An Integrated Approach to Multi-Robot Systems for Tactical Urban Missions

Response to DARPA BAA 98-08
Phase I, Part A

Principal Investigator:
Professor Shankar Sastry
Co-Principal Investigators:
Professor Randy Katz
Professor Jitendra Malik
Professor Christos Papadimitriou
Professor Stuart Russell
Co-Principal Investigators at Army Research Laboratories:
Dr. Jose Salinas
Key Researchers:
Dr. Jana Kosecka

Department of Electrical Engineering and Computer Science, and
Electronics Research Laboratory
University of California Berkeley,
Berkeley, CA 94720-1774.

Contents

1 Executive Summary

U.S. Special Forces are trained to carry out missions in hostile urban environments, including reconnaisance and clearing of urban targets. In the years to come it will be possible to reduce casualties on such missions dramatically by deploying a team of reconfigurable autonomous mobile agents (either ground-based mobile robots or airborne autonomous robots, or a combination of the two). Individual robotic agents should be able to operate autonomously or semi-autonomously (that is, in a teleoperated mode), gather information about the environment using multiple sensors, build maps of unknown and possibly time varying buildings or terrain, and coordinate their actions by communicating with each other, with a "mother ship" (which may be an airborne robot), or with human agents or commanders. While there is extensive doctrine about the conduct of special operations today, we feel that the use of autonomous and semi-autonomous unmanned agents offers possibilities for defining some new capabilities and missions. In this project we identify certain key research problems and tasks whose solution must be an important part of the development of such capabilities, such as: integrated perception of the environment; association and matching of objects perceived by different agents; decentralized map-making; team identification of evasive moving targets in a partially known environment; navigation of partially known and changing terrain; hierarchical mission planning and decentralized control. We propose to put together a research team that is uniquely qualified for this important and difficult research project, in that *each of its members has been involved in the successful solution of the previous generation of at least one* of the research problems above. The proposed research team involves people with expertise in robotics, computational vision, computational geometry, artificial intelligence, wireless communication and networking, robot control, and decentralized decision making, as well as engineers and researchers at the Army Research Laboratory with special expertise in intelligent systems, terrain building, mobile robots and Army doctrine.

In this project we will develop tools, software, and test simulations and experiments to provide these capabilities. Our research is organized along the following seven thrust directions, with roughly the first three to be completed in Year 1, and the others in Year 2:

1. **PER**. An Integrated Approach to Multi-Sensor Environmental Sensing.
2. **MAP**. Decentralized Map Building, Communication and Navigation.
3. **MULTI**. Multi-agent coordination, decentralized decision making and planning, fault handling.
4. **MAPPER**. Multi-sensor, multi-agent based map building.
5. **DEPLOY**: Terrain, urban environment based military (Army) protocols for deployment.
6. **MULTIMAP**. Multi-agent based map building, gaming strategies to deal with hostile agents.
7. **PERMULTI**. Hierachical control of perception, degraded modes of operation, hierarchical mission planning based on sensory data.

2 Program Plan

The deployment and operation of robotic agents in hostile urban environments is an important and difficult technological task, whose research underpinnings are at this time not fully mature. Such agents will be mobile, be equipped with multiple sensors and communication devices, and will operate either autonomously or semi-autonomously. They will be deployed by special operations forces in urban areas or other terrains that need to be explored and cleared of hostile elements. The deployment decisions will be based on a combination of prior intelligence and on-line information gathering by the special operations forces or by airborne robotic agents. The deployed agents must create models and maps of the unknown environment and identify human targets; to do so they must navigate and search the environment, communicate to reconcile their individual views and maps, and coordinate their planning, decisions, and actions. There are four large research problem areas that are essential to bring the capabilities of multiple mobile robot systems, each consisting of a series of research questions that will be considered at this time. These four research thrusts and the questions to be answered are as follows: (they are listed from the physical layer up through the decision layer)

1. **Integrated Perception of the Environment.** How are the images captured by the robots' sensors decomposed into objects such as walls, windows, and humans? How are key attributes of each to be extracted? How do the robots solve the problem of associating objects seen by each as being the same object?

2. **Decentralized Map-Building and Navigation.** How can a robot carry out its map-making task expeditiously, stealthily, and accurately? (What are the trade-offs between these three aspects?) How can the robots partition and dynamically reallocate the map-making task? How can a team of robots identify reliably all enemies present in a partially known environment? How can navigation be compiled down into a sequence of simpler positioning tasks?

3. **Multi-Agent Cooperation, Planning, and Decentralized Control.** What are appropriate architectures for planning and decentralized control by a team of distributed agents in an unknown environment? What are the trade-offs between communication (minimized for stealth and security) and accuracy of reconaissance? How does the team reconfigure in the face of agent loss?

4. **Deployment Decisions** How is prior intelligence on an urban environment reconciled with real-time aerial observations? What exterior information about particular buildings is necessary and useful to the reconaissance robotic team?

The basic project is partitioned into these four tasks. The first three tasks have partial overlap, since successful accomplishment of their individual subtasks is directly related to the urban missions in indoor environments (for example, building clearing, reconaissance) mentioned in the BAA98-08. The research directions will concentrate on perception and autonomy and demonstrate them using mobile robot bases. We will not focus on the locomotion capabilities on the ground needed to demonstrate the capabilities that we develop to accomplish our goals. The fourth task addresses the issues relevant to outdoors monitoring and exploration by airborne platforms.

2.1 Scenario

The basic perceptual operation that this robot will perform is to produce a *topo-sketch* of its environment. Good generic image segmentation techniques need to be employed in order to successfully identify different objects. The advances in algorithms for recognition and grouping of the information acquired by visual sensing make it possible to employ vision both in hazard detection and detection and classification of familiar and unfamiliar objects (e.g. other members of the (friendly) robot team, enemies, obstacles). The vision techniques for identifying and tracking human targets, which are currently done at the level of individual joints, can be used to further identify their intent or state (for example alive, moving, wounded, etc.). In order to assess their location and numbers, re-recognizing people seen before is an important issue (the association problem).

The topo-sketch is a dynamic representation over space and time. The sensing robot could change its position as it moves around, and even for a stationary robot, the environment would change as the dynamic agents move about. We envision that each robot would have a private copy of a data structure representing its topo-sketch with as accurate absolute location and time stamp information carried along as available. The location information would be derived from inertial navigation and is therefore expected to become unreliable over time (due to sensor drift) unless upgraded in quality from fixes from known landmarks. A key representational question is that of the nature of the topo-sketch: is it 3D or a collection of multiple 2D views? We expect to use a hybrid approach — a crude 3D structure to which are linked a set of 2D views. On obtaining a topo-sketch in static or dynamic environment, the agent needs to choose particular exploration strategy. This choice will depend on the available knowledge of the environment (or stored intelligence), the available time constraints (so-called "anytime" algorithms) and the presence of foe or friendly forces. We will develop so-called "master algorithms" where constraints are not an issue and scaled down versions of them to fit an any time scenario (when stealth is an issue) or a scenario where the robot has to move or take cover to prevent exposure to hostile fire.

Different robots would have different private topo-sketches. They could communicate with each other to acquire common knowledge about a larger part of the environment than each one has privately. Communication could be encrypted and authenticated using standard techniques to ensure trustworthiness. The chief tradeoffs between individual exploration and communication to benefit from the exploration of others are those of trading

1. Communication time and reconciliation time (forming a common representation from the topo-sketches of multiple robots) vs. time spent in continued individual exploration.
2. Increased probability of detection with greater communication time. In different contexts, this might impose more or less severe penalties.

The reconciliation process has two aspects:

1. Static objects that can be matched across multiple topo-sketches help provide fixes to the drift in the Inertial Navigation Systems (INS) or odometers of the individual robots.
2. Dynamic agents pose interesting challenges — one should expect to see some of the agents reappear in different topo-sketches as they move about, but there also some

that are only found in individual topo-sketches. The matching process has to be probabilistic and what one wants to obtain are the best maximum likelihood or Bayesian estimates of variables such as the numbers of particular kinds of agents in various partitions of space-time.

When the robots collaborate the key means of collaboration is communication, the kinds of hierarchies that can be proposed are:

1. Communicate with nearby robots, This has the advantage that there might be landmarks seen by both robots that permit a more accurate combined topo-sketch. That also leads to a more compressed representation to be sent on in a tree-structured hierarchy. There are relatively few compression possibilities in the disjoint topo-sketches of robots that are far apart in an environment.
2. Communicate directly with the "mother ship" and have it transmit the reconciled information (or subsets of it) back.

Of course these are two end points in the design space. We would expect to study heterarchical information structures, with intermediate communication architectures.

2.2 Team Expertise

Our research team is uniquely qualified to take on and make large strides in these difficult research problems. In fact, each member has already been involved with the successful solution of at least one, in some cases several, of the previous generations of these problems. Randy Katz is an expert in multi-media, reconfigurable wireless networks for indoor environments. He has developed compression and communication protocols for flexible networking in urban areas. Jitendra Malik is a leader in low-level and intermediate vision, with recent work in the crucial aspects of image segmentation, association, grouping, and attribute evaluation ([23, 24, 3, 26, 5, 17]). Christos Papadimitriou works on algorithms and complexity, and was involved in some of the first work on on-line terrain exploration and navigation [20], as well as searching for evasive targets; he is also investigating the communication/quality of decision trade-offs [29]. Stuart Russell is an expert in AI with expertise in decision making under uncertainty, learning as well as well as with Bayesian techniques for association and matching of objects observed by different agents at different places and times [8]. Jose Salinas has expertise in military doctrine for the deployment of multiple mobile robots at the Army Research Laboratory. He has also developed the Combat Information Processor for use in battlefield awareness on personalized soldier displays. Shankar Sastry specializes in the decentralized control of distributed systems [27] and hybrid design and verification techniques [19]. He has applied his methods to problems in automated highway systems, air traffic management systems and robotics. Jana Kosecka has previous experience and expertise in sensing for navigation and map building [14], control strategies for basic navigational capabilities [16] and has been recently involved in project on high-speed vision based automobiles [10]. More generally, our team is unique in its collaboration with the Information Sciences and Technology Directorate of the Army Research Laboratory (ARL) with major thrusts in the area of intelligent systems, information processing and mobile computing and communications. ARL is the Army's corporate research laboratory for technology generation and transfer of technology to deployable military systems. We have had a long standing collaboration with several personnel there and in particular with Dr. Chandra, the Deputy Director of Information Sciences and Technology and Director of Atmospheric Research. He will be making available his services at no-cost to this project.

3 Technical Approach

3.1 An Integrated Approach to Environmental Perception

Following our bottom up approach to the integrated architecture of multi-robot systems, we will first elaborate on perception component. The effort that we wish to undertake within this task will concentrate primarily on visual sensing. By visual sensing we include not only optical cameras, but also infra red cameras, and other two dimensional sensors. We will elaborate in later sections on the use of other sensing modalities (e.g. INS, odometry, sonar), which will be closely tied with the vision capabilities during the navigation and map building. For the perception component, more specifically the techniques and algorithms which need to be developed fall into the following categories.

Segmentation

We [23, 24] have developed a novel approach to the problem of image segmentation by modeling it as the process of finding a partition of the image into regions such that there is high similarity within a region and low similarity across regions. Rather than focusing on local features and their consistencies in the image data, our approach aims at extracting the global impression of an image. This is made precise as the "normalized cut" criterion which can be optimized by solving a generalized eigenvalue problem. The resulting eigenvectors provide a hierarchical partitioning of the image into regions ordered according to salience. Brightness, color, texture, motion similarity, proximity and good continuation can all be encoded into this framework. Results on complex images of natural scenes demonstrate the significant superiority of this technique over classical approaches such as those based on edge detection, Markov Random Fields etc. Figure (1) shows results on segmentation of a challenging natural scene. Note that we are able to find regions, which correspond directly to surfaces or parts of objects in the scene. These directly enable us to determine obstacles, open spaces and the like instead of having to perform complicated inference on edge outputs. For many years, lack of reliable image segmentation has been a bottleneck in the effective use of computer vision; we believe that our results represent a breakthrough and make possible a new generation of mobile robotic perception strategies. We can only include a few examples here due to space limitations. More can be found on the web site `http://HTTP.CS.Berkeley.EDU/~jshi/Grouping/overview.html` We are also able to segment scenes using multiple view analysis–see Figures (2,3).

Further Grouping Using Object Class Models

The proto-regions extracted by the normalized cuts algorithm need not correspond in a one to one fashion with the objects of interest. In the motion example above, the person is broken up into regions corresponding to the face and the torso. Based on low-level measurements of color, texture or motion, this is indeed the right decision. High-level knowledge of object

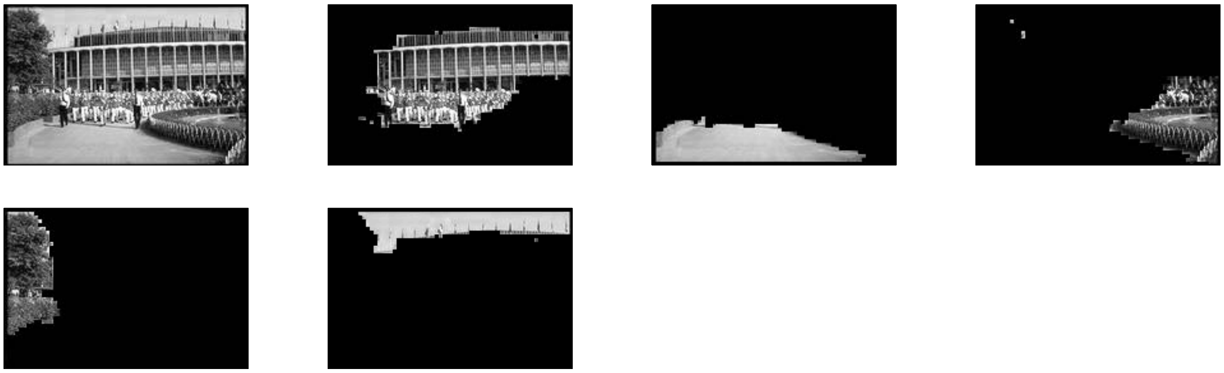

Figure 1: Segmentation of a building scene. The original image is on the top left.

classes, such as humans, walls, windows, doors, etc. can be used to group together proto-regions. For geometric scene elements such as walls or floors, the problem is one of fitting consistent geometric models. For any of these planar elements, monocular analysis (e.g. from vanishing points) or multiple view analysis (e.g. from textured patches which permit stereopsis or structure from motion reasoning) can be used to determine the plane and coplanar continuation used to group the elements together. The idea generalizes to other object classes with well-characterized geometry.

Tracking Dynamic Agents

We are interested in tracking dynamic agents — these could be other robots, vehicles, or people. In contrast, tracking rigid objects, such as vehicles, is essentially a solved problem. At Berkeley, we have demonstrated a real-time system [2] for tracking cars, that deals with significant occlusion due to traffic congestion and operates over a wide variety of lighting and environmental conditions. This system has been tested over many hours of video and its performance compared with ground truth on more than 40,000 vehicles. The more challenging problem is that of tracking people. When they are far away, one can model them as approximately rigid and track them in a standard way. When they are closer to the camera, the possibility of tracking individual joints becomes available permitting inference of behaviors and status for example, walking or running, healthy or wounded.

Tracking people at the joint level is a difficult problem for several reasons:

- Low contrast between the different parts of the body,
- Significant amounts of self-occlusion,
- Large angular velocities which result in significant movement from frame to frame, and
- Cluttered environments.

We have developed a novel approach [3] to this problem that has given us the ability to track individual joints and the global 3D pose at high accuracy with just a single camera. (Approaches such as those using a sea of sensors seem inapplicable in the current setting.)

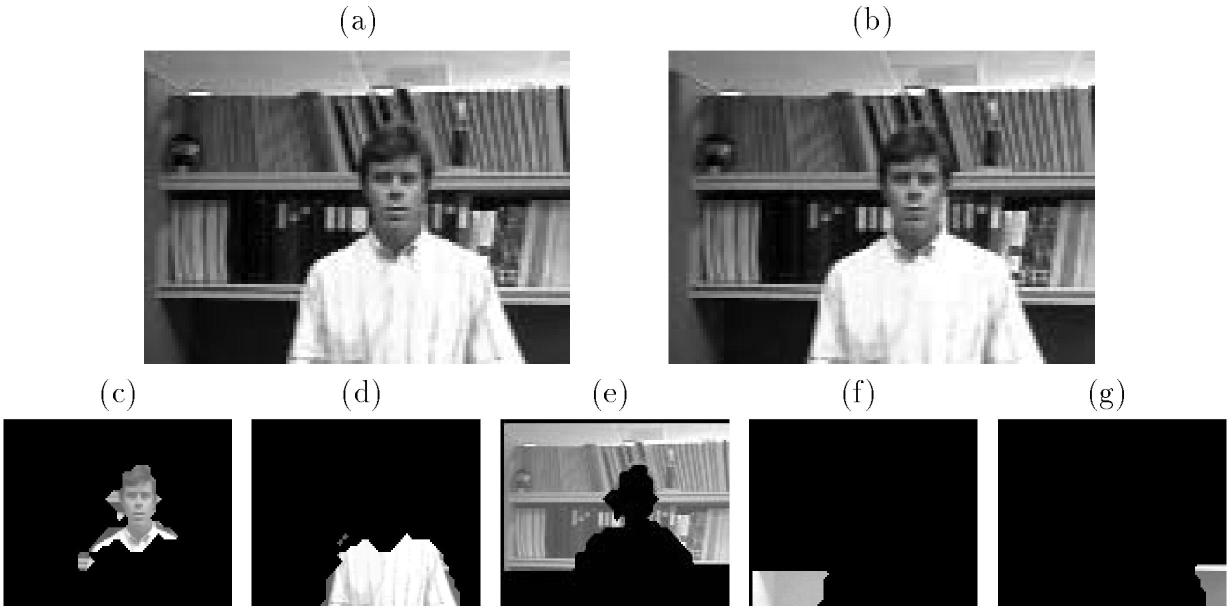

Figure 2: Subimage (a) and (b) show the left and right image of a stereo image pair, and segmentation is shown in subimages (c) to (g).

The idea is to exploit the mathematical formalism of *twists* and the *exponential map* developed by one of us ([22]), for representing and reasoning about kinematic chains. This can be integrated into a region-based tracking technique and leads to the problem of updating the joint angles and pose becoming a simple *linear* problem. The image velocity of a chain segment is a linear combination of the global pose change, and angular velocities of the kinematic chain. The kinematic model itself can be estimated during this process. We show some example results of tracking in Figure (4). Note that we are able to do joint-level tracking which enables us to infer behavior — walking, running, wounded, for example. It is important to note that our methods do not rely on the subjects having special clothes or markers — techniques which rely on such features are common in computer vision literature, but are obviously irrelevant in a special operations theatre.

3.2 Decentralized Map Building, Communication and Navigation

Data Association

The techniques described in Section 3.1 allow each robot to detect and track a variety of objects such as humans, "landmarks" in buildings, and other robots. Over time and with multiple robots, a large set of detected objects can be accumulated. *Data association* is the problem of deciding which of these objects are, in fact, the same object, and which are distinct objects. Consider the following cases:

- Robot A explores several complex, interconnected hallways in a building and detects, at various times, 3 stairwells. Depending on the building topology (possibly unknown), the odometry and compass error (large), and the appearance of the stairwells and

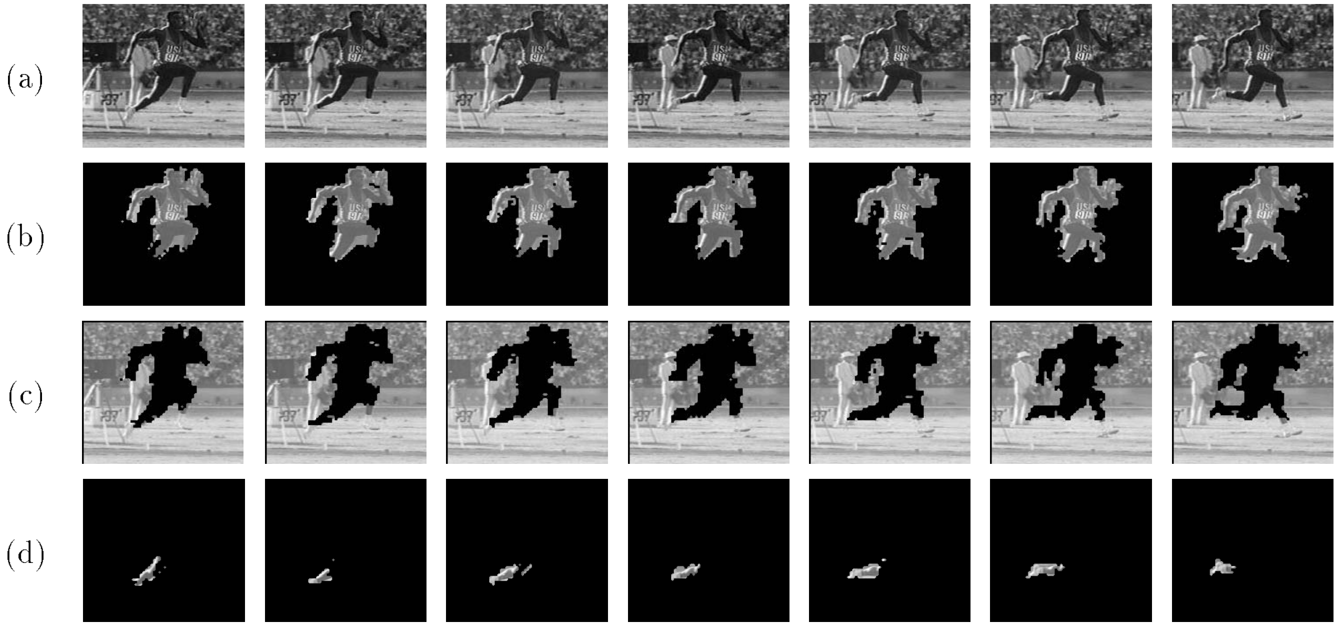

Figure 3: Row (a) shows an image sequence of Carl Lewis running. Notice that the background is moving to the left as the camera is panning to keep the runner in the center of the image. Consequently, background subtraction would not work as an image segmentation technique. Row (b) to (d) show the motion segmentation produced by our algorithm. Note these regions found corresponds the runner in (b), moving background in (c), and the left lower leg in (d). The left lower leg is segmented from the runner because it undergoes significant upward rotation in these seven image frames. By recursive cuts the other moving limbs can be found.

neighboring landmarks (possibly poorly characterized, and possibly changing over time in a combat situation), the robot must decide if these 3 locations are distinct or not in order to obtain a reliable building map.

- Robots B and C make, at various times, a total of 12 and 17 observations, respectively, of humans in the building. At no time does either robot observe more than 2 humans simultaneously. Depending on the humans' appearance, their estimated mobility, building geometry, and so on, there may be anywhere between 2 and 29 humans in the building (in addition to some unknown number of as-yet-undetected humans).

Traditionally, data association has been used in situations such as radar tracking of multiple missiles, where observations are fairly continuous, target descriptions are low-dimensional vectors, motion of targets is fairly predictable (locally linearizable), sensor models are constant, and targets redetected after a significant interval of nondetection can safely be treated as new targets [1]. None of these assumptions are valid in the current problem, and novel techniques are needed. A start on Bayesian techniques for association and matching of objects observed by different agents at different places and times has been made in our IJCAI best paper prize work [8].

Figure 4: Example configurations of the estimated kinematic structure during one walk cycle of a person seen from an oblique view.

Decentralized Map-Building

Building the map of an unknown 2-D scene from visual data by a moving agent is a well-studied problem in on-line algorithms, computational geometry, and robotics. The goal is typically to build the map as fast as possible, with very little motion and computation when the environment is simple, but allowing more motion and computation when the environment is complex. There is an extensive literature on this problem after it was posed. The present problem, however, is further complicated by several factors. In the list below, we describe how problems in task partitioning, survivability and stealth issues, dynamic environment and communication costs will influence the development of our new algorithms.

1. **Task Partitioning**
   The map-building task must be partitioned between the agents, and reassigned dynamically. The agents start with a default partition of the environment as they "fan out," each following its own, deliberately differentiated, motion plan. For example, one robot may start by tracing the left wall, another the right, a third heading to the second opening from the left, etc. Each robot engages in map-building of the part of the environment it perceives through its sensors and the integrated perception mechanism, following one of the known effective algorithms for this task (such as the heuristic "head for the furthest clockwise point that has ever been visible" known to always be within a factor of two of the ideal optimum map-building plan [25]). Each robot will resolve the association problem regarding to objects it sees at different times, and robots may communicate for resolving their intra-agent association problem. However, the robots may have to communicate for another reason: The parts of the environment assigned to each robot by their default partition may differ dramatically in their complexity, as well as in the degrees to which they overlap. When such imbalances are discovered in the course of decentralized map-building and communication, the map-building task must be re-partitioned for balance and overall effectiveness.

2. **Stealth and Survivability Issues**
   Navigation of the environment for the purpose of map-building may be undesirable for reasons of survivability and stealth. Each elementary cell in the terrain to be navigated can be assigned a *cost* commensurate with the undesirability of navigating in it. For example, the cost could be low if it is currently known that the cell is not under enemy surveillance, moderate if it is not known to be free of enemy surveillance, and highest

if it is known to be under enemy surveillance. The cost may also vary with the tactical desirability of stealth and safety at the time. The robot must therefore solve a version of the optimum map-building problem in which the terrain is weighted with possibly time-varying costs, and we must minimize the total weight of the path traversed, as in [18].

3. **Dynamic Environment**
   Building maps of static elements of the environment is the easier task. It is much harder to identify all mobile parts, especially if they are trying to evade such detection. This more complex identification task is closely related to *clearing* the environment of enemy presence. It is known that a single robot will not in general suffice, and the smallest required size of a team depends crucially and subtly on the topological properties of the environment [4]. A weak team strategy may have as a result *recontamination*, that is, the opening up of the possibility that one or more hostile agents have moved to an area that was previously cleared and known to be free of hostile agents; this has grave consequences on the time required for the search. There are algorithms known for optimally clearing complex environments, without recontamination [4]. However, the problem of clearing of a geometric environment in the absence of a complete map has not been addressed in the literature.

4. **Communication Costs**
   Communication between the robots during the mission may have to be minimized for reasons of stealth. In all three subtasks identified so far (association, decentralized map-building, and environment clearing) there seem to be intricate trade-offs between the allowed communication and the quality and efficiency of the task (quality and reliability of the map built, efficiency of the environment clearing strategy, completeness of the association). Since communication may be undesirable in varying degrees depending on tactical considerations, we will have to study many points of this trade-off. Communication–quality of decision trade-offs have been studied extensively in other contexts, see for example [29].

Wireless Communications and Networking in Urban and Indoor Environments

When there are multiple robots which are operating covertly, they will need to communicate to achieve coordination. However, the level and quantity of the communication with each other as well as with the “mother ship” needs to be minimized for stealth. We will work on architectures for configuring the communication between the robots using flexible wireless networks, which are reconfigured based on mission specifications. An important area of research in the multi-robot system is wireless communications in the urban battle space is to provide network connectivity for robots and mobile sensors. We propose to develop new communications protocols for mobile sensor platforms that are both power efficient and can operate in a stealth mode.

Reliable transport layer protocols, such as the TCP protocol suite, have been developed to insure reliable communications over a wide diversity of access technologies. It has long been recognized that TCP does not perform well in an environment characterized by high loss rates, such as found in wireless environments. A variety of techniques have been developed to improve TCP performance over wireless links, but it remains that such techniques require power-consuming acknowledgements and frequent retransmissions. An issue for military operations is that existing protocols based on acknowledgements have the weakness that the mobile node must radiate energy in order to complete the reliable transport layer protocol. Not only does this consume precious energy, it also reveals the location of the mobile node, which may be lethal on the battlefield.

We propose to develop a radical alternative to traditional transport layer protocols, based on protocols that are application- and power-sensitive. In particular, we propose to develop pseudo-reliable transmission schemes that tradeoff bandwidth in the form of multiple (re)transmissions of a packet against acknowledgement-based explicit retransmission requests. The idea borrows from pager networks, which improve the probability of successful packet delivery by retransmitting packets more than once.

Situation reports, which are frequently updated, can be sent in this pseudo-reliable fashion. If the status packets are lost, a new report will be generated in the near future. The receiver simply receives the broadcast report without explicit acknowledgement. Command messages, such as "take the hill", targeted for a specific mobile node, must be acknowledged. Thus, it is important to understand the applications semantics to determine whether pseudo-reliable transport is appropriate or not. It is inherent in the pseudo reliable scheme that there will never be a 100 % guarantee that a transmitted message will have been successfully received.

Simply sending a transmission multiple times is a rather naive solution to the problem. More sophisticated techniques will be investigated. One idea is borrowed from RAID. An 'n packet' packet stream can be appended with an additional $n+1$ parity protection stream. Should one packet be lost, its content can be reconstructed from the $n-1$ received packets as well as the parity packet. Within this context, it will be important to develop the protocol suite so that it is adaptive to the nature of errors on the underlying communications channels (burst errors vs. continuous errors), as well as sensitive to the requirements of the overarching sensor application.

Visual Servoing for Navigation

Navigation capability is an integral part of both map making and reconaissance missions given a map. In the map making context, there is a need for implementation/development of motion plans which would reliably position the robot between visibility regions associated with explored places, while maintaining accurate pose estimates of the robot relative to the signatures associated with the places. Similarly in the second stage, after obtaining a "topo-sketch" of the environment, navigation between neighboring places needs to prone to errors in odometry. This can be accomplished by incorporating additional sensing capability (in addition to the odometry) in a control loop of the mobile robot.

We propose a novel approach to global navigation utilizing the idea of visual servoing. By choosing particular servoing primitives, represented as a part of the qualitative map of the environment we try to formulate the global navigation problem as a sequence of relative positioning tasks and design stable closed feedback loop control strategies for various relative positioning tasks. Then assuming availability of a topological model of the environment represented in the form of a directional place graph and with the arcs as the control strategies which need to be invoked in order to go from one place to another, more robust plans can be generated and directly executed. The visual servoing framework can then combine the geometry of the environment and the control of the mobile robot in such a manner that global navigation can be done in a robust and reliable way.

Pose Estimation

A mobile robot requires both a means of determining its position and controlling the vehicle position. For control purposes and map generation task, it is desirable to know the position and additional vehicle state information (i.e. attitude, velocity, angular rates, linear acceleration) at high sampling rates. Odometer systems frequently used for determining position of mobile robots, can not provide the full-state information (i.e. accelerations, attitude), and suffer from severe drift errors. Also, odometer systems can not provide enough information in order to determine accurately such parameters as pathway inclination, turning radius and direction.

The Global Positioning System (GPS) is the most convenient and accurate method for determining vehicle position in a global coordinate system. The accuracy of GPS position estimates is on the order of 50 meters. Increased accuracy on the order of 10 centimeters can be achieved by using Differential GPS (DGPS). Moreover, DGPS system with carrier phase measurement capability can provide accuracies of 1 to 3 centimeters. However, GPS is limited by the update rate and the occlusion of satellite signals by natural and man-made features.

Inertial Navigation Systems (INS) have been developed to provide complete vehicle state information at rates suitable for accurate vehicle control. An INS system integrates the differential equation describing the system dynamics for a short period of time. During the integration step, the variance of the state estimates increases. Therefore, INS states have to be periodically corrected.

The traditional approach to designing integrated navigation system is based on the combination of (1) determining absolute navigation parameters (coordinates and velocity), which are obtained from the GPS, and (2) determining relative navigation parameters (relative changes) in coordinates and velocity), which are obtained from the INS. At Berkeley we have developed and tested algorithms for an inexpensive gyro-free INS system in combination with accurate DGPS system, which with effective sensor fusion and fault detection techniques can provide accuracy, speed, and stability required by most outdoor autonomous mobile vehicle applications (when DGPS signals are available).

Taken independently, each of the above position sensing techniques (INS, GPS, and vision) has advantages and disadvantages. The strongest criticism of any of the individual techniques is that it is susceptible to a single point failure. No single reference system can supply adequate reliability and availability to ensure safe control. However, used jointly the overall system performance and reliability can be improved. It is important to note that sensor fusion is accomplished by using the different kinds of sensor data in different ways.

It is anticipated that representative data will be made available to verify the capabilities of the system developed under this effort. We shall demonstrate the performance of the techniques in an experimental setting, and shall address issues related to the speed of processing. The techniques developed to fulfill the requirements of this topic will provide a flexible solution to many problems in surveillance, robotic navigation, targeting, and multi-sensor robot exploration. The developed system will provide capabilities to assemble, update, and maintain geometrical/geographical information databases using imagery and navigation information collated from different sensor systems.

3.3 Agent Coordination, Planning and Decentralized Control

The issues addressed in the previous tasks were focused on the sensing, communications, tracking, recognition and map building. These are the enabling blocks for the team of robots, to be able to deal effectively with issues of how to plan the coordination of the multiple robots. This is addressed in this section. We shall explore the use and choice of general design principles for complex distributed multi-agent systems and instantiate them on the problems which need to be solved along in order to organize the team of special op robots. The findings of the research conducted here then in turn provide some guidelines, and possibly introduce additional constraints into the design process for the *utilities* that are developed for sensing, communications, recognition, tracking and map building for the mission organization.

Distributed Decision Making

We wish to explore new design paradigms for distributed control and decision making which would avoid the high communication and computation costs of centralized control, and would allow the individual participants of the mission to operate much closer to their performance limits and achieve greater efficiency of resource use. We believe that such a challenge can be met by organizing the distributed control functions in a hierarchical architecture that makes those functions relatively autonomous (which permits using all the tools of central control), while introducing enough coordination and supervision to ensure the harmony of the distributed controllers necessary for high performance. In this case, each agent is trying to optimize its own usage of the resource and coordinates with “neighboring” agents in case there is a conflict of objectives. Methods from hybrid systems (combining discrete protocols, continuous dynamics, and hierarchies of sensing and action) analysis and control, along with work in distributed algorithms, plan merging and consistent update are needed to guarantee safety (survivability in a military context) of individual physical agents, and optimality (mission completion) of the composite system. Such a scheme also insures that while the physical agents can have a high degree of autonomy, they operate only within the freedoms and bounds set for them by their human soldier operators.

Hierarchical Mission Planning and Execution

Our approach to planning addresses the two main deficiencies of AI planning algorithms:

1. They do not account for the uncertainty inherent in real world problems.
2. They do not reason hierarchically.

Traditional AI planners construct plans under the assumption that the effects of actions are deterministic. Extensions to this planning paradigm permit hierarchical planning in which abstract, high-level actions like moving from one building to another are mixed with low-level actions like moving forward a single step. The abstract and low-level actions then can be combined to form an overall plan. The problem with this approach is that the plans it produces are difficult to use in the real world, when actions often can fail to have the desired outcome or can inadvertently undo the effects of other actions.

The Markov Decision Process (MDP) framework permits optimal planning under conditions of uncertainty. Unlike traditional planning methods, the MDP framework has not yet been extended to permit hierarchical reasoning. MDP algorithms are forced to deal exclusively with low-level actions, a limitation that has hampered efforts to scale planning under uncertainty to large problems. The Hierarchical Abstract Machines (HAM) approach [21] makes this leap by introducing a language that permits the introduction of hierarchically structured constraints into the MDP framework. For example, the system can be told that the overall strategy must involve a sequence of movements from building to building and that the next building should not be visited until an objective is achieved in the current one. In [21], both offline and online (reinforcement learning) algorithms are given. However, some technical advances are needed to allow these methods to scale up to the kinds of problems faced here. First, we need to extend the HAM language and algorithms to handle *partially observable* environments. The main requirement here is that each agent maintain some representation of its uncertainty about the current state of the environment. We will also need to incorporate generalized value function representations that enable computed (or learned) substrategies to be applied across a wide variety of situations. For example, when an agent has devised a method for scanning a rectangular room, it should be able to apply this method in future to any other rectangular room. Strategies for obtaining smaller objectives within buildings can be encoded directly into the HAM framework at varying levels of specificity. However, partial plans or strategies with pieces that are incomplete can also be included. Our algorithms will find the optimal strategy that is consistent with the provided structure. The more structure the system is given, the faster it will run. Constraints like rules of engagement actually will speed the system up.

Architectures for Decentralized Control

In distributed, decentralized control applications of the kind found in C^4I systems, it is important to be able to evaluate hierarchies and heterarchies of control architectures for the following reasons:

- Degree of centralized vs. decentralized decision making
- Safe performance of individual agents to be reconciled with optimal performance of the composite system.
- Failure mode detection, fault handling and survivability of the system after faulted modes.

To elaborate on these points, consider that to achieve the common optimum we should ideally have a centralized control scheme that computes the global optimum and commands the agents accordingly. A solution like this may be undesirable, however, for several reasons: (1) it is likely to be very computationally intensive, as a large centralized computer is needed to make all the decisions; (2) it may be less reliable, as the consequences may be catastrophic if the centralized controller is disabled; (3) the information that needs to be exchanged may be too demanding of communication resources; and (4) the number of agents may be large and/or dynamically changing.

If the performance degradation of a completely decentralized solution is unacceptable and a completely centralized solution is prohibitively complex or expensive, a compromise will have to be found. Such a compromise will feature semi-autonomous agent operation. In this case, each agent is trying to optimize its own usage of the resource and coordinates with "neighboring" agents and a base station in case there is a conflict of objectives. It should be noted that *semi-autonomous agent control is naturally suited for hybrid designs.* At the continuous level, each agent chooses its own optimal strategy, while discrete coordination is used to resolve conflicts. Thus, the class of hybrid systems that we will be most interested in are semi-autonomous multi-agent systems, where the hybrid dynamics arise from the interaction between continuous single agent "optimal" strategies and discrete conflict resolution or coordination protocols.

The need for rapid and varied agent and inter-agent reconfiguration, the uncertainty of information, its localization and the inter-agent or agent-base communication required to optimize global performance in the face of such localization, result in emergent behaviors of extraordinary complexity. Once again, our experience with the design of such systems suggests that the management of complexity demands that a theory of multi-agent control embed within it a theory of hierarchical control and a corresponding theory of hierarchical abstraction of information. Needless to say, the theory must locate and manipulate within its discourse both planning and decision making layers designed by soft computing techniques as also layers that quickly and decisively translate decisions or plans into deterministic sequences of agent and inter-agent actions that are formally designed.

There is a continuum of design choices for system decomposition, ranging from strict hierarchical control to a fully distributed, multi-agent system. Furthermore, different choices may be appropriate at different levels of abstraction, ranging from the (typically continuous-domain) low-level control systems concerned with safety and smooth execution to the (typically symbolic/discrete) strategic levels concerned with optimization and planning for high-level goals. We will investigate theoretical and design issues involved in the choice of system architecture, and methods for interfacing elements of the resulting hybrid system.

4 Statement of Work

The primary deliverables for this research project will be in development of algorithms, software for simulation and visualization of the results and algorithms and software for demonstration of the experimental results.

4.1 Year-One Program

The basic plan is to have the development of tools, algorithms, simulations and experiments in the first year of the grant separately in the following basic thrust areas:

PER Integrated Approaches to Environmental Perception.

MAP Decentralized Map Building, Communication and Navigation.

MULTI Multi-agent Coordination, Planning and Decentralized Control, including decentralized control for individual survivability, and centralized control for overall mission success.

DEPLOY We will develop enhancements of the ARL Combat Information Server to add to it a topographical, terrain feature, and urban feature server,

The detailed statement of deliverables in these three thrust areas labeled PER, MAP and MULTI follow in Section 5.

4.2 Year-Two Program

In the second year, we will have three separate efforts involving tools, algorithms, simulations and experiments combining the thrust areas pair by pair as follows:

MAPPER We will integrate the algorithms for sensing the topology of the environment using vision, ultrasonic and other sensors with those for building and representing maps. The maps will also have facilities for association and tracking of multiple moving targets.

MULTIMAP We will integrate map building and navigation with the gaming strategies such as Stackelberg, Nash to enable exploration of the new terrain while guarding old terrain (i.e. updating estimates of what it contains). Networking protocols will also be tested. We will build hierarchical dynamic representations of the environment capturing the geometric and topological properties of the static entities and maintaining the consistent representation of the dynamic targets.

PERMULTI Especially germane here will be the development of algorithms for transmission of the minimal amount of information for decision making. We will investigate these issues in the context of visual sensing and map information. We would test the fault handling algorithms for the loss of sensors or individual agents. We would give probabilistic figures of merit of the composite functioning of the system involving perception, multi-agent planning, fault handling and mission reconfiguration.

A brief summary of the statement of work in all the thrust areas follows. More detailed description of individual subtasks is in Section 2.

4.3 Deliverables

Task PER

Subtask 1.1 *Segmentation.* Development of the segmentation algorithm into regions of coherent brightness, color texture and disparity and motion algorithm. Call these “proto-surfaces”. The algorithm will be tested on real image sequences acquired by mobile robot in the indoor environment.

Subtask 1.2 *Grouping.* Development of grouping criteria and algorithm which will taken into account some generic knowledge of the floors, walls, doors and windows and associate properties with these regions, such as slant, depth and spatial extent. The algorithm will tested on real image sequences obtained from the segmentation step.

Subtask 1.3 *People Identification.* Algorithms for identification of people. The algorithm would be employed on image streams during the exploration phase and would work in case of low resolution images and instances where people may be seen from distance. The testing would be performed under ordinary lighting conditions.

Subtask 1.4 *Target Tracking.* We will develop algorithms for real time tracking of people, using the new approaches for tracking linked rigid bodies.

Subtask 1.5 *People Tracking and State Assessment.* Development of an additional layer on the tracking algorithm which would be for identification of additional properties of the tracked target (for e.g., walking, running, unusual activities, female, male, etc.).

Subtask 1.6 *Feature Detection for Data Association.* Experimentation with the choice of suitably robust features for data association in a Bayesian framework, and integration of the association algorithms into visual tracking algorithms.

All the techniques and algorithms developed as a part of this task will be tested on real image sequences of a dynamic environment acquired either by a moving or static observer and tested on the mobile robot bases.

Task MAP

The Decentralized Map Building, Communication and Navigation task concentrates mainly of the geometric and tactical aspect of cooperative map building process and basic navigational capabilities. We plan to develop the following techniques:

Subtask 2.1 *Data Association.* The data association problem presented by the decentralized map-building task shares many characteristics with a vehicle identification problem, but also presents novel difficulties. We shall extend the vehicle association work to allow for the mobility of camera platforms (robots), and the much richer motion behaviors exhibited by humans, and complex environment geometries. This involves some deep technical issues, such as inferring *absence* of targets in some region based on nonobservation over time and space even when targets may be mobile. It will also require the use of dynamic probabilistic networks, which are a generalization of Kalman filters, in order to handle track complex behaviors (see [7]). For this domain, we will need to extend our real-time approximate inference algorithm ER/SOF [12] to handle the combined use of discrete and continuous variables required in this domain.

Subtask 2.2 *Decentralized Map-Building.* We shall develop and test algorithms for estimating the local complexity of an environment, for the purpose of load balancing the map-building process; such algorithms will have to partly rely on prior stochastic understanding of the characteristics of the environment being explored. We shall modify our efficient map-building algorithms taking into account the novel objective of identifying complex parts of the environment early and reliably. We shall develop algorithms for optimally partitioning a partially explored environment among agents already in the scene.

Subtask 2.3 *Weighted Terrain.* We shall develop models for assigning weights on cells of the terrain We shall extend the map-building algorithms to the case of non-uniform terrain costs.

Subtask 2.4 *Clearing the Environment.* We shall combine our understanding of the problems of exploring an unknown environment and clearing a known environment, to develop algorithms for the more complex problem of clearing an unknown environment.

Subtask 2.5 *Communication–Quality Trade-offs.* We shall re-examine the three problems of object association, decentralized map-building, and environment clearing from the point of view of communication requirements. We shall develop a theoretical understanding of the trade-offs between intensity of communication (which includes possibly encoded information) and performance in these three subtasks. We shall develop new algorithms for these subtasks, for a fairly comprehensive range of increasingly stringent regimes of communication.

Subtask 2.6 *Visual Servoing.* We will design control strategies for mobile robot relative positioning with respect to planar objects (e.g. walls, doors) in a static setting and corridors as well as strategies for vision based following of another dynamic target (foe or friendly). The designed strategies will be verified in simulation and experimentally on a mobile robot platform.

Subtask 2.7 *Pose Estimation.* We propose to integrate vision system into the GPS/INS based navigation system described in Section 3.2 for control and map generation purposes. The objective is to develop techniques for establishing the precise geometric relationship between the image sources and other position sensors. The development of algorithms to perform automated positioning of different data types with respect to a single coordinate system is a new and challenging task.

Subtask 2.8 *Protocol Design for Wireless Networks.* Protocol design and simulation modeling using the ns simulator. Protocol implementation; integration with applications being developed for the mobile sensor network. Initial deployment in testbed.

Task MULTI

The main trust of this task is investigation and instantiation of some basic design and planning principles in the problem domain of mobile robots involved in tactical urban missions. This will be done in the scope of following subtasks:

Subtask 3.1 *HAM.* Using doctrine and the results of our algorithmic analysis of the problem of exploring/clearing unknonw environments, we will devise HAM mission plans for a set of typical scenarios, and develop optimal offline and online solutions translating these mission plans into concrete behaviors. As part of this task, we will propose and implement schemes for handling partial observability and across-subtask generalization.

Subtask 3.2 *Architectures for Decentralized Control.* We will develop a discrete event model of a decentralized multi-agent system engaged in a clearing mission and formulate the distributed observation problem for fault detection, with fault events being unobservable by the other agents. The answers to the observability issues will guide the design of a communication scheme between multiple agents engaged in the mission.

Subtask 3.3 *Fault Handling and Reconfiguration of Mission.* When one of the agents or some of the sensors on board a platform is lost, an assessment has to be made as to how much of the mission can be accomplished or whether the mission should be abandoned for reasons of survivability. An architecture for fault management at the sensor or robot level will be developed.

Subtask 3.4 *Probabilistic Verification.* When a multi-agent system, with various protocols for communication, sensing, and map building and architectures for fault handling are put together, it is important to assess the overall performance. We will use a combination of probabilistic formal methods and methods developed for HAMs for this purpose.

Task DEPLOY

Subtask 4.1 *Topographical Elevation Server.* ARL's Combat Information Processor (CIP) has a topographical elevation server which provides a user with a map of a given geographic area. ARL will undertake the development of algorithms and make necessary modifications to the code so that the results are available for robotic implementation.

Subtask 4.2 *Terrain Feature Server.* The CIP also has a terrain feature server which supplies the users the features in a given geographic area. ARL will undertake the development of algorithms and make necessary modifications to the code for use in robotic implementation.

Subtask 4.3 *Urban Feature Server.* ARL will create a new "urban feature" server which will be used for querying the urban areas for the mission planning computations including inside building and city information.

Subtask 4.4 *Military Mission Planning.* ARL has previously developed single route planning algorithms for single robots in a special forces environment with detailed descriptions of the appropriate protocols. These will be modified to take for multiple robots with planning capabilities for urban city and inside building war fighting.

Task MAPPER

Subtask 5.1 *Integration of Decentralized Map Building algorithms with Perception Routines.* Evaluation of Map building algorithms with input from the perception followed by refinement and redesign.

Subtask 5.2 *Integration of Decentralized Map Building with Communication Protocols.* Evaluation of protocol design in testbed, followed by refinement and possible redesign.

Subtask 5.3 *Integration of Human and Other Visual Tracking Algorithms with Mapping Routines.* Implementation of dynamic map update for objects as well as mobile human targets.

Task MULTIMAP

Subtask 6.1 *Integration of gaming strategies with mapping routines.* Here we will update maps of the environment, using game theoretic strategies to "guard' territory while exploring new territory.

Subtask 6.2 *Hierarchical Abstract Models for representation of maps.* We will develop algorithms for storing prospected maps at various levels of granularity.

Task PERMULTI

Subtask 7.1 *Fault Handling.* When some of the individual sensors or robots are blinded or lost, we will test the algorithms for map update, mission replanning based on probabilistic assessments of data that we do not have access to.

Subtask 7.2 *Tradeoffs between Information and Stealth.* We will study the tradeoff between having more sources of information and the potential loss of stealth when transmitting. We will also study "semantically" based measures of information to be transmitted from rich sensors such as vision sensors for map building.

5 Milestone Chart

The deliverables listed below are classified as SW (algorithms and software design tools), or SIM (for simulation demonstrations) or EXP (for experiments). Year 1 experiments will be performed on mobile robot platforms that are recommended to us by DARPA or by platforms that are currently at ARL. At Berkeley, the current mobile robot platforms are helicopters (aerial mobile robots). While the importance of developing algorithms for aerial mobile robots is clear and in Year 2, we will perform experiments using one of the aerial mobile robots as the "mother ship", we will use terrestrial wheeled robots for the Year 1 experiments.

Task	Symbol	Duration	Deliverables
Integrated Approach to Perception	**PER**		
Segmentation, Grouping, Identification	**1.1-1.6**	4/98 - 4/99	SW, SIM
Tracking, Association		1/99 - 6/99	EXP
Map Building, Comm., Nav.	**MAP**		
Association, Map Building,	**2.1 - 2.8**	4/98 - 4/99	SW, SIM
Tracking, Communication		6/98 - 6/99	EXP
Multi-Agent Coordination	**MULTI**		
Architecture Evaluation	**3.1-3.4**	4/98 - 10/99	SW, SIM
HAMs, Degraded Modes		6/98 - 6/99	EXP
Combat Development	**DEPLOY**		
Terrain, Building,	**4.1-4.4**	4/98 - 4/99	SW, SIM
Army Mission Protocols		5/99 -3/00	EXP
Sensor Based Mapping	**MAPPER**		
Communication + Sensing	**5.1 - 5.3**	4/99 - 10/99	SIM
based Map Building		6/99 - 3/00	EXP
Multi Agent Mapping	**MULTIMAP**		
Game based cooperation	**6.1 - 6.2**	4/98 - 6/99	SW, SIM
HAM representation,		1/99 - 3/00	EXP
Multi Agent Perception	**PERMULTI**		
Fault Handling,	**7.1 - 7.2**	9/98 - 6/99	SW, EXP
Stealth + Communication		6/99 - 3/00	EXP

6 Facilities and Equipment

Most of the simulations and algorithms will be developed on Ulta Sparc 2 workstations and Silicon Graphics Indigo 2 and Octane workstations which we currently have in our laboratories. Visualization of the simulation for various mission executions will be written in Open GL or an extension of our X-windows based simulator. We are also planning to experiment with combination of odometry and INS for maintaining better pose estimates of the robot. We have some experience in this are through experimentation with an aerial robot based on a Kyosho helicopter model. Further facilities which will be used within a project are: CCD cameras, 4 processor Pentiums, Crossbow accelerometers, compact B/W CCD camera, Video Transmitter, Imagination Image Grabber, Futaba transmitter/receiver servo, some 200Mhz Pentium Notebooks. The real-time design software will be made available to us from other projects. Facilities available at ARL include 3 experimental lab bays which include the following equipment: 5 Sun Sparc 20 Workstations, 12 Ultra 2 Sun Sparc Stations, 3 Ultra 2 Sun Sparc Stations with enhanced 3D creator graphics, 2 Data Max Video 200 real time image processors, 1 Mercury Parallel SHARC VME system with 12 processors, 1 Mercury Parallel PowerPC VME system with 4 processors, 2 HelpMate Indoor robotic platforms with: 200 lbs payload capacity, pan/tilt platform, 24 ultrasonic range sensors, 3 CCD cameras, 2 Video Capture CCD cameras and assorted lenses, 2 IR CCDs, 3 Silicon Graphics Octane workstations with 4 processors and 512 Mbytes of memory, 1 Silicon Graphics Infinite Reality Engine with 4 processors and 256 Mbytes of memory, 1 Silicon Graphics Reality Engine workstation, 1 Sand Table horizontal display device, and AMX resolution multiplexed large projection system with dual displays. Development software at ARL includes: OpenGL development package for SGI and SUN workstations, Imagine Systems terrain processing software, Mercury SHARC, PowerPC development software, and Mutigen Modeling software,

We have included a modest budget for two of these on this proposal and we will purchase others on supporting Army Research Office grants. However, both team members have tremendous expertise in mobile robot navigation, vision, terrain building, etc. At Berkeley, we have just completed some experiments involving high speed (close to 100 m.p.h.) robotically driven cars in close platoon formation using vision in the feedback loop. At ARL, we have just completed experiments in a combat information processor and display using stored terrain and other data, We also have extensive experience and have access to mobile robot platforms with large payloads and CCD cameras and ultrasonic sensors onboard. As a part of this proposal, we are also planning to increase the level of sophistication of the techniques for maintaining the correct pose estimate of the mobile agent. Owing to wheel slipage, the estimate of the mobile base pose maintained by odometry degrades severely over a short period of time (for example, after fifteen minutes of exploration, the pose estimate has been shown to be more then 10 meters off, and is essentially useless). Maintaining an accurate estimate of the robots pose is especially important during the map making process where external references to the environment are sparse. We are planning to develop and implement estimation schemes for combining odometric readings with estimates obtained from inertial sensors for indoor setting and GPS for outdoor navigation. To this end, we are planning to purchase the following items (equally spaced over the two years of the project).

Equipment

Qty	Item	Price/Unit	Total
12	Crossbow accelerometers	179	5,370
6	Compact B/W CCD Cameras in Die-Cast Case	240	1,440
4	Video Transmitter	199	796
6	Imagination Image grabber	895	5,370
12	Freewave radio transceiver with 5dB Omni antenna	1,505	18,060
2	200 MHz Pentium 64 Mb RAM, 4.0 Gb disk	2,300	4,600
4	300 MHz Dual PentiumII 256 Mb RAM, 4.0 Gb disk	6,100	24,400
4	Notebook PC 32MB RAM, 1.4GB disk	3,645	13,580
4	4 Nomad Scout Mobile Robots	8,000	32,000
4	Novatel RT-20 DGPS system	8,297	33,188
16	Senix Ultrasensor ULTRA-U	321	5,136
	TOTAL		143,940

The quote for the Dual Pentiun was obtained from Bret Collard, at 1800-274-7799, ext. 62330. The PC Notebook computer quote is from Computer Warehouse (800) 664-4239. The accelerometer quote is from Crossbow (408)324-4830. The camera quote is from Polaris Industries (800)752-3571. The Video-link quote is from AEGIS research (604)224-0416. Imagination Image grabber is quoted by Imagenation Co. (503)641-7408. The radio modem is quoted by Freewave Technologies, Inc. (303)444-3862. The mobile robot quote is from John Slater from Nomadic Technologies of Palo Alto at (650)-988-7200. The mobile robot platform includes batteries, onboard computing and a camera. The GPS quote is from Novatel Inc. at (403)295-4900. The Ultra-sensor quote is from Senix at (802)453-5522.

7 Relevant Prior Work

The use of visual sensing for navigation and map building has been studied in several contexts. Attempts at constructing purely metric maps were prone to inaccuracies both in sensing and dead-reckoning of the mobile robots [13]. Various landmark based approaches has been widely pursued and successful in the enviroments where the landmarks can be placed artificially or are generic enough and can be readily recognized (outdoors) [15]. Without having any prior information about the environment some attempts to build appearance based topological models has been successful [6]. However the qualitative representations of the environment obtained from by these techniques were limited to merely capturing the neigborhood relationship between individual places. The representaion of the environment which we envision would still remain qualitative with some metric information associated with it, but in addition to the purely appearance-based representation of a "place" we would like to associate some geometric properties to the description of a "place" as well as associate some visual events (e.g. due to occlusions, appearance of an edge, surface) with the transitions between places. In this line of work we can draw from our previous experience in visually guided navigation using topological maps [14] as well aspect representations of objects/scenes [17]. Malik and his group have worked both on tracking of humans at the joint level [3], and the tracking of vehicles [2]. The image segmentation techniques for initial processing of the video input were developed in [23, 24]. Work on object recognition techniques based on color and texture attributes is described in [5]

In recent work [9], we have developed a rigorous probabilistic framework that allows assumptions made in the "classical" data association literature (see for example, [1]) to be relaxed. We have developed and tested this framework for tracking vehicles in freeway traffic using multiple, geographically separated cameras under a wide variety of time-varying conditions. Our system is completely extendable in that arbitrary new object features (including additional sensor types such as infra-red) can be added to the detection system to improve data association reliability.

In [20] exploration and navigation in an unknown environment was formulated for the first time within the framework of on-line algorithms. The map-building problem was first considered within the same framework in [25], and in a topological context in [28]. In [4] the problem of clearing a known environment by a team of agents was first considered algorithmically (it had been posed as a graph-theoretic problem by Parsons in the 1970s), and largely solved. Starting with [29], and in several articles thereafter, Papadimitriou has been studying trade-offs between the amount of communication and the quality of decisions made by a team of decentralized agents.

We have begun a systematic investigation of the problem of designing *intelligent control architectures for distributed systems.* The successful completion of our investigations has resulted in an architectural design theory applicable to semi-autonomous multi-agent systems designed to execute coordinated surveillance or combat missions. We view a control design problem as one that aims to design a safe and efficient control on the basis of a given observation structure. An architecture design problem is concerned with the design of both the observation and control. An architecture design problem for a distributed system begins with specified safety and efficiency objectives and aims to characterize communication, observation, and control. The conceptual separation between observation and communication is a distinction between local and remote observation in a spatially extended environment.

This crucial separation of communication and observation in a distributed system is essential for architectural research.

We have investigated the intelligent control architecture design problem in two formalisms. They are the *intrinsic model supervisory control of discrete event systems* and the *hybrid system formalism.* In order to understand the purely combinatorial aspects of architecture design we have begun our investigation in the context of discrete event systems. By starting with this type of model we are able to take advantage of a literature characterizing the relationship between distributed control and observation. During the past few months we have developed insights into the communication requirements of intelligent distributed systems. The hybrid system literature on multi-agent distributed control is small [27, 11]. Therefore, in parallel, we have started to formulate the first distributed hybrid control problems. Fortunately, the research team has executed some distributed hybrid control designs for automated vehicle control and coordination problems. This past experience together with research on vision-guided navigation and control being done under this grant is proving to be invaluable in guiding our formulation of hybrid distributed control problems.

Our investigations of discrete event models have yielded the following insights to date. We have formulated a multi-agent decentralized observation problem where each agent observes some events that occur in the system and environment and aims to detect the occurrence of a few distinguished events. We think of these distinguished events as failure events. Therefore, we state the problem as a decentralized diagnosis problem. In general, the observations of the agent are not rich enough to infer the events of interest, and therefore, the agents are connected by an inter-agent communication bus that they use to exchange messages. We have a result that characterizes decentralized observation problems that cannot be solved by communication. For problems that can be solved, we can algorithmically synthesize a communication scheme from a system model without communication.

Design and Verification Tools. We are developing a new approach to probabilistic verification. The heart of the approach is to not verify that every run of the hybrid system satisfies certain safety or liveness parameters, rather to check that the properties are satisfied with a certain probability, given uncertainties of actuation and sensing. In this sense, this is a "softening" of the notion of verification and represents a rapproachment between stochastic control, Bayesian decision networks and soft computing. We are hopeful that this new softer approach to verification will push the decidability barriers that cloud deterministic hybrid verification problems.

Past Experience with Automated Highway Systems, Air Transportation Systems and Unmanned Combat Air Vehicles We have investigated some tradeoffs between centralized and distributed multi-agent architectures within different application testbeds. While in case of air-traffic control (work done for NASA, FAA) with web address
`http://robotics.eecs.berkeley.edu/~gpappas/atms`
each individual agent is trying to achieve its own objective and maximize the performance the main concern is to guarantee safety of individual agents and minimize total delays and fuel efficiency which become both global and individual performance measures of a particular strategy. In case of intelligent highway systems (done for Caltrans and FHA, DoT web address: `http://path.berkeley.edu`) the global measure becomes maximize throughput capacity while guaranteeing safety of individual vehicles. We have also worked on architectures to coordinate the deployment of multiple unmanned combat air vehicles in a project for ONR.

8 Management Plan

The project has Sastry as PI at Berkeley, with co-PIs Katz, Malik, Papadimitriou and Russell at Berkeley and Salinas at the Army Research Laboratory. Dr. Kosecka, a research scientist at Berkeley has full time devoted to this project. She will act as a technical project manager, to coordinate the collaborative arrangements and maintenance of the test bed at Berkeley. Salinas will be responsible with his colleagues at ARL for maintaining the test bed at ARL. Both Sastry and Kosecka have had contacts over the last year with several personnel at the ARL and have a good working relationship with them. Coordination will be maintained by visits, and electronic means of collaboration.

Many of the project subtasks have a software (SW), simulation (SIM) phase followed by an experimental validation (EXP) phase. While we will begin with the hardware configuration immediately after the start of the project, it is of importance to debug and *verify* the algorithms extensively in realistic simulation environments prior to *validating* them experimentally. Two sets of test beds will be maintained, one at Berkeley and the other at ARL. The ARL test bed will be especially involved in transitioning the algorithms that we develop to realistic Army special ops protocols and obtain user feedback. During the first year, the mobile robots will be essentially ground based, though we will already have been working on our "aerial" mobile robot (a helicopter which is currently on site at Berkeley fully equipped with INS/GPS, computation onboard, cameras, and transmitters). During the second year, we will use the helicopter as the mother ship in outdoor tests (the web address for this project is `http://robotics.eecs.berkeley/~koo/dv8.html`). This will be a challenging test for the whole system, since the environment will not be a controlled laboratory environment (the helicopter is too large to fly indoors, in addition) and will serve as a final demonstration experiment.

The demonstration experiments have not been described in detail for reasons of space, but we have extracted from ARL's experience with realistic scenarios for searching and warfighting in a building or exploring an urban area (a few city blocks), some benchmark experiments for each year. In the final experiment, the helicopter will be used to provide aerial 3-D views of the scenes being explored for a search and rescue operation.

9 Description of Proprietary Data Rights

All software developed at Berkeley is the property of the Regents of the University of California. However, it is generally speaking made available in the public domain to other researchers, government agencies, and for other non-profit uses, with the caveat that it is not to be further distributed. No classified work is generally undertaken at the University of California. Army Research Laboratory personnel working on joint projects with University personnel are cognizant of this.

The Combat Information Processor and further software developed for this project at the Army Research Laboratories is the property of the U.S. Government and is classified. Software developed at Berkeley will be made available to the Army Research Laboratory for the task **DEPLOY**.

If patents are to be applied for any inventions or discoveries on this project made by Berkeley investigators, they will be sought through the University Office of Technology Licensing after consultation with DARPA and ARL. Joint inventions will also treated similarly.

References

[1] Yaakov Bar-Shalom and Thomas E. Fortmann. *Tracking and Data Association.* Academic Press, New York, 1988.

[2] D. Beymer, P. McLauchlan, B. Coiffman, and J. Malik. A real-time computer vision system for measuring traffic parameters. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pages 495–501, June 1997.

[3] C. Bregler and J. Malik. Tracking people using twist kinematics. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 1998. Submitted.

[4] N. Megiddo et. al. The complexity of searching a a graph. In *1981 FOCS Conference*, 1981.

[5] D. Forsyth, J. Malik, and R. Wilensky. Searching for digital pictures. *Scientific American*, 276(6):88–93, June 1997.

[6] M. Franz, B. Schölkopf, H. Mallot, and H. H. Bültfhoff. Learning view graphs for robot navigation. *to appear in Autonomous Robots*, 199?

[7] T. Huang, D. Koller, J. Malik, G. Ogasawara, B. Rao, S. Russell, and J. Weber. Automatic symbolic traffic scene analysis using belief networks. In *Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI-94)*, pages 966–972, Seattle, Washington, August 1994. AAAI Press.

[8] T. Huang and S. Russell. Object identification on a bayesian context. In *Proceedings of IJCAI 97*, Nagoya, Japan, 1997.

[9] Tim Huang and Stuart Russell. Object identification in a Bayesian context. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence (IJCAI-97)*, Nagoya, Japan, August 1997. Morgan Kaufmann.

[10] C. J. Taylor J. Kosecka, R. Blasi and J. Malik. Vision-based lateral control of vehicles. In *Intelligent Transportation Systems 1997*, 1997.

[11] Datta N. Godbole John Lygeros and Shankar S. Sastry. A verified hybrid controller for automated vehicles. In *CDC 1996*, 1996.

[12] Keiji Kanazawa, Daphne Koller, and Stuart Russell. Stochastic simulation algorithms for dynamic probabilistic networks. In *Uncertainty in Artificial Intelligence: Proceedings of the Eleventh Conference*, pages 346–351, Montreal, Canada, August 1995. Morgan Kaufmann.

[13] A. Kosaka and A. C. Kak. Fast vision-guided mobile robot navigation using model-based reasoning and prediction of uncertainities. *CVGIP*, 56(3):271 – 329, November 1992.

[14] J. Košecká. Visually guided navigation. *Robotics and Autonomous Systems*, 21(1):37–51, July 1997.

[15] A. Lazanas and J. C. Latombe. Landmark-based robot navigation. *Algorithmica*, 13:472–501, 1995.

[16] Y. Ma, J. Košecká, and S. S. Sastry. Visual servoing for a nonholonomic mobile robot. In *IEEE Conference on Decision and Control*, 1997.

[17] J. Malik and Z.Gigus. Computing the aspect graph for line drawings of polyhedral objects. *IEEE Transactions of PAMI*, 12(2):113–122x, February 1990.

[18] Mitchell and C. H. Papadimitriou. The weighted region problem: Finding shortest paths through a weighted planar subdivision. In *JACM, 38, 1*, 1991.

[19] A. Nerode P. Antsaklis, W. Kohn and S. Sastry. *Hybrid Systems IV*, volume 999. Springer Verlag, Lecture Notes in Computer Science, 1997.

[20] C. H. Papadimitriou and M. Yannakakis. Shortest paths without a map. In Springer-Verlag, editor, *Proc. of the 1989 ICALP Conference*, 1989.

[21] Ron Parr and Stuart Russell. NIPS. In *Conference on Neural Information Processing Systems*, 1997.

[22] Z. Li R. Murray and S. Sastry. *A Mathematical Introduction to Robotic Manipulation*. CRC Press, 1994.

[23] J. Shi and J. Malik. Normalized cuts and image segmentation. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pages 731–737, 1997.

[24] J. Shi and J. Malik. Motion segmentation and tracking using normalized cuts. 1998.

[25] C. H. Papadimitriou T. Kameda, X. Deng. How to learn an unknown environment I: The rectilinear case. *Journal of the ACM*, 1995.

[26] C.J. Taylor, J. Malik, and J. Weber. A real-time approach to stereopsis and lane-finding. In *Proc. of Intelligent Vehicles 96*, pages 207–212, September 1996.

[27] C. Tomlin, G. J. Pappas, J. Lygeros, D. Godbole, and S. Sastry. A next generation architecture for air traffic management systems. In Springer-Verlag, editor, *Hybrid Systems IV, Lecture Notes in Computer Science*, October 1996.

[28] C. H. Papadimitriou X. Deng. Searching an unknown graph. In *Proc. 1990 FOCS*, 1990.

[29] C. H. Papadimitriou X. Deng. Competitive distributed decision-making and the value of information. *Algorithmica*, 16(2):133–50, 1996.

10 Resumes

Principal Investigator: Shankar Sastry

Professor of Electrical Engineering and Computer Sciences
and Director, Electronics Research Laboratory
269 M Cory Hall
University of California, Berkeley CA 94720
Tel: (510)642-1857. Fax: (510) 642-1341.
E-mail: sastry@robotics.eecs.berkeley.edu

Field of Specialization

Robotics, Complex systems and hybrid control, Simulation and Visualization

Education

B.Tech. (1977), Indian Institute of Technology, Bombay, India. M.S. EECS (1979), University of California, Berkeley. M.A. Mathematics (1980), University of California, Berkeley. Ph.D. EECS (1981), University of California, Berkeley.

Experience

Professor, University of California, Berkeley, 1988 - present, Gordon Mc Kay Professor of Electrical Engineering and Computer Sciences, Harvard University, 1994, Visiting Vinton Hayes Professor of Electrical Engineering, MIT, Fall 1992, Directeur Recherche, Center Nationale Recherche Scientifique (CNRS), Toulouse, France, Summer 1991, Professore A Contratto, Universita di Roma. Summer 1990, 1991, Associate Professor, University of California, Berkeley, 1984 - 1988, Visiting Fellow, Australian National University, Canberra, Summer 1985, Assistant Professor, University of California, Berkeley, 1983 - 1984, Assistant Professor, MIT, Cambridge, 1980 - 1982.

Associate Editor of:

IMA Journal of Mathematical Control and Information, *Journal of Mathematical Systems, Estimation and Control*, *International Journal of Adaptive and Optimal Control.*

Past Associate Editor of: *IEEE Transactions on Circuits and Systems*, *IEEE Transactions on Automatic Control*, *IEEE Control Magazine*, *Large Scale Systems.*

Awards: President of India Medal, 1977, IBM Faculty Development Grant, 1983, NSF Presidential Young Investigator Award, 1985, IEEE Student Best Paper Award, 1977, Eckmann Award of the American Control Council, 1990. M. A. Arts and Sciences, (honorary, 1994), Harvard University, Cambridge, Fellow IEEE (1995).

Books Published

S.S. Sastry and M. Bodson, *Adaptive Control: Stability, Convergence and Robustness*, Prentice Hall, 1989.

R. Murray, Z. Li and S. Sastry *A Mathematical Introduction to Robotic Manipulation*, CRC Press, 1994.

P. Antsaklis, W. Kohn, A. Nerode and S. Sastry (editors), *Hybrid Systems II*, Springer Verlag, Lecture Notes in Computer Science, Vol. 999, 1995.

P. Antsaklis, W. Kohn, A. Nerode and S. Sastry (editors), *Hybrid Systems IV*, Springer Verlag, Lecture Notes in Computer Science, Vol. 1293, 1997.

Publications

Here are a selection from over 200 papers:

1. D. C. Deno, R. Murray, K. Pister and S. Sastry, "Control Primitives for Robot Systems," *IEEE Trans. on Systems, Man and Cybernetics*, Vol. 22, (1992), pp. 183-193.
2. M. B. Cohn, L. S. Crawford, J. M. Wendlandt and S. S. Sastry, "Surgical Application of Milli-robots", Journal of Robotic Systems, Vol. 12, No. 6, June 1995, pg. 401-416.
3. S. S. Sastry, M. B. Cohn and F. Tendick, "Millirobotics for Remote, Minimally Invasive Surgery", Proceedings of the International Symposium on Intelligent Robotics (ISIR), Singapore, October 1995, also *to appear in Journal of Robotic Systems, 1997.*
4. J. Hauser, S. Sastry and G. Meyer, "Nonlinear Control Design for Slightly Nonminimum Phase Systems: Application to V/STOL aircraft", *Automatica*, Vol. 28 (1992), pp. 665-679.
5. R. M. Murray and S. S. Sastry, "Nonholonomic Motion Planning – Steering using sinusoids", *IEEE Transactions on Automatic Control*, Vol. 38 (1993), pp. 700-716.
6. D. Tilbury, R. Murray and S. Sastry, "Trajectory Generation for the N trailer system using the Goursat Normal Form", *IEEE Transactions on Automatic Control*, Vol. 40, (1995), pp. 802-819.
7. G. Walsh and S. Sastry, "On Reorienting linked rigid bodies using internal motions", *IEEE Transactions on Robotics and Automation*, Vol. 11 (1995), pp. 139-146.
8. D. Godbole and S. Sastry, "Approximate Decoupling and asymptotic tracking for MIMO systems", *IEEE Transactions on Automatic Control*, Vol. 40 (1995), pp. 441-450.
9. J. Lygeros, D. Godbole and S. Sastry, "A verified hybrid controller for automated vehicles", Proceedings of the 35th IEEE Conference on Decision and Control, Kobe, Dec. 1996, pp. 2289-94, also *to appear IEEE Transactions on Automatic Control, Special Issue in Hybrid Systems, April 1998.*
10. C. Tomlin, G. Pappas and S. Sastry, "Conflict Resolution in Multi-Agent Hybrid Systems", Proceedings of the 35th IEEE Conference on Decision and Control, Kobe, Dec. 1996, pp. 1184-89, also *to appear IEEE Transactions on Automatic Control, Special Issue in Hybrid Systems, April 1998.*

Faculty Co Principal Investigator: Randy Katz

Randy H. Katz received his A.B. degree, with highest honors, in Computer Science and Mathematics from Cornell University in 1976. He received the M.S. and Ph.D. degrees in Computer Science from the University of California, Berkeley in 1978 and 1980 respectively. He joined the Berkeley faculty in 1983, where he now holds the United Microelectronics Corporation Distinguished Professorship and serves as the Department Chair of the Electrical Engineering and Computer Science.

Professor Katz is a leading researcher in computer system design and implementation. Under DARPA sponsorship, he led the implementation of the SPUR multiprocessor memory system, the first such system to integrate coherent multiprocessor cache memories with efficient virtual memory management, and was responsible for developing the concept of Redundant Arrays of Inexpensive Disks (RAID). Katz's recent research has focused on wireless communications, mobile computing applications, collaboration technology, and video archive systems.

From January 1993 through December 1994, Katz was a program manager and deputy director of the Computing Systems Technology Office (now the Information Technology Office) of DARPA. One of his key achievements during his Washington service was to oversee the connection of the White House to the Internet, and to establish electronic mail accounts for the President and the Vice President at WhiteHouse.Gov.

Five Relevant Publications:

1. R. H. Katz, "Adaptation and Mobility in Wireless Information Systems," *IEEE Personal Communications Magazine*, V 1, N 1, (First Quarter, 1994), pp. 6-17.
2. H. Balakrishnan, S. Seshan, R. H. Katz, "Reliable Transport and Handoff Protocols for Cellular Wireless Networks," *ACM Wireless Networks Journal*, V 1, N 3, (December 1995), pp. 469-482.
3. S. Narayanaswamy, S. Seshan, E. Brewer, R. Brodersen, F. Burghardt, A. Burstein, Y-C Chang, A. Fox, J. Gilbert, R. Han, R. Katz, A. Long, D. Messerschmitt, J. Rabaey, "Application and Network Support for InfoPad," *IEEE Personal Communications Magazine*, V 3, N 2, (April 1996), pp. 4-17.
4. S. Seshan, H. Balakrishnan, R. H. Katz, "Handoffs in Cellular Wireless Networks: The Daedalus Implementation and Experience," *Wireless Personal Communications*, Kluwer Academic Publishers, V 4, N 2, (March 1997), pp. 141-162.
5. M. Stemm, R. H. Katz, "Measuring and Reducing Energy Consumption of Network Interfaces in Hand-Held Devices," *IEICE Transactions on Fundamentals of Electronics, Communications, and Computer Science, Special Issue on Mobile Computing, V.* E 80-B, No. 8, (August 1997), pp. 1125-1131.

Faculty Co Principal Investigator: Jitendra Malik

Professor, Computer Science Division, Department of EECS
University of California, Berkeley, CA 94720

Field of Specialization

Computer Vision, Computational Modeling of Human Vision

Education

Ph.D. in Computer Science, Stanford University, December 1985, B.S. in Electrical Engineering, Indian Institute of Technology, Kanpur, 1980.

Experience

Vice-Chair for Graduate Matters, EECS, UC Berkeley, 1995-continuing, Assistant/Associate or Full Professor EECS, UC Berkeley, Jan 1986-continuing,, Member, Groups on Cognitive science and Vision Science, UC Berkeley

Selected Honors and Awards

Rosenbaum Fellow, Isaac Newton Institute, University of Cambridge, 1993., Presidential

Young Investigator Award 1989., IBM Faculty Development Award, 1986.
Keynote or invited speaker at various meetings including NAS/NRC workshop (Irvine, 1990), ESPRIT Insight (Nice, France,1991), NATO workshop (York, Canada, 1992), CIBA foundation (London, 1993), British Machine Vision Conf (1995)
Best Graduating Student in Electrical Engineering, IIT Kanpur 1980., One of the top ten students in the Indian School Certificate Examination 1974.
Ten Selected publications (from more than 75):

1. J. Malik and P.Perona, "Preattentive texture discrimination with early vision mechanisms," *Journal of Optical Society of America A*, **7** (2), May 1990, pp. 923-932.

2. P. Perona and J. Malik, "Scale space and edge detection using anisotropic diffusion," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, **12** (7), July 1990, pp. 629-639.

3. D. Jones and J. Malik, " Computational framework for determining stereo correspondence from a set of linear spatial filters," *Image and Vision Computing* **10**(10), December 1992, pp. 699-708.

4. J. Weber and J. Malik, "Robust computation of optical flow in a multi-scale differential framework," *International Journal of Computer Vision*, 14(1), Jan 1995, pp. 67-81.

5. J. Malik, J. Weber, Q.T. Luong and D. Koller, "Smart Cars and Smart Roads," Proc. of British Machine Vision Conference, September 1995, pp.367-382.(Keynote lecture).

6. J. Weber and J. Malik, "Rigid Body Segmentation and Shape Description from dense optical flow under weak perspective," *IEEE PAMI*, February 1997,pp. 139-143.

7. J. Malik and R. Rosenholtz, "Computing local surface orientation and shape from texture for curved surfaces," *International Journal of Computer Vision*, 23(2), June 1997.

8. D. Forsyth, J. Malik and R. Wilensky, "Searching for Digital Pictures," *Scientific American*, 276(6), June 1997, pp. 88-93.

9. J. Shi and J. Malik, "Normalized Cuts and Image Segmentation," *Proc. of IEEE CVPR, Puerto Rico*, June 1997.

10. D. Beymer, P. McLauchlan, B. Coiffman and J. Malik, "A Real-time Computer Vision System for Measuring Traffic Parameters," *Proc. of IEEE CVPR, Puerto Rico*, June 1997.

Faculty Co Principal Investigator: Christos Papadimitriou

Professor, Computer Science Division, Department of EECS
University of California, Berkeley, CA 94720
Education:
October 1972: Diploma in Electrical Engineering, National Technical University of Athens.
August 1976: Ph.D., Princeton University. November 1997: D.T.S (honoris causae), ETH,

Zürich, Switzerland.

Professional Activities:

On the Editorial board of *J.ACM* (1983-1986), *J. of Computer and Systems Sciences*, *Combinatorica*, *Journal of Computer Science and Technology* (Beijing, China), *Algorithmica*, *Information and Computation*, *J. of Operations Research and Computer Science*, *Journal of AI Research*, *SIAM J. on Discrete Mathematics*, *Annals of Mathematics and Artificial Intelligence*, and *Mobile Networks and Applications*, Chair of the Program Committee for the 1989 STOC. Chair of the organizing committee for the 12th ICALP (1985, Nafplion, Greece). Chair of the IEEE Computer Society's Technical Committee on the Mathematical Foundations of Computing (1986-89).

Positions held:

July 1976–August 1978: Gordon McKay Assistant Professor of Computer Science, Harvard University.

September 1978–January 1979: Miller Fellow for Science, University of California, Berkeley.

February 1979–March 1983: Assistant Professor of Computer Science, Massachusetts Institute of Technology. Associate Professor, as of July 1981. June 1976–present: Consultant with Bell Labs, Murray Hill. December 1981– January 1988: Professor of Computer Science, National Technical University of Athens. April 1983–1988: Professor of Computer Science and Operations Research, Stanford University. January 1988–present: Irwin Mark and Joan Klein Jacobs Professor of Computer Science and Engineering, University of California at San Diego. January 1996–present: McKay Professor of Computer Science (since 1997: C. Lester Hogan Professor of Electrical Engineering and Computer Science), University of California Berkeley.

Books

1. H.R. Lewis, C. H. Papadimitriou *Elements of the Theory of Computation*, Prentice-Hall, 1981.
2. C. H. Papadimitriou, K. Steiglitz *Combinatorial Optimization: Algorithms and Complexity*, Prentice-Hall, 1982.
3. C. H. Papadimitriou *The Theory of Database Concurrency Control*, Computer Science Press, 1986.
4. C. H. Papadimitriou *Computational Complexity*, Addison-Wesley, 1993.
5. H.R. Lewis, C. H. Papadimitriou *Elements of the Theory of Computation*, second edition, Prentice-Hall, 1997.

Publications (a selection from over 200 papers)

1. X. Markenscoff, L. Ni, C. H. Papadimitriou "The Geometry of Grasping", *International J. of Robotics Research*, 1990.
2. J. Mitchell, C. Papadimitriou "The Weighted Regions Problem", *J.ACM* 1990.
3. C. H. Papadimitriou, M. Yannakakis "Optimization, Approximation, and Complexity Classes," special issue of *JCSS* for the 1988 STOC Conference.

4. C. H. Papadimitriou, M. Yannakakis "Towards an Architecture-Independent Analysis of Parallel Algorithms," *SIAM J. on Computing*, 1989.

5. C. H. Papadimitriou "On Games Played by Automata with a Bounded Number of States," *J. Games and Econ. Behavior, 4*, 1, pp. 122-131, 1992.

6. C. H. Papadimitriou, M. Yannakakis "Shortest Paths without a Map," special issue of *Theor. Computer Science* for the 1989 ICALP Conference, 1991.

7. X. Deng, C. H. Papadimitriou "Searching an Unknown Graph," submitted to *SIAM J. Computing.*

8. C. H. Papadimitriou "On the Complexity of the Parity Argument and other Inefficient Proofs of Existence" *JCSS, 48*, 3, 498–532, 1994.

9. T. Kameda, X. Deng, C. H. Papadimitriou "How to Learn an Unknown Environment I: The Rectilinear Case," to appear in *J.ACM.*

10. X. Deng, C. H. Papadimitriou "Competitive Distributed Decision-Making and the Value of Information," *Algorithmica*, 1994.

11. E. Koutsoupias, C. H. Papadimitriou "On the k-server Conjecture," *J.ACM*, 1995.

Faculty Co Principal Investigator: Stuart Russell

Stuart Russell is a Professor of Computer Science at UC Berkeley. He is leading researcher in the field of Artificial Intelligence, winner of the principal AI research award (Computers and Thought) and author of the leading AI textbook. His research interests include machine learning, agent architectures, real-time decision making, reasoning and planning under uncertainty, autonomous vehicles, search, game playing and foundations of intelligent systems.

Graduate education

Stanford University (Sept 1982 - Dec 1986), PhD program, Dept. of Computer Science

Honours, Awards: 1st Scholar, St. Paul's School, London, 1974-8., Distinction, British Mathematical Olympiad, 1977, Wadham College Major Scholar 1979-82, NATO Scholar, 1982-85, NSF Presidential Young Investigator, 1990-95, UK SERC Visiting Professorship, 1992, Computers and Thought Award, 1995 (awarded biennially by the International Joint Committee on Artificial Intelligence for the most significant contributions to artificial intelligence), Miller Professorship, UC Berkeley, 1996, Fellow of the American Association for Artificial Intelligence, 1997.

Selected Publications

1. Stuart Russell *The Use of Knowledge in Analogy and Induction.* London: Pitman, 1989.

2. Stuart Russell and Eric H. Wefald *Do the Right Thing: Studies in Limited Rationality.* Cambridge, MA: MIT Press, 1991.

3. Stuart Russell and Peter Norvig *Artificial Intelligence: A Modern Approach.* Englewood Cliffs, NJ: Prentice Hall, 1995.

4. Stuart Russell and Devika Subramanian "Provably bounded-optimal agents." *Journal of Artificial Intelligence Research*, **2**, 1995.
5. Shlomo Zilberstein and Stuart Russell "Optimal composition of real-time systems." *Artificial Intelligence*, **82**, 181–213, 1996.
6. Jeff Forbes, Tim Huang, Keiji Kanazawa, and Stuart Russell, "The BATmobile: Towards a Bayesian Automated Taxi." In *Proc. Fourteenth International Joint Conference on Artificial Intelligence*, Montreal, Canada, 1995.
7. Timothy Huang and Stuart Russell, "Object identification in a Bayesian context." Distinguished Paper Prize, in *Proc. Fifteenth International Joint Conference on Artificial Intelligence*, Nagoya, Japan, 1997.
8. Nir Friedman, Stuart Russell, "Image Segmentation in Video Sequences¡/A¿." In *Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence*, Providence, Rhode Island: Morgan Kaufmann, 1997.
9. Ron Parr and Stuart Russell, "Reinforcement Learning with Hierarchies of Machines." In *NIPS '97: Neural Information Processing Systems*, Denver, 1997.

ARL co-Principal Investigator: Dr. Jose Salinas

Dr. Jose Salinas is a Computer Scientist within the Battlefield Visualization and Processing Branch of the U.S. Army Research Laboratory (ARL). ARL has participated strongly in all of the OSD robotics demonstrations to date. Our role has varied from systems integrator to DARPA agent. Several types of robotic platforms have been developed with a variety of sensor and application modes, and have been used in several demonstrations/exercises at Fort Hood (robotic scout mission) as well as for major demonstrations at our own extensive facilities at Aberdeen. Dr. Salinas is currently working on the Mission and Path Planning Algorithms for the Demo III project funded from OSD. In addition, he is also developing a 2D/3D Operator Control Unit (OCU) for the project.

Dr. Salinas is a graduate from Texas A&M University where he earned a Bachelor of Science in 1988, a Master of Science in 1991, and a Doctor of Philosophy in Computer Science in 1994. While working as a graduate student he was a member of the Computer Systems and Architecture Development Group within the department where he performed research in Fault-Tolerant Systems Design, Digital Systems Testing, and Wafer-Scale Integration Techniques. After receiving his Ph.D., Dr. Salinas joined the faculty of the Computer Science and Mathematics Department at Fairleigh Dickinson University in Madison, New Jersey as associate professor. He joined the Naval Undersea Warfare Center in 1995 where he was part of the Systems Software and Architecture Development Branch. Research projects in this branch included developing new algorithms for real-time sonar beam forming in submarine sonar systems using COTS hardware, in addition to develop new approaches for implementing Virtual Reality technologies into submarine systems.

Dr. Salinas has published 21 articles in conference proceedings and refereed journals and is currently a member of several honor and professional societies.

Research Engineer: Jana Kosecka

Education

1990-95, University of Pennsylvania, Philadelphia, PA.
Ph.D., Computer and Information Science, Spring 1996.
1983–1988, Slovak Technical University, Bratislava, Slovakia.
M.S.E. in Computer and Information Science, with honors.

Professional Experience and Activities

1996-97, Postdoctoral fellow, Department of Electrical Engineering and Computer Sci ence, University of California, Berkeley and PATH.
1990-96, Research fellow, Department of Computer and Information Science, University of Pennsylvania, Philadelphia.
1993-1994, Reviewer for *International Journal of Computer Vision, Computer Vision, Graphics and Image Processing, IEEE Transactions on Robotics and Automation, Autonomous Systems.*
1996 **Program Committee Member:** *Conceptual descriptions from images, ECCV'96 Workshop.*, *AAAI 97*

Fellowships and Awards

1990–1995, Research Fellow, University of Pennsylvania.
1988 Deans award for distinguished M.S.E Thesis, Slovak Technical University.

Selected Publications

1. J. Košecká. Visually Guided Navigation. *To appear: in Journal of Robotics and Autonomous Systems*, 1997. Also to appear as a book chapter in Modelling and Planning for Sensor-based Intelligent Robot Systems, eds. Robert C. Bolles, Horts Bunke and Hartmut Noltemeier

2. J. Košecká and H. I. Christensen. Experiments in Behavior Composition. *Journal of Robotics and Autonomous Systems*, 1996

3. J. Košecká, R. Bajcsy and H. I. Christensen. Discrete Event Modeling of Visually Guided Behaviors. *International Journal on Computer Vision. Special Issue on Qualitative Vision*, Vol. 12(3): 295–316, 1995

4. J. Košecká and R. Bajcsy. Discrete Event Systems for Autonomous Mobile Agents. *Journal of Robotics and Autonomous Systems*, Vol. 12:187–198, 1994

5. 1994 J. Košecká, R. Bajcsy and M. Mintz. Control of Visually Guided Behaviors, in C. Brown and D. Terzoupoulos (eds.), *Real-Time Computer Vision, Cambridge University Press*, 1994