both in spirit and in method, is that of Ramanathan and Rosales-Hain [23], which 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, our continuous optimization naturally lends itself to *distributed online* solution.

Topology control is also a closely related area of research, and there is a significant body of research relevant to our own. We refer mainly to the work of Wattenhofer, Li, Bahl, and Wang *et al.* [31], and Salonidis, Bhaghwat, and Tassiulas [25] which both present distributed algorithms for constructing connected networks. The latter is potentially quite relevant to our own, as we envision Bluetooth-like technologies playing a significant role in multi-vehicle systems. We also refer to Gupta and Kumar [12] which, although not strictly addressed at topology control, gives an analytical treatment of the broadcast power required to achieve asymptotic connectedness in a wireless network.

Finally, we discuss location-aware and mobility-aware networking. Two high-level discussions which introduce these areas are Tseng, Wu, Liao, and Chao [30] and Lee, Su, and Gerla [15]. Most of this work is based on using position data (perhaps from GPS) to assist the mobile-networking protocol. For an early and representative development in this area, see Navas and Imielinski [19]. Other work of particular relevance to multi-vehicle systems includes Ko and Vaidya [13] and Saverese, Rabaey, and Beutel [26], which both address location-aided networking

issues in the absence of a centralized positioning system.

2 Main Results

The main body of the paper presents the notion of geometric connection robustness, and shows how a certain local connectedness criterion can be defined in terms of this quantity. We then formulate an optimization problem subject to local connectedness constraints, and observe that a standard barrier-type formulation of the optimization problem lends itself to solution by a distributed steepest-descent algorithm. We also present a simple distributed heuristic for constructing sparse connected information flows using the geometric position robustness data.

2.1 Notation and Assumptions

We begin with a set V of n nodes arranged in the plane. To each node i we associate a position vector $q_i \in \mathbf{R}^2$, and we denote the distance between two nodes i and j, $\|q_i - q_j\|$, by d_{ij}. We will occasionally make reference to the stacked vector of positions $\mathbf{q} \in \mathbf{R}^{2n}$.

We suppose that each node is equipped with an omnidirectional radio antenna, and that the transmitter can control the broadcast range subject to a maximum-range constraint, $r_i \leq r_{max}$ for all $i \in V$. Again, we will denote the stacked vector of broadcast ranges by $\mathbf{r} \in \mathbf{R}^n$.

We assume that two nodes form a connection if and only if each node is within broadcast range of the other. Thus, a bidirectional link (we only consider bidirectional links) exists between nodes i and j if and only if $\min\{r_i, r_j\} \geq d_{ij}$. Under this assumption, a choice of the positions q_i and a choice of the broadcast ranges r_i induces a graph $G_{\mathbf{q},\mathbf{r}} = (V, E_{\mathbf{q},\mathbf{r}})$, which we call the *communication network*. We again draw attention to the fact that this graph is distinct from the *information flow*.

For the purposes of optimization, we suppose that each node also has a cost function associated with its broadcast range, $c_i : [0, r_{max}] \to \mathbf{R}^+$. We assume these functions to be strictly convex, increasing, and twice differentiable. One physically relevant choice for these functions is the power required to broadcast packets to that range at a particular level of received power, which would motivate a function like $r_i{}^\alpha$. Typical choices of the parameter α in the wireless networking research range between two and four, but we do not devote any additional attention to the choice of cost function.

One final assumption we make is that, if a link exists between two nodes, each node can know the other's broadcast range. This can be inferred physically by measuring the received power and knowing the radio capabilities of the sender, or by attaching a small header to all packets indicating their intended broadcast range. Relative to the continuous data-exchange overhead associated with multi-vehicle control algorithms, we feel this should be a negligible addition to the required throughput.

2.2 Geometric Connection Robustness

The notion of robustness is inherent to any engineering design. We would like any wireless network that we construct for a multi-vehicle system to be robust to the kinds of variability that could potentially disconnect or otherwise impair the network. Wireless networks are known to suffer from many such sources of variability, including uncertainty of the wireless channel, obstacles, and interference.

Here we concentrate on the difficulty imposed intrinsically by the fact that the network is operating over a collection of moving vehicles, i.e., the fact that the transmitters move relative to each other. We thus call our notions of robustness *geometric* to emphasize that we do not consider robustness to electromagnetic, information-theoretic, or computational sources of variability. However, for brevity, we will suppress the explicit use of the word geometric, and simply refer to the robustness of the network or link.

As a first stab at quantifying the robustness of network links, we consider a single link between two nodes. We are interested in the minimal perturbation, either to the relative positions of the nodes, or their respective broadcast ranges, which will result in destruction of the link. We refer to this as the *pair-wise connection robustness*, and denote it by

$$R_{ij}^{PW}(\mathbf{q}, \mathbf{r}) \equiv \min\{r_i - d_{ij}(\mathbf{q}, \mathbf{r}), r_j - d_{ij}(\mathbf{q}, \mathbf{r})\}.$$

Here we explicitly indicate the dependence on $\mathbf{q}$ and $\mathbf{r}$, but we suppress it in subsequent formulae in order to lighten the already cumbersome notation.

Note that, under the aforementioned assumptions, this quantity can be computed locally by each node whenever it is positive, i.e., whenever the two nodes are connected. We also observe the following trivial (and equivalent) facts about the graph $G_{\mathbf{q},\mathbf{r}}$:

$$
\begin{aligned}
R_{ij}^{PW} \geq 0 &\Leftrightarrow ij \in E_{\mathbf{q},\mathbf{r}} \\
N_i^{G_{\mathbf{q},\mathbf{r}}} &= \{j \in V | R_{ij}^{PW} \geq 0\}.
\end{aligned}
$$

Here $N_i^{G_{\mathbf{q},\mathbf{r}}}$ denotes the graph-theoretic neighborhood of node i relative to the graph $G_{\mathbf{q},\mathbf{r}}$. We will also come to discuss neighbors of node i relative to the information flow, and we thus require this additional specificity. Where there is no danger of confusion, we will suppress the superscript.

Having defined the robustness of a single link, we are naturally inclined to ask about the robustness of a path. For intuitive reasons, the robustness of a path must be the robustness of its least robust link. This in fact coincides exactly with the minimal perturbation required to destroy a particular path. For the case of a two-link path with node k acting as a router for nodes i and j, we define the *one-hop path-wise connection robustness*,

$$R_{ikj}^{P,1} \equiv \min\{R_{ik}^{PW}, R_{kj}^{PW}\}.$$

Robustness definitions for longer paths, $R_{iklj}^{P,2}$, and so on, follow suit, but we do not use them in this paper. Note that the above quantity is non-negative whenever it is possible for node i and j to exchange messages through node k, and negative otherwise. Further, since both pair-wise robustness quantities are available to the router, k, nodes i and j can have local knowledge of this quantity by polling their mutual neighbor.

Clearly, node i may have more than one routing option for sending messages to j, including the option of not using a router at all. We wish to quantify the smallest perturbation that will destroy *all* of i's messaging options to j. We thus define the *one-hop routing connection robustness*,

$$R_{ij}^{R,1} \equiv \max_{k \in N_i \cup \{i\}} R_{ikj}^{P,1}.$$

We include i in the set over which the maximum is taken because it is possible that direct messaging from i to j is the most robust option available, and we naturally want the one-hop routing robustness to be at least as large as the pair-wise (or zero-hop) robustness.

Again, we see a trivial, but useful graph-theoretic property:

$$R_{ij}^{R,1} \geq 0 \quad \Leftrightarrow \quad N_i \cap N_j \neq \emptyset. \tag{1}$$

We also see that, since each term $R_{ikj}^{P,1}$ can be known locally, $R_{ij}^{R,1}$ can also be found locally.

Finally, we mention that we can recursively define higher hop-count robustness quantities by

$$R_{ij}^{R,m} \equiv \max_{k \in N_i} \min\{R_{ik}^{PW}, R_{kj}^{R,m-1}\}.$$

Each additional hop allowed in the routing will require additional message passing in order to be calculated locally. We expect there to be a strong tradeoff between reducing energy expenditure (suggesting small ranges and large hop-count) and minimizing messaging complexity (suggesting large ranges small hop-count), but we do not attempt to quantify it.

Finally, we wish to discuss the robustness of a particular information flow under a choice of **q** and **r**. We recall that, by an information flow, we mean a graph $F = (V, E)$ where each edge represents the exchange of data between two nodes for the purposes of some distributed algorithm. We would like to know the size of the minimal perturbation such that each information flow edge can no longer be implemented by at most a m-hop data path in the communication network. We thus define the *m-hop information flow robustness of i relative to F*,

$$R_i^{F,m} \equiv \min_{j \in N_i^F} R_{ij}^{R,m}.$$

Here we have again used the notation N_i^F to denote the neighborhood of i in the graph F. We also observe that, as before, this quantity can be computed locally, with no more message passing than is required to construct each individual $R_{ij}^{R,m}$ term.

Having defined information flow robustness, we now define some additional notation which will make analysis of this function easier and more intuitive.

We first define, for each node i, the set of m-hop F-limiting information-neighbors (this set will typically be a singleton, but we formulate it as a set for completeness),

$$N_{i,L}^{F,m} \equiv \{j \in N_i^F | R_{ij}^{F,m} = R_i^{F,m}\}.$$

Here the additional subscript L signifies the fact that these are the *limiting* information neighbors. Thus, in order to increase the m-hop robustness of i relative to F, the routing robustness to each member of $N_{i,L}^{F,m}$ must be increased. Conversely, decreasing $R_{ij}^{F,m}$ for *any* $j \in N_{i,L}^{F,m}$ will decrease $R_i^{F,m}$.

At this point we are faced with a notational dilemma: we must choose between formulating the notion of limiting routers in full generality (with respect to hop-count), and incurring the additional notational burden of hypergraphs, or restricting ourselves to a particular number of hops and avoiding this complication. We choose to err on the side of comprehensibility, and present the remaining definitions for 1-hop quantities. We feel that the forthcoming concepts are sufficiently transparent that the reader will be able to construct the analogous higher-hop quantities with little difficulty.

Now, we define the set of 1-hop F-limiting routing-neighbors of i,

$$N_{i,L}^{G,1} \equiv \{k \in N_i^G | R_{ikj}^{P,1} = R_i^{F,1} \text{for some } j \text{ in } N_i^F\}.$$

These are the neighbors which provide the maximally robust message-routing service for i to one of its F-limiting information-neighbors. This set may include i, as direct messaging may be the maximally robust route.

We now define the set of one-hop *limiting edges* of i in G relative to F. A limiting edge is one that participates in a path from i to a member of $N_{i,L}^{G,1}$ through a limiting routing-neigbor. We denote this set by $E_{i,L}^{F,1}$.

Finally, we note that for any edge, the pair-wise robustness associated with that edge is limited by the broadcast range of one or both of the participating transmitters. We denote the set of such limiting transmitters by T_i, and observe the following useful property:

$$R_i^{F,1} = r_t - d_{tl},$$

where t and l are members of T_i and $N_{i,L}^{F,1}$ respectively. As a consequence of this equality, we say that the variable r_t is *active* for node i. We denote the set of nodes for which r_i is active by A_i, and note that each member of the active set of i is known to a neighbor of i, and hence
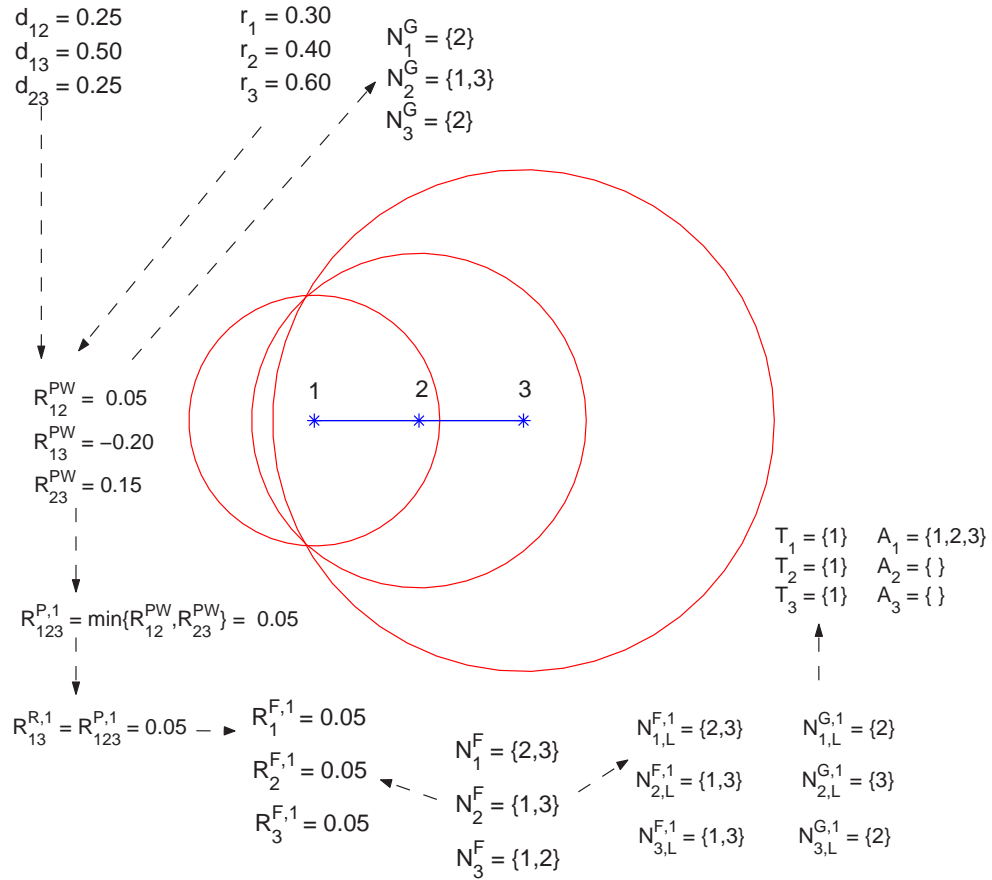
$d_{12} = 0.25$
$d_{13} = 0.50$
$d_{23} = 0.25$

$r_1 = 0.30$
$r_2 = 0.40$
$r_3 = 0.60$

$N_1^G = \{2\}$
$N_2^G = \{1,3\}$
$N_3^G = \{2\}$

1   2   3

$R_{12}^{PW} = 0.05$
$R_{13}^{PW} = -0.20$
$R_{23}^{PW} = 0.15$

$R_{123}^{P,1} = \min\{R_{12}^{PW}, R_{23}^{PW}\} = 0.05$

$T_1 = \{1\}$ $A_1 = \{1,2,3\}$
$T_2 = \{1\}$ $A_2 = \{\ \}$
$T_3 = \{1\}$ $A_3 = \{\ \}$

$R_{13}^{R,1} = R_{123}^{P,1} = 0.05$ ➝ $R_1^{F,1} = 0.05$

$R_2^{F,1} = 0.05$

$R_3^{F,1} = 0.05$

$N_1^F = \{2,3\}$

$N_2^F = \{1,3\}$

$N_3^F = \{1,2\}$

$N_{1,L}^{F,1} = \{2,3\}$ $N_{1,L}^{G,1} = \{2\}$
$N_{2,L}^{F,1} = \{1,3\}$ $N_{2,L}^{G,1} = \{3\}$
$N_{3,L}^{F,1} = \{1,3\}$ $N_{3,L}^{G,1} = \{2\}$

Figure 2: An illustration of the various robustness quantities we have defined. For the information flow F, we require all nodes to communicate with one another. In the communication network G, only two links are available, so the information flow from 1 to 3 is routed through 2. Note that all three nodes have the same information flow robustness relative to F because they are all limited by the 0.05 robustness link between 1 and 2.

can be constructed at i locally by polling neighbors. These active variables, and hence the sets A_i will prove to characterize the differential properties of the robustness function quite nicely.

We now explore a few properties of information flow robustness, and discuss its differentiability. Throughout this discussion we assume F is a connected information flow.

We first state the obvious limitation of the above robustness quantities: for $m < n - 1$, positivity of $R_i^{F,m}$ for all i is not necessary for connectedness of $G_{\mathbf{q},\mathbf{r}}$ (unless $m \geq n - 1$, where the necessity is trivial). However, $R_i^{F,m} \geq 0$ for all i is *sufficient* for connectedness of $G_{\mathbf{q},\mathbf{r}}$ for *any* m. We show this in the following simple proposition:

Proposition 1. *Let F be a connected information flow. If $R_i^{F,m} > 0$ for all i, then $G_{\mathbf{q},\mathbf{r}}$ is connected.*

Proof. Consider two nodes i and j. Since F is connected, there exists a sequence of edges $\{kl\}_t$ in F linking i to j. Let kl be any edge in this sequence. Then, by definition, $k \in N_l^F$. Now, since each node has positive routing robustness to each of its neighbors in F, there exists a sequence of at most $m+1$ edges in $G_{\mathbf{q},\mathbf{r}}$ between k and l. Now, we can construct a sequence of edges in $G_{\mathbf{q},\mathbf{r}}$ by substituting, for each kl in the sequence $\{kl\}_t$, the $m+1$-edge path in which is guaranteed to exist by the information flow robustness. This path connects i to j. □

The importance of this simple result is that we now have a *distributed* numerical assessment of the connectedness of $G_{\mathbf{q},\mathbf{r}}$. At this point, the reader may feel that we have slightly cheated, since we have in essence pushed back the difficulty of assessing connectedness to that of having, *a priori*, a connected information flow F. This is in part true, but it is also possible to cheaply construct information flows which are "as connected as physically possible" by simply setting $F = G_{\mathbf{q},\mathbf{r}_{\mathbf{max}}}$, which can be done on-line. We will also show a simple heuristic for constructing *sparse* connected information flows using $G_{\mathbf{q},\mathbf{r}_{\mathbf{max}}}$ and the robustness quantities above.

We now state another simple connectedness result:

Proposition 2. *Let $\mathbf{q}$ be a given configuration, and F any connected information flow. Suppose $R_i^{F,m} > 0$ for all i. Now, let $\tilde{\mathbf{q}}$ be any other configuration satisfying $\max_i \|q_i - \tilde{q}_i\| < \frac{R}{2}$. Then $G_{\tilde{\mathbf{q}},\mathbf{r}}$ is connected.*

Proof. Since $R_i^{F,m} > 0$, $G_{\mathbf{q},\mathbf{r}}$ is connected, from the previous proposition. From the definition of robustness, every edge in $G_{\mathbf{q},\mathbf{r}}$, say ij persists under a perturbation in d_{ij} of size R. Since the perturbation to each q_i is less than $\frac{R}{2}$, the perturbation to each d_{ij} is less than R by the triangle inequality. Thus, each edge in $G_{\mathbf{q},\mathbf{r}}$ is also in $G_{\tilde{\mathbf{q}},\mathbf{r}}$, and so the latter is connected. □

This is really a very simple result which, unfortunately, is partly obscured by our notation. The intuitive idea is that this gives something like a $\|\cdot\|_\infty$ neighborhood of $\mathbf{q}$ in which all possible configurations result in a connected communication network.

We are now interested in discussing differential properties of information flow robustness. We will first show that it is continuous wherever it is positive.

Proposition 3. *Let $\mathbf{q}$ and $\mathbf{r}$ be given, and let F be such that $R_i^{F,m}(\mathbf{q},\mathbf{r}) > 0$. Then, $R_i^{F,m}$ is continuous at $(\mathbf{q},\mathbf{r})$.*

Proof. By the construction of the robustness function, and the fact that the robustness is positive, a perturbation Δ to $(\mathbf{q},\mathbf{r})$ such that $\|\Delta\|_\infty < \delta$ induces a change in $R_i^{F,m}$ of less than 2δ (see proposition 2). This implies continuity in $\|\cdot\|_\infty$, and by equivalence of norms in finite-dimensional vector spaces, continuity in all norms. □

Corollary 1. *$R_i^{F,m}$ is Lipschitz. In $\|\cdot\|_\infty$, its minimal Lipschitz constant is 2.*

It may seem that we have not utilized the positivity of the robustness function, but we have tacitly embedded it in the 2δ bound. Were the robustness negative, an arbitrarily small perturbation could yield an additional neighbor with larger (negative) routing robustness. This situation is eliminated in the positive robustness case because any neighbor acquired through

an arbitrarily small perturbation will provide routing robustness arbitrarily close to zero (and hence, smaller than the maximal routing robustness defining $R_i^{F,m}$). Conversely, any neighbor lost with an arbitrarily small perturbation cannot provide positive routing robustness, and so is not the limiting router.

Now, it is clear that $R_i^{F,m}$ is piecewise affine in r, as it is obtained through max and min operations over families of affine functions (of the form $r - d$). It is thus differentiable almost everywhere (i.e. anywhere except the interfaces between the piece-wise affine regions). We now present a generalized gradient (with respect to $\mathbf{r}$) for this function (which reduces to the gradient where it exists).

Proposition 4. *Wherever $R_i^{F,m}$ is differentiable, we have:*

$$\frac{\partial R_i^{F,m}}{\partial r_j} = \begin{cases} 1 \ iff \ j \in T_i \\ 0 \ else. \end{cases}$$

Proof. We again write $R_i^{F,m} = r_t - d_{tl}$. If T_i is a singleton, this equation holds in an open set, and hence the partial derivatives of $R_i^{F,m}$ must match those of $r_t - d_{tl}$. We have thus proved the proposition in this case. We will now show that this is the only case where the function is differentiable. Suppose there are at least two elements in T_i, t_1 and t_2. There are two cases:

I: t_1 and t_2 route to the same information-neighbor.
   Then, increasing r_{t_1} will (locally) increase $R_i^{F,m}$ linearly, since the latter is defined by the maximally robust path. However, decreasing r_{t_1} will have no effect on routing robustness for the same reason. This implies that the partial derivative with respect to r_{t_1} (or r_{t_2}) does not exist.

II: t_1 and t_2 route to distinct information-neighbors.
   Then, decreasing r_{t_1} will (locally) decrease $R_i^{F,m}$ linearly, while increasing it will have no effect, since it is determined by the maximally robust path to the *minimally robust* information neighbor. This again implies that the associated partial derivative does not exist.

Thus, whenever T_i is not a singleton, $R_i^{F,m}$ is not differentiable. This proves the proposition. $\square$

Corollary 2. *The above formula defines a maximum-increase direction for $R_i^{F,m}$.*

2.2.1 Using Robustness to Construct Sparse Connected Information Flows

Here we address the issue of constructing connected information flows using robustness information. We are particularly interested in constructing *sparse* information flows in order to limit the maximal eigenvalues of a certain matrix associated with the graph, the Laplacian. 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 [11] explores the relevance of the eigenvalues of the Laplacian to stability and performance issues in multi-vehicle systems.

We suppose that the configuration $\mathbf{q}$ is a random variable, uniformly distributed on some bounded domain. We will now present a simple application of robustness to produce, almost

surely, an information flow with maximum degree 5. The proof is essentially a consequence of basic plane trigonometry, but we include it for clarity.

Proposition 5. *Let $F_1 = G_{\mathbf{q}, \mathbf{r}_{max}}$. Define $F_2^m = (V, E_2^m)$ as follows:*

$$E_2^m = \{ij \in E_{\mathbf{q}, \mathbf{r}_{max}} | R_{ij}^{PW} = R_{ij}^{F,m}\}.$$

Then, almost surely, every node has at most 5 neighbors in F_2^m.

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 (this is unique almost surely). 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 $d_{ik} > d_{jk}$. Applying the definition of $R_{ijk}^{R,1}$, we see that $R_{ik}^{F,m} \geq R_{ijk}^{R,1} > R_{ik}^{PW}$. Thus, $ik \notin E_2^m$. So, i is connected to at most one neighbor in this sector.

We can cover the entire circle with six such sectors, and hence bound the degree of i in F_2^m 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, which will occur with zero probability. Thus, almost surely, the degree of i is at most 5. $\qquad\square$

Corollary 3. *The maximum eigenvalue of the Laplacian matrix associated with F_2 is at most 10, almost surely.*

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

This corollary provides a simple (albeit conservative) bound for analyzing stability and performance issues of the resulting information flows. For more information on this issue, see [11].

We now show that this information flow will also be connected, if physically possible.

Proposition 6. *Suppose F_1, as defined above, is connected. Then F_2^m is also connected, almost surely.*

Proof. Consider a path between any two nodes i and j in F_1, and let kl be an edge in this path which is not in F_2. By construction of F_2, $R_{kl}^{F_1,m} > R_{kl}^{PW}$. Thus, there is a path in F_1 from k to l beginning with an edge which is also in F_2, say kh. Reapplying the previous argument on the new path (from h to l), we can construct a path in F_1 from i to j beginning with *two* edges in F_2, and so on (the new edge is distinct from the previous almost surely). The finiteness of the graphs implies that this process will eventually yield a path in F_2 from i to j. $\qquad\square$

We direct the reader to Section 2.4 for an example of a dense communication network "pruned" into a very sparse information flow.

2.3   Optimization Problem

We are now in a position to formulate the optimization problem for a static arrangement of nodes. Our aim is to minimize an aggregate cost function subject to local guarantees that a connected information flow can be implemented with a (prescribed) bounded number of hops. We view the latter constraint as a geometrical proxy for quality-of-service and delay-time issues which are impossible to quantify within our simplistic radio model.

Upon formulating the problem and discussing pertinent optimization-theoretic issues, we present a barrier formulation and an associated distributed algorithm for solving this problem on-line. We deliberately refrain from calling the power-control algorithm a protocol, as both it

and the underlying mathematical model lack certain features key to a practical implementation. However, we note that most communications protocols operate in several phases (startup, link formation, tuning, termination, restart, etc.), and we believe that this algorithm can form a theoretical basis for the tuning/optimization phase of a realistic protocol.

2.3.1 Problem Formulation

We will attempt to minimize the aggregate cost incurred in the network, i.e. $\sum_{i \in V} c_i(r_i)$. We choose this as an objective function due to the widespread use of additively separable objectives in distributed optimization. Despite this convenience, we must point out some caveats.

First, for non-physical choices of cost functions, there is no intrinsic interpretation to be assigned to this aggregate cost, especially if the cost functions are not selected to be identical. One may question the utility of distinct cost functions, but we envision the possibility of multi-vehicle applications in which the communications capabilities of each vehicle differ. In particular, if some members of the group play the role of "mobile routers", i.e. they move so as to maintain information flow feasibility in the formation, they would probably be designed with superior radio equipment.

A second drawback of the aggregate cost is that, de facto, it does not penalize the exploitation of a single transmitter for the benefit of the aggregate, except insofar as such a penalty is provided by convexity of the individual cost functions. This can cause the power supply of the "exploited" transmitter to be depleted much faster than that of others, and hence result in a shorter network lifespan. This is, of course, an important consideration, but we expect that using a battery-dependent cost function could potentially resolve this issue.

With these limitations in mind, we will consider the following optimization problem, which we denote $P(F, \mathbf{q}, m)$:

$$P(F, \mathbf{q}, m): \qquad \min \sum c_i(r_i)$$
$$\text{s.t. } R_i^{F,m}(\mathbf{q}, \mathbf{r}) \geq b \text{ for all } i.$$

Here, b is a non-negative scalar which defines the demanded information flow robustness for the network. We do not discuss choices of b, except to refer the reader to the consequences of position robustness in the previous section. We will also see in discussion of the barrier problem that the algorithm will actually produce slightly larger robustness values than demanded, due to the inherent conservativeness of barrier methods.

2.3.2 Characterizing the Optimization Problem

We now turn to characterizing the optimization problem P, essentially for the sake of showing that the problem is well-posed despite the unusual constraint function. We begin by discussing the feasible set.

Proposition 7. *The feasible set of P is compact and connected.*

Proof. It is clear that the feasible set is bounded, since it is a subset of $[0, r_{max}]^n$. Since demanded robustness b is positive, the constraint function is continuous in the feasible set, by Proposition 3. Continuous inequality constraints define closed sets, and so the feasible set is a closed and bounded subset of $\mathbf{R}^n$, thus compact.

To prove connectedness, consider any $\mathbf{r}$ within the feasible region. Linear interpolation between this $\mathbf{r}$ and $\mathbf{r}_{max}$ only increases the individual broadcast ranges. Thus, since $\mathbf{r}$ is feasible, so is $\mathbf{r}_{max}$. Hence, a continuous arc exists between any feasible point and $\mathbf{r}_{max}$. By concatenating the paths of these arcs, we can obtain a continuous arc between any two feasible points. Hence, the feasible set is connected. $\square$

Corollary 4. *P has a global minimum, and there exists a feasible continuous arc linking this global minimum to any other feasible point.*

This set is not convex in general, owing to the possibility of different routing topologies. We thus cannot assert any uniqueness properties for the global minimum, and in fact we can easily construct examples with local minima in addition to the global minimum. We will now show that all minima occur on the boundary.

Proposition 8. *At any local minimum of P, all constraints are active.*

Proof. We will use a contrapositive argument. Let $\mathbf{r}$ be any feasible point such that $R_i^{F,m} > k$ for some i. Then, i must have at least this much robustness to each of its neighbors in $G_{\mathbf{q},\mathbf{r}}$. One can thus reduce r_i by up to $R_i^{F,m} - b$ without reducing the robustness in the network below b. Our assumptions on c_i imply that this will strictly reduce the value of the objective function. Thus, $\mathbf{r}$ cannot be locally optimal. $\qquad\square$

2.3.3 Barrier Formulation and Distributed Steepest-Descent Algorithm

We now present a modified optimization problem in which the constraints are appended to the objective as barrier terms. Before doing so, we comment on the lack of useful alternatives for distributed solution for constrained optimization.

The primary approach taken in distributed constrained optimization is that of duality, i.e. explicitly including Lagrange multipliers as optimization variables. This formulation often results in algorithms which can be implemented in a decentralized fashion. Indeed, the dual problem of our optimization demonstrates similar structure. However, we do not use dual methods because of the special physical significance of our constraint functions, i.e. connectedness of the network. Allowing excursions into infeasible areas could cause disconnection of the network, and suspension of connectivity services for higher-level functionalities such as motion planning.

We thus feel that barrier methods are quite appropriate, as they work strictly with *feasible* points. The resulting (unconstrained) optimization problem is:

$$\min_{[0,r_{max}]^n} \sum_i c_i(r_i) + \mu \sum_i \log(R_i - b).$$

where $\mu > 0$ is a barrier parameter, presumed small relative to the scale of the variables in question (see [2] for details). For sufficiently small choices of μ, the optima of this barrier problem will approach the optima of P arbitrarily closely. Further, for any non-zero choice of the barrier parameter, the barrier objective function rises rapidly to infinity near the boundary of the feasible region.

Now, we consider the gradient of this objective function with respect to the broadcast ranges (for convenience, we denote the objective function f):

$$\frac{\partial f}{\partial r_j} = \frac{dc_j}{dr_j} - \mu \sum_l \frac{1}{(R_l - b)} \frac{\partial R_l}{\partial r_j} = \frac{dc_j}{dr_j} - \mu \sum_{l \in A_j} \frac{1}{(R_l - b)}.$$

The latter formula exhibits the desired distributed structure of the gradient. The derivative of the individual cost function can certainly be computed locally, and the sum over the active set can be accomplished by exchanging constraint-surplus information with communication neighbors (as opposed to *information* neighbors).

We are thus motivated to use the following distributed steepest-descent algorithm (we deliberately say steepest-descent rather than gradient because of non-differentiability):

$$r_i \leftarrow r_i - \gamma \left(\frac{dc_i}{dr_i} - \mu \sum_{l \in A_i} \frac{1}{(R_l - b)} \right).$$

Here γ is a parameter defining the step-size. We discuss γ in the next section.

This is thus a synchronous distributed algorithm for approximately finding a local minimum of P. The intuitive structure of the algorithm is apparent: when the node's variable is not active for any other nodes, it uses a gradient algorithm to minimize its own cost function. When A_i is not empty, the node modifies its own selfish action to prevent disconnection according to a hyperbolic repulsion.

2.3.4 Convergence Issues

The convergence of the steepest-descent algorithm is tricky to characterize. The usual convergence proof for these constant stepsize methods requires globally Lipschitz derivatives, which we do not have in our case because of potential non-smoothness and because of the hyperbolic terms. The problem of non-smoothness can be addressed using generalized gradient arguments as in [4]. It is sufficient to use a generalized gradient which is a maximum-ascent direction to recover the local behavior of a smooth gradient algorithm, and in this regard we have no problem. Unfortunately, the hyperbolic terms will yield to no such easy answer.

It is a standard result that in regions where the derivatives are locally Lipschitz, there is a non-zero γ such that the algorithm will reduce the objective function at each iteration, but it is impossible to patch this result into something global because of the barrier terms. Indeed, by choosing an initial point sufficiently close to the boundary of the feasible region, we can generate an arbitrarily large initial step for any non-zero γ, and so cannot hope to find a stepsize for which this algorithm will converge from arbitrary initial point.

We do not have an answer to this problem, except to say that barrier methods such as these are used regularly in optimization with dynamic stepsize methods without explicit treatment of convergence (see, for example, [20]). We conjecture that, for each initial point, there is a sufficiently small γ such that the algorithm converges to a local minimum, but we have found no way to prove this, nor have we found any related results in the optimization literature.

On the other hand, numerical simulations strongly suggest that initializing the algorithm from $\mathbf{r}_{max}$ and using the *the objective function of P* to determine γ as in [2] will produce a sequence converging to the *global* minimum.

2.3.5 The Problem of Asynchronous Implementation

While we do not claim this algorithm to be sufficiently practical for real-world implementation, we do concern ourselves with the restrictive requirement of synchronous operation (i.e. that all nodes update simultaneously). Synchronization, though possible, is costly in terms of delay, and one would like to avoid this in practice.

We are again presented with a deeply unsatisfying situation because of the barrier terms: in their absence, we would be able to directly apply a bounded-delay asynchronous convergence result from [3] for distributed gradient algorithms.

The numerical results of the synchronous implementation suggest to us that the barrier terms, though theoretically problematic, do not cause the kind of trouble that would arise from a more general function lacking a global Lipschitz constant. This is, of course, difficult to quantify, and we mention it only because we suspect that the kinds of asynchronous results obtained in [3] could also be obtained using this algorithm. We leave this as a conjecture, and emphasize that any practical implementation of such an algorithm will necessarily have to face the problem of non-synchronization.

2.4 Simulations

In this section we first show an application of the distributed heuristic for constructing a sparse information flow. Fifty nodes were distributed uniformly on the unit square, and their maximum broadcast range was set to 0.5. Figure 3 shows the initial information flow F_1 and the resulting sparse information flow F_2^1. We also show bar graphs of the node degrees of F_1 and F_2^1 to verify the degree bound presented in Section 2.2.1.
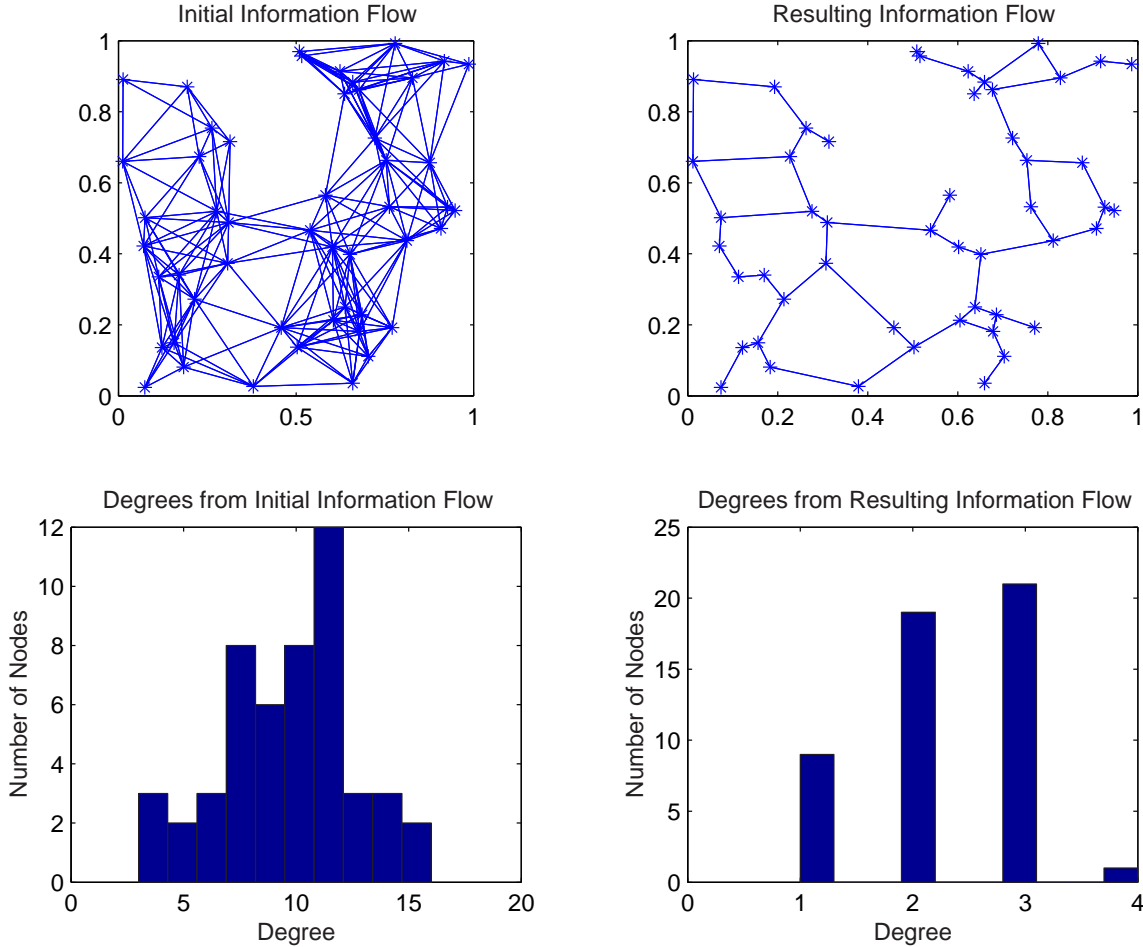
Figure 3: An application of the distributed heuristic of Section 2.2.1. We indeed verify the bound on node degree using this scheme. Although no node occurs with degree 5 in this example, it is easy to construct examples which achieve the bound (and which occur with non-zero probability)

It is evident that the graph created is vastly sparser than the original graph, and maintains connectedness.

We present a sample run from the optimization algorithm in Figure 4. The cost functions were chosen to be r_i^2, and the demanded robustness was set at 0.05. The step-size was 0.1 and the barrier parameter was 0.01.

The "chattering" behavior is a consequence of the non-smoothness. In this run, it is clearly not a source of major difficulty. We have found the chattering to become a problem only if the step-size is set too large (approaching the bound based on P and [2], which in this case is 0.25).

3 Conclusions and Future Work

This paper has presented a continuous-variable constrained optimization problem for the allocation of broadcast power in wireless networks, and a distributed steepest-descent algorithm for the solution of this problem on-line. We have presented a notion of geometric robustness for an information flow on a wireless network, and shown some simple properties of this robustness function. We believe that this geometrical viewpoint, though simplistic as a radio communications model, can provide useful insight into combining wireless networking with dynamic
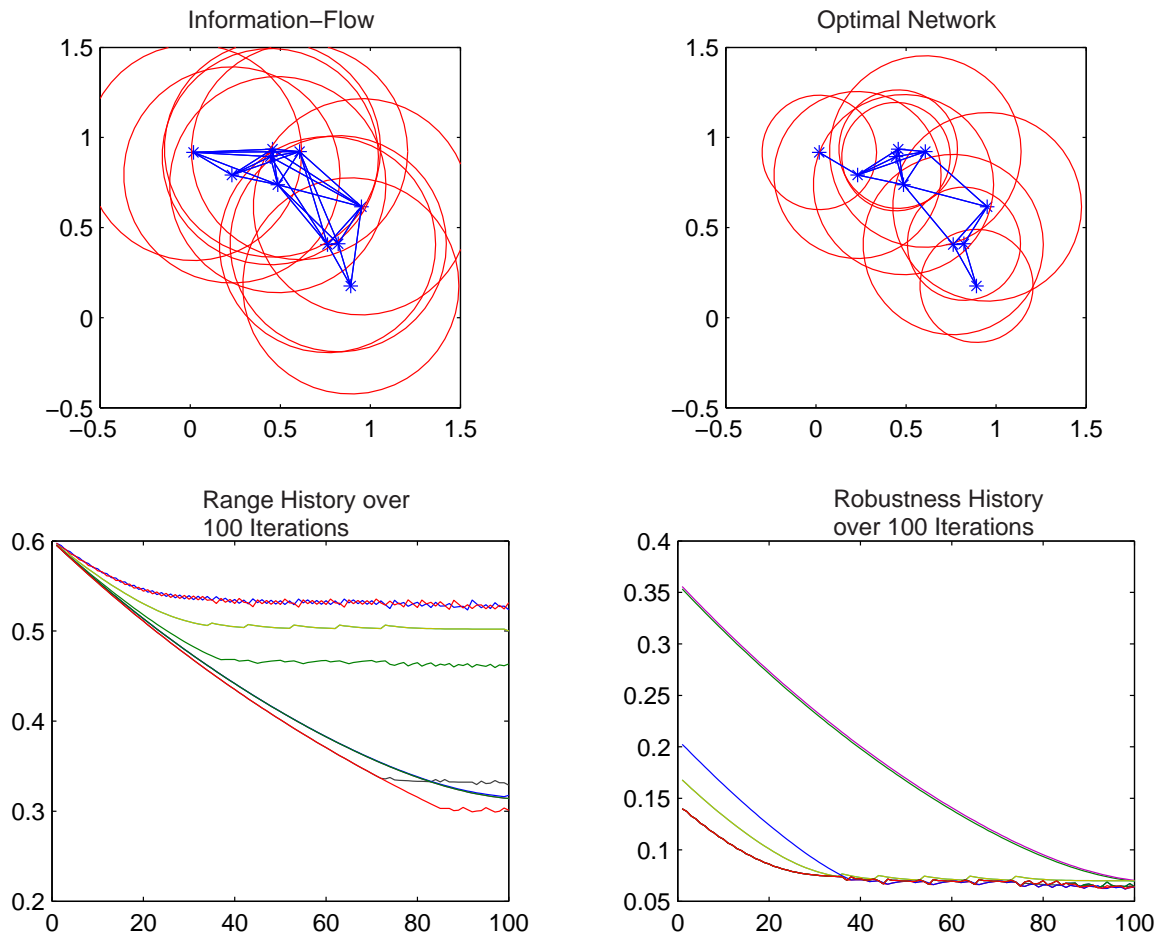
14

Figure 4: A sample run of the optimization algorithm. Note the slight conservativeness of the solution, which is typical of barrier methods. This solution has been shown (by exhaustive enumeration of the possible topologies) to be the global minimum.

phenomena such as motion-control.

We have also presented a distributed heuristic for constructing sparse connected information flows which are feasible for a given wireless network. We feel this may have applications in flocking of automated agents, where the number of individuals makes manual construction of the information flow impractical.

A natural extension of this work would be to formulate a dynamic optimization problem with these constraint functions. Preliminary numerical results suggest that, for fast enough updates, the algorithm we presented can handle mobile transmitters. However, we have not been able to quantify this behavior, and believe that an intrinsically dynamic optimization framework would provide the appropriate vehicle to do so.

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] F. Clarke, *Methods of Dynamic and Nonsmooth Optimization*, SIAM Publishing, 1989

[6] L. Cremean, W. B. Dunbar, D. van Gogh, J. Hickey, E. Klavins, J. Meltzer, and R. M. Murray, *The Caltech Multi-Vehicle Wireless Testbed*, Proceedings of IEEE Conference on Decision and Control, 2002

[7] R. D'Andrea and R. Murray, *The RoboFlag Competition*, Proc. of the American Control Conference, 2003

[8] R. D'Andrea and R. Murray, *The RoboFlag Testbed*, Proc. of the American Control Conference, 2003

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

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

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

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

[13] Y. Ko and N. Vaidya, *Using Location Information in Wireless Ad Hoc Networks*, Proceedings of the 49th IEEE Vehicular Technology Conference, Vol. 3, pp. 1952

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

[15] J. Lee, W. Su, and M. Gerla, *Ad hoc Wireless Multicast with Mobility Prediction*, Proceedings of IEEE ICCCN'99, 1999

[16] 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

[17] N.A. Lynch, *Distribted Algorithms*, Morgan Kauffmann Publishers, 1996

[18] J. Monks, V. Bharghavan, and W. Hwu, *Transmission Power Control for Multiple Access Wireless Packet Networks*, 25th Annual IEEE Conference on Local Computer Networks, 2000

[19] J.C. Navas and T. Imielinski, *GeoCast - Geographic Addressing and Routing*, Proceedings of Third ACM/IEEE International Conference on Mobile Computing and Networking, 1997

[20] J. Nocedal and S. J. Wright, *Numerical Optimization*, Springer-Verlag, 1999

[21] R. Olfati-Saber, R. M. Murray, *Distributed Cooperative Control of Multiple Vehilce Formations Using Structural Potential Functions* Proceedings of IFAC World Congress, 2002

[22] 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

[23] 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

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

[25] T. Salonidis, P. Bhagwat, L. Tassiulas, and R. Lamaire, *Distributed Topology Construction of Bluetooth Personal Area Networks*, Proceedings of IEEE Infocom 2001

[26] C. Saverese, J.M. Rabaey, J. Beutel, *Locationing in Distributed Ad Hoc Wireless Sensor Networks*, International Conference on Acoustics, Speech, and Signal Processing, 2001

[27] A. Stubbs and G. Dullerud, *Networked Control of Distributed Systems: A Testbed*, Proceedings of IMECE 200, 2001

[28] W. Su, S. Lee, and M. Gerla, *Mobility Prediction in Wireless Networks*, Proceedings of the IEEE Military Communications Conference (MILCOM), 2000

[29] C.K. Toh, *Maximum Battery Life Routing to Support Ubiquitous Mobile Computing in Wireless Ad Hoc Networks*, IEEE Communications Magazine, June 2001

[30] Y. Tseng, S. Wu, W. Liao, and C. Chao, *Location Awareness in Ad Hoc Wireless Mobile Networks*, IEEE Computer Magazine, June 2001

[31] 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