Adaptive Coordinated Control in the Multi-Agent 3D Dynamic Battlefield: Adaptive Coordinated ControL of Intelligent Multi-Agent TEams (ACCLIMATE)

S. S. Sastry, PI [*], V. Kumar, co-PI [†], H. Choset, co-PI [‡]

---

[*] and investigators L. El Ghaoui, M. Jordan, J. Malik, S. Russell, and P. Varaiya, Department of Electrical Engineering and Computer Science, University of California, Berkeley

[†] and investigators R. Bajcsy, K. Danillidis, J. Ostrowski, G. Pappas, and C. J. Taylor, GRASP Laboratory, Departments of Computer and Information Science, Electrical Engineering, and Mechanical Engineering, University of Pennsylvania, Philadelphia

[‡] and investigators A. Rizzi and C. Thorpe, The Robotics Institute and Department of Mechanical Engineering, Carnegie Mellon University, Pittsburgh

1 Abstract (1 Page)

We propose the design and evaluation of the adaptive hierarchical control of mixed autonomous and human operated semi-autonomous teams that deliver high levels of mission reliability despite uncertainty arising from rapidly evolving environments and malicious interference from an intelligent adversary. The design of architectures combining both hierarchical and heterarchical elements, the analytical foundations of interacting hybrid systems, the design of controllers for such systems that are robust against uncertainty, the management of rich sensory information from networked sensors among distributed and mobile teams; and the incorporation of human intervention in a mixed-initiative system are all key areas of our work. Additionaly, the novelty of our approach is to explicitly take into account the need to adaptively replan missions to take into account environmental uncertainties and the deliberate malicious actions of a determined adversary. Our approach builds on the following research thrusts:

Thrust I: Architecture Design and Analysis for Dynamic, Adaptive Planning. The architectures that we design will organically incorporate human intervention at all levels of planning and execution. Architectural design begins with an overall hierarchy featuring flexible team formation, task specification, pre-mission evaluation, and changes in goals, team composition, and communications during mission execution. In the rapid adaptive, dynamic replanning, it is an absolute necessity to have modules which can be composed interoperably on the fly when warranted by the actions of the adversary. A main drawback of traditional approaches of *hybrid systems* has been their *extreme conservativeness of compositionality* when designing *intrinsically complex architectures.* We will address these issues through work on: (1) abstractions for perception and action, and (2) assume-guarantee reasoning and interface theory for compositionality.

Thrust II: Integration of Rich Multi-sensor Information into Virtual Environments for Incorporating Human Intervention in Mission Planning and Execution. A key difficulty in the use of unmanned aerial vehicles (UAVs), unmanned ground vehicles (UGVs) and unmanned underwater vehicles (UUVs) is the difficulty in acquiring the rich sensory data gathered by the vehicles contributing to sensory-overload where teams of up to four warfighters are required to control a single robotic asset (either a UAV or a flight control system). We propose to handle the fusion of rich multi-sensor information over an unreliable network by developing new classes of algorithms combining recent work in omni directional vision, the extraction of graphical models from video sequences, and the joint rendering of simulated (synthetic) environments with multi-sensor (real) data. The research directions are: (1) adaptive hierarchical networks for acquiring and providing information, (2) extraction of 3D models from distributed video and other sensors networks, and (3) environments for human intervention and decision making.

Thrust III: Handling Uncertainty and Adversarial Intent in Adaptive Planning. Two types of uncertainty pervade mission planning and execution: (1) probabilistic uncertainty having to do with environmental unknowns, such as weather, terrain data uncertainty, the probabilistic nature of failures of hardware or software, *information attack*, (2) adversarial uncertainty having to do with systematic attempts by an intelligent adversary (red-force) to defeat the mission. A key mathematical framework for the modeling of adversarial actions comes from the theory of games, and partially observable Markov decision processes and games. We will develop methods for; *learning of adversarial strategies.* We will develop teaming and game strategies to allow for defeating a *dynamic adversary*, that is one who changes his strategy, cost function, information patterns during the course of an engagement.

Finally the strategy for the integration of the research of the three University teams is through a set of two or more scenario-based challenge problems involving intelligent adversaries on the extensive testbeds at the three partner institutions.

Technical Approach

We are being called upon to protect our national security interests in progressively more complex and hostile environments. Even though we are the dominant superpower today, major threats arise from asymmetric threats such as terrorism, guerilla attack, and other unconventional methods of warfare. The technology challenge for dealing with these asymmetric and extremely rapidly adapting adversaries in the battlefield are many, and of course the battlefield itself is in a wide variety of terrains, and in urban environments and in some cases also the homeland. These in turn need us to develop adaptive, intelligent, multi-agent cooperative control technologies which are responsive to the needs of our project title, ACCLIMATE:

1. Control of the 3D Digital Battlefield: the need to use the 3rd dimension: aerial forces, robotic and mixed initiative with untethered communication channels between them.
2. Adaptive Coordinate Control of Multiple Agents: the ability to reconfigure teams of our own assets dynamically to take into account communication patterns, as well as the responses of the adversary during the course of an engagement.
3. Intelligent Coordination of Multiple Agents: the ability to determine the intent and objectives of an enemy during the course of an engagement to adaptively reconfigure our own strategies to defeat an intelligent adversary with time varying objectives.

However, it is important to note that concerns over the reliability, adaptability, robust and fault tolerance of complex systems such as those encountered on modern battlefields or in homeland defense increase dramatically with the nation's reliance on these systems. While functionality, speed, and availability dominate the headlines about new IT systems, the success of complex adaptive systems will ultimately depend on reliability, including safety, predictability, fault tolerance, and their ability to interact with hard real-time constraints, and their ability to reconfigure after failure. We believe that this program of research needs three important components:

Thrust I: Architecture Design for Adaptive, Dynamic Planning

The architectures that we design will organically incorporate human intervention at all levels of planning and execution. Architectural design begins with a hierarchical overall system architecture featuring flexible team formation, task specification, pre-mission evaluation, and changes in goals, team composition, and communications during mission execution. The first "weak link" in current design practice stems from the intrinsic "complexity" of modern battlefield systems. While "complexity" in science usually refers to the understanding of complex systems that occur in nature (such as weather prediction), we submit that a different kind of complexity arises in systems that are designed by humans, and that if properly understood, this complexity can be controlled. The complexity of modern battlefield systems arises from the large number of distributed but interacting components, the heterogeneous nature of components (digital computers as well as analog devices), the many levels of abstraction used in design (such as physical, system/protocol, and application layers), the many different aspects of system performance (such as functionality, timing, fault tolerance), and the many, often unpredictable ways in which the environment (sensors, users, failures, attackers) can influence the behavior of the system. Digital battlefields involve complex digital systems (computers and networks) interacting with the physical world through sensors and actuators. In the rapid adaptive, dynamic replanning, it is absolutely necessary to have **modules which can be composed interoperably** on the fly when warranted by the actions of the adversary. This in turn calls for the need to deviate from strictly hierarchical architectures to ones which allow for greater agility of decision making and replanning, either because of failure modes caused by the actions of the adversary or the environment or by losses to the friendly forces. A key drawback of traditional approaches of *hybrid systems* which has resulted in their exclusive use for hierarchical systems has been their *extreme conservativeness of compositionality* when designing *intrinsically complex architectures.* We will address two key issues to develop the interoperability and compositionality of planning modules: (1) abstractions for perception and action. We will develop systematic procedures for consistent and non conservative abstraction of models of reasoning and computation at multiple levels of a hierarchy. (2) assume-guarantee reasoning and interface theory for compositionality. Existing formal design methodologies are either optimistic when they are top down or pessimistic when they are bottom up. The centerpiece of our approach is the development of component interfaces which are much more expressive in that they specify not only what a component does, but what it expects the environment to do. Such "assume-guarantee" interfaces with *multiple aspects* such as safety, fault tolerance security will allow for both an optimistic view of those aspects which are well understood or identified during operation and a pessimistic view of those that are unknown or unpredictable.

Thrust II: Integration of Rich Multi-sensor Information into Virtual Environments for Incorporating Human Intervention in Mission Planning and Execution. A key difficulty in the use of robotic UAVs, UGVs and UUVs is the difficulty in acquiring the rich sensory data gathered by networked sensors contributing to sensory-overload where teams of up to four warfighters are required to control a single robotic asset (either a UAV or a FCS platform). We propose to handle the integration of rich multi-sensor information over an unreliable network by developing new classes of algorithms combining recent work in omni directional vision, the extraction of graphical models from video sequences, and the joint rendering of simulated (synthetic) environments with multi-sensor (real) data. The research directions are: (1) adaptive hierarchical networks for acquiring and providing information. In environments like battlefields, where radio spectrum is at a premium and stealth is a desirable feature for survivability of networked assets, it is important that data-rich sensors producing high bandwidth streaming be tasked with reporting on their sensory input at a level of granularity set by a query. We will develop tools at the interface of control, sensing and communication based on random graphs, nonlinear control and estimation theory, and dynamics. (2) extraction of 3D models from distributed video sensors. For the purposes of integrating real data with virtual data it is important to build 3D graphical models (with animation) of both objects as well as linked rigid bodies (humans) from streamed data coming from either surveillance or scout cameras. (3) environments for human intervention and decision making. The primary function of robotic UAVs, UGVs and UUVs is to operate autonomously on specific tasks until a requested intervention arrives. Assessments of the effectiveness of our methods will be performed using cognitive models of the decision maker as well as in experiments.

Thrust III: Handling Uncertainty and Adversarial Intent in Adaptive Planning. Two types of uncertainty pervade mission planning and execution: (1) Probabilistic uncertainty related with environmental unknowns, such as weather, terrain data uncertainty, the probabilistic nature of failures of hardware or software, *information attack*, measures of uncertainty about the effectiveness of assets will be assessed using dynamic probabilistic networks using graphical models. (2) Adversarial uncertainty related with systematic attempts by an intelligent adversary (red-force) to defeat your mission. A key mathematical framework for the modeling of adversarial actions comes from the theory of games, and partially observable Markov decision processes. Two separate sub-themes will be addressed here: *computation of strategies for Nash, Stackelberg and other kinds of dynamic games with hybrid models of computation.* Here we will hit the complexity barrier for NP-hard or even undecidable problems very soon, hence we will need to develop good sub-optimal greedy strategies and evaluate their effectiveness. *learning of adversarial strategies.* In our opinion a big drawback in current techniques of teaming and game strategies is their inability to allow for a *dynamic adversary*, who changes his strategy, cost function, information patterns during the course of an engagement by using a strategy such as *feinting.* An engagement can be preceded by a *learning phase* when a number of scout UAVs/UGVs are sent out to probe and learn about adversary reactions for use in an engagement, using new graphical learning techniques.

2 Research Program

In order to motivate the proposed research thrusts we will develop two different scenarios. These scenarios will illustrate our vision for research over the next five years and help define experiments, demonstrations and milestones over the course of the project. In each case, it is important to note that we will be planning in the face of an unknown environment and a hostile and intelligent adversary.

Reconnaissance and intelligence: Robotic ranger force. A platoon of warfighters is charged with the responsibility of scouting a 100 mile^2 area of moderately undulating terrain in two hours. They deploy a team of five UAVs and ten UGVs. The UAVs take off in formation, and break formation over the hostile terrain for exploration. They cover the specified area while avoiding previously identified areas of threat, and new pop-up threats. They identify target areas that require closer investigation and automatically deploy UAVs and UGVs to further explore these target areas. The team reconfigures so that one UAV flies low while others provide coverage. The remotely located warfighter is able to immerse himself in a virtual environment simulating the battlefield and monitor the operation. This environment is reconstructed from the information obtained by 3D computer vision systems on-board the UAVs. When necessary, the UAVs engage the UGVs to obtain finer ground level information while the UAVs provide protection against threats. During ground operations, the sensors on-board the UGVs provide additional information for the reconstruction and the ability to locate ground-level pop-up targets. The network of UAVs hovering above and UGVs patrolling the battlefield below provide the warfighter with real-time dynamic updates and allow the warfighter to task other robotic vehicles on tactical missions.

Mixed initiative engagement. In this scenario, the warfighter deploys UGVs and UAVs to engage a platoon of enemies. The terrain here is to be chosen to be urban. The warfighter uses a 3D immersive environment to command, monitor and visualize the 3D battlespace. As in the reconnaissance scenario, the sensors from the unmanned vehicles are used to dynamically update the 3D environment. Specially equipped UGVs are also able to launch micro-UAVs, e.g., quad-rotor helicopters, for local scouting missions. After collecting data on enemy locations and broadcasting it to team members, the micro-UAVs return to their home UGV to continue the missions. In engaging the enemy, some vehicles assume the role of decoys, while others respond to the enemy's actions. The remote user is able to task the team to intercept enemy vehicles and take out multiple targets[1]. The warfighter introduces new pieces of information, re-tasks the team, and changes the mode of operation from clandestine to operation with communication. However, he never has to specifically task individual robots.

The next three sections describe the detailed research objectives. While there is no formal statement of milestones, each section begins with a discussion of the key areas to be fully addressed in the five-year cycle of the MURI program. We expect this project to be evaluated with respect to progress on these objectives and the scenario demonstrations at each of the partner Universities.

2.1 Thrust I: Hierarchical Architecture Design and Analysis for Dynamic, Adaptive Planning.

Faculty Lead: Howie Choset, Vijay Kumar, James Ostrowski, George Pappas, Alfred Rizzi, Shankar Sastry, Pravin Varaiya.

The rapid progress in embedded hardware and software makes plausible ever more ambitious distributed, multi-layer, multi-objective, adaptive control systems. However, adequate design methodologies and design support lag far behind. Consequently, today most of the cost in system development is spent on ad-hoc, prohibitively expensive systems integration and validation techniques that rely almost exclusively on testing more or less complete versions of the entire system. Our project addresses this bottleneck by focusing on the *systematic design* of hierarchical architectures and the design of controllers for individual agents at all levels of the hierarchy. We say "hierarchical" though we really mean an architecture which has a hierarchical backbone (nominal architecture), with the ability to add information patterns, data provisioning and other links which cross layers of the hierarchy for dynamic replanning in alert or emergency modes. We will discuss how to deviate from hierarchical architectures for ease of decision making, robustness to attack and low latency adaptive replanning. These are the metrics that will be used, with tools to evaluate architectures in the normal, replanning or emergency mode of operation. We also propose to continue our efforts in building a solid analytical foundation based on *hybrid systems*, and a practical set of *software design tools* that support the construction, integration, safety and performance analysis, on-line adaptation and off-line functional evolution of multi-agent hierarchical control systems—for a swarm of our rotorcraft UAVs, UGVs and other real-time systems for the real-world validation of our techniques.

The following research topics will be addressed:

1. Architecture for dynamic, adaptive, multi-agent systems:
    - Centralized vs. Decentralized: How to determine the degree of centralization/decentralization
    - High-confidence architectures: Latency, fault tolerance, ease of verification of design and compositionality
    - Adaptability & flexibility for replanning: Transition mechanisms from normal modes to degraded modes of operation.
2. Tools for the analysis of hybrid systems:
    - Abstractions of lower layers for higher level actions: How to determine what should be abstracted. How to adapt this abstraction as circumstances change?
    - Analysis & controller design: Continuous and discrete variables. Deterministic, probabilistic
    - Compositionality: Assume-guarantee reasoning for reducing the pessimism of compositional tools. Probabilistic and deterministic.

---

[1] In our laboratory experiments, a take-out will be defined as a planned, successful collision with an enemy ground based vehicle, or the dropping of an object over the vehicle.

Our efforts have been centered, for several years, and with considerable success, around the mathematical foundation of hybrid systems. "Hybrid" refers to the tight coupling and interaction of discrete with continuous phenomena. The hybrid characteristic of embedded control systems arises from several sources. First, the high-level, abstract protocol layers of hierarchical control designs are discrete so as to make it easier to manage system complexity and to accommodate linguistic and qualitative information. The low-level, concrete physics-based control laws are naturally continuous. Second, while individual feedback control scenarios are naturally modeled as interconnections of modules characterized by their continuous input/output behavior, *multi-modal* control naturally suggests a state-machine-based view, with states representing discrete control modes. Software-based control systems are an integrated mixture of both types. Third, *every* digital hardware/software implementation of a controller is ultimately a discrete approximation that interacts through sensors and actuators with a continuous physical environment.

Abstractions of Hybrid Systems for Architecture Design and Validation

The overall command and control framework has a hierarchy of abstractions: high-level planning, team strategization and trajectory planning. A key to this hierarchical decomposition is the ability for higher levels to have *consistent abstractions* of the lower layers. Our previous work [12] has been on the development of hierarchically consistent abstractions of continuous control systems and the composition of a special class of controllers with guarantees [29, 30]. We propose to extend this research in two key directions:

1. **Hierarchical abstractions of hybrid systems.** Consistent abstractions of hybrid systems are useful since for decision making at the high levels we need discrete state approximations of continuous state trajectory plans. However, the granularity of the discrete abstraction should not be so coarse as to destroy features of the plan to be implemented at the lower levels such as asking for conflicting or unflyable trajectories for UAVs. Thus, while there is an infinity of abstractions of a continuous system, the abstraction driven by specific requirements on the functionality needs to be developed.
2. **Assume-guarantee reasoning for abstractions.** Existing formal design methodologies are either optimistic or pessimistic. The optimistic approach advocates strictly top-down, stepwise refinement, with the design team in full control of the complete design. It does not allow for some parts or aspects of the design to be unknown, unpredictable, or evolving. The pessimistic approach advocates a strictly bottom-up, component-based design, where some components may be preexisting, and an environment that behaves in an arbitrary, possibly adversarial way must be considered in the design of each component. It does not allow for individual components to be designed under assumptions about how the environment or the other components of the system behave. We submit that neither the fully optimistic ("all parts work together by design") nor the fully pessimistic ("each part is on its own") paradigm is realistic for complex, heterogeneous designs. We propose to develop the foundations of a formal approach which we refer to as *an interface theory* that permits a synthesis of both paradigms.

Central to our approach are component interfaces that are much more expressive than traditional interfaces used in software or hardware designs. First, the interfaces we envision not only specify, on some abstract level, what a component does, but also what the component expects the environment to do. Such "assume-guarantee interfaces" allow a component designer to adopt an optimistic view about some aspects of the other components, as if those aspects were under the designer's control, and at the same time adopt a pessimistic view about unknwon or unpredictable aspects of the other components and the environment. Second, the interfaces we envision specify not only aspects that are traditionally specified in interfaces, such as the number and types of the arguments of a procedure, but also a variety of different aspects, for example, that the call of the specified procedure must always be preceded by the call of another (e.g., initialization) procedure. There has been considerable work on functional interface languages, little on timing and security, and virtually none on other system aspects such as resource management, performance, and reliability. This lack of multi-aspect interface formalisms has forced designers to address timing, security, performance, and reliability issues at all levels of the implementation in order to attain the desired properties for the overall system. We propose to develop a composition theory for multi-aspect interfaces, which expose resource properties, such as real-time assumptions and guarantees, and algorithms and tools for checking the consistency and compatibility of the multi-aspect interfaces.

The technical approach of the proposed research agenda will focus on formalizing notions of abstraction and composition for control systems. This will be considered independent of whether the systems are inherently continuous, discrete, or hybrid systems. In order to extract more optimistic and higher performance hierarchies, we shall develop a theory for *context or environment-dependent abstractions* of control systems. Context-dependent abstractions of continuous and hybrid systems will critically depend previous characterizations of property, preserving abstractions for linear systems [12, 36], nonlinear systems [14], and hybrid systems [32, 37]. In particular, we shall develop a particular form of assume-guarantee theory for compositional abstraction, where the one component gets refined or abstracted by another, under specific assumptions about the environment. The abstraction of the system will continuously adapt to the environmental changes, allowing the dynamic reconfiguration and adaptation of the control hierarchy, and thus increasing the overall performance.

Control of Hybrid Systems

Approaches to hybrid system control design and their limitations can be summarized as follows:

1. Algorithmic "model checking" approaches of theoretical computer science, which have led to great successes in hardware verification, grow in computational complexity very rapidly, and problems are infeasible or even undecidable for all but the simplest classes of hybrid systems (referred to as linear hybrid automata).
2. Deductive approaches involving "theorem proving" techniques are ad-hoc and require considerable inspiration for deriving invariants associated with these methods.
3. Synthesis approaches, which we have pioneered, derive "pre-verified" hybrid systems, face similar obstacles as model checking and are, so far, restricted to the worst-case safe-case.

To alleviated these difficulties, we will develop constructive methods and tools for the (semi-)automatic synthesis of hybrid controllers. In the continuous case, optimal control laws may be derived as solutions of the Hamilton-Jacobi-Bellman equation, while discrete controllers can be synthesized by solving games on finite automata. Both methods can be seen as special cases of a generic game-theoretic approach to safety control [6]. This insight can likely be extended to more general hybrid control objectives. In order to enable the simulation and implementation of the resulting controller, we will put special emphasis on constructivity.

We have shown the importance of a unified approach to the treatment of methods from control, and computer science in [8]. We were able to show that synthesis problems for hybrid conditions with so-called *safety specifications* can be solved under tremendously generous conditions for nonlinear hybrid automata by using Hamilton-Jacobi equations, and an algorithmic construction for the *maximal safe sets* and *least restrictive control laws.* The Hamilton-Jacobi equation (and its discrete counterpart) is a partial differential equation of the form $\frac{\partial J}{\partial t}(x,t) = -H^*(x, \frac{\partial J}{\partial x}(x,t))$ where H^* is the so-called Hamiltonian determined from the appropriate target sets and the game between the controllable and uncontrollable actions (or between the software and the data). The following open problems need to be addressed to convert this general approach into an *automated synthesis* procedure:

Numerical solutions of Hamilton-Jacobi equations. Since the synthesis procedure described above is general, its utility depends on the availability of efficient numerical tools to compute the solutions. *Approximation techniques* are a first step, where we will use ellipsoid, projection, fast wavefront methods, etc. In addition, we will use quantitative computing methods when the equation has *shocks*, corresponding to changes in the gaming strategy. This draws from the modern mathematical theory of *wavefront propagation* and *viscosity solutions* of Hamilton-Jacobi equations. Some algorithms have been developed in [4, 8, 7]. The ellipsoidal methods for the approximation of safe sets from the inside and outside has been presented in [15]. Methods using semi-definite programming and linear hyperplane methods have also been proposed by us in [13].

Hierarchical solutions of synthesis procedures. For systems with high dimensional state spaces or for many agents, a hierarchical application of the approach is necessary to facilitate the computations associated with safety and least restrictive control computations. A promising avenue is the work begun in [9]. .....write one more.....

Liveness and other acceptance conditions. Liveness, fairness and other acceptance conditions such as $\Box\Diamond$. $\Diamond\Box, \ldots$ are important to model in the framework described for safety games, not just for mathematical completeness, but also to establish the fairness of certain synthesized solutions for multi-agent problems, because they are requirements for the eventual completion of a mission (always eventually complete and eventually always complete). We believe that this is possible not by nested Hamilton-Jacobi solutions, but by a single new Hamiltonian.

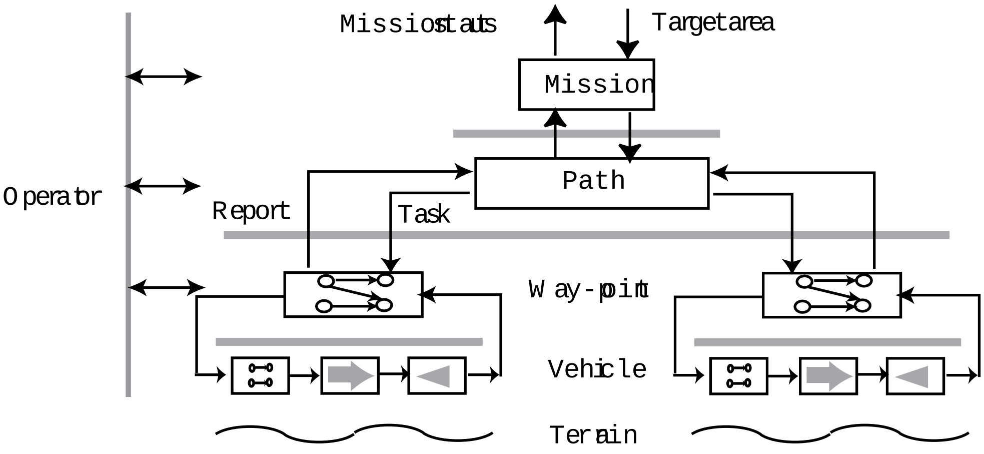

Figure 1: Illustration of surveillance mission organized in a four-level hierarchy.

One important factor in the synthesis procedure is decidability boundaries. Recently, we extended the decidability frontier to capture classes of hybrid automata with linear dynamics in each discrete location. In particular, the very recent notion of *O-minimality* from model theory is used to define a class of hybrid systems, called *O-minimal hybrid systems.* In [10] it is shown that all O-minimal hybrid systems admit finite bisimulations. The main computational tool for symbolic set manipulation in this context is *quantifier elimination* for decidable theories. This immediately leads to new decidability results, which will allow to model significant disturbances as well as provide us with a framework for symbolic controller synthesis. The implementation of the above methodology by a computational tool whose kernel is a quantifier elimination engine (REDLOG) will be benchmarked and integrated with existing tools, see for example [11].

Controller Libraries

One of our primary goals is to build libraries of controller models for teams to explicitly articulate the capabilities of each team agent and that of the team as a whole. The libraries will include (see Figure 1) the family of networking protocols, terrain models, vehicle/sensor/actuator models, real-time controllers for vehicles, waypoint controllers, path planners, and mission command. The libraries will also include entire mission scenarios and mission fragments obtained by combining models of terrain, vehicles, functional agents (including those representing human intervention), and communications. Each mission fragment representing a high-level capability, such as surveillance, may be combined to create a multi-capability mission. Each model in the library will be a SHIFT component, which corresponds to a class in object-oriented programming languages. Each model will have a well-defined interface to facilitate combining them into entire missions or mission fragments. Figure 2 illustrates these models. The box on the top left summarizes SHIFT syntax for a component (class). A component description includes a data model (inputs, outputs, and internal state), and hybrid dynamics (discrete state or mode changes and differential equations. SHIFT permits creation and destruction of objects as the simulation proceeds. The Lisp-based extension of SHIFT, Lambda-SHIFT (see www.gigascale.org/shift), permits changing class definitions as the simulation proceeds.

Hierarchy Semantics

Another major goal of our research is to provide connections between the multiple world of declarative semantics and the single world of imperative semantics. We will pursue the following directions.

Ideal compilation. Instead of syntactic translation followed by semantic interpretation at lower level as in the one-world semantics, a higher-level expression is compiled into an idealized lower-level expression and then interpreted.

Invariants. These may be used to show that higher-level truth-claims now become conditional lower-level truth-claims: higher-level truth-claims thereby become necessary conditions. By contrast, higher-level claims are sufficient conditions in one-world semantics (cf. abstraction in verification). An example of an invariant associated with a 'lift' mode might be: the vehicle cannot go above 3000 feet.

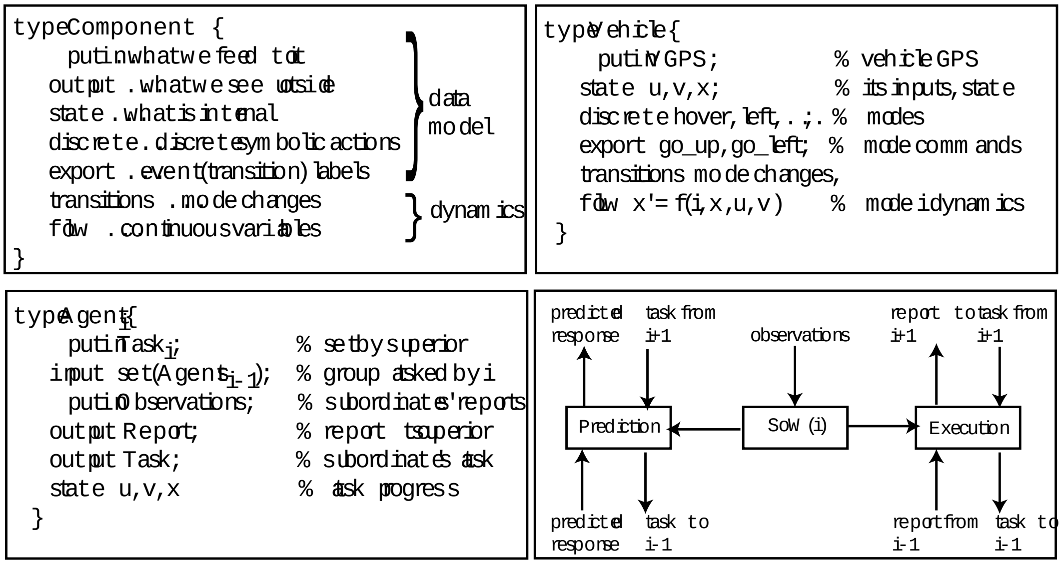

```
type Component {
    input . what we feed to it
   output . what we see outside
   state . what is internal
   discrete . discrete symbolic actions
   export . event (transition) labels
   transitions . mode changes
   flow . continuous variables
}
```

```
type Vehicle {
    input GPS;               % vehicle GPS
   state u,v,x;              % its inputs, state
   discrete hover, left, ...; % modes
   export go_up, go_left;   % mode commands
   transitions mode changes,
   flow x'= f(i,x,u,v)      % mode i dynamics
 }
```

```
type Agent_i {
    input Task_i;             % set by superior
   input set(Agents_i-1);    % group tasked by i
    input Observations;      % subordinates' reports
   output Report;            % report to superior
   output Task;              % subordinates' task
   state u,v,x               % task progress
 }
```

Figure 2: Top left summarizes SHIFT syntax; top right is a skeleton of a vehicle component; bottom left is a generic agent: there are sensor inputs that give the control inputs, the continuous/discrete states representing modes, mode transitions, and continuous dynamics in each mode. The 'exported' events allow an external agent (such as the way-point controller in Figure 1) to select the modes. The box on the bottom left represents a generic agent at level i in the hierarchy. Its inputs consist of tasks from a superior agent at level i+1, and reports from those subordinate agents at level i-1. Its state summarizes the progress of its own task, and its outputs consist of tasks assigned to its subordinates and reports to its superior.

Modal decomposition. Suppose the one-world interpretation leads to falsification of higher-level claims that are true in multi-world semantics. This can happen because lower level allows faults that are not accounted for in higher-level descriptions. Then, multi-world semantics must be "split" into multiple-frameworks, each dealing with identified faults. for example, if the waypoint controller had assumed that there was no obstacle in tasking a vehicle to go to a particular waypoint, but the obstacle was noted in a simulation, then a new mode should be created that incorporates the actions to be taken if an obstacle is encountered.

Exceptions. All these directions are subsumed under a theory of exceptions that we intend to elaborate. Exception means a situation that arises during execution which the agent (program) is not designed to handle. In the case of hierarchical control, the system should be designed to report the exception to a superior agent or a human operator. Modal decomposition, invariants, etc. would be considered within the theory as means to detect the cause of exceptions or methods to handle them.

Team and Task Allocation

A third major goal is the allocation of resources to tasks. In order for an agent to carry out a certain task, the agent requires certain inputs from authority on task assignments. The importance of the task assignments increases when the mission requires a team of agents. The decision to alllocate resources in order to carry out a particular task, or to modify the task to operate within the limits of the available resources, can itself be formulated as a higher-level problem of resource allocation. We intend to pursue this formulation and to develop suitable algorithms to resolve it. Logic programming provides an approach where a high-level goal is to be met by combining lower-level resources. Numerical algorithms provide another approach that explicitly formulates trade-offs between different costs and corresponding outcomes.

2.2 Thrust II: Integration of Rich Multi-sensor Information into Virtual Environments for Incorporating Human Intervention in Mission Planning and Execution.

Faculty Lead: Ruzena Bajcsy, Kostas Danillidis, Vijay Kumar, Camillo Taylor, Laurent El Ghaoui, Jitendra Malik, Stuart Russell, Charles Thorpe.

With recent advances in information technology, sensors, energy storage devices, networking infrastructure, and control technology, and the falling price-to-performance ratio of the technology associated with personal computing, we have the necessary technological infrastructure to develop and field truly intelligent systems that will allow the military to fight wars with minimal loss of life. Now, for the first time, the main obstacle is our lack of understanding of the fundamental issues underlying cooperation, human-robot interaction, the control, command and coordination of a large number of human and robotic agents, and the theoretical and computational tools to support this activity. Thrust II brings together a team of researchers to address the following three critical areas:

1. Adaptive hierarchical networks for acquiring and providing information
   - Networked sensors, communications and traffic flow based on stealth considerations
   - Bandwidth utilization and sharing
2. Extraction of 3D models from distributed sensors
   - 3D Models from video data for humans and mobile vehicles
   - Integration of real and virtual environments
3. Environments for human intervention and decision making
   - Situational awareness for humans when called upon to intercede
   - Triaging of data for decision making
   - Display of uncertain data

Adaptive Hierarchical Networks

Figure 3 a team of agents, characterized by spatial and communication links between agents. Spatial links represent physical interaction between agents and are enabled by the sensing and actuation capabilities of individual agents. For example, sensors allow agents to localize themselves, estimate the relative positions and orientations of neighboring agents, obtain information about their environments, and assess local physical constraints. Actuators allow agents to change their relative positions or work together to accomplish a physical task. Communication links between agents are used to exchange information. We use these links to associate network structures to groups of agents. These networks must be established, hierarchically organized, and continuously adapted to changes in the environment and commands from human warfighters (see Figure 4).

Modeling, analysis, and adaptation of ad-hoc networks Specifically, the team of agents is built on three different networks: a *physical network* that captures the physical constraints on the dynamics, control and sensing of each agent; a *communication network* that describes the information flow between agents; and a *computational network* that describes the computational resources available to each agent. We model each network by a graph with n nodes, one node for each agent. R is a finite set of nodes, $R_1, R_2, \ldots, R_n$. The physical network is a directed graph, $G_p = (R, E_p)$, where E_p consists of edges each of which represent the flow of sensory information (relative state). $G_c = (R, E_c)$ is an undirected graph. The edge set consists of pairs of agents that can communicate with each other. (We assume omnidirectional transmitters and receivers on each agent.) G_p and G_c are determined by the hardware limitations, the physical distribution of the agents, and the characteristics of the environment. The key goal is to design a computational network, which is modeled by a directed graph $H = (R, E)$. E consists of edges in the edge set $E_p \times E_c$. The design of the graph H is based on the task. For example, if we want a team of robots to transport a payload through a field of obstacles to a destination, a control-centric point of view leads to the assignment of edges with the goal of maximizing performance metrics such as stability, robustness, and time-optimality. On the other hand, if the task is to explore the environment and build a 3D reconstruction of the environment, H must be designed with different perception-centric performance measures. In tasks involving target detection, pursuit and evasion of threats, the performance measures are more intertwined.

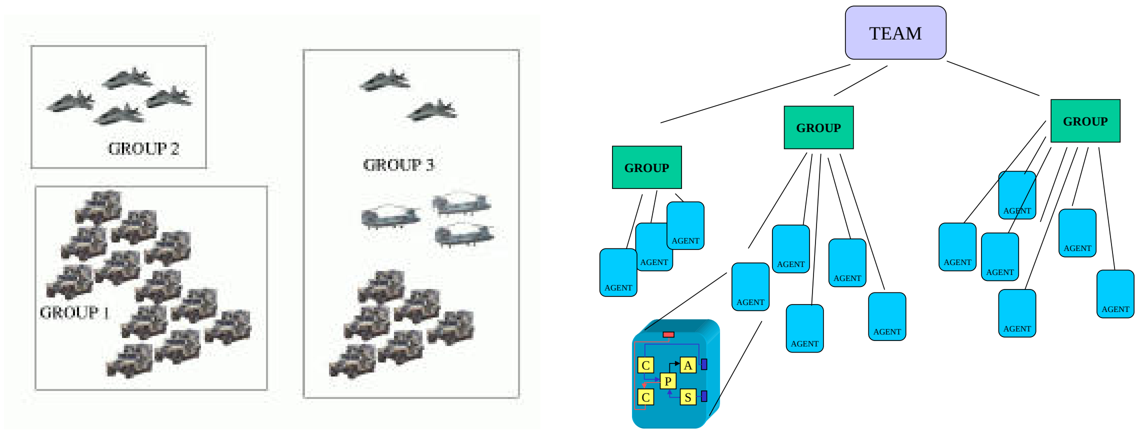

Figure 3: We envision cooperation between heterogeneous agents (left). Each agent is a collection of processors (P), sensors (S), actuators (A), and communication channels (C). The reconfiguration of the team into groups and the dynamic hierarchical organization (right) is a key to adaptation.

Communication for control and sensing Wireless network involves the use of untethered devices with a very small fixed infrastructure, and shared used of the spectrum are key to the battlefields of the future. One of the most fundamental design challenges is the unpredictable and highly constrained wireless communication channel. In particular, the radio spectrum is a scarce resource which directly limits the throughput of the wireless channel. Moreover, as a signal propagates through a wireless channel, it experiences random power fluctuations over time due to changing reflections and attenuation. These power fluctuations cause intermittent connectivity, which affects link transmissions as well as routing protocols. In *ad-hoc* topologies where the network does not have a fixed structurthe nodes self-configure into a network, and all control and routing functions are shared among the terminals. Since there is no centralized controller, packets sent by different terminals may collide, in which case the packets must be retransmitted, thereby incurring delay and sometimes loss. Moreover, the finite resources of the network, including power, bandwidth, and rate, must be allocated in a decentralized and efficient manner to meet the requirements of the underlying application.

Ad hoc networks for control Our description of the physical network will be based on two elements: the position and orientation of the team in space denoted by g; and the shape of the formation denoted by r. g is an element of the motion group (usually the special Euclidean group such as $SE(2)$ or $SE(3)$), while $r \in \Re^m$ describes the distribution of the agents in g.

The optimal design of the computational network depends on the control policies used by each agent. Consider the case in which E_c is empty or the agents must operate in a clandestine mission (see Figure 4). In [18] we show that if linear feedback of relative state information is used for fully-actuated holonomic vehicles with linear dynamics, and this can be extended to fully actuated systems with nonlinear dynamics. H is acyclic and connected is necessary and sufficient for asymptotic stability of the formation. When edges in E_c are included in H, feedforward information is possible. For the case with linear dynamics, we can design H to guarantee global, exponential stability [18, 23]. Performance measures that incorporate notions of input-state stability allow us to characterize the optimality of a graph H [26]. The lengths of paths of information flow are now known to be important design parameters. Similarly, extensions to nonholonomic, underactuated systems show that performance of the system degrades as the length of the longest path in H increases [24].

It is very difficult to guarantee quality of service (QoS) in the face of the random varying end-to-end network performance. There have been significant research efforts to address variability at different levels of the network protocol stack. We will study the tradeoffs and specific characteristics of the network data rate, delay, and loss, and the impact of these parameters on stability and performance degradation. Further, we will pursue algorithms that will allow the network to adapt and vary these tradeoffs according to controller requirements and the state of the underlying communication channel.

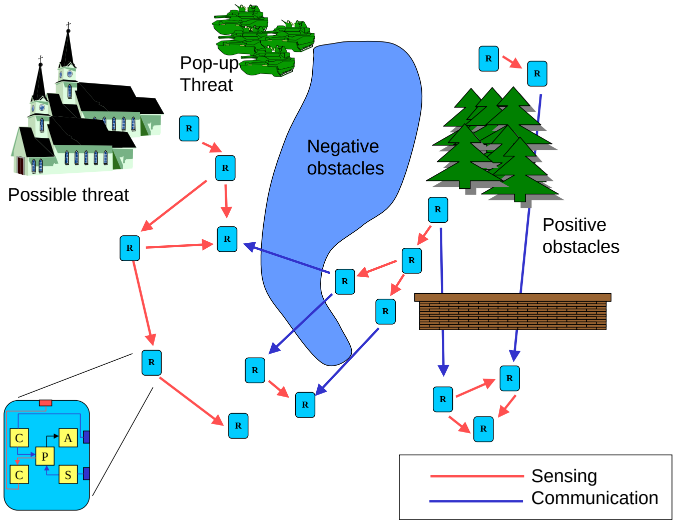

Figure 4: An *Ad Hoc* network of agents tasked with reconnaisance and surveillance. Communication links are disabled in areas where stealth is necessary. The performance of the team in the navigation and exploration subtasks directly depends on the acyclic graph.

Ad hoc networks for sensing and perception The paradigm of *active perception* [17] involves data fusion from both homogeneous and heterogeneous dynamic sensors to improve performance, including ranging accuracy and completeness of data. Our framework can be viewed as the distributed version of active perception, where surveillance of dynamic environments and the recovery of their 3D structured is made possible or enhanced by the use of networked, mobile video sensors. In our preliminary work in tracking and estimating moving features, a quality function that captures how expected errors in the estimates of the feature depend on the configuration of the team and the shape of the formation guides the selection of the graph H [27]. We propose to develop the theoretical foundation and the algorithms for establishing and maintaining sensor networks.

3D Models from Distributed Sensors

In our scenario [28], a commander can deploy a team of robotic vehicles on a specific floor of a building (Figure 5). Robots with omnidirectional cameras were used to provide immersive visualization of a remote area. The information gathered by the deployed sensor network can be presented as a panoramic view from one of the agents as well as a view of the entire floor. Then the commander can be immersed at every possible position into this floor and view it from the correct 3D perspective (Figure 6). Omnidirectional vision sensors provide a robust recovery of the relative 3D pose between the sensors and thus a correct registration for rendering. Our recent results on the geometry of multiple views [25] show a powerful projective geometric feature representation in omnidirectional imagery that enables a more robust 3D registration with less views than with conventional perspective cameras, using the algorithms that blend sensory information with reconstructed models into a 3D virtual model. As illustrated in the figures, the immersive visualization system can provide several levels of representation and abstraction in the interaction with an operator. The system recovers the positions and orientations of all the cameras from correspondences of linear features using distributed sensor fusion. Then, a view-independent 3D layout of the environment is recovered from multiple cameras

Figure 5: A panoramic view (right) obtained from the omnidirectional camera mounted on the robot (left).

Figure 6: An overhead view of the reconstructed floor (left) and immersed 3D views of the environment (middle and right).

with known relative pose.

In this project, we aim to construct the theoretical foundations and the software tools for deploying networked cameras and synthesizing 3D virtual environments, along the two thrusts outlined below.

1. Anytime, anyplace immersion for warfighters. We will develop a virtual environment that will enable one or more remotely located human commanders to be tele-immersed in the virtual battlefield for a mixed-initiative engagement (see Section 2). This *anytime, anyplace immersion* will require the close coupling between the immersion algorithms to distributed implementation and the computational network architecture, and a systematic approach to model abstractions of spatio-temporal data at different levels of resolution in the control hierarchy. For example, it will be possible to provide many human agents with a global view of the operations, while allowing a single human agent to get a detailed perspective of the data and models for a single robotic agent. The anyplace immersion of a human operator requires the real-time rendering capability in any visual direction and from any viewpoint. Since network constraints will not allow a continuous 3D-information flow, we will pursue a view-independent scene acquisition that is independent of the display and rendering rates used by different viewers at different sites. The 3D acquisition process will be modeled as a dynamical system pursuing a trajectory in the real 3D-space with different scales of detail. The dynamics will explicitly incorporate the sensor characteristics, the measurement process, the estimation rate, and the possible latencies, in order to enable a scalable, tractable approach to immersion.

2. Task-oriented sensor network characterization. We will address the characterization of sensor networks and algorithms for configuring such networks. Video signals require the square of any 1D-signal throughput and a delay affordable in human/computer interaction, and challenging constraints are posed on sensor placement and tuning. Free-space optical (FSO) communication is likely to be applicable in situations with high visual communication demands, such as in urban environments. It is necessary for the efficient use of FSO that the *ad hoc* network maintain physical constraints, requiring *control for communication.* In addition, there is the issue of controlling the position and viewing angle to recover 3D geometry. These questions are grounded in projective geometry and estimation theory, and we will specifically address (1) the recovery of a partially unknown static environment and the appropriate visualization and rendering; and (2) the continuous monitoring of spatiotemporal changes, including monitoring of agents in a team, identifying and tracking agents in peer teams, and the detection and tracking of unmodeled adversarial agents.

Environments for Human Intervention and Decision Making

The special case where one human controls a single machine is the commonly studied human-machine interaction problem. We are interested in mixed-initiative engagement where many humans must control many machines and human warfighters and machines must deconflict adaptively.

Real-time immersion Our main goal is to develop a system for tactical planning in mixed initiative engagement with multiple cooperating unmanned vehicles, allowing multiple generals to be simultaneously immersed in the battlefield. We propose to accomplish a real-time immersive visualization using the concept of tele-immersion, which we have been studying for collaboration between remote offices as illustrated in Figure 7. The tele-immersion [31]

Figure 7: A user in Chapel Hill (left) wearing polarized glasses and an optical tracker communicates with two remote users from Philadelphia and Armonk. The stereoscopically displayed remote 3D-scenes are composed from incoming streams of textured 3D data. A user in Chapel Hill (right) wearing the same devices interacts in a mini office model which is virtually shared with the remote user in Armonk (Courtesy of UNC-UPenn)

enables remotely-located users to collaborate in a shared space that mixes the local with the remote realities. The user is immersed in a rendered 3D-world that is acquired with computer vision at a remote site and transmitted over the network. In this project, our work will be directed to develop (1) the similar tele-immersion technology for outdoor, unstructured 3D environments; and (2) visualization system for reconstruction, based on the algorithms on active and distributed sensing. Algorithms for target localization and local reconstruction will run in a distributed fashion.

Sensor abstractions for hierarchical control systems The collection and fusion of sensory data across multiple agents makes it necessary to develop abstractions at the team level, at the group level, at the agent level, and within the agent (Figure 3). It is necessary for one or more commanders to be able to command the team as a whole, view groups as supervisors, view robotic agents as peers, or "enter" agents for diagnosis, debugging or reprogramming.

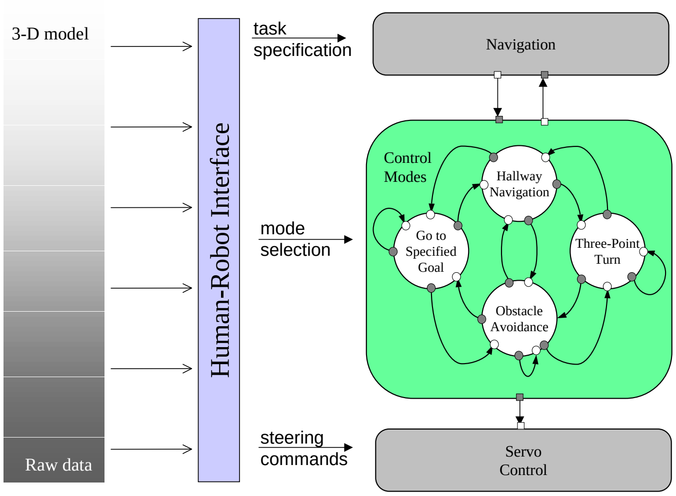

Figure 8: The human-agent interaction must allow the human to go down the hierarchy in a top-down fashion and interact at the lowest level, if necessary. The three-level hierarchy in the figure illustrates the control of an indoor, unmanned system by a remote user. The user can immerse herself in a 3D reconstruction at the top navigation level, and also view real-time imagery at the lower-most level for directly piloting the vehicle.

It is difficult, and sometimes impossible, to decouple the choice, design, and processing of sensors from the given control task. The mission or control objectives (at any level of the hierarchy) generally dictate the necessary sensory inputs, although the converse, where restrictions of sensors cause changes to the control schemes, can also occur. With this in mind, we seek to develop a sensing architecture that is compatible with the control architecture in the sense that sensory objectives at each level of the hierarchy need to be compatible with control objectives at similar levels. For example, at low levels of control, raw acceleration or rate information may be important for maintaining stability of fast processes. At slightly higher levels, visual servoing techniques could perform control using minimally processed visual data, while at higher levels, multi-sensor information might be fused to build terrain maps for mission planning and navigation. Furthermore, there is a significant advantage to be gained when sensory information is combined between agents, e.g., large-baseline triangulation of targets. By appropriately abstracting the shared data, one can provide formal, "object-oriented" interfaces between agents and also reduce the bandwidth necessary to share information.

One important research question is the extent to which the control hierarchy (addressed in Thrust I) helps to define and restrict the sensing hierarchy. In particular, to what extent does the complexity reduction that comes from abstracting the dynamics also arise in the sensor abstractions. It is clear that sensory data can and should be abstracted into higher-level representations, but to what extent this can be done in a hierarchical fashion is unclear. Standard tools, such as (extended Kalman) filtering, naturally lend themselves to use here, particularly in reducing and assimilating low-level sensory data. However, generating higher-level abstractions, such as target recognition, map-building, and fusing of information from distributed, heterogeneous, and perhaps dynamically configured sources, and representations that allow humans to interact with them will require new tools and algorithms.

Architecture for human intervention and decision making. Complex software-enabled control systems, and particularly unmanned systems, are generally constructed in a "bottom-up" manner, whereby one starts with the full sensing and actuation complexity and develops approaches to reduce the complexity. Such systems are generally constructed by layering increasingly sophisticated modes on top of each other. Although the high-level goal guides this process, the focus is guided by those lower-level tasks that can be most easily performed by an autonomous control system.

Human users, on the other hand, often take a "top-down" approach, starting with the goal, say searching for a target using multiple agents, and refining it into increasingly simpler tasks, for example, localized search regions, until finally the lowest level of implementation is reached, e.g., controlling actuator for mobile robots. In this case, the refinement of the task structure would appear to be motivated by avoiding tasks that require complex controls at the lowest level, or processing large amounts of low-level data.

Thus, the goal is to seek not only methods for generating hierarchical control and sensing architectures, but also those that best strike a balance between the advantages of doing low-level autonomous controls and high-level human intervention.

Explicit design of task/mode constraints. It is necessary to place appropriate constraints on the types of high-level modes that are made available to the human user. Doing so helps to restrict the types and complexity of controllers to be developed, while also relying on the human input at a level for which he is well-suited– choosing when to switch between well-organized and focused control modes. This also allows formal performance guarantees using, for example, the assume-guarantee paradigm for compositionality described in Thrust I.

Since a human agent can usually focus his attention on only one agent at a time, we propose a query-based system where robot agents can ask humans for assistance. We arbitrate among the queries so that the human is always presented with the one which is most urgent in terms of safety and timeliness. This effectively reduces the level of attention and control the operator must dedicate to each robot, and thus, makes the human intervention more manageable from the human's stand point and more predictable from the stand-point of the system.

2.3 Thrust III: Handling Uncertainty and Adversarial Intent in Adaptive Planning

Faculty Lead: Michael Jordan, Stuart Russell, Anant Sahai, Shankar Sastry.

Two types of uncertainty pervade mission planning and execution: (1) uncertainty related with environmental unknowns, such as weather, terrain data uncertainty, the probabilistic nature of failures of hardware or software, and *information attack*. Measures of uncertainty about the effectiveness of assets will be assessed using dynamic probabilistic networks and graphical models. (2) uncertainty related with systematic attempts by an intelligent adversary (red-force) to defeat a mission. A key mathematical framework for the modeling of adversarial actions comes from the theory of games, and partially observable Markov decision processes and games.

Topics to be Addressed in this thrust include:

1. Models of Uncertainty
    - Environmental: non-deterministic and probabilistic
    - Adversarial
2. Guarantees of Success of Plans in the Presence of Uncertainty
    - Analysis techniques
    - Decision making in the presence of uncertainty
3. Learning of Adversarial Strategy
    - Probing strategies
    - Games, partial information, solution concepts
    - Adaptation to changing utility functions of the adversary

Models of Uncertainty

Uncertainty rises for two different reasons: (1) The possibility of mission plans having to be redone because of environmental factors, or to compensate for losses suffered by friendly (blue) forces in mission completion. (2) The possible hostile evasive actions of the opposing (red) forces to evade the pursuers. We will pursue two styles of modeling uncertainty, and develop corresponding approaches to confronting uncertainty.

The first model of uncertainty is probabilistic: the actual input and the actual model is selected from the known sets according to a probability distribution. The control design is now evaluated as expected or average behavior over the set of possible inputs and models. Statistical decision theory and stochastic control theory address this probabilistic uncertainty.

The second model of uncertainty is adversarial. It assumes that the unknown input or unknown model is selected by an adversary whose goal is to degrade as much as possible the execution of the mission. This leads to a game-theoretic formulation, and the design criteria that emerge are associated with the names of game theorists such as Nash, Shapley, and Stackelberg. Two separate sub-themes will be addressed here: (1) *The computation of strategies for Nash, Stackelberg and other kinds of dynamic games with hybrid models of computation.* Here we will hit the complexity barrier for NP-hard or even undecidable problems very soon, hence we will need to develop good sub-optimal greedy strategies and evaluate their effectiveness. (2) *Learning of Adversarial Strategies.* In our opinion a big drawback of current techniques of teaming and game strategies is their inability to allow for a dynamic adversary, who changes his strategy, cost function, and information patterns during the course of an engagement by using a strategy such as *feinting.* An engagement can be preceded by a *learning phase* when a number of scout UAVs/UGVs are sent out to probe and learn about adversary reactions for use in an engagement, using new graphical learning techniques.

Plan Completion in the Presence of Uncertainty

In this project we propose a new approach to verification of the accurate functioning of "safety critical systems". The heart of the approach is not to verify that every run of the hybrid system satisfies certain safety or liveness properties, but rather to check that the properties are satisfied *with a certain probability*, given uncertainties in actuation, sensing and the actions of the environment. In this sense, this is a "softening" of the strict logical notion of verification, which is typically unachievable due to environmental and adversarial uncertainty. Also, in a game-theoretic context, optimal blue-team strategies may themselves be partially randomized (hence unpredictable by red-team). Since the systems are safety critical, we will be interested in guarantees of performance which are high in the unfaulted mode and slightly lower in the faulted modes.

Performance tradeoffs. Control algorithms need to be designed to not compromise safety, and yet allow efficient functioning. Such safety-performance tradeoffs are conveniently characterized in a probabilistic setting. For example, to facilitate mission accomplishment for a groups of UAVs, one would like to avoid executing special collision avoidance maneuvers unless the probability of a collision is fairly high.

Mode switching. Stochastic effects such as sensing and actuation noise that invariably enter any realistic system may result in erroneous mode switching.

Fault tolerance. Fault handling routines have to be proven to work, with high probability in both detecting the fault and then providing the sequence of mode changes as well as tactical and possibly strategic redeployment so as to make for reliable operation.

Malicious environment. The UAV system will have to operate in a potentially malicious environment, for example enemy forces and adverse weather. In this case it is sometimes fruitful to characterize the actions of the environment in terms of probabilistic strategies.

Design mode verification. UAVs will need to have multiple modes of operation, including hover, take off, land, track, etc. It will be important to prove that control algorithms that switch between these modes based on high level commands and vision data do not cause the aircraft to enter unstable or unsafe states. Control algorithms need to be designed not to compromise safety, but to allow for efficient functioning with multiple agents for wide area coverage.

Faulted mode verification. The UAV will need to have fault detection and handling routines, in order to maintain integrity of the aircraft and safety with possible gradual degradation in the performance of the functioning of the system. Fault handing routines have to be proven to work, with high probability in both detecting the fault and then providing the sequence of mode changes as well as tactical and possibly strategic redeployment so as to make for reliable operation.

Probabilistic verification. In the conceptual underpinnings of the research, we will attempt a rapprochement between Markov decision networks and AI-based Bayesian decision networks to come up with a “soft” version of verification for hybrid systems. Conceptually this method will be like the change in computational complexity theory between studying the worst-case and mean behavior of an algorithm.

Probabilistic Hybrid Systems As part of this proposal we propose to develop a comprehensive framework for dealing with probabilistic phenomena in the hybrid domain. This will involve:

- **Probabilistic modeling.** We will extend the current hybrid system modeling formalisms to formally allow the introduction of probabilistic phenomena. We will draw on the experience of the members of our team with probabilistic automata and stochastic control. For large-scale probabilistic modeling, we will require expressive formal languages that combine relational logic with probability.
- **Probabilistic control.** Deterministic control strategies are vulnerable to attacks that exploit their regularity and predictability. To achieve a higher degree of robustness, and to meet mission objectives that cannot be met with deterministic control strategies, we propose the use of randomized control strategies. To achieve the control objectives, the control strategies will compete with randomized strategies modeling probabilistic disturbances and faults. We will give probabilistic performance bounds on their performance using viscosity solutions for the resulting games. We also intend to generalize recent results on reachability objectives in discrete multi-agent games to the liveness case and the hybrid case and their combination.
- **Probabilistic analysis.** When it is not possible to meet specifications deterministically, we will give probabilistic bounds on the performance of controllers and the reliability of the system. Specifically, for multi-modal and multi-agent systems, algorithms and tools will be developed to provide probabilistic estimates of safe behavior (envelope protection, mission completion, etc.). We will develop tools for the analysis of the reliability and performance of distributed multi-agent systems operating in probabilistic and malicious environments. The tools we envision rely on formal-method and probabilistic techniques to enable the analysis both of mature systems and of early conceptual designs and prototypes. Hence, in addition to validating proposed designs, the tools will help with the process of selecting appropriate control and communication architectures capable of meeting the performance and reliability requirements of the system. In particular, we will extend to hybrid, multi-agent systems our techniques for the reliability and performance validation of timed systems. We will also develop fault detection and handling tools, with guarantees of performance after the onset of a fault, as well as

Markov decision processes and Bayesian decision network based tools for incompletely observed modules. In general, we advocate a shift from worst-case behavior to mean behavior estimates of control algorithms.

Intelligent Adversarial Strategies: Modeling and Learning of Adversary Strategies

The second source of uncertainty arises from the deliberate evasive actions of an adversary. In this part of the proposal we will focus on the modeling and learning of adversarial strategies. A key mathematical framework for the modeling of adversarial actions comes from the theory of games, and partially observable Markov decision processes and games. Pursuit-evasion games model a class of zero sum games where the two adversaries have complimentary goals. Each adversary is allowed to execute teaming strategies among groups of its UAVs/UGVs and in addition to make changes in its cost function (feinting) to confuse or misinform the adversary. We review our work (theoretical and experimental) to date in this area in the next subsection. Here, we discuss how we will organize our proposed research:

Partially observable Markov games. It can be shown that either Nash or Stackelberg strategies exist and are Markov for completely observable decision processes and games. However, when the games and decision processes are not fully observable, the exact nature of the information patterns, extent of sharing, misinformation, etc. make the determination of answers very complicated. We will study existence of solutions in the presence of different information patterns and patterns of concealment under which no regret for Nash solutions. We will single out Nash over Stackelberg mainly because we consider the overall Nash strategies more robust.

Sub-optimal greedy approximations. A major difficulty in solving Partially Observable Markov Decision Processes (POMDPs) and Games (POMGs) is that the solutions of the Bellman equations are normally intractable since the size of the solution space grows exponentially in the size of the discretized state space and number of agents and the time period of the solution (because the optimal policies are not Markov). However, one or two or N-step greedy approximations are computationally more attractive and have bounded degradation in performance. We will study the quality of the sub-optimal approximants and their relative performance. Key features of the solutions that we will study are the dependence on how the speed and sophistication of the adversarial actions and information patterns will thwart the completion of the mission. We expect to have the *mathematical equivalent of war gaming* built into our solutions.

Policy search for large POMDPs. To solve large POMDPs, We also apply the PEGASUS method [43], which gives an efficient algorithm for finding good control policies even in very large POMDPs. The key in PEGASUS lies in its searching for policies by evaluating each of them on a small number of "representative scenarios." As a concrete example, if we are learning to quickly locate mines randomly buried in a road, a "scenario" might be a specific placement of the mines. As another example, if we are trying to track and intercept an unguided missile, the scenario might be the specific (random) trajectory taken by the missile. In many situations, it is not clear what a representative "scenario" is, but we show they can be automatically defined and generated for *any* POMDP. This stems from the rather surprising mathematical fact that any stochastic POMDP problem can be reduced to one with only deterministic transition dynamics (but still a random initial state). We also proved that in order to find good policies, the number of scenarios needed is small—generally only a low-order polynomial of the dimension of the state space. This should be contrasted with methods that discretize the state space or that attempt exact solutions to the Bellman equations, which suffer from the curse of dimensionality and hence are totally inapplicable to even moderate-sized problems. We also note that PEGASUS applies straightforwardly in hybrid control settings and for learning distributed policies for agents that may need to act independently, but in a way that their actions remain maximally cooperative/coordinated.

Robust and model predictive games. As pointed out above, Nash solutions are less brittle than Stackelberg or leader follower solutions. In more general terms, it is important to choose robust strategies for dealing with adversarial intent. One way to build robustness into a strategy is to not have strategies for infinite horizon cost functions, but rather to make them finite receding horizons with recalculations during operations. In the case of control for a single agent these control strategies are referred to as model predictive control. We propose to develop a theory of model predictive games to deal with variable planning horizons for the adversary and the blue team. We will explore strategies for defeating a forward thinking (long horizon) adversary who is slow (since the calculations are likely to be many) or a short sighted adversary who is fast (with fewer calculations) and explore the tradeoffs between "intelligence" and "speed".

Learning of adversarial strategies. In our opinion a big drawback of current techniques of teaming and game strategies is their inability to allow for a *dynamic adversary*, that is one who changes his strategy, cost function, information patterns during the course of an engagement. We will address this problem head on by exploiting ideas of dual control. An engagement can be preceded by a *learning phase* when a number of scout UAVs/UGVs are sent out to probe and learn about adversary reactions for use in an engagement. We will use Q-learning and other reinforcement learning techniques to develop the learning of the adversarial intent. During an engagement, we would also propose an "observer" UAV to obtain data about the adversarial strategy to provide an "outer learning loop" which slowly adjust blue team tactics to compensate for changing adversarial intent and information patterns. The main concern here is to not have too many nested levels of learning of adversarial action, since counteraction by the adversary will need to be counteracted and so on, ad inf.

While we previously described the application of PEGASUS to problems with stochastic uncertainty, we have, by building on the work of [45], recently further generalized our algorithm to apply to problems with adversarial uncertainty. Whereas we had previously pointed out that computing Nash equilibria in large games such as pursuit-evasion games tends to be completely intractable, these methods will allow us to compute *approximate* Nash equilibria. Thus, this will enable us to compute policies that are robust to any actions that an adversary may take to try to thwart our actions.

Probabilistic Pursuit-Evasion Games

Multi-agent pursuit-evasion games is a promising application for cooperative multi-robot systems in which a team of agents acting as *pursuers* attempts to capture a group of *evaders* within a bounded but unknown environment (see Figure 9).

The pursuit-evasion game is a mathematical model of the two scenarios which were suggested at the begin of the proposal. It is also important to note that the pursuit-evasion game can be easily adapted to model several scenarios of interest to the Army and DoD including suppression of enemy air defenses (SEAD), precison strike of time critical targets, or for that matter many other scenarios of theater air defense, and hostage rescue operations.

A probabilistic framework is used to model the locations of the pursuers on the ground as well as obstacles. The basic setup considers multiple pursuers trying to capture a single evader undergoing random motion. We have used a simple sensor model based on the probability of false positives $p \in [0,1]$ and false negatives $q \in [0,1]$ of an agent detecting an evader in adjacent locations. Also, we assume that pursuers have perfect knowledge of their own locations, that is $\mathbf{v}(t) = \mathcal{X}_p(t)$. It is possible to have one pursuer collect all the measurements, build the maps and broadcast its findings to the other pursuers (centralized case), or each pursuer builds its own map and shares its measurements (decentralized case). We have assumed that pursuers are able to identify each evader separately. Therefore, pursuers keep one map for each evader and one map for the obstacles. Whenever an evader is captured, that evader is removed from the game and its map is no longer updated. *Capture* is defined as follows: Let $x_{p_k}(t) \in \mathbf{v}(t)$ and $x_{e_i}(t) \in \mathbf{e}(t)$ be the estimated positions of pursuer k and evader i at time t, respectively. We say that evader i is captured by pursuer k at time t if $x_{e_i}(t) \in V_{p_k}(t)$ and $d(x_{p_k}(t), x_{e_i}(t)) \leq d_m$ where $d(\cdot,\cdot)$ is a metric in $\mathcal{X}$ and d_m is a pre-specified *capture distance.*

A hierarchical system architecture shown in Figure 10 is employed for real-time implementation of the pursuit-evasion games on the BEAR platforms. This architecture segments the control of each agent into a number of different layers of abstraction, dynamically changing conditions in the environment are perceived by various sensor, and the sensed data is processed by map building and interpreted by strategy planning to determine proper actions. The sensing and action coordination occurs frequently than strategy planning so that each agent can have the capability to avoid the immediate danger, e.g. collision, in real-time.

The *strategic planner* handles the selection and control of tasks at the highest level. It maintains a state-space of the system useful for mission planning and tasks the agents according to mission objectives, and handles inter-agent communication. State information maintained by the strategic planner is used by the *tactical planner* for motion control of the agent. Each agent will make observations of the environment using sensors and through communication with other agents, and then decide a course of action (to map the environment or attempt to capture an evader,

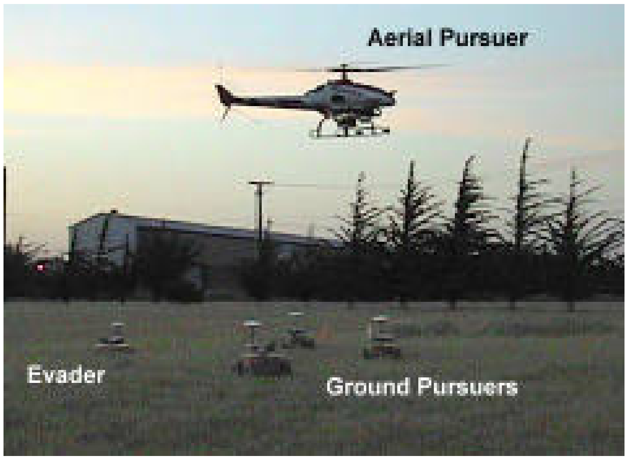

Figure 9: The Berkeley AeRobot testbed for pursuit-evasion games.

depending on the scenario). The *trajectory planner* is responsible for the design of a realizable trajectory for each agent, based on a detailed dynamic model of the vehicle and the set of way-points given by the tactical planner. It is at this level safety routines, such as obstacle avoidance, also reside. The trajectory planner provides a set of way-points to the *regulation layer.* The regulation layer uses various control techniques to guide the agent to the desired way-point and send the tracking error back to the trajectory planner in case rescheduling is necessary. The Pegasus method has been successfully applied by the Carnegie-Mellon University autonomous helicopter group to learning a policy for hovering and for arbitrary trajectory following. [44] We propose to apply Pegasus to learning UAV policies for increasingly challenging maneuvers, such as autonomous landing (where the ground effect makes the dynamics noisy and non-linear, so that more traditional linear control design methods are completely inapplicable) and difficult combat maneuvers. Given the inherent danger in flying UAVs, an important part of any implementation must be guarantees on the performance of its policies, such as are traditionally given by verification methods. Extending the work of [43], we will also prove strong conditions under which the UAV policies learned by Pegasus can be rigorously guaranteed to meet performance levels required for a successful combat operation.

We propose to prove the probabilities of success and the robustness of some strategies in complete end to end designs of the different levels of the hierarchy: is there a *persistent* pursuit policy that guarantees that the evader can be captured in finite time with probability one. We extended the basic scenario to consider supervisory agents, such as a helicopter, that can estimate the position of the evader but not capture it. We also extended the approach to multiple evaders, under the assumption that each one can be identified separately. The case where the evaders actively avoid the pursuers has also been investigated, where a dynamic programming solution to a Stackelberg equilibrium of a partial information Markov process is proposed. This solution is computationally infeasible in realistic scenarios, which inspires the search for efficient sub-optimal solutions with good performance. Thus far, we have validated the proposed pursuit policies on numerous realistic simulations and experiments, and have studied the effects of the agents' physical limitations (speed, visibility region, etc.) on their behavior under the different pursuit policies. As a typical example, Figure 11 shows the evolution of the probabilistic pursuit-evasion game between a team of two ground and one aerial pursuers and one ground evader. Ground pursuers have a trapezoidal visibility region, while aerial pursuers have a rectangular one. Also, while the ground robots cannot realize arbitrary trajectories due to the non-holonomic constraints on the conventional wheel-based drive mechanism, the aerial agents are free to move in any direction. On this MURI we propose the following lines of research to be validated on the experimental testbed on the two scenarios laid out at the start of this proposal:

Robust strategies for mission completion. There may be many strategies for planning a mission, but some are more robust than others: for a simple example Nash strategies are more robust than Stackelberg ones. We will prove the success probabilities not only for a single strategy but for tubes of outcomes around a given initial condition and parameters. We will verify this in experiments. Experiments provide an important way of validating and debugging the technologies developed here.

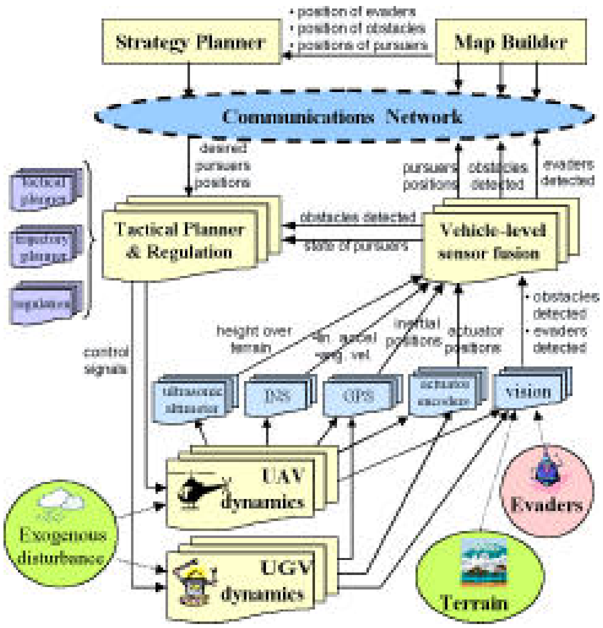

Figure 10: Blue-force architecture for pursuit-evasion game

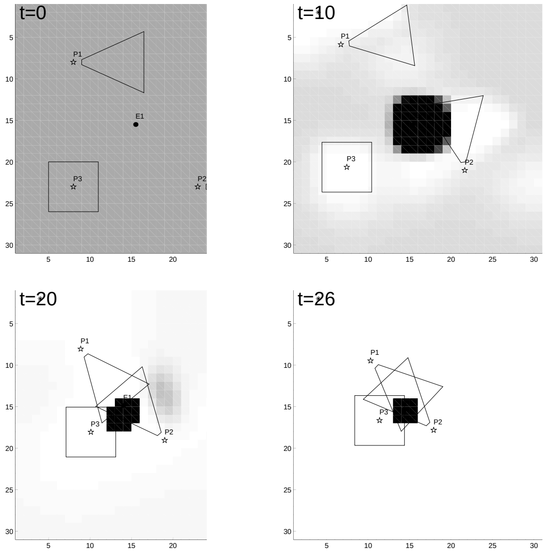

Figure 11: one randomly moving evader is captured by two ground pursuers and one aerial pursuer using the global maximum policy. Background color of the $30m \times 30m$ arena indicates the value of $p_e(x, t+1|Y_t)$, dark meaning the high probability of containing an evader, t denotes time in seconds.

Evaluation of reflective versus agile strategies. A thoughtful strategy may be globally optimal but may take too long. Agile strategies on the other hand may be outsmarted by a clever opponent. We will explore in experiment the tradeoffs of when to spend more computation to reflect and when to rapidly take decisions and to make mixed strategies embodying a combination of the strategies to be robust against red-teaming.

Red-teaming strategies. One of the key features of our UAV/UGV testbed is to allow for the possibility of the red team being a "human expert" team opposing the blue "robotic or mixed initiative automateam". We propose to use the experimental testbed to bring in seasoned military commanders to red-team our strategies with full knowledge of their algorithms.

3 Facilities: Testbeds for Experiments and Research Integration

For this project, we propose hardware verification and validation of our methods on a UAV/UGV platform developed at Berkeley with fully autonomous and semi-autonomous behavior for scenarios like scout missions for time-critical target acquisition and tracking of an intelligent evasive target. A testbed for mapping of urban areas for hostage rescue scenarios will be developed using UGVs at UPenn for experimentation with algorithms for fusing multiple omni-directional camera data to combine with prior simulated views of complex urban environments for location of foes and friendly forces. The CMU testbed will focus on unexploded ordinance detonation and de-mining of an intelligent adversary's minefield with groups of robots.

Furthermore, at Berkeley, we have ready four types of testbeds for experimental and theoretical work which have some connections with the style of intelligent control research proposed here. These testbeds are supported by other funds, but provide some background information and knowledge that is relevant to this project:

- Intelligent Vehicle Highway Systems testbed: we have three fully instrumented cars with sensors, an extensive simulation facility, and surveillance testbeds.
- Air Traffic Management Systems testbed: we have a simulation testbed in addition have access through the FAA and Boeing to simulation models of commercial and some civilian aircraft.

- Intelligent Battlefield TeleMedicine testbed: we have a telesurgery testbed with cooperating robots equipped with tissue modeling and tactile display tools. We also have access to surgical video data from UCSF.
- Mobile Offshore Base testbed: we have rapidly deployable offshore landing platforms with multiple units that can be assembled and controlled for landing large aircraft in high sea states.

3.1 Berkeley Testbed

Unmanned Robotic UAV/UGV testbed The BErkeley AeRobot (BEAR) project [40] is a research effort at UC Berkeley that encompasses the disciplines of hybrid systems theory, navigation, control, computer vision, communication, and multi-agent coordination, since 1996. We currently have eleven instrumented model-scale helicopters equipped with GPS/INS, camera, and other sensors on board, which we have been using to validate our control systems design algorithms for UAVs. Information about this is available through our web site at `http://robotics.eecs.berkeley.edu/~koo/bear/`. In addition, we also have four mobile ground-based robots for pursuit-evasion games between the ground based robots and UAVs. No funding is requested for purchasing this hardware. In addition, the robotics and intelligent machines lab at Berkeley has an extensive network of workstations, Suns, SGIs, and other graphics, simulation and visualization workstations. Also, the SHIFT group which pools the acquisition and maintenance of computing resources owns some 60 Sun and DEC workstations, an equal number of Intel PC's, 10 major cpu servers, including Sun Sparc, DEC Alpha, and DEC Alpha multiprocessor architectures.

With a view towards the two scenarios which will act as unifying ways of harvesting the best of breed technologies for each of the university partners, we have proposed distributed hierarchical hybrid system architecture which emphasizes the autonomy of each agent yet allows for coordinated team efforts. Our hierarchical system design (previously shown in Figure 10 for the pursuit-evasion game) was inspired by the hierarchical architectures of Automated Highway Systems [1, 2], Air Traffic Management Systems [3]. This architecture has been implemented on a fleet of UGVs and UAVs [41, 42], including components such as high level pursuit policy computation, inter-agent communication, navigation, sensing, and control. We have studied the effectiveness of the proposed hierarchical architecture, which has been demonstrated on the application of probabilistic pursuit-evasion games between teams of UAVs and UGVs.

Our UGV fleet consists of five ActivMedia Pioneer 2-AT all-terrain ground robots (see the ground robots in Figure 9). These rugged UGVs are four-wheel drive, differential skid-steering robots designed for all-terrain operations. The regulation layer is responsible for low-level control of the robot, managing the the motors, position encoders and user accessories such as sonars, compasses, grippers, etc. The tactical/trajectory planners are implemented by the vision computer with either the Ayllu or Saphira software. The vision computer (client) communicates with the micro-controller (server) through a serial connection. The vision computer receives the current state of the robot (position, heading, translational velocity, rotational velocity, sonar readings, etc.) and sends motion commands (move forward, rotate, go to a certain position, etc.).

Our UAV frames are integrated with navigation sensors based on an integrated INS/GPS and localizing sensors such as ultrasonic altimeters and vision-based sensors, flight control computers and wireless communication devices, allowing the UAV to achieve autonomous flight by following dynamically uploaded commands from an external source. The flight control system we designed [38] is capable of performing autonomous hover, turns, low-speed forward flight and longitudinal-lateral flight with fixed heading. The implemented autonomy can be sequenced or dynamically requested through the use of a novel framework called Vehicle Control Language (VCL) [39]. VCL is a script-style language that specifies a sequence of flight-mode associated destination coordinates and optional parameters. VCL is seamlessly integrated into the hierarchical structure shown in Figure 12, providing a nice abstraction between the high-level mission guidance layer and the low-level vehicle stabilization and control layer.

3.2 UPenn testbed

The General Robotics Automation Sensing and Perception (GRASP) Laboratory is an interdisciplinary laboratory dedicated to research and education in robotics and automation, with eight Faculty and thirty students. In relation to this proposal, the GRASP lab provides laboratory space in which experiments will be run, and provides significant computing power to do computational modeling and testing.

The GRASP Laboratory occupies 11,000 ft.^{2} on a single floor of a modern office building. Its expansive computing facilities include 15 Sun/Solaris and 32 Dell/Win32 workstations, X-Terminal, Linux, and SGI hosts. The Lab's four-

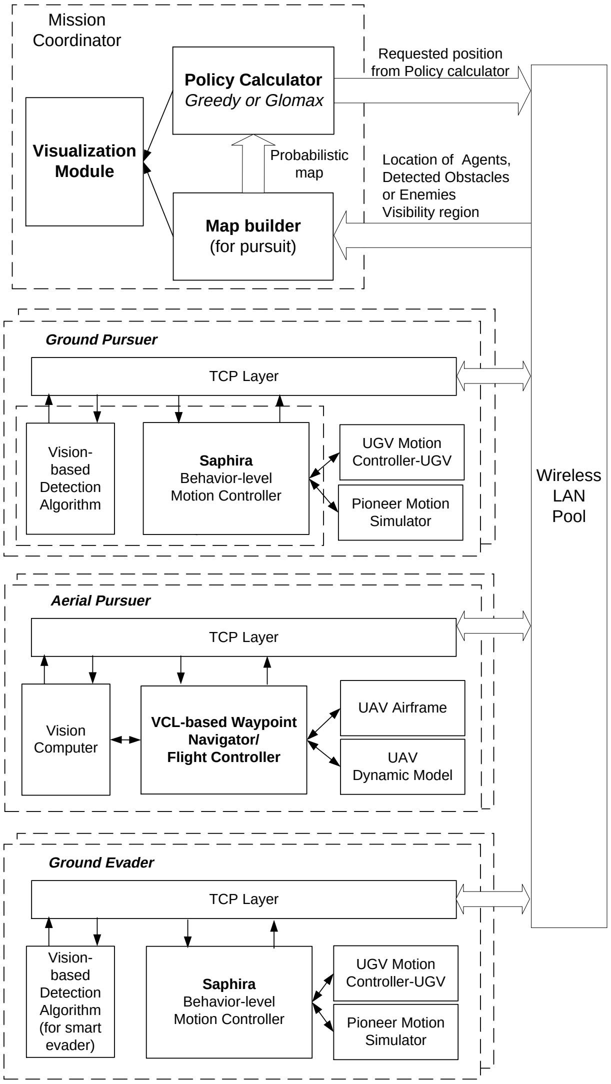

Figure 12: Implementation of hierarchical network architecture on the Berkeley AeRobot testbeds.

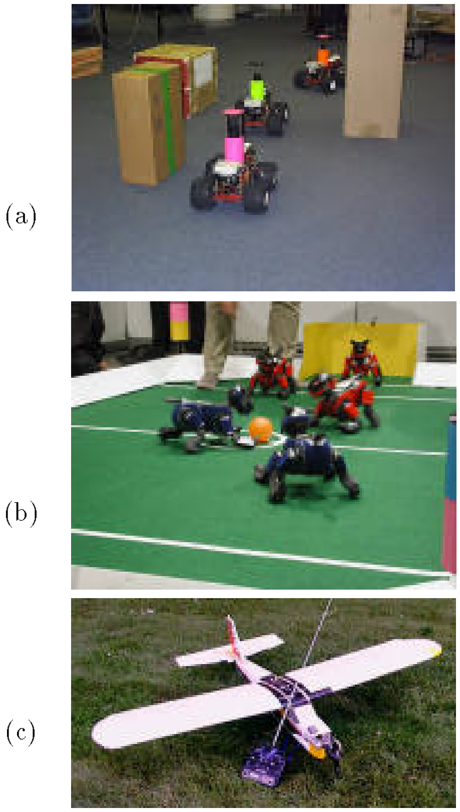

Figure 13: The GRASP/UPenn multi-vehicle testbed includes (a) seven wheeled platforms equipped with omnidirectional cameras; (b) legged robots with visual sensing; and (c) three UAVs, including quadrotor helicopters with cameras for indoor operation (not shown).

processor Sun Ultra-450 file and application server integrates these workstations to provide a coherent, data-centric computing environment for researchers. A common network, with both wired and wireless access, ties all machines together, and this network connects to Internet-2's Abilene backbone. The Lab has many ground based and aerial robots, robot manipulators, and sophisticated sensor systems in its diverse projects. It includes an optics laboratory for developing, omnidirectional, and hyperspectral (IR and UV) sensors.

3.3 CMU testbed

The Robotics Institute of Carnegie Mellon University is the world's largest academic organization for robotics research. In the 21 years since its founding, the Robotics Institute has constructed hundreds of mobile robots, ranging from centimeter-sized crawlers to full-sized robot backhoes, buses, and mining machines. Thus multitude of unique research equipment is available in over a dozen laboratories spanning basic areas of perception, cognition, and manipulation to mobile robot systems, advanced manipulators and manufacturing. Three labs will be conducting work from the Robotics Institute: The Navigation Lab, the Microdynamics Systems Lab, and the Robotics Sensor Based Planning Lab.

The Navigation Laboratory (Navlab) project produced a series of 10 autonomous vehicles (included two HMMWVs, two transit buses, a panel truck, two minivans, and a three-passenger vehicle), integrating vision, radar, ladar, sonar, GPS and laser scanners for both on road and off road applications. Technologies developed under the Navlab umbrella include laser scanners for terrain mapping, 2-, 3-, and 5-eyed stereo vision systems for obstacle detection; 77 GHz phased-array radar for obstacle detection (with no moving parts); and algorithms for obstacle avoidance, road following, route following. NavLab houses a 1,400 square feet indoor facility assembly area containing test equipment and parts for building and servicing electronic control systems for autonomous navigation and sensor integration. The Robotics Institute is currently working to automate Navlab 11, a Jeep Wrangler that will serve as a testbed vehicle capable of both on and off road autonomous navigation. A portion of the proposed research will be performed within the Microdynamic Systems Laboratory which is part of the Robotics Institute. Systems currently under development within the Microdynamic Systems Laboratory include precision high-speed vision based guidance systems for legged robots, a modular distributed precision automation system capable of performing sensor guided micron scale assembly tasks, and a high-fidelity haptic interface device. The Department of Mechanical Engineering has given Professor Choset a 1600 square foot lab space and secretarial support. The lab houses several computers, two Nomadic Technologies Mobile robots, two custom built outdoor demining robots, and the JPL Serpentine Robot.

4 Training of Students

We believe that the research proposed here will have a profound effect on training of a new generation of graduate students and scientists on the use of technology. We believe that advances in technology including simulation, visualization, and virtual environments can play a fundamental role in addressing this issue. We are seeking funds to give students and researchers an interdisciplinary environment to integrate sensing, control, multi-sensor fusion and real-time computing. We will create a resource for students to learn about both algorithmic aspects, hardware as well as software and integration and validation of the multi-sensor multi-agent platforms. This is critical to being able to understand the full autonomous and mixed initiative high confidence systems design process.

In addition to semantic integration of the models used by CAD tools such as SHIFT, HyTech, CHARON, Ptolemy, a common interface must be designed to enable the hand-off of control specifications between the various tools. For experimentation and evaluation of the entire design flow, from architecture conception to real-time code. Funds are requested in this proposal to expand and further develop autonomous vehicle testbeds which will bring together the group of researchers on this project to develop the entire suite of software tools from high level architectural design to real-time control law implementation (mirroring the research thrusts of the proposal), we will be implementing the control law design on a set of rotorcraft robotic UAVs and UGVs. The research proposed here will be used to prototype control laws for multiple mobile robot task force for specialization in urban environments with complex geometries and real world images of dynamical agents. This is not yet another mobile robot testbed proposal since it has some complex multi-agent problems built into it as well as cluttered real world environments. Not mentioned in this proposal, but something that we will do in the context of the on-going ARO activity is to interface this multiple mobile robot testbed with the helicopter testbed, with the helicopter hovering over the scene and playing the role of the mother ship.

The University of California, Berkeley, University of Pennsylvania and Carnegie Mellon University are some of the best engineering institutions based on a wide variety of national rankings. Our programs in Electrical Engineering and Computer Sciences, Mechanical Engineering and the Robotics Institute are unique in that they stress the interplay of automata theory and control theory from the sophomore level through the graduate programs. To our knowledge, we are in the only program concentrating so much of our teaching and curricular resources to this "third bubble" spanning computer engineering and traditional electrical engineering systems. The autonomous vehicle testbeds that will be created with this equipment will train a generation of graduate students with cross disciplinary expertise of this unique flavor. Judging from the needs of industry, this is the kind of engineer that will be most in demand in the future. Silicon Valley is ready to pay high premiums for graduates who have real-time control and embedded systems expertise with verification and program validation techniques in their domain of expertise. From our own understanding of DoD priorities, we feel that UAVs/UGVs and also UUVs (all Unmanned Autonomous Vehicles) will be critical to the Army, Navy and Air Force in a variety of different contexts: urban, open terrain, air and sea warfare. The UAV-UGV laboratory at Berkeley will support all of these DoD developmental efforts.

We are committed to the recruitment of women and minorities to participate in our laboratory activities. As proof of this, we are proud to have a list of distinguished women alumnae from our laboratory now in many positions of importance all over the country and also in France and Italy. Several African American graduate students from our laboratories have compl program for involving undergraduates from HI UPERB.

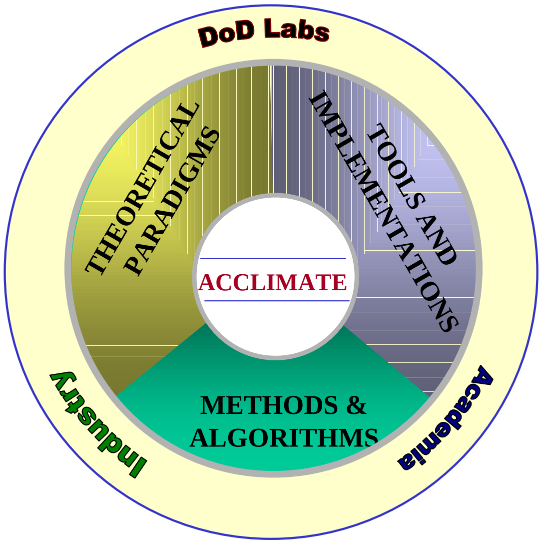

Figure 14:

5 Proposed Team, Subawards, and Management

The project will be jointly managed by Sastry, Kumar, and Choset. The primary administration will be done at Berkeley with subcontractual arrangements to UPenn and CMU. There will be *one annual project reviews* to which Army and other DoD personnel will be invited: its location will move between the partner campuses. The meetings will give exposure of the MURI projects to a academic community, as well as our industrial sponsors and testbed participants. Projects on the MURI will be evaluated after each semi-annual review both for scientific content and relevance to the MURI and DoD every year with recommendations made for changes. Post doctoral scholars will be be jointly supervised by several MURI PIs and they will be encouraged to travel between the MURI locations. All software developed on the project will be placed in the public domain.

The team is well-integrated with a history of long collaborations, not only in papers but in exchanges of students, postdocs, philosophies of research and national security:

Research Team UC Berkeley: Profs. Shankar Sastry (PI), Stuart Russell, Laurent ElGhaoui, Michael Jordan, Anant Sahai, Jitendra Malik, Pravin Varaiya
UPenn: Profs. Vijay Kumar (co-PI), G. Pappas, Kostas Danillidis, Ruzena Bajcsy, James Ostrowski.
CMU: Profs. Howie Choset (co-PI), Alfred Rizzi, Charles Thorpe. The disbursal of requested funds on the MURI is in proportion to the effort levels at Berkeley, UPenn and CMU $ 650K to Berkeley, $ 330K to UPenn and $ 220K to CMU. In addition of the MURI functions, Berkeley will be primarily responsible for the integration of the technology onto the testbeds, and industry liaison, UPenn primarily for the DoD, governmental liaison and CMU for the dissemination into the academic robotics and control communities.

6 Relevance to Army Needs

The design of new 3D adaptive digital battlefields modeling, simulation, and validation problems are beyond those of traditional weapon system design because of the diversity of interacting physical and logical systems which have to be represented and optimized. Weapon systems, both robotic and mixed initiative are characteristically described by evolution of solutions of ordinary and partial differential equations with parameters to allow for adaptation to changing circumstances. Automated command and control systems can be formulated as logical systems in classical logic plus logics developed for AI and common sense reasoning. Evolutionary and adaptive behavior of knowledge and belief is represented in command and control by evolving knowledge bases of rules and facts, which rely on automated deduction to recommend decisions. Command and control, human behavior, and weapon systems interact with each other and the environment as a single feedback system. Many systems being developed have feedback loops which control the systems in which some of the controllers are human agents using decision aids, some are automated decision software, and some are conventional physical controllers. The research proposed on this grant will address the issue of design of hierarchies and heterarchies for the design of an effective management of a digital battlefield without cognitive overload.

Interactions with Army Labs Each of the principal PIs as well as a number of the faculty investigators are long time investigators for the Army Research Office and have built up extensive relationships with Army labs, such as Picatinny Arsennal, ARDEC, TACOM, CAA, ARL, etc. We have had extensive discussions with Dr. Coleman of Picatinny Arsennal on hybrid systems in helicopter fire control, researchers at TACOM on platooning and coordination of HumVees, researchers at Natick on analysis and simulation of load bearing GIs, and with the former Army hospital at Fort Ord, Ca on telemedicine. We will continue a pattern of visits and exchanges and would welcome longer term visits of Army personnel to Berkeley, UPenn and CMU. In addition both Berkeley and UPenn have recently been good performance on ARO MURIs which have just ended. They both have many direct ties with Army organizations, will use these contacts under the current MURI to bring AI-Hybrid system technology into modeling and design of as many Army systems as possible. For instance, Bajcsy and Sastry have been on the ARO Board of Visitors. Kumar, Malik, Russell, Danillidis, Sastry, and Bajcsy have had extensive interactions with ARL groups and with several Army sponsored meetings in such areas as distributed interactive simulation, semi-automated forces, linguistic command and control interfaces, distributed command and control, etc. Sastry has just finished a term as the Director of the Information Technology Office at DARPA with strong participation in the creation of the Army-DARPA FCS program. Former, ITO Deputy Director, Col. (US Army retired) Mark Swinson, now at Sandia Laboratories running their robotics program will be a key military consultant to the project. Brigadier General (USMC retired) Keith Holcomb will be a consultant to the project on war fighting in urban terrain.

Technology transfer

One of the key characteristics of our project is that we will be synergistic with the larger world-wide research community working on “Software for Distributed Multi-Agent Control Systems”. We have been involved as the leaders in the last 5 years in fostering and growing a world-wide research community in this area with yearly conferences and workshops in the areas of “Hybrid Systems”, “Intelligent Control”, and Multi-Agent Systems. Meetings on this topic have been held at Berkeley, Philadelphia and Pittsburgh from 1996 onwards under the auspices of the Army Research Office, DARPA and the Office of Naval Research in addition to scholarly meetings like the IEEE Conference on Decision and Control, IEEE Conference on Robotics and Automation, IEEE Conference on Computer Aided Design. Berkeley, UPenn, CMU have become known as centers of a world-wide activity in “Multi-Agent Systems: Hybrid Control, Communication and Computation”. Some of our colleagues in this area who are not explicitly listed on this proposal will collaborated with us on this, such as Henzinger at Berkeley, Alur at UPenn and Krogh at CMU. We are well connected (with frequent visitors and seminar exchanges) to research efforts as collaborators (on other projects) at Cornell, Stanford, Yale, MIT, Harvard, CNRS, Grenoble, INRIA, Rennes and INRIA, Grenoble, Royal Institute of Technology, Stockholm, University of Porto, Portugal, the University of Rome, Pisa. We have special collaborations with several government laboratories: Air Force Research Laboratory, Wright Patterson with whom we collaborate in the area of SEAD missions and hybrid systems (groups of Dr. Siva Banda, Mr. Ray Bortner), ARL with the groups of Phil Emmerman and others, NASA Ames Research Center, with whom we collaborate in the area of flight management systems, and air traffic control, Army Research Laboratories, with whom we collaborate in the area of multi-agent autonomous intelligent agents, Picatinny Arsennal (US Army ARDEC) with whom we work on hybrid control.

In terms of industrial laboratories, we have a long history of collaborating with some of the best industrial control research laboratories in the USA: Boeing in the area of embedded and autonomous software, Raytheon in the area of distributed air traffic control and unmanned aerial vehicels, Honeywell Technology Center, with whom we have a large and ongoing program of collaboration in the area of Air Traffic Management Systems and Flight Control

Systems Software, United Technologies Research Center, with whom we collaborate in the area of jet engine control systems, SRI a non-profit with interests in communication, robotics and control systems design and high confidence real-time software, Cadence a CAD company with interests in Embedded Control hardware, Scientific Systems, Inc. a small business in Massachusetts with whom we collaborate with on mobile offshore base and unmanned combat air vehicles research, Xerox PARC with whom we collaborate on formal methods and controller software for copiers, and Hughes (Raytheon), General Motors and Ford.

The products of this research will be software and methodology. Methodology results will be reported in conference presentations, papers, and reports that will be posted on the World Wide Web. Fundamental results will also be implemented in software that will be freely distributed. At Berkeley, we have a long standing tradition of disseminating many of key research results by distributing high-quality software. The software is distributed with a very liberal copyright that permits commercialization and requires only attribution and disclaimer of liability. This mechanism has proven effective for technology transfer. We also produce high-quality documentation, which we make available on-line. We plan to continue this high calibre and professionalism. We believe that it greatly enhances the impact of the software, and hence of our fundamental research results as well.

6.1 Current and Pending Support of PI

7 Personnel

8 Biography of Investigators

8.1 Berkeley PI and Investigators

- **Shankar Sastry, PI,** is Professor of Electrical Engineering and Computer Sciences and Bioengineering, as well as the chairmen of the Department of Electrical Engineering and Computer Sciences at Berkeley. He has just completed a term as Director of the Information Technology Office at DARPA from November 1999 to March 2001. Dr. Sastry received his Ph. D. in Electrical Engineering from the University of California, Berkeley in 1981 and a Master of Arts (honoris causa) from Harvard University in 1994. He taught at MIT in 1981-1982 and has been at Berkeley since 1983. He was a Gordon McKay Professor of Electrical Engineering and Computer Sciences at the Division of Applied Sciences, Harvard University in 1994, and a visiting Vinton Hayes fellow at MIT in the fall of 1992. He was awarded the President of India medal in 1977, the NSF Presidential Young Investigator award in 1985, the Eckman Award of the American Control Council in 1990, and the David Marr Best Paper Prize for the best paper in the International Conference on Computer Vision in 1999. He is a Fellow of the IEEE and a Member of the National Academy of Engineering. He is the author/co-author of three books, one on adaptive control (with Marc Bodson, Prentice Hall, 1989), and on robotics (with Richard Murray and Zexiang Li, CRC Press, 1994), and the latest book "Nonlinear Systems: Analysis, Stability and Control", Springer Verlag, 1999. He has co-edited five books on hybrid systems and control and published over 250 papers. His recent research is in the areas of hybrid control systems; distributed control of multi-agent control systems including air traffic management and road transportation systems; sensor fusion, fault handling and coordinated control for UAV/UGVs; millimeter-scale robotics for surgery and simulation and visualization techniques for training surgeons.
- **Laurent El Ghaoui** is an Associate Professor in the Department of Electrical Engineering and Computer Sciences at UC Berkeley. He received his Masters' degree from Ecole Polytechnique (France) in 1985 and his PhD in Aeronautics and Astronautics from Stanford University in 1990. He taught at Ecole Nationale Superieure de Techniques Avancees and Ecole Polytechnique in France before joining the faculty of the University of California at Berkeley in 1999. He received the bronze Medal of the Centre National de la Recherche Scientifique in 1998, the NSF Career Award in 2000, and the award from the Okawa foundation for Information and Telecommunications in 2001. He is an associate editor of the SIAM Journal of Matrix Analysis and its Applications. His research interests include optimization with uncertain data, robust statistics, robust control, data mining, air traffic management and circuit design. He has published over 80 papers and two books, "Linear Matrix Inequalities in Systems and Control" (with S. Boyd, V. Balakrishnan and E. Feron, SIAM, 1994) and "Recent Advances in Linear Matrix Inequalities in Control" (edited book with S. Niculescu, SIAM, 1999).
- **Michael Jordan** is Professor in the Department of Electrical Engineering and Computer Science and the Department of Statistics at the University of California at Berkeley. He received his Masters in Mathematics from Arizona State University, and earned his PhD in Cognitive Science from the University of California, San Diego. He was a professor at the Massachusetts Institute of Technology from 1988 to 1998. He received an NSF Presidential Young Investigator Award (1991-1996) and Best Paper Awards at the American Control Conference (1991) and the Conference on Uncertainty in Artificial Intelligence (1996). His research interests focus on the interface between computer science and statistics, where he has published over 140 research papers. He has worked on a number of topics in machine learning, including neural networks, decision trees, hidden Markov models, general graphical models and reinforcement learning. He has recently focused on variational methods for approximate probabilistic inference for graphical models. He has given invited plenary lectures at the International Conference on the Mathematical Theory of Networks and Systems, the American Association for Artificial Intelligence, the International Joint Conference on Neural Networks, the ACM Conference on Computational Learning Theory, and the Conference on Uncertainty in Artificial Intelligence.
- **Jitendra Malik** is Professor in the Department of Electrical Engineering and Computer Science at University of California at Berkeley. He graduated from IIT Kanpur in 1980 with the gold medal for the best graduating student in Electrical Engineering, and earned his PhD in Computer Science from Stanford University in 1985. He received a Presidential Young Investigator Award in 1989, the Rosenbaum fellowship at the Newton Institute of Mathematical Sciences, University of Cambridge in 1993, and a Miller Professorship at UC Berkeley

for 2001. He received the Diane S. McEntyre Award for Excellence in Teaching from UC Berkeley, in 2000. He is Co-Editor-in-Chief of the International Journal of Computer Vision. His research interests span a wide range of problems in computer vision and computational modeling of human vision. He has made fundamental contributions to the areas of image segmentation and grouping, differential equations for image analysis, stereopsis, texture, line drawing interpretation, and object recognition. His research has found applications in computer graphics, content based image querying, and intelligent vehicle highway systems. He has authored or co-authored more than 100 publications on these topics, including an invited article in Scientific American.

- **Stuart Russell** is a Professor of Computer Science at UC Berkeley, where he holds the Smith-Zadeh Chair in Engineering. He received his B.A. with first-class honours in Physics from Oxford University in 1982, and his Ph.D. in Computer Science from Stanford in 1986. He then joined the faculty of the University of California at Berkeley. In 1990 he received the Presidential Young Investigator Award of the National Science Foundation, and in 1995 he was co-winner of the Computers and Thought Award, the highest international award in the field of artificial intelligence. He was a 1996 Miller Professor of the University of California and was appointed to a Chancellor's Professorship in 2000. His research interests cover many areas of artificial intelligence, including reasoning and planning under uncertainty, real-time decision-making, machine learning, intelligent agent architectures, autonomous vehicles, search, game-playing, and commonsense knowledge representation. He has published over 120 papers and three books, "The Use of Knowledge in Analogy and Induction" (Pitman, 1989), "Do the Right Thing: Studies in Limited Rationality" (with Eric Wefald, MIT Press, 1991), and most recently "Artificial Intelligence: A Modern Approach" (with Peter Norvig, Prentice Hall, 1995), which is the leading textbook in the field. In 1998 he gave the Forsythe Memorial Lectures at Stanford University. He is a Fellow and former Executive Council member of the American Association for Artificial Intelligence, Associate Editor of the Journal of the ACM, the Journal of AI Research, and the Journal of Machine Learning Research.

- **Pravin Varaiya** is Nortel Networks Distinguished Professor in the Department of Electrical Engineering and Computer Sciences at the University of California, Berkeley. His areas of research are control of transportation systems, hybrid systems, and communication networks. From 1994 to 1997 he was Director of Califonia PATH, a multi-university program of research in Intelligent Transportation Systems. From 1975 to 1992 he was also Professor of Economics at Berkeley. Professor Varaiya has held a Guggenheim Fellowship and a Miller Research Professorship. He is a Fellow of IEEE, and a Member of the National Academy of Engineers. He received an honorary doctorate from the L'Institut National Polytechnique de Toulouse. He is on the editorial board of several journals. He has co-authored four books and more than 200 technical papers, including "High-Performance Communication Networks," (second edition, Morgan-Kaufmann, 2000) with Jean Walrand, and "Structure and Interpretation of Signals and Systems," (Addison-Wesley, 2001) with Edward A. Lee.

8.2 U Penn PI and investigators

- **Vijay Kumar, PI,** received his M.Sc. and Ph.D. in Mechanical Engineering from The Ohio State University in 1985 and 1987 respectively. He has been on the Faculty in the Department of Mechanical Engineering and Applied Mechanics at the University of Pennsylvania since 1987. He is currently a Professor and also holds a secondary appointment in the Department of Computer and Information Science. He is the Deputy Dean of the School of Engineering and Applied Science and also the Directory of the General Robotics, Automation, Sensing, and Perception (GRASP) Laboratory. He is a member of the American Society of Mechanical Engineers, Institution of Electrical and Electronic Engineers, and Robotics International, Society of Manufacturing Engineers. He has served on the editorial board of the IEEE Transactions on Robotics and Automation, Editorial Board of the Journal of Franklin Institute and the ASME Journal of Mechanical Design. He is the recipient of the 1991 National Science Foundation Presidential Young Investigator award and the 1997 Freudenstein Award for significant accomplishments in mechanisms and robotics. His research interests include robotics, dynamics, control, design, and biomechanics.

- **Ruzena Bajcsy** is currently a Professor of Computer and Information Science at the University of Pennsylvania, with a secondary appointment in the Department of Mechanical Engineering and Applied Mechanics. From December, 1998 to July, 2001, she served as the Assistant Director for the Computer Information Science and Engineering Directorate (CISE). She has also been the director of the General Robotics, Automation, Sensing, and Perception (GRASP) Laboratory, which she founded in 1978. Dr. Bajcsy received her master's and Ph.D. degrees in electrical engineering from Slovak Technical University in 1957 and 1967, respectively. She received a Ph.D. in computer science in 1972 from Stanford University, and since that time has been teaching and doing research at UPenn's Department of Computer and Information Science. Dr. Bajcsy has done research in the areas of human-centered computer control, cognitive science, robotics, computerized radiological/medical image processing and artificial vision. She is a member of the National Academy of Engineering as well as the Institute of Medicine.

- **Kostas Daniilidis** is Assistant Professor of Computer and Information Science, University of Pennsylvania, affiliated with the interdisciplinary GRASP laboratory. He obtained his PhD in Computer Science from the University Karlsruhe, 1992 and a Diploma in Electrical Engineering from the National Technical University of Athens, 1986. Prior to his current appointment he was with the Cognitive Systems Group, University of Kiel. Nowadays, his research centers on omni-directional vision and vision techniques for tele-immersion and augmented reality. Daniilidis was the chair of the IEEE Workshop on Omnidirectional Vision 2000. He is the co-chair of the computer vision TC of the Robotics and Automation Society. He has served as Program Co-Chair of the CAIP-97 conference, as member of the Program Committee at ECCV-02, CVPR-01, CVPR-98, ICPR-96, ICPR-00, SMILE-00, VAA-01, VMV-01, and he is reviewer in multiple journals. His research on tele-immersion was featured in Scientific American (April 2001). He is the 2001 recipient of the Ford Motor Company Award for the Best Penn Engineering Faculty Advisor.

- **Jim Ostrowski** is an Assistant Professor in the Department of Mechanical Engineering and Applied Mechanics at the University of Pennsylvania and holds a secondary appointment in the Department of Computer and Information Science. He obtained his undergraduate degree from Brown University (Sc.B., 1986) and graduate degrees from the California Institute of Technology (M.S., 1991; Ph.D., 1996). His expertise is in the areas of nonlinear dynamics and control for robotics, with a particular emphasis on the mechanics and control of robotic locomotion systems and vision-based control of mobile robots. He has studied a wide variety of systems, including wheeled vehicles, snakes, eel-like swimming robots, and blimps. He is currently an Associate Editor for the Conference Editorial Board of the Control Systems Society, has served on the National Organizing Committee for the 2000 IFAC Workshop on Lagrangian and Hamiltonian Methods in Nonlinear Control and on the Program Committee for WAFR-2000 and IROS-2001. He is the recipient of a 1998 NSF CAREER award.

- **George J. Pappas** is currently an Assistant Professor of Electrical Engineering at the University of Pennsylvania, where he also holds a Secondary Appointment in the Department of Computer and Information Sciences. He received the B.S. degree in Computer and Systems Engineering in 1991, the M.S. degree in Computer and Systems Engineering in 1992, both from Rensselaer Polytechnic Institute, Troy, NY. In 1994, he was a Graduate Fellow at the Division of Engineering Science of Harvard University. In 1998, he received the Ph.D degree from the Department of Electrical Engineering and Computer Sciences at the University of California at Berkeley. He was a postdoctoral researcher at the University of California at Berkeley and the University of Pennsylvania. Dr. Pappas is the recipient of the 1999 Eliahu Jury Award for Excellence in Systems Research from the Department of Electrical Engineering and Computer Sciences at the University of California at Berkeley. His research interests include embedded hybrid systems, hierarchical control systems, nonlinear control systems, geometric control theory, flight and air traffic management systems, robotics, and unmanned aerial vehicles.

- **C.J. Taylor** is currently an Assistant Professor in the Computer and Information Science Dept at the University Pennsylvania. He has carried out research on several problems in Computer Vision and Robotics including: reconstruction of 3D models from images, automatic control of vision-guided motor vehicles, mobile robot navigation and multi-robot coordination. Dr. Taylor received his A.B. degree in Electrical Computer and Systems Engineering from Harvard College in 1988. He received his M.S. and Ph.D. degrees from Yale University in 1990 and 1994 respectively. Dr. Taylor was the Jamaica Scholar in 1984, a member of the Harvard chapter of Phi Beta Kappa and held a Harvard College Scholarship from 1986-1988. From 1994 to 1997 Dr. Taylor was a postdoctoral researcher and lecturer with the Department of Electrical Engineering and Computer Science at U.C. Berkeley. He joined the faculty of the Computer and Information Science Department at the University of Pennsylvania in September 1997. He received an NSF CAREER award in 1998 and the Lindback Minority Junior Faculty Award in 2001.

8.3 CMU PI and Investigators

- **Howie Choset, PI,** is an Assistant Professor of Mechanical Engineering and Robotics at Carnegie Mellon

University where he conducts research in motion planning and design of serpentine mechanisms, coverage path planning for de-mining and painting, mobile robot sensor based exploration ofunknown spaces, distributed manipulation with macroscopic arrays, and education with robotics. In 1997, the National Science Foundation awarded Professor Choset its Career Award to continue the work in the underlying fundamentals of roadmaps for arbitrarily shaped objects; the long-term goal of this work is to define roadmaps for highly articulated robots. Recently, the Office of Naval Research started supporting Professor Choset through its Young Investigator Program to develop strategies to search for land and sea mines and to construct a land-mine-search robot. Professor Choset co-chairs the IEEE Technical Committee on Mobile Robots co-chairs the SPIE Mobile Robots Conference each year with Doug Gage. In 1999, he co-chaired with Dr. John Bares the Field and Service Robotics conference and has co-organized with Professor Karl Bohringer a workshop on distributed manipulation; Professors Bohringer and Choset edited a book on the subject. Finally, Professor Choset directs the Undergraduate Robotics Minor at Carnegie Mellon and teaches an overview course on Robotics. Recently, he developed a series of Lego Labs to complement the course work. Professor Choset received his Ph.D. from the California Institute of Technology under the direction of Joel Burdick in 1996.

- **Alfred Rizzi** has served as a Research Scientist in The Robotics Institute at Carnegie Mellon University since January of 1998. He is a member of the Microdynamic Systems Laboratory and is involved in a number of research projects focused on hybrid and sensor-based control of distributed systems. Prior to joining Carnegie Mellon's faculty he served as an NSF postdoctoral fellow within the Robotics Institute under the supervision of Dr. Ralph Hollis from January 1996. From January 1995 through December 1996 he was a postdoctoral fellow with the Center for Display Technology and Manufacturing at the University of Michigan, under the supervision of Prof. Daniel Koditschek. He was employed as an electrical design engineer at the Electro-Mechanical Division of the Northrop Corporation from July 1986 through July 1988. Alfred Rizzi is currently serving as a member of the editorial board for the International Journal of Robotics Research, and has helped guide the introduction of electronic multimedia publication in the journal. He has received a Ph.D. and M.S. in Electrical Engineering, from Yale University in December 1994 and June, 1990 respectively, and an Sc.B. in Electrical Engineering from the Massachusetts Institute of Technology in June 1986.

- **Chuck Thorpe** is the director of the Robotics Institute, and is a Principal Research Scientist. His research interests include computer vision, planning and architectures for outdoor robot vehicles. Since 1984, his Navlab research group has built a series of 11 robotic cars, trucks, and busses for military and civilian research. The Navlab group has pioneered new methods in stereo vision, laser rangefinding, 3D terrain modeling, neural networks for perception, route planning, driver performance modeling, traffic simulation, teleoperation, vehicle control on rough terrain and system architectures. Dr. Thorpe has also been involved with automated helicopters, walking robots and underwater robots. Dr. Thorpe founded the Master's Degree program in Robotics at Carnegie Mellon. Dr. Thorpe received his doctor's degree in Computer Science from Carnegie Mellon in 1984. He earned his undergraduate degree in natural science from North Park College in Chicago, Ill. in 1979. He is a Fellow of the American Association for Artificial Intelligence.

9 Budget

References

[1] J. Lygeros. *Hierarchical, Hybrid Control of Large Scale Systems.* PhD thesis, UC Berkeley, 1996.

[2] J. Lygeros, D.N. Godbole, and S. Sastry. Verified hybrid controllers for automated vehicles. *IEEE Trans. Automatic Control*, 43(4):522–539, 1998.

[3] G. Pappas, C. Tomlin, J. Lygeros, D. Godbole, and S. Sastry. A next generation architecture for air traffic management systems. In *Proc. of 36th IEEE CDC*, pp. 2405–2410, 1997.

[4] C. Tomlin. *Hybrid Control of Air Traffic Management Systems.* PhD thesis, UC Berkeley, 1998.

[5] C. Tomlin, G. Pappas, J. Lygeros, D. Godbole, and S. Sastry. Hybrid control models of next generation air traffic management. In *Proc. of Hybrid Systems: 4th Intl. Conf.*, 1996.

[6] C.J. Tomlin, J.L. Lygeros and S.S. Sastry, "A game theoretic approach to controller design for hybrid systems", *IEEE Proc.*, Vol. 88, pp. 949–970, 2000.

[7] Hybrid Systems IV, edited by P. Antsaklis, W. Kohn, A. Nerode and S. Sastry, Springer Verlag Lecture Notes in Computer Science, 1997, Vol. 1273.

[8] C. J. Tomlin, J. L. Lygeros and S. S. Sastry, "Synthesizing Controllers for Nonlinear Hybrid Systems", *Hybrid Systems, Computation and Control*, Springer Verlag Lecture Notes in Computer Science, editor T. Henzinger and S. Sastry, 1998, pp. 360-373.

[9] T. J. Koo. Hierarchical system architecture for multi-agent multi-modal systems. In *Proc. of 40th IEEE IEEE Conf. on Decision and Control*, 2001. Submitted.

[10] G. Lafferiere, G. Pappas and S. Sastry, "O-minimal Hybrid Systems", *Mathematics of Control, Signals and Systems*, Vol. 13, pp. 1–21, 2000.

[11] O. Shakernia, G. Pappas, and S. Sastry, Semidecidable synthesis for Triangular Hybrid Systems, *Hybrid Systems: Computation and Control: 4th Intl. Workshop, HSCC2001, Rome, Italy,* editors M. Di Benedetto and A. Sangiovanni-Vincentelli.

[12] G. Pappas, G. Lafferiere and S. Sastry, Hierarchical Consistent Control Systems, *IEEE Trans. on Automatic Control*, 45(6):1144–1160, 2000.

[13] R. Vidal, S. Shaffert, J. Lygeros, and S. Sastry, "Controller Invariance of discrete Time Systems," *Hybrid Systems: Computation and Control: 3rd Intl. Workshop, HSCC2000, Pittsburgh, PA'*, edited by B. Krogh and N. Lynch, pp. 437-450.

[14] A. Puri and S. Tripakis, "QoS routing," available at http://wow.eecs.berkeley.edu/papers.html.

[15] A. Kurzhanksi and P. Varaiya, Dynamic optimization for reachability problems, *J. Optimization Theory and Applications*, 108(2): 227–251.

[16] A. Deshpande, A. Gollu and P. Variaya, "SHIFT: a formalism and a programming environment for dynamic networks of hybrid automata", in *Hybrid Systems IV*, Springer Verlag Lecture Notes in Computer Science, editors P. Antsaklis, W. Kohn, A. Nerode and S. Sastry, 1997, pp. 113-133.

[17] R. Bajcsy. Active perception. *Proc. of the IEEE*, 76:996–1005, 1988.

[18] A. Das, R. Fierro, V. Kumar, J. Southall, J. Spletzer, and C. J. Taylor. Real-time vision based control of a nonholonomic mobile robot. In *IEEE Intl. Conf. Robot. Automat.,* pp. 157–162, Seoul, Korea, May 2001.

[19] A. Das, J. Esposito, R. Fierro, G. Grudic, V. Kumar, J. Ostrowski, J. Southall, J. Spletzer, and C. Taylor. A framework and architecture for multirobot coordination. In *Seventh Intl. Symposium on Experimental Robotics*, Honolulu, Hawaii, 2000.

[20] Jaydev P. Desai, Vijay Kumar, and James P. Ostrowski. Control of changes in formation for a team of mobile robots. In *IEEE Intl. Conf. Robot. and Automat.*, Detroit, MI, April 1999.

[21] Jaydev P. Desai, James P. Ostrowski, and Vijay Kumar. Controlling formations of multiple mobile robots. In *IEEE Intl. Conf. Robot. Automat.*, pp. 2864–2869, Leuven, Belgium, 1998.

[22] J. Esposito, V. Kumar, and G. Pappas. Accurate event detection for simulating hybrid systems. Hybrid Systems: Computation and Control 2001.

[23] R. Fierro, A. Das, V. Kumar, and J. P. Ostrowski. Hybrid control of formations of robots. In *IEEE Intl. Conf. Automat.,* pp. 157–162, 2001.

[24] A. Das and V. Kumar and R. Fierro. Stability and Performance of Networked Teams of Robots. Submitted to *IEEE Intl. Conf. Robot. Automat.,* 2002.

[25] C. Geyer and K. Daniilidis. Structure and motion from uncalibrated catadioptric views. In *IEEE Conf. Computer Vision and Pattern Recognition.* Hawaii, 2001.

[26] H. Tanner, V. Kumar and G. Pappas. Input-State Stability in Teams of Robots. Submitted to *IEEE Intl. Conf. Robot. Automat.,* 2002.

[27] J. Spletzer and C. J. Taylor. Cooperative sensing, estimation and localizaton in multi-robot teams. Submitted to*IEEE Intl. Conf. Robot. Automat.,* 2002.

[28] C.J. Taylor. Video plus. In *IEEE Workshop on Omnidirectional Vision,* Hilton Head, SC, pp. 3–10, 2000.

[29] R. R. Burridge, A. A. Rizzi, and D. E. Koditschek. Sequential composition of dynamically dexterous robot behaviors. *Intl. Journal of Robotics Research*, 18(6):534–555, 1999.

[30] A.A. Rizzi. Hybrid control as a method for robot motion programming. In *IEEE Intl. Conf. Robot. Automat.*, pp. 832–837, Leuven, Belgium, 1998.

[31] J. Mulligan and K. Daniilidis. View-independent scene acquisition for tele-presence. In *Proc. Intl. Symposium on Augmented Reality*, pp. 105–110, Munich, Germany, 2000.

[32] R. Alur, T. Henzinger, G. Lafferriere, G.J. Pappas, *Discrete Abstractions of Hybrid Systems*, Proc. of the IEEE, 88(2):971–984, 2000.

[33] P. Antsaklis, W. Kohn, A. Nerode, and S. Sastry, editors. *Hybrid Systems II*, volume 999 of *Lecture Notes in Computer Science.* Springer-Verlag, 1995.

[34] T. Henzinger and S. Sastry, editors. *Hybrid Systems : Computation and Control*, volume 1386 of *Lecture Notes in Computer Science.* Springer-Verlag, 1998.

[35] H. Bronnimann and M.T. Goodrich, Almost Optimal Set Covers in Finite VC-Dimension. *Discrete and Computational Geometry.* Vol. 14, pp. 463–479, 1995.

[36] G.J. Pappas, Bisimilar Linear Systems, Submitted to *IEEE Trans. on Automatic Control.*, 2001.

[37] P. Tabuada and G.J. Pappas, *Hybrid abstractions that preserve timed languages*, Hybrid Systems : Computation and Control, Lecture Notes in Computer Science, vol. 2034, Springer, 2001.

[38] H. Shim, H.J. Kim, S. Sastry, Control system design for rotorcraft-based unmanned aerial vehicles using time-domain system identification, *Proc. of IEEE Intl. Conf. on Control Applications,* Anchorage, 2000.

[39] H. Shim, H.J. Kim, and S. Sastry. Hierarchical control system synthesis for rotorcraft-based unmanned aerial vehicles. In *Proc. of AIAA Conf. on Guidance, Navigation and Control*, Denver, 2000.

[40] BErkeley Aerial Robot (BEAR) Project homepage. `http://robotics.eecs.berkeley.edu/bear`.

[41] R. Vidal, S. Rashid, C. Sharp, O. Shakernia, J. Kim, and S. Sastry. Pursuit-evasion games with unmanned ground and aerial vehicles. In *Proc. of IEEE Conf. Robot. Automat.*, Seoul, Korea, 2001.

[42] H.J. Kim, R. Vidal, H. Shim, O. Shakernia, , and S. Sastry. A hierarchical approach to probabilistic pursuit-evasion games with unmanned ground and aerial vehicles. In *40th IEEE Conf. on Decision and Control*, 2001.

[43] A.Y. Ng and Michael I. Jordan, PEGASUS: A policy search method for large MDPs and POMDPs, *Uncertainty in Artificial Intelligence*, pp. 406–415, 2000.

[44] J. Bagnell and J. Schneider, Autonomous Helicopter Control using Reinforcement Learning Policy Search Methods, In *Proc. of the Intl. Conf. Robot. Automat.*, Seoul, Korea, 2001.

[45] M. Kearns and Y. Mansour and A.Y. Ng and S. Singh, Game Theory with Restricted Strategies, Unpublished manuscript.