

Field Results for Tactical Mobile Robot Missions

Thomas R. Collins, Ronald C. Arkin, Michael J. Cramer, and Yoichiro Endo

Mobile Robot Laboratory, College of Computing,
Georgia Institute of Technology

Tel: (404) 894-8209

Fax: (404) 894-0957

Email: tom.collins@gtri.gatech.edu, arkin@cc.gatech.edu, mjcramer@cc.gatech.edu, and
endo@cc.gatech.edu

URL: <http://www.cc.gatech.edu/aimosaic/robot-lab/index.html>

ABSTRACT

In 1999, Georgia Tech conducted two field experiments to determine the performance of its mission specification system. The experiments were developed for the DARPA Tactical Mobile Robotics (TMR) Program and were conducted at Fort Sam Houston, Texas. The goal of the TMR Program is to develop robotic tools that can perform useful tasks on future military missions involving complex obstacle negotiation, autonomous indoor navigation, and robust machine perception for urban environments. As a part of the program, Georgia Tech has been developing fault-tolerant multi-robot behaviors and a reusable mission-specification/user-interface system.

Pioneer-AT robots were integrated with vision and sonar sensors, infrared proximity sensors, and differential global positioning system (DGPS) to achieve the goals of approaching and conducting an interior search of a hospital. The emphasis of these particular experiments was the practical implementation of schema-based behavioral control with the mission specification system when designed with the novice user in mind. This paper details the results obtained and lessons learned during those preliminary field trials.

1. INTRODUCTION

The purpose of the DARPA Tactical Mobile Robotics (TMR) Program [1] is to develop tools for autonomous robots in urban terrain and other high-risk war scenarios. The robot will encounter complex obstacles, need to perform autonomous indoor navigation, and require robust machine perception. One scenario, for example, would be a biohazard detection mission, where a robot or team of robots would be deployed to enter and navigate an unknown building and search and report the existence of any hazardous materials. Deploying robots in such scenarios is expected to reduce the risk of losing human lives.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2006		2. REPORT TYPE		3. DATES COVERED 00-00-2006 to 00-00-2006	
4. TITLE AND SUBTITLE Field Results for Tactical Mobile Robot Missions				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Defense Advanced Research Projects Agency, 3701 North Fairfax Drive, Arlington, VA, 22203-1714				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT see report					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 15	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

In order to assess the status and reliability of our intermediate demonstrations that will eventually evolve into an end-to-end demonstration of urban warfare scenarios, two experiments were conducted, in July 1999 and October 1999. A vacant military hospital at Fort Sam Houston in San Antonio, Texas, was chosen for the demonstration site. Both demonstrations utilized the *MissionLab* mission specification system, with the emphasis not on the robot hardware, sensors, or sensor processing, but rather on the practical implementation of schema-based behavioral control with a mission specification system designed with the novice user in mind [2, 3]. The *MissionLab* system has previously been described in the context of similar laboratory experiments [4].

The field experiments were divided into two parts: an indoor mission and an outdoor mission. All missions were configured using the *MissionLab* toolkit.

- 1) **Building Approach:** An outdoor mission. A robot approaches the hospital from a known position, following waypoints, using differential global positioning system (DGPS) integrated into the robot hardware.
- 2) **Indoor Assessment:** An indoor mission. A robot navigates a corridor inside the hospital, looking for an open room, which may contain a potential biohazard. Two phases of the assessment were tested. In the first phase, the approximate location of the target room that contained the biohazard was given to the robot before launching the mission. In the second phase, the robot had no such prior knowledge, and had to find a biohazard sign, which was marked on the wall next to the room that contained a potential biohazard, to decide which room to be inspected.

2. MISSION SPECIFICATION AND HARDWARE SETUP

2.1 Building Approach Mission

The outdoor mission was created with *CfgEdit*, a *MissionLab* tool that can allow users to assemble a robot mission using the robot behavior library, and the experiment was conducted outside the

hospital. The task of the robot was to approach the hospital autonomously from a remote parking lot by following the waypoints (Figure 1). The distance between the parking lot and hospital was approximately 100 meters. The waypoints were selected on site to create a mission that would force the robot to leave the paved road at a low curb, cross varied vegetation and small rubble, and return to the roadway at another low curb before continuing to the hospital loading dock. High curbs, especially at an oblique orientation, were avoided. Because of the long distance that the robot had to travel, and the slippage that was likely to occur on the grass, dead-reckoning errors were assumed to be substantial. Thus, differential global positioning system (DGPS) was implemented in order for the robot to acquire its current position.




Figure 1: Satellite image with the overlaid path shown (white line). Dark areas are primarily the building shadows. Significant features are shown in the graphical legend.




Figure 2: CfgEdit is displaying the Finite State Acceptor (FSA) representation of the building approach mission.

2.1.1 Mission Specification

The building approach mission was assembled with seven states and six triggers (three types of states and two types of triggers). The finite state acceptor (FSA) diagram of the mission is shown in Figure 2, and the states and triggers that were used in the mission are listed in Table 1. The first four pairs of *GoTo* state and *AtGoal* trigger were used for the robot to go through the selected waypoints,

and the last pair of *GoTo* state and *AtGoal* trigger was used to approach the final destination, the hospital loading dock.

Agent Name	Type	Description
<i>Start</i>	State	The robot starts a mission.
<i>GoTo</i>	State	The robot moves to the location specified.
<i>Stop</i>	State	The robot stops moving.
<i>Immediate</i>	Trigger	The transition occurs immediately. (Label “FirstTime” in Figure 2.)
<i>AtGoal</i>	Trigger	The transition occurs when the robot is at the specified point.

Table 1: States and triggers that were employed in the building approach mission.

2.1.2 Hardware Setup

The Pioneer-AT from ActivMedia (Peterborough, NH) was selected as the robot platform for its ability to meet basic locomotion requirements. To supplement the measurement of current position using the embedded shaft-encoders, a NovAtel Communications (Calgary, Canada) RT-20 carrier-phase differential GPS was integrated into the robot. The base station was set up on the roof of the hospital to provide the differential datalink.

During the mission, the robot was controlled by *MissionLab*, which runs on an onboard Dell Latitude laptop computer with an Intel Pentium microprocessor (233 MHz). Even though *MissionLab* allows the robot to reactively avoid obstacles that are detected using the robot's onboard sonar sensors, the position of these sensors was so low that the medium-to-high grass, whose height was often higher than the sonar sensors, would be considered to be obstacles. Thus during this mission, the obstacle avoidance capability was turned off to permit the robot to be able to pass through the grassy area. During the traversal on the pavement across the parking lot, however, the obstacle avoidance was enabled.

2.2 Indoor Assessment Mission

An indoor mission where the robot assesses the interior of a hospital to locate the presence of a biohazard was created and examined. In this experiment colored hazard markers were used to substitute for actual biohazard materials. The experiment was conducted in two phases. For Phase I,

the second floor in the west wing of the hospital was chosen as the test site (Figure 3). The target room was approximately 15 meters distant from the room where the robot was deployed. *CfgEdit* was used to generate an autonomous robot mission that would make the robot depart its original room, find the hallway, navigate along the hallway, find a doorway for the second room on the right hand side, then enter the room and test for the presence of a biohazard.




Figure 3: Site maps for the indoor assessment mission: for Phase I (left) and for Phase II (right).

In the Phase II mission, which was conducted at the first floor in the west wing of the hospital, the robot was deployed from a hallway instead of from inside a room (Figure 3). The task of this robot was to align with the hallway, navigate the hallway, find a biohazard sign on the wall and send an alert to the console, find a doorway near the biohazard, enter the room, then test for a biohazard and alert the console if one is found. The main difference between the Phase I mission and the Phase II mission was in the first case, the robot had knowledge of the approximate location of the room which contained the biohazard, as “the second room at the right hand side.” In the latter case, the robot had to find the biohazard sign on the wall to identify which room to be inspected and had no a priori knowledge of its whereabouts.

2.2.1 Mission Specification

As shown in Figure 4, the Phase I mission was constructed with two types of FSAs, a motor FSA and a camera FSA, to reflect the two types of actuators utilized in the mission. A set of ten finite states

and ten triggers (five types of states and six types of triggers), which had already been coded prior to our arrival at the hospital site, were assembled in the motor FSA (Table 2). The camera FSA was assembled with four states and five triggers (three different types of states and three types of triggers) as listed in Table 3.




Figure 4: The Finite State Acceptor (FSA) representation of the indoor assessment mission (Phase I): the motor FSA (left) and the camera FSA (right).

The Phase II mission also utilized two types of FSAs: the motor FSA and the camera FSA (Figure 5). The motor FSA was assembled with fifteen states and fourteen triggers (nine different types of states and six triggers), and the camera FSA was assembled with six states and five triggers (three different types of states and three types of triggers). The list of the states and triggers utilized in each FSA is described in Table 4 and Table 5, respectively.

Agent Name	Type	Description
<i>Start</i>	State	The robot starts a mission.
<i>GoTo</i>	State	The robot moves to the location specified.
<i>Spin</i>	State	The robot rotates.
<i>MoveToward</i>	State	The robot moves to the object specified.
<i>Terminate</i>	State	The robot terminates its process.
<i>Immediate</i>	Trigger	The transition occurs immediately. (Labeled “FirstTime” in Figure 4.)
<i>AtDoorway</i>	Trigger	The transition occurs when the robot is at the front of a door.
<i>MovedDistance</i>	Trigger	The transition occurs when the robot moved for a specified distance.
<i>Detect</i>	Trigger	The transition occurs when the robot detects a specified object.
<i>NotDetected</i>	Trigger	The transition occurs when the robot is not detecting a specified object. (Labeled “UnDetect” in Figure 4.)
<i>HasTurned</i>	Trigger	The transition occurs when the robot rotates specified angle.

Table 2: States and triggers that were employed in the motor FSA of the indoor assessment Phase I mission.

Agent Name	Type	Description
<i>Start</i>	State	The robot starts a mission.
<i>LookTo</i>	State	The camera of the robot faces to the direction specified.
<i>LookAt</i>	State	The camera of the robot faces to the object specified.
<i>Immediate</i>	Trigger	The transition occurs immediately. (Labeled “FirstTime” in Figure 4.)
<i>Detect</i>	Trigger	The transition occurs when the robot detects a specified object.
<i>NotDetected</i>	Trigger	The transition occurs when the robot is not detecting a specified object. (Labeled “UnDetect” in Figure 4.)

Table 3: States and triggers that were employed in the camera FSA of the indoor assessment Phase II mission.




Figure 5: The Finite State Acceptor (FSA) representation of the indoor assessment mission (Phase II): the motor FSA (left) and the camera FSA (right).

Agent Name	Type	Description
<i>Start</i>	State	The robot starts a mission.
<i>AlignWithHallway</i>	State	The robot aligns with a hallway.
<i>Localize</i>	State	The robot localizes the current position as it is specified. (Labeled “SetXYTheta” in Figure 5.)
<i>GoTo</i>	State	The robot moves to the location specified.
<i>Alert</i>	State	The robot alerts to the console and sends a message, along with a captured image from the camera, via email, to the address specified.
<i>AlignWithDoorway</i>	State	The robot aligns with a doorway to enter.
<i>Spin</i>	State	The robot rotates.
<i>Stop</i>	State	The robot stops moving.
<i>Terminate</i>	State	The robot terminates its process.
<i>Immediate</i>	Trigger	The transition occurs immediately.
<i>HasStopped</i>	Trigger	The transition occurs when the robot stops moving.
<i>Detect</i>	Trigger	The transition occurs when the robot detects a specified object.
<i>Alerted</i>	Trigger	The transition occurs when the robot sends the alert message.
<i>AtDoorway</i>	Trigger	The transition occurs when the robot is at the front of a door.
<i>MovedDistance</i>	Trigger	The transition occurs when the robot moved for a specified distance.

Table 4: States and triggers that were employed in the motor FSA of the indoor assessment Phase II mission.

Agent Name	Type	Description
<i>Start</i>	State	The robot starts a mission.
<i>Scan</i>	State	The camera of the robot swings left and right.
<i>TestSample</i>	State	The camera of the robot faces to the left shoulder of robot to check the color of the paper that, hypothetically, reacts to the specified chemical/biological agent.
<i>Immediate</i>	Trigger	The transition occurs immediately.
<i>MovedDistance</i>	Trigger	The transition occurs when the robot moved for a specified distance.
<i>Detect</i>	Trigger	The transition occurs when the robot detects a specified object.

Table 5: States and triggers that were employed in the camera FSA of the indoor assessment Phase II mission.

2.2.2 Hardware Setup

For both Phase I and Phase II missions the Pioneer-AT was again chosen as the robot platform. A Sony EVI-D31 pan/tilt/zoom camera and a Newton Cognachrome vision system were mounted on the robot to provide the capability of identifying colored objects. The default configuration of the embedded sonar sensors was disabled, since it did not provide an adequate spread for sufficient

obstacle detection. Modification was made to the platform, rearranging the mount positions of the sonar sensors to produce an even 360-degree sonar spread. For the Phase II mission two infrared proximity sensors were also added to the front of the robot to reinforce frontal obstacle avoidance.

Two laptop computers were utilized to carry out this indoor mission. A Dell Latitude, the same type that was used in the building approach mission (Section 2.1.2), was used to run the *MissionLab* console, which issues commands to the robot to start or abort the missions as well as to monitor the progress of the experiment. In the Phase I mission, the Dell laptop was set up in the room where the robot was deployed. In the Phase II mission, it was set up in one of the rooms at the site, approximately 7 meters away from the starting position. The second laptop computer, a Toshiba Satellite with an Intel Pentium microprocessor (75 MHz), was mounted on the robot to serve as the onboard controller for the robot. Among other functions, the executable program on the Toshiba laptop maintains serial communications on three datalinks: with the robot for gathering sensory data and for issuing move commands, with the camera for real-time position control, and with the *MissionLab* console for reporting position information and any other data. Wireless communication between the Dell laptop and Toshiba laptop was achieved through the use of two wireless data transceivers (FreeWave Technologies), one onboard the robot and the other at the *MissionLab* console. Prior to execution, a PPP link was initiated between the console and the robot. To speed the download of the robot executable after compilation on the console to the onboard laptop, while the robot was “docked” a regular Ethernet connection was used in lieu of the slower PPP connection. While the *MissionLab* console and the “docked” robots shared a local network, this network was completely standalone and required no external connectivity.

3. RESULTS

3.1 Building Approach Mission

The robot was repeatedly able to reach the final destination by following the specified waypoints. With an average speed of approximately 0.45 m/sec, the robot took about 5 minutes to travel the approximately 140 meters of the path. It successfully negotiated small curbs (about 2 inches) and shallow puddles, both of which were deliberately incorporated into the chosen path. Still frames from the demonstration video are shown in Figure 6. As mentioned in Section 2.1.2, the obstacle avoidance had to be turned off on the grass area because the medium-to-high grass covering the embedded sonar sensors would otherwise cause the robot to take a needlessly undulating path.



Figure 6: Sequence of images from the building approach mission; (a) the robot is moving out from the parking lot; (b) the robot is crossing over a curb; (c), (d) the robot is navigating a grass and rubble area; (e) the robot is approaching the final destination.

It was also found that satellite images, such as the one shown in Figure 1, were sufficient for specifying relative waypoints under certain conditions. These conditions include:

- Robots must be able to either detect-and-avoid or traverse obstacle features not seen on satellite image (most notably, curbs and small shrubs).
- The satellite image must be recent enough to include any buildings or other large obstacles that would be non-trivial to circumnavigate.
- Important features must not be in shadow.

This experiment indicated that DGPS performance was adequate for the TMR behaviors. Over 95% of the time during cross-country traversal, the robot appeared to maintain positional knowledge within ± 10 cm. The DGPS base station was available 100% of the time, and the datalink to the robot appeared to be 100% reliable, as well. There were isolated instances of satellite signal loss to the robot

GPS unit, usually at one particular waypoint near a tree. During the signal loss from the satellite, the shaft-encoder-based navigation degraded the heading of the robot with each instance of turning or slippage. Once satellites were reacquired (normally within 5-10 seconds), a noticeable course correction occurred, and the robot moved more precisely toward the current waypoint again.

3.2 Indoor Assessment Mission

In the indoor assessment Phase I mission, eight trials were recorded for analysis. Out of the eight trials, the robot successfully completed its mission six times, for a 75% success rate. The sequence of still frames from one of the successful trials is shown in Figure 7. With an average speed of approximately 0.3 m/sec, it took about 2 minutes for the robot to complete the successful mission. Object detection was relatively reliable during our test runs. The typical cause of an unsuccessful mission occurred when the robot did not detect the second doorway, and the mission had no recourse for recovery from this situation. Another failure cause was when the robot simply did not detect the visual cue. Inconsistencies in lighting condition and slow reactivity of the vision algorithm were assumed to be the cause of the problem.



Figure 7: Sequence of images from the indoor assessment Phase I mission; (a) the robot is moving out from the room; (b) the robot is entering a corridor; (c), (d) the robot is navigating the corridor; (e) the robot is entering the target room.

Twelve trials were recorded for analysis during the experiment on the Phase II mission. Out of the twelve trials, the robot successfully completed the mission five times, for 42% success rate. On average, it took over 2 minutes for the robot to successfully complete the mission when it was moving with an average speed of 0.4 m/sec. Still images from a successful trial are shown in Figure 8, and two sets of the alert email message and the captured image attached to the email that were sent by the robot during the mission are shown in Figure 9. In this experiment, half of the unsuccessful trials were

caused by human errors, such as minor bugs in the source code and incorrect parameter setting in a robot behavior. It was also found that failing to align with the corridor at the beginning of the mission would critically affect the rest of the mission.



Figure 8: Sequence of images from the indoor assessment Phase II mission; (a) the robot is aligning with the corridor; (b) the robot is navigating the corridor; (c) the robot is recognizing the biohazard sign; (d) the robot is entering the target room; (e) the robot is in the target room.




Figure 9: Email messages and attached images that were sent during the Phase II mission. (Due to the incorrect setting of the clock, the date shown in the message does not reflect the actual time they were sent.)

4. CONCLUSIONS

Relative to the outdoor Building Approach Mission, the results show that it is straightforward to specify the mission so that it is reliably executed, as long as accurate position information is available and the terrain is traversable. This is indeed what was expected, and recent related usability studies have confirmed that even novice users are in fact able to configure robots to perform such missions in short order. Such waypoint-following missions are relatively limited with respect to the degree of “intelligence” being exhibited, but they often form the underlying backbone of a mission in which other behaviors are exhibited concurrently.

Most of the issues involving the capability of the Pioneer robot traversing outdoor obstacles would vanish with the use of a platform like the Urban Robot, developed by iRobot (Somerville, Massachusetts) for DARPA. The problems associated with spurious sonar readings in high grass could be resolved by either remounting the sensors at a higher position or relying on range sensors that provide higher angular resolution.

The conventional DGPS setup used in the experiments may not be suitable as the primary means of navigation in a hostile situation because of the difficulties associated with setting up a differential base station and maintaining a base-to-robot datalink. There are potentially a variety of ways in which DGPS technology can be applied to TMR-type scenarios. One of these is the use of *approximate* DGPS base stations that can provide highly accurate *relative* position data in a given locality. If each robot in a multi-agent team has the dual capability of being either a base station or a remote user, then it becomes possible for the base stations to be both mobile and redundant. At somewhat higher cost, it would even be possible for stationary robots to act as pseudolites, providing GPS signals directly and eliminating the dependency on the orbiting satellite constellation. Stationary robots can act as base stations (or pseudolites) while moving robots “leapfrog” past them and take on the stationary role. Tactically, these robots can simultaneously perform a military maneuver known as “bounding overwatch,” with the stationary base station robots monitoring the progress of the moving robots with vision or other appropriate sensors.

In the indoor assessment mission, unsuccessful trials of the mission were largely produced by poor performance of integrated sensing devices. For example, slow and unreliable readings from sonar sensors impaired the certain perceptual triggers, such as *AtDoorway*, *AlignWithDoorway*, and *AlignWithHallway*. This problem could be resolved by developing more elaborate missions to allow for alternate actions to recover from the errors, as well as by integrating more reliable sensors. A laser scanner would, for instance, provide faster and more reliable information of the environment.

The sensor inadequacies and the inconsistencies of the simulation environment relative to the real world adversely affected the setup times for our experimental runs. They not only impact the ability to perform the task itself, but also have the effect of requiring more initial “pre-runs” to ensure that the prescribed behavior was appropriate for a particular task. Due to this restriction, the setup time for the mission was much higher than expected in these novel surroundings. However, with improved robustness and further experimental analysis, this should be reduced to an acceptable level.

5. SUMMARY AND FUTURE WORK

Mobile robotic technology continues to advance, particularly with respect to novel locomotion schemes and improved sensory capabilities. Behavioral architectures are a significant research area in their own right, but often with little consideration given to usability. We have demonstrated, through an initial series of experimental trials, that the *MissionLab* system provides a suitable toolset for an end user to create useful robot behaviors in tactical applications.

In real-world robotics, it is not possible to perform research in usability without significant effort in system integration, configuration, and maintenance, both with regard to hardware and to software. Our work therefore also reflects the results of the initial integration of a minimalist TMR platform, including the inherent limitations of the sonar and visual sensors, as well as the terrain restrictions of the wheels. We continue to assess and improve the usability of the software tools while simultaneously removing some of the platform limitations by integrating emerging COTS technology. Some of the areas for improvement have been identified from the experiments described here.

Ongoing usability experiments, in which test subjects attempt to construct mission to perform various tasks, will provide insight into how easily untrained operators can create the TMR missions with the *MissionLab* system. Moreover, a real-time assessment capability for *MissionLab*, which was developed by the project team members at Honeywell Technology Center, will be integrated fully into our experimental configuration, and additional research at Honeywell will address hard real-time

scheduling. Finally, upcoming experiments will provide additional results with improved hardware platforms and *MissionLab* enhancements.

ACKNOWLEDGEMENTS

This research is funded by DARPA under contract DAAE07-98-C-L038. The authors would like to thank Sunundo Ghosh, Karen Haigh, David Musliner (all from Honeywell), Jonathan Diaz, William Halliburton, David Johnson, and Donghua Xu for their participation on this project.

REFERENCES

- [1] Blitch, J., "Tactical Mobile Robots for Complex Urban Environments," *Mobile Robots XIV*, Boston, MA, September 1999, pp.116-128.
- [2] MacKenzie, D. and Arkin, R., "Evaluating the Usability of Robot Programming Toolsets," *International Journal of Robotics Research*, Vol. 4, No. 7, April 1998, pp 381-401.
- [3] MacKenzie, D., Arkin, R.C., and Cameron, J., "Multiagent Mission Specification and Execution," *Autonomous Robots*, Vol. 4, No. 1, Jan. 1997, pp. 29-52.
- [4] Arkin, R.C, Collins, T.R., and Endo, Y., "Tactical Mobile Robot Mission Specification and Execution," *Mobile Robots XIV*, Boston, MA, September 1999, pp. 150-163.