Information Flow in Cooperative Control of Multi-Vehicle Systems

Richard M. Murray (PI) and J. Alex Fax

Control and Dynamical Systems
California Institute of Technology
Pasadena, CA 91125

AFOSR BAA 2001-02, Theme 1: Cooperative Control
January 24, 2001

Abstract

As the Air Force moves toward its vision of a mixed force of inhabited and uninhabited air and space vehicles, the role of integrated command, control, communications, and computation will continue to increase in importance. A key element in all future military systems will be the role of communications networks and the increased use of information flow as an integral part of military operations. In many applications, such as coordinated control of unmanned aerial vehicles (UAVs), the presence of a central or supervisory controller cannot be assumed or is undesirable, and the presence and reliability of inter-vehicle communication channels is itself subject to change. As such, there is a need for decentralized control and decision making strategies which are robust to changes in the underlying inter-vehicle communication topology.

This proposal is focused on developing the underlying theory required to achieve integration of information flow into control analysis and design for cooperative tasks of multi-vehicle systems. By making use of tools from control theory, dynamical systems, and graph theory, we will develop a framework for analyzing the effects of information and sensor topology on feedback systems and develop tools for designing information flow and control algorithms that build on this framework. We will apply these ideas to several test problems involving real-time, distributed control of a set of multiple vehicles performing cooperative tasks. In addition to computational exploration through simulation, we plan to implement our algorithms on a multi-vehicle, wireless testbed for networked control, communications, and computing that is being developed at Caltech under separate funding.

1 Introduction

As the Air Force moves toward its vision of a mixed force of inhabited and uninhabited air and space vehicles, the role of integrated command, control, communications, and computation will continue to increase in importance. In addition to ongoing challenges in traditional areas of flight control, new technical challenges will be introduced in areas of man-machine systems, vehicle management systems, distributed control across multiple-platforms, and management of uncertainty and complexity. A key element in all future military systems will be the role of communications networks and the increased use of information flow as an integral part of military operations.

Several recent reports have identified the need for basic research in cooperative control of uninhabited aerial vehicles (UAVs). The Air Force Scientific Advisory Board highlighted the role of UAVs in their New World Vistas study [2] and noted that UAVs must be considered as revolutionary, rather than evolutionary, technology. In 1997, an AFOSR panel recommended that the Air Force support research in networked control systems for tightly-couple, multi-aircraft systems [1] and articulated the need for work in the areas of distributed control architectures, disturbance propagation and string stability, failure modes and health monitoring, and control over noisy communication channels. These recommendations are consistent with the high level needs described DoD's Joint Vision 2020 [12], where technological innovation and information superiority are identified as enablers for key operational goals.

This proposal is focused on developing the underlying theory required to achieve integration of information flow into control analysis and design for cooperative tasks of multi-vehicle systems. By making use of tools from control theory, dynamical systems, and graph theory, we will develop a framework for analyzing the effects of information and sensor topology on feedback systems and develop tools for designing information flow and control algorithms that build on this framework. We will apply these ideas to several test problems involving real-time, distributed control of a set of multiple vehicles performing cooperative tasks. In addition to computational exploration through simulation, we plan to implement our algorithms on a multi-vehicle, wireless testbed for networked control, communications, and computing that is being developed at Caltech under separate funding.

Current tools for analysis and design of flight control systems are focused on control of individual vehicles or small numbers of coordinated vehicles. While these tools allow sensing and communications channels to be included, there is not yet a deep understanding of the role that information flow plays within a networked control system, especially as the number of vehicles becomes large. There has been some recent progress in this area for networked control systems (see, for example, Walsh *et al* [14]), but there has not yet been an exploration of the fundamental limits imposed by sensor and communication topologies (e.g., what is the minimal information flow between a set of vehicles in order to maintain stable formations in the presence of uncertainty). Early work on string stability in longitudinal control of platoons of automobiles showed that this information flow was important and provided insights for the one dimensional case [15]. This situation becomes much more complicated for aerial vehicles operating in dynamic, uncertain and adversarial environments.

Without this understanding, it will be difficult to develop the scalable approaches to control of multi-vehicle systems that are required for future applications. It is also likely that such applications will need to operate in an environment where the sensor and network topology is dynamic (due to motion of the cluster as well as failures of sensing and communication channels) and computation is highly distributed, with no well-defined, central controller. The approach that we propose will help develop the basic theory required to understand such systems and design the information flow as well as the underlying control algorithms needed to develop robust solutions.

We believe that such tools are an essential element of the control toolbox for future uninhabited flight systems. By making use of well-designed information flow, clusters of vehicles will be able to carry out tasks in a manner that provides high confidence operation in noisy, dynamic, and adversarial environments. These systems will be *highly interconnected*, providing robustness to component failure as well as increased capability, *dynamically reconfigurable* based on changes in mission, condition, and environment, and *fault-tolerant* to allow continued operation in the presence of individual failures. While some of these features are present in current vehicle systems, the combination of these features in a large-scale, heterogeneous, multi-vehicle system would provide new capabilities and approaches to motion control for military, scientific and commercial systems. The tools that we will develop under this proposal will help develop the rigorous understanding required to develop systematic techniques for analysis and design of such systems.

In addition to its potential impact on Air Force and other DoD applications, we believe the research proposed here will be important for many scientific and commercial applications as well. Problems ranging from control of the power grid, to spacecraft formation flight for planetary exploration, to air traffic control, to supply chain management are all examples of systems in which networked controls, communications, and computation are playing an increasingly integral role. In all of these problems, distributed computation, cooperative planning, and uncertainty management present challenges for which the existing theory leaves many important questions unaddressed.

While our eventual goal is clearly very ambitious, we believe that our initial focus on information flow in feedback systems is the proper starting point for a broader activity. It builds on our previous results in motion control and dynamical systems, and complements current and proposed work at Caltech in optimization-based, receding horizon control for cooperative systems. By making use of our multi-vehicle testbed, we will be able to leverage these other programs and demonstrate the role of information flow in multi-vehicle systems.

This proposal is organized as follows. In Section 2, we outline some existing challenges in coordinated control of vehicle formations, including some background information on motion control, networked control systems, and information flow in interconnected systems. In Section 3, we outline our research goals, including a test problem which we will use to evaluate the performance of different control and information flow architectures. In Section 4, we present a research methodology which is well suited to achieving those goals and which capitalizes on our previous research experience. In Section 5, we present a year-by-year program plan. In Section 6 we describe existing research and technology transfer efforts which complement and leverage the activities proposed here, including a description of the multi-vehicle testbed that will be used to implement our results. Finally, in Section 7 we provide resumes for key personnel and a program budget.

2 Challenges in Coordinated Control of Vehicle Formations

Unlike many application areas for decentralized control, vehicles in a formation are, as a rule, *dynamically decoupled*, meaning the motion of one vehicle does not directly affect any of the other vehicles.[1] Instead, the vehicles are coupled through the *task* which they are jointly asked to accomplish. Tasks of this nature include requiring a formation to fly in a specific pattern, distribute itself evenly over a specified area, or arrive simultaneously at specified endpoints. Other tasks include the assignment of roles to individual vehicles within a formation which enable the entire formation to accomplish a higher-level task.

When the formation is dynamically coupled, that coupling constrains, or at least naturally suggests, what information must be available to each component of the decentralized controller. In the case of cooperative vehicle control, no such architecture is necessarily suggested. In some cases, such as one-dimensional vehicle following [14], the task may suggest which vehicles ought to have information about one another, but for many tasks this is not the case. As such, central to any discussion of cooperative control of vehicles is a determination of the nature of the *information flow* throughout the formation. We will distinguish between two types of information flow: *sensed information*, meaning the ability of a single vehicle to sense some information (e.g. relative position) about another vehicle in a way which involves no action on the part of that vehicle, and *transmitted information*, meaning transfer of information between two vehicles which requires some action on the part of both the transmitter and recipient of that information. Sensing and transmission, or "seeing" and "hearing," as we refer to them colloquially, together are the means by which each vehicle acquires the information necessary to perform its task within the formation.

Several observations about the information flow within a formation make clear the need to consider its impact on the formation performance. The first is that as a rule, no vehicle will be able to see or hear the entire formation. Thus, having each vehicle simultaneously solve a centralized control problem using complete information is infeasible, thus necessitating some form of decentralized control. The second observation is the sensing and transmitted information topologies are themselves *dynamic*, meaning they are subject to change due to changes in the formation itself or due to outside influences. As such, a control law which is optimized for one topology may exhibit poor performance, or even instability, for another topology.

One possible approach to vehicle formation control is to implement a centralized controller or decision maker, and to overcome the limitations in the information flow topologies by having each vehicle transmit all information it possesses, thereby allowing all information to eventually arrive at one vehicle, which can compute the centralized control law on behalf

[1]There are exceptions, of course, such as drag reduction via formation flight of airplanes, but we are not considering those in this proposal.

of the formation. While this architecture may be appropriate in certain cases, it possesses certain deficiencies which render it infeasible in a dynamic and adversarial environment.. This architecture, which requires maximal information flow, is necessary slow and expensive in terms of bandwidth. It is fragile, in that it depends on the *reliable* transmission of large amounts of data over potentially *unreliable* and dynamic communication channels. Instead, we intend to research the implementation of decentralized control laws augmented by *minimal* information flow. A minimal information flow paradigm has the potential to balance the performance improvements achieved through information sharing with the requirements of reliability and stealth.

Vehicle formations could be required to perform multiple tasks, only some of which traditionally fall under the domain of controls. Traditional controls tasks include:

1. **Stabilization / Disturbance Rejection** The most basic needs of the formation are stabilization of individual vehicles and formation maintenance in the face of external disturbances. Ideally, one would be able to design controllers locally which guarantee some level of performance of the formation as a whole. While this concept has been explored in the case of one-dimensional vehicle following, leading to the notion of string stability [15], we wish to generalize these concepts to arbitrary formations. Simple examples have shown that receiving information from other vehicles, rather than simply treating their motion as a disturbance, leads to vastly improved performance. These concepts have yet to be explored in the setting of more complex formations and information flow topologies.

2. **Trajectory Generation** Another traditional task is the generation of trajectories for individual vehicles to follow. In the case of formations, this could include computing a trajectory for the formation as a whole to follow, or for individual vehicles to follow to reconfigure the formation. Trajectory generation for formations in a decentralized setting poses interesting challenges. If the desired location of each vehicle is predetermined, then the trajectory generation problem for the formation decouples to a trajectory generation problem for each individual vehicle. However, in many cases the formation reconfiguration problem possesses *symmetries* which induce coupling. For example, if the vehicles in the formation are require to acquire a pattern relative to one another, then the absolute final position of each vehicle is not uniquely defined. The optimal trajectory for each vehicle to follow is therefore a function of the current state of each vehicle in the formation. A related issue is the resolution of *discrete* symmetries. The task may require that the vehicles arrive at certain locations, but which vehicle arrives at which location is unimportant. Again, the question arises: what is the minimal amount of information each vehicle must possess to compute its trajectory in some reasonably optimal fashion? How does each vehicle acquire that information in a decentralized architecture?

Other, nontraditional tasks include

3. **Formation Self-Assessment** In the absence of a centralized controller, the formation will be required to *self-assess*, meaning determine its own properties. Examples of

self-assessment include health monitoring and role management. When the information flow topologies are themselves uncertain and dynamic, a crucial component of formation self-assessment is the determination of those topologies to enable implementing other distributed algorithms which act along those communication networks.

4. **Environmental Awareness** The formation will, of necessity, be required to sense its environment in order to react accordingly. Examples include obstacle detection, target acquisition, and threat identification. When the sensing and computational capacities of the formation are distributed across multiple vehicles, environmental awareness can only be effected by through the flow of sensed information to a location or locations where it can be synthesized and interpreted. Again, the need arises for efficient, reliable information flow.

5. **Role Selection** Accomplishing the formation's mission often requires the assignment of different roles to vehicles in the formation. These roles could be physical (e.g. attack, scan for threats), or computational. The formation will be required to determine what roles are needed at any given time as a function of the current environmental conditions, and to assign those roles based on each vehicles current capabilities. Again, effecting this requires methodologies for distributed role assignment and re-assignment based on changes in the formation status, including changes in the information flow topologies.

3 Research Goals

What the above challenges have in common is that they involve managing the interaction of an intelligent, dynamic networked system with a dynamic, uncertain, and adversarial environment in a setting where the means to manage that interaction are constrained by limitations in the flow of information within the network.

Accordingly, we have identified the following areas to pursue to in our research:

3.1 Fundamental limits of formation control with limited sensing

We propose to investigate the relationship between the sensed information topology and the formation behavior. This represents a specific subproblem within the general topic of decentralized control which merits systematic study. As we discuss above, this problem is highly relevant to coordinated control of vehicle formations.

Questions we wish to address on this topic include: how does the sensed information topology affect formation controllability and observability? For a given controller or class of controllers, how does the sensing topology affect the stability and performance of the formation? How do changes in the sensing topology affect the stability/performance of a given controller? What are appropriate characterizations of performance for vehicle formations? Can the notion of string stability be generalized, or are there other, more appropriate concepts? Thorough understanding of the effects of limited sensing on formation performance should lay the foundation for identifying what information should be passed between vehicles to improve formation performance.

3.2 Information Management for Cooperative Control

As discussed above, any discussion of formation control of vehicles leads naturally to the issue of information management. We intend to research methods by which information management can be systematized to produce improved formation performance or to extend formation functionality. As discussed above, there is a tradeoff between the quantity of information flow and its reliability. In dynamic, uncertain, and adversarial environments such as those combat vehicle formations are likely to face, it is our opinion that the minimal information paradigm possess advantages which merit consideration, although further research is needed. In the sections below, we divide this area of research into three (highly interrelated) subtopics:

Minimal Information Representation

The range of information which could be available to an individual vehicle lies on a continuum from full information, leading to centralized control algorithms, to having only sensed information. The first question one must address in using minimal information flow is: how does one minimally represent the relevant information in a way which is still useful to the individual vehicle. Steps in this process include quantification of the task and establishing some performance index by which to measure the utility of any information. Once that is done, it would then be possible to discuss the tradeoff between information size and system performance. We intend to research methods by which the quantification of a task and the minimization of the associated information can be systematized and analyzed.

Distributed Information Computation

Once the information which each vehicle needs has been determined, a method by which that information arrives at each vehicle must be devised. In many cases, the information that each vehicle must possess must be computed from other information resident on multiple vehicles. This naturally leads to the issue of distributed information computation, meaning methods by which the minimal information can be synthesized without resorting to collecting all information at a single vehicle (and thereby violating the minimal information paradigm). In particular, any realistic method must take into account the existing transmitted information flow topology, and be robust to possible changes in that topology. We intend to investigate distributed computation algorithms and their interaction with the information flow topologies.

Performance Augmentation

The final area of research is an investigation of how information is used to augment the formation performance. In Section 2, we outlined several areas where inter-vehicle communication can improve formation performance and facilitate intelligent formation behavior. We intend to research methods by which those performance improvements can be achieved. Our initial emphasis will be on the extension of existing single-vehicle control methods to vehicle formations (see, for example, Section 6.3), with a goal of understanding what role information sharing plays in facilitating their implementation.

3.3 Test Problem

A sample test problem, which we hope to implement experimentally, highlights these issues. Consider a fleet of vehicles, tasked to fly in a predetermined pattern, which confronts an obstacle in its path as shown in Figure 1. The formation must circumvent the obstacle and, if

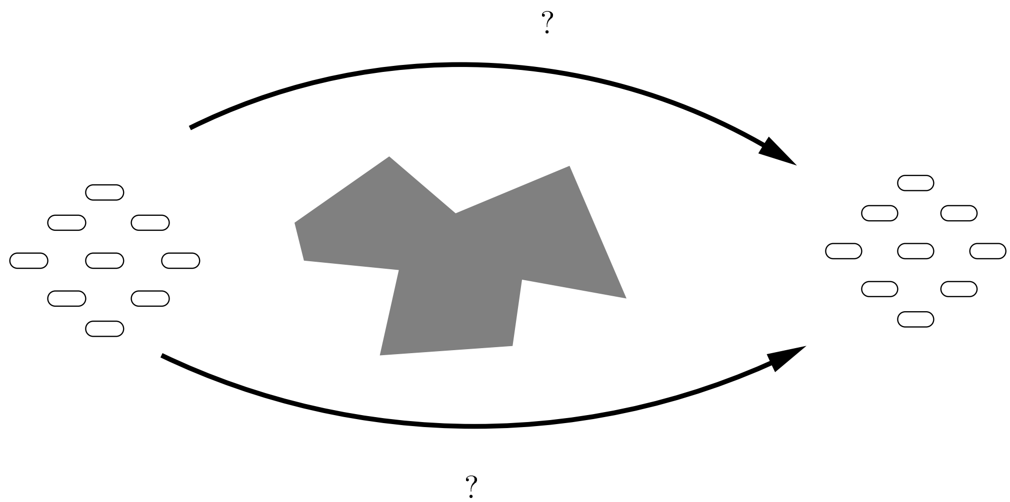

Figure 1: Formation Obstacle Avoidance Problem

necessary, re-establish the formation pattern on the far side of the obstacle. To achieve this, the formation must accomplish:

- Recognize the presence and shape of the obstacle and communicate this information to the formation. This may involve synthesis of data from multiple vehicles which sense the presence of the obstacle.
- Determine a path to circumvent the obstacle. Each vehicle must determine which direction to follow around the obstacle, and it cannot be assumed that the vehicles will (or should) each make the same choice. As such, the formation may be asked to decompose into two smaller formations, and to recompose into one formation on the far side of the obstacle.
- Implement inter-vehicle collision avoidance during the reconfiguration.

The realities of vehicle dynamics, formation speed, and communication bandwidth will play a large role in the difficulty of this problem. Conversely, the ability of the formation to rapidly assess the obstacle, and formulate a reaction strategy will determine the maximum allowable operating speed of the formation. The sample problem can be implemented in two ways: with predetermined flow topologies, to test their relative merits, or with flow topologies which change according to deterministic (e.g. inter-vehicle distance, line-of-sight) formation properties and random effects, to test the robustness of the formation to variations in the flow topologies.

3.4 Concluding Remarks

Taken together, these different research goals address aspects of cooperative control of multi-vehicle systems which we consider to be crucial in enabling this technology. Our overarching goal is the implementation of a control architecture and control laws which are capable of achieving mission objectives and reacting intelligently to the environment in a way which is robust to uncertainty in the environment and its effects on the ability of the formation to communicate with itself. We intend to proceed incrementally toward this admittedly ambitious goal, emphasizing the need to clearly articulate a framework through which these questions can be analyzed, and leveraging our research experience in dynamical systems theory and control of individual vehicles wherever possible.

4 Research Methodology

In this section, we outline a conceptual framework through which we can address the research areas discussed above.

4.1 The formation as a graph

Graph theory possess a rich language for conceptualizing a set of vehicles and their associated communication links. A graph is a set V of nodes or vertices v_i, along with a set E of edges $e_{ij} = (v_i, v_j)$, representing a link between from node i to node j. If we assign a node to each vehicle, we can represent each communication topology as a graph on that set of nodes. In this case, we do not wish to assume bidirectional communication; it is possible that one vehicle can hear another without the reverse being true. As such, we wish to consider directed graphs (digraphs), as a suitable model for a formation and its network topologies.

Graph theory possesses many methods of characterizing the structure of graphs, and we intend to determine which of those methods are relevant to the problem of coordinated control. Graph-theoretic characterizations may be a source of minimal characterizations of the graph structure which would be of use in information flow. One area which appears promising is spectral graph theory, in which the graph is represented as a matrix, and properties of that matrix (e.g. its eigenstructure) are related to graph-theoretic properties such as connectivity, diameter, girth, etc. [6, 5, 4, 9, 7]

4.2 Information flow as a dynamic system

As information is transmitted throughout the formation, each vehicle receives some information set and performs computations using that information along with other information, such as sensed information or data stored from earlier measurements or transmissions. Each vehicle then transmits data based on those computations. Regardless of the specific nature of the computations, we can regard the computations together with their transmissions as a dynamical system whose states are the information computed at each vehicle. Specifically, let x_i denote the state of the ith vehicle, and let us model the dynamics of each vehicle (in discrete time, for simplicity) as

$$x_i(k+1) = f(x_i(k), u_i(k)) \tag{1}$$

Suppose further that each vehicle can sense some subset of the vehicles in the formation, denoted S_i, and receives transmissions from another subset of vehicles in the formation, denoted T_i. Letting y denote the sensed information, $c(\cdot)$ the function which returns y, z the transmitted information, and $h(\cdot)$ the computation performed at each vehicle, we can represent the computation in the following way:

$$y_i(k) = c(x_i(k), x_{S_i}(k)) \tag{2}$$

$$z_i(k) = h(y_i(k), z_k(k-1), z_{T_i}(k-1)) \tag{3}$$

Note that $z_i(k)$ is a function of information from previous time steps, representing a delay in transmission. These equations therefore define a dynamical system in which z evolves as a function of its current and past values as well as the state x of the formation. If, in turn, the control u is computed as a function of z, i.e.

$$u_i(k) = g(y_i(k), z_i(k)) \tag{4}$$

then the system is fully coupled. Of course, many variations on this could be devised, such as variable time delays and computations based on information from multiple past time periods.

4.3 Putting it all together

Within this framework, one can conceptualize the multiple roles information flow might take on. For example, if $h(\cdot)$ were independent of the measurements y, the computations might represent formation self-assessment, e.g. determining the structure of the current communication topology. If u were computed using z, then the computation might represent a distributed trajectory generation algorithm or a decentralized stabilizing controller. Alternatively, if the control law g is chosen from within a set based on the value of z, then the computation might represent a role selection, where each choice of g represents a different role. In each case, assessing the performance of the algorithm amounts to analysis of a dynamical system whose structure is intimately tied to the graph structure of the information flows. Furthermore, design of the algorithms will require a synthesis of ideas and tools from control theory, dynamical systems theory, and graph theory.

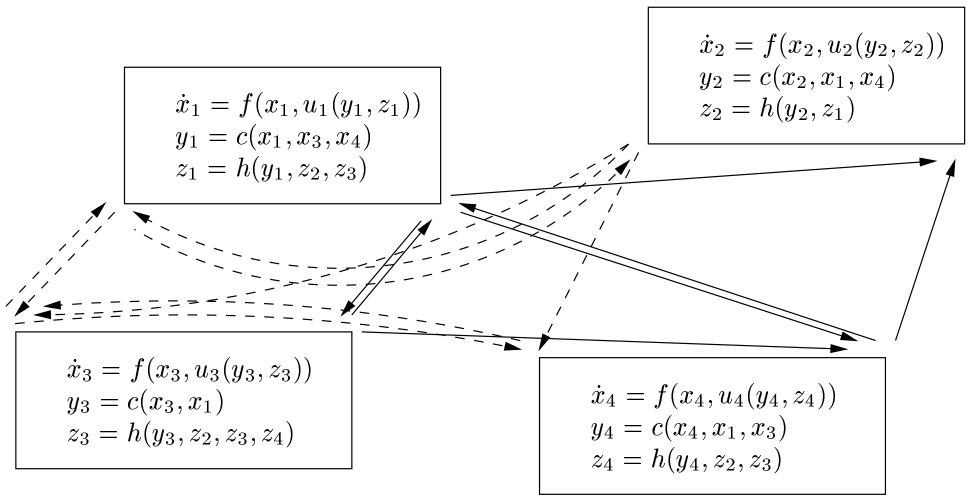

Figure 2: Formation Equation Schematic

Figure 2 shows a schematic of the equations with the associated graphs for a formation of four vehicles. The solid lines represent the flow of *sensed* information from vehicle to vehicle, and the dashed lines represent the flow of *transmitted* information. The graph structure of the information flow is readily apparent in the diagram.

Mathematically, the graph structure can be shown to arise in the equations of motion as well. Consider a set of vehicles with linear system dynamics and (for simplicity) no information flow:

$$\dot{x}_i = Ax_i + Bu_i \tag{5}$$

$$y_i = C\left(x_i - \frac{1}{d_i}\sum_{j \in S_i} x_j\right) \tag{6}$$

$$u_i = -Ky_i \tag{7}$$

where d_i is the *indegree* of the ith vehicle (the number of edges directed to that node). Note that each vehicle collects its sensed information by averaging them. If we wish to collect these equations together, we can do so efficiently using notation from graph theory. The equations can be restated as

$$\dot{x} = \left(\hat{A} - \hat{B}\hat{K}\hat{C}D^{-1}L\right)x, \tag{8}$$

where the hat superscript indicates block diagonal stacking of the associated matrix n times. The matrix D is the the matrix with the indegree of each vertex along the diagonal, and L is the *Laplacian matrix* or *admittance matrix* [9], defined by placing the indegree along the diagonal, and setting the (i, j) element to -1 if an edge leads from vertex i to vertex j, and is 0 otherwise. For the sensed information graph in Figure 2, the indegree matrix is diag(2, 2, 1, 2), and the Laplacian is

$$L = \begin{pmatrix} 2 & 0 & -1 & -1 \\ -1 & 2 & 0 & -1 \\ -1 & 0 & 1 & 0 \\ -1 & 0 & -1 & 2 \end{pmatrix}. \tag{9}$$

4.4 An Example

Preliminary investigation has shown that the Laplacian useful not only notationally, but for assessing stability properties of the formation. For example, consider a SISO second-order plant of the form

$$A = \begin{pmatrix} 0 & 1 \\ -2\zeta\omega & -\omega^2 \end{pmatrix}, \quad B = \begin{pmatrix} 0 \\ \omega^2 \end{pmatrix} \quad C = \begin{pmatrix} 1 & 0 \end{pmatrix}, \quad K = 1 \tag{10}$$

and ask the question: for a given graph (i.e. Laplacian), what is the range of ζ, ω for which the formation is stable. The result turns out to be that stability is independent of ω (assuming it is positive), but that the minimum value of ζ for which the system is stable varies, and is closely correlated with the maximum argument of the eigenvalues of the (normalized) Laplacian $D^{-1}L$. Table 4.3 shows the set of all nonisomorphic digraphs of minimum indegree 1 (meaning each vehicle can see at least one other vehicle) for a formation of three vehicles, the minimum stable value of ζ, and the maximum argument of the eigenvalues. Note that the naive information management paradigm "use all available information" appears not to be the appropriate one in terms of stability: the fourth graph, with a minimum ζ of 0.16, can improve its stability margin to $\zeta = 0$ if the lower right vehicle *ignores* the information the lower left vehicle, thereby rendering itself isomorphic to the sixth graph.

This correlation is seen in formations with larger numbers of vehicles as well. Figure 4.4 shows ζ_{min} versus Laplacian eigenvalue argument for a large selection of formations with between three and seven vehicles. As is seen in that figure, stability margins and the position of the Laplacian eigenvalues are clearly correlated. This correlation was determined through examination of the data: the graph-theoretic significance of the eigenvalue argument is not yet known to us. Nonetheless, that the eigenstructure of the Laplacian plays a significant role is clear, and will form one of our first avenues of investigation.

Graph	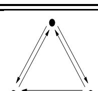	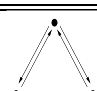	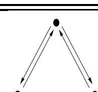	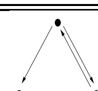	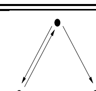	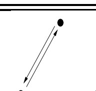	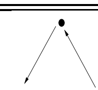
ζ_{min}	0	0	0	0.16	0	0	0.27
$\max \angle\lambda(L)$	0	0	0	18.44°	0	0	30.00°

Table 1: Graphs, Stability Margins, and Laplacian Eigenvalues

One existing research effort within spectral graph theory is the relationship between locations of the graph spectra and structural properties of the graph. For many applications, the graphs under consideration are undirected, which implies symmetry of the Laplacian. In this application, the graph cannot be assumed to be undirected, and therefore another natural avenue of research is the extension of existing results to the case of digraphs. The existing literature provides a good foundation from which to determine which structural properties of the graph influence the formation behavior, and as such are natural candidates for calculation in a minimal information flow setting.

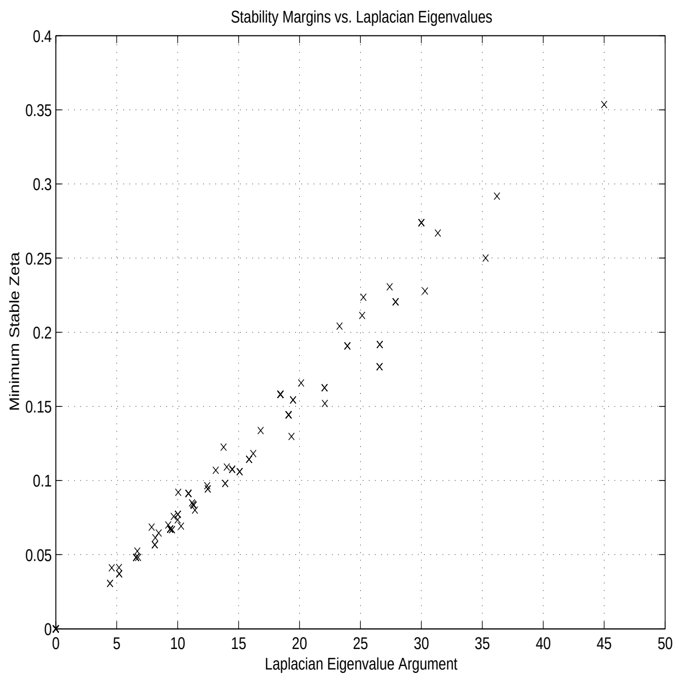

Figure 3: Minimum stabilizing ζ vs. maximum argument of Laplacian eigenvalues

4.5 Conclusions

Dynamical systems theory can pose as a unifying framework for many issues in coordinated control of vehicle clusters. Graph theory has already developed methods for representing the flow of information which are amenable to inclusion in this framework. Furthermore, preliminary investigation has revealed that graph-theoretic concepts play a role in characterizing formation behavior. Finally, this framework naturally allows the inclusion of ideas from control theory to achieve both formation performance objectives and also in the design of information flow algorithms. The synthesis of ideas from dynamical systems theory, control theory, and graph theory has the potential to form the theoretical foundation for addressing many of the challenges facing the Air Force in implementing coordinated multi-vehicle formations.

5 Program Plan

In this section, we outline a year-by-year plan for our proposed research program.

Year 1:

- Investigate the relationship between graph theory and control theory, focusing on identification of graph-theoretic parameters which are of use in relating individual vehicle behavior to formation behavior.
- Research the current state of the art in techniques for multi-agent role assignment, focusing on minimal representations of vehicle and information state.
- Research extensions of optimization based control to decentralized multi-vehicle control.
- Develop simulation tools for testing control architectures for the obstacle avoidance problem concurrently with testbed development.

Year 2:

- Continue research into the relationship between graph theory and control theory, focusing on identification of control strategies which are robust to changes in the information topology, and what information is needed locally to facilitate implementation.
- Identify methodologies for multi-agent role assignment, and investigate means by which it can be implemented in a decentralized, information-minimal fashion. Focus on roles which are necessary to implement obstacle avoidance (e.g. distributed sensing, obstacle identification, formation splitting and merging).
- Evaluate, via simulation, the performance of different control strategies and role assignment for accomplishing obstacle avoidance.

Year 3:

- Continue research into relationship between graph theory and control theory.
- Implement obstacle avoidance on the multi-vehicle testbed. Evaluate, via experimentation, the performance of different control architecture. Quantify, for this example, the effects of communication bandwidth and topology on formation performance.
- Identify open problems and avenues for future research.

6 Research Interactions

Caltech has several ongoing research efforts that we intend to leverage in our research in cooperative control of multi-vehicle formations.

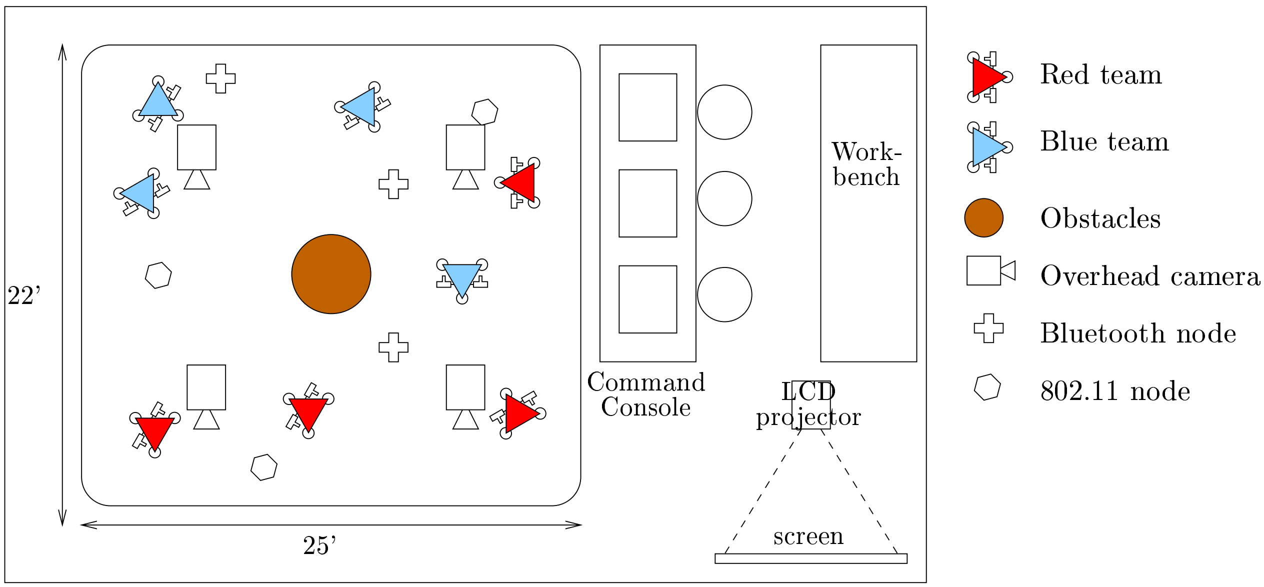

Figure 4: Multi-vehicle testbed layout.

6.1 Multi-Vehicle Wireless Testbed

We will implement the research described above on a multi-vehicle testbed that is under development at Caltech as part of an newly approved DURIP effort. Although this platform is not being developed under the proposed program, we give a short description since it will be used to implement and test the proposed deliverables from this program.

Figure 4 shows the layout of the eventual system that we plan to construct. The testbed consists of

- 8–10 mobile vehicles with integrated computing and communications, including wireless Ethernet (802.11), Bluetooth, and infrared communications capability
- 2-4 fixed communications nodes, capable of broadcasting on multiple channels, connected to fixed computing platforms
- A set of overhead cameras that can be used to provide position information to the vehicles (perhaps by simulating GPS information across the wireless Ethernet network)
- A command console consisting of computing and communications nodes

A unique feature of the testbed is the use of vehicles that have second order dynamics, requiring real-time feedback algorithms to stabilize the system while performing cooperative tasks.

The testbed is capable of operating in a variety of different modes. For the purposes of this proposal, we will concentrate on the use of distributed computation across a communications network to execute a cooperative task. The system can be commanded to drop communications links or an individual vehicle can have a real or simulated failure, allowing robustness of the approach to be verified. While the testbed is capable of being used in an adversarial environment (4-5 vehicles per team), we anticipate initial usage under this program to be as a single (larger) cluster of vehicles.

The individual vehicles will consist of four subsystems: power, propulsion, computing and communications. To minimize cost and complexity, we intend to use commercial parts whenever possible. The propulsion system will consist of two electrically driven ducted fans, which are readily available through hobby dealers. These fans are capable of producing approximately 7N of force each and are similar to those that are employed on the Caltech ducted fan [10]. The fans will be used in a differential mode to vector the thrust on the vehicle, allowing rotation. The motors will be powered by batteries and we expect to be able to run the system for 30–60 min between charges. The location of the vehicles will be provided by overhead cameras, augmented by local sensors, if needed. We intend to use a system similar to that used in the RoboCup competition,[2] where markers are placed on each vehicle to simplify the localization problem.

Computation on the vehicle will be performed on either ruggedized, laptop computers or PC-104 based, single board computers. The computers will be responsible for running the vehicle real-time control systems, performing guidance and navigation tasks, and managing the communications system. The real-time control system will use commercial data acquisition and control hardware to interface with the propulsion system.

6.2 Coordinated Control of Microsatellite Clusters

We are currently in the final year of a three-year contract from AFOSR to research coordinated control of microsatellite clusters. Microsatellite clusters are an example of a vehicle formation in the sense defined in this proposal. In the case of microsatellite clusters, the joint task is the maintenance of some pattern in space over a period of time to implement sparse aperture radar or a similar interferometric task. An exemplary mission of this type is the Air Force's TechSat21 mission[3]. Many of the ideas presented in this proposal were developed during our research into microsatellite clusters.

One of the mission drivers is the minimization of fuel consumption, which leads to the need for formation control strategies and the need to understand how control laws implemented by each vehicle affect overall formation motion. Our initial avenue of research was understanding the dynamics of near-satellite orbital motion, and develop methods for finding fuel-efficient natural orbits and control laws about those orbits [3]. One method we have used to address this is optimization-based control (see the next section), in which the optimal trajectories for all vehicles are computed centrally [11]. Extending that idea to computing each trajectory in a decentralized fashion will form one of the initial avenues of research of this proposal.

[2] See `http://www.robocup.org` for more information.

[3] See `http://www.vs.afrl.af.mil/FactSheets/techsat21.html` for more information

6.3 Real-Time, Optimization Based Control

As part of a DARPA-sponsored program in Software Enabled Control (SEC) and a proposed MURI on Cooperative Control, we are developing techniques for real-time control of flight vehicles using online optimization and coordinated control of multi-vehicle systems. In our ongoing SEC program, we are developing techniques that exploit the dynamics of the vehicle to minimize the amount of computation required. Optimal trajectories are computed the respect the vehicle's dynamics, as well as state and input constraints. In the proposed MURI activity, these motion control primitives are used as the basis for multi-vehicle planning and coordination.

The basic framework we have developed is a receding horizon approach, where an optimal trajectory is computed over a finite time that is shifted forward in time at each sampling instant. To insure stability in the closed loop, we make use of a control Lyapunov function approached developed at Caltech by Doyle and his co-workers [8, 13]. The key idea in this approach is to use a control Lyapunov function as a terminal cost for a receding horizon optimization problem. This approach can be show to have good theoretical stability properties and initial implementations on sample problems indicate that even very short horizons give near optimal performance. This approach works in part because the control Lyapunov function is an approximation for the optimal "cost to go" and hence using it as a terminal constraint allows a simplification of the optimization without unduly sacrificing performance.

There are many advantages over this approach for low-level control over other, more traditional approaches. The key advantage of optimization-based control is reconfigurability in the presence of changes in the system, mission, or environment. Because optimization-based techniques rely explicitly on an internal model of the plan, performance objective, and constraints, the low-level control can automatically reconfigure its operations appropriately as these change during a mission.

As part of the SEC program, we are extending our work to consider multiple vehicles, but with no communications constraints. The theoretical basis for this extension will be considered in the proposed MURI program, along with development of a formal specification and control language for generation of high confidence, embedded systems (joint with Prof. Jason Hickey, Caltech CS).

The current proposal complements our existing work and focuses on the role that *information flow* plays in optimization-based control of multi-vehicle systems. Through collaborations between students, faculty, and postdocs, we expect significant interaction between these different programs, with significant leveraging of activities by virtue of the collaborative nature of the research environment at Caltech. The multi-vehicle testbed will be used for all three programs, providing a convenient mechanism for integrated the different techniques being explored.

6.4 Industrial Transitions

Over the past five years, Caltech has developed several relationships with aerospace companies that can serve as mechanisms for transfer of this technology to application. Due to the long range nature of the proposed research, we anticipate that such interactions with be through collaboration on 6.2 programs in the area of multi-vehicle flight control. Specific companies that have ongoing interactions with Caltech CDS include:

- **Northrop Grumman** As part of the AFOSR PRET Center on "Robust Nonlinear Control Theory with Applications to Aerospace Vehicles", Caltech has built an interaction with Northrop, primarily through contact with Dr. Frank George, of the Advanced Systems Development group. Northrop is also a participant in the DARPA Software Enabled Control (SEC) program, where they are looking at control of UAV swarms. We anticipate that the proposed research would be of high interest to Northrop after sufficient development and demonstration.

- **Boeing** Researchers at Boeing are responsible for the development of the Open Control Platform (OCP) for the DARPA SEC program and there have been frequent interactions with Boeing in this contact (primarily through John Hauser and Gary Balas, who are co-investigators on the Caltech PRET center and the SEC program. We intend to implement OCP on the multi-vehicle testbed and hence the results of this program may also be implemented under the OCP. This would enable Boeing and other researchers to transfer algorithms for information flow developed under this program to other systems making use of the OCP.

- **Honeywell Technology Center** Caltech has had a long-standing relationship with Honeywell, who was the primary partner for the AFOSR PRET Center. This interaction includes frequent visits by Honeywell researchers to Caltech (particularly Tierno and Glavaski, former Caltech CDS students) as well as participation in the DARPA SEC program.

- **Lockheed Martin Aerospace** Caltech is currently working with Lockheed on a proposed activity involving the use of real-time, optimization-based control for reconfigurable control. Through this activity, we hope to develop a stronger relationship with Lockheed and explore possible future collaborations on multi-vehicle problems are possible. Dr. James Buffington, a former researcher for the AFRL Center of Excellence in Control Science, is the primary contact. We have also established preliminary contact with the guidance and control group for the X-33 program, based in Palmdale, CA.

As part of this program, we intend to keep our industrial partners and collaborators aware of the theory, software, and demonstrations developed under this program and look for opportunities to participate in collaborative programs. We are particularly interested in the possibility of participating as a subcontractor in 6.2 programs that provide industrial funding to explore advanced concepts (such as we are currently doing with Lockheed).

References

[1] S. Banda, J. C. Doyle, R. M. Murray, J. Paduano, J. Speyer, and G. Stein. Research needs in dynamics and control for uninhabbited aerial vehicles. Panel Report, November 1997. `http://www.cds.caltech.edu/~murray/notes/uav-nov97.html`.

[2] Air Force Scientific Advisory Board. New world vistas: Air and space power for the 21st century. Summary Volume, December 1995. `http://afosr.sciencewise.com/afr/-sab/edu/text/any/afrtnwv.htmd`.

[3] Dong-Eui Chang, David Chichka, and Jerrold E. Marsden. Lyapunov transfer between elliptic orbits. Submitted, *ASME Journal of Guidance, Navigation and Control.*

[4] Fan R. K. Chung. *Spectral Graph Theory*, volume 92 of *Regional Conference Series in Mathematics.* American Mathematical Society, 1997.

[5] Dragos Cvetkovic, Peter Rowlinson, and Slobodan Simic. *Eigenspaces of Graphs*, volume 66 of *Encyclopedia of Mathematics and Applications.* Cambridge University Press, 1997.

[6] Dragos M. Cvetkovic, Michael Doob, and Horst Sachs. *Spectra of Graphs*, volume 87 of *Pure and Applied Mathematics.* Academic Press, Inc., 1980.

[7] Miroslav Fiedler. Laplacian of graphs and algebraic connectivity. In *Combinatorics and Graph Theory*, volume 25 of *Banach Center Publications*, pages 57–70. PWN-Polish Scientific Publishers, 1989.

[8] Ali Jadbabaie, Jie Yu, and John Hauser. Unconstrained receding horizon control of nonlinear systems. to appear, *IEEE Transactions on Automatic Control*, 2001.

[9] Russell Merris. Laplacian matrics of graphs: A survey. *Linear Algebra and its Applications*, 197,198:143–176, 1994.

[10] Mark Milam and Richard M. Murray. A testbed for nonlinear flight control techniques: The caltech ducted fan. In *Proceedings of IEEE Conference on Controls Applications*, 1999.

[11] Mark Milam, Nicolas Petit, and Richard M. Murray. Constrained trajectory generation for microsatellite formation flying. In preparation.

[12] Department of Defense. Joint Vision 2020. `http://www.dtic.mil/jv2020/`.

[13] Jim Primbs, Vesna Nevistic, and John Doyle. Nonlinear optimal control: a control Lyapunov function and receding horizon perspective. *Asian Journal of Control*, 1(1):1–11, 1999.

[14] G.C. Walsh, H. Ye, and L. Bushnell. Stability analysis of networked control systems. In *Proceedings of the American Control Conference*, pages 2876–2880, June 1999.

[15] Diana Yanakiev and Ioannis Kanellakopoulos. A simplified framework for string stability analysis in AHS. In *Proceedings of 13th IFAC World Congress*, volume Q, pages 177–182, San Francisco, CA, 1996.

7 Key Personnel and Budget

7.1 Personnel

The key personnel for this proposal will be Professor Richard Murray and Alex Fax.

Richard Murray Richard M. Murray received the B.S. degree in Electrical Engineering from California Institute of Technology in 1985 and the M.S. and Ph.D. degrees in Electrical Engineering and Computer Sciences from the University of California, Berkeley, in 1988 and 1991, respectively. He is currently a Professor of Mechanical Engineering and the Chair of the Division of Engineering and Applied Science at the California Institute of Technology, Pasadena. In 1998-99, Professor Murray was the Director of Mechatronic Systems at the United Technologies Research Center. His research interests include nonlinear control of mechanical systems with applications to aerospace vehicles and robotic locomotion, active control of fluids with applications to propulsion systems, and nonlinear dynamical systems theory. Murray is the recipient of an NSF Career Award, an ONR Young Investigator Award, and the 1997 Donald P. Eckman award from the American Automatic Control Council.

Professor Murray is an expert in nonlinear control of mechanical systems and real-time trajectory generation for flight control systems. His work in differential flatness has led to new algorithms for rapid generation of trajectories that satisfy the equations of motion for flight systems. Murray has substantial experience in developing controls experiments and implementing real-time control algorithms for trajectory generation and tracking. Murray will act as Principal Investigator for this research effort. A *curriculum vitae* for Prof. Murray is attached at the end of this proposal.

Alex Fax Alex Fax received his B.S.E. *magna cum laude* in Mechanical and Aerospace Engineering from Princeton University in 1993. He is currently completing a Ph.D. in Control and Dynamical Systems at Caltech under the direction of Prof. Murray. He is the recipient of an NSF Graduate Research Fellowship. His research interests include nonlinear control of mechanical systems, optimal control, orbital mechanics and control of microsatellite formations. Fax will participate in this research effort as a postdoctoral scholar.