

Reconnaissance and Autonomy for Small Robots (RASR)

Final Report

Solicitation Number : FA5209-11-T-0116
NSN : AJ11-AO-ARD-MAGI
Purchase Request Number : MIPR1DAS04D101
Item Number : 0001AC

Government Point of Contact:
Mr. John Harrison
john.harrison@yokota.af.mil

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 29 JUN 2012		2. REPORT TYPE Final		3. DATES COVERED 25-05-2011 to 24-05-2012	
4. TITLE AND SUBTITLE Multi-Robot Teaming - MAGIC 2010 Third Place-Reconnaissance and Autonomy for Small Robots (RASR)				5a. CONTRACT NUMBER FA520911P0169	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Alberto Lacaze				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Robotic Research, LLC, 555 Quince Orchard Road, Suite 300, Gaithersburg, MD 20878, United States, NA, NA				8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AOARD, UNIT 45002, APO, AP, 96338-5002				10. SPONSOR/MONITOR'S ACRONYM(S) AOARD	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) AOARD-114025	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT The Reconnaissance and Autonomy for Small Robots (RASR) team developed a system for the coordination of groups of unmanned ground vehicles (UGVs) that can execute a variety of military relevant missions in dynamic urban environments. This entry to the MAGIC-2010 competition provides robust autonomous mobility for small UGVs as well as an innovative coordination strategy capable of dealing with communication losses. This effort involved the development of a system that used 1) a relevant deployable platform; 2) a minimum set of relatively inexpensive navigation and LADAR sensors; 3) an expandable and modular control system with innovative software algorithms to minimize computing footprint; and that minimized 4) required communications bandwidth to handle communication losses; and 5) additional power requirements to maximize battery life and mission duration.					
15. SUBJECT TERMS Multi-Agent Technology, Unmanned Vehicles, Robotics					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 28	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

FINAL REPORT
FOR
DEPARTMENT OF THE AIR FORCE
AIR FORCE MATERIAL COMMAND
ASIAN OFFICE OF AEROSPACE RESEARCH AND DEVELOPMENT
(AOARD)
RECONNAISSANCE AND AUTONOMY FOR SMALL ROBOTS (RASR)

The *Reconnaissance and Autonomy for Small Robots* (RASR) team developed a system for the coordination of groups of unmanned ground vehicles (UGVs) that can execute a variety of military relevant missions in dynamic urban environments. Historically, UGV operations have been primarily performed via tele-operation, requiring at least one dedicated operator per robot, and requiring substantial real-time bandwidth to accomplish those missions. Our team goal for entering the MAGIC 2010 competition was to develop a system that can provide practical long-term value to the warfighter. To that end, we self-imposed a set of constraints that would force us to develop technology that could readily be used by the military in the near term:

- Use a relevant (deployed) platform
- Use low-cost, reliable sensors
- Develop an expandable and modular control system with innovative software algorithms to minimize the computing footprint required
- Minimize required communications bandwidth and handle communication losses
- Minimize additional power requirements to maximize battery life and mission duration

Introduction

The use of small unmanned ground vehicles (UGVs) saves lives in Iraq and Afghanistan by distancing humans from dangerous areas. These robots are instrumental in EOD applications, including improvised explosive device (IED) detection and neutralization. However, current controllers require at least one operator for each robot. An operator must tele-operate a single UGV to the suspect object where the operator remotely manipulates and deactivates the IED. All of these operations are performed using remote video, requiring the complete attention of the operator.

To help break this 1:1 ratio and to promote autonomous control of small Unmanned Ground Vehicles (UGV), the Defence Science & Technology Organisation (DSTO) in Australia and the United States Army's Research Development & Engineering Command (RDECOM) took the lead in organizing MAGIC 2010. This challenge required multi-vehicle robotic teams that could execute an intelligence, surveillance and reconnaissance mission in a dynamic urban environment. To complete the challenge, competitors were required to: (i) accurately and completely explore and map the challenge area; (ii) correctly locate, classify and recognize all simulated threats; and (iii) complete all phases within 3.5 hours. The challenge event was conducted in Australia during November 2010.

The challenge was mindful that, at the current state of autonomy, operators still need to provide oversight of the UGVs. Therefore, this challenge also forced developers to design a Human Machine Interface (HMI) that minimized operator workload and increased overall effectiveness.

The RASR team believes that the challenge was a realistic step toward the next generation of small UGV operation and that our solution introduced a new philosophy of distributed control that is well suited for the platform size and significantly reduced communication bandwidth. Figure 1 shows the RASR entries.



Figure 1: Eight RASR-Bots at the beginning of a training run.

MAGIC 2010

The Multi-Autonomous Ground-robotic International Challenge (MAGIC 2010) was a challenge designed to draw cutting-edge proposals for fully autonomous teams of ground vehicles. The capabilities of the unmanned ground vehicles were required to meet situations applicable to being deployed quickly and effectively during either a military operation or a civilian emergency, such as a hurricane. This international challenge was open to industry and academia with an elimination process that reviewed initial proposals and selected 12 teams. This was further reduced to 6 semi-finalists after on-site inspections during the summer 2010.

The final Challenge was held at the Royal Adelaide Showground in Adelaide, South Australia. The larger central area of the showground's racetrack was used to host a ground control station and command center as well as three service zones. A mix of temporary and permanent boundaries was used to contain the UGVs to the desired challenge area. The test area design and testing mythology was overseen by representatives from the National Institute of Standards and Technology (NIST) with input from the team sponsors.

Teams were allowed a total time of three and a half hours to complete the competition. The phases all increased in complexity over time. A mock urban environment approximately 500m x 500m was used for the challenge. This environment contained obstacles and features that would be encountered in the real world, including, but not limited to: buildings, grass, sand, holes, curbs, fences, and humans. The operators were not in line of sight view of the UGVs, as they were isolated in the control tent for the duration of the competition. While GPS was available outdoors, it was not available indoors and was also subject to the normal interruptions encountered when using GPS. Some a priori knowledge was provided to teams, such as the location, number, and area of buildings. During the competition, objects of interest (OOI) were required to be located, identified, and neutralized. The OOI were both static and mobile and were located randomly in the challenge area. OOIs included humans who may be hostile combatants or non-combatants (wearing jump suits of different colors for identification) as well as static objects, such as specifically colored and shaped canisters. In addition, a real-time data feed was used to simulate the information that would typically be relayed by an Unmanned Aerial System (UAS).

Teams were required to have a minimum of three robots and a maximum of two operators during the challenge. There were two different designations for the unmanned vehicles: disruptor-bots and sensor-bots. The disruptors could neutralize OOIs after they were identified by the "sensor" bots. Sensor-bots were to explore and map the area as well as identifying OOI. In order to complete the

challenge, teams had to completely explore and accurately map the challenge area and accurately identify, classify, and neutralize all of the hostile OOI within a three and a half hour period without any accidental neutralizations of non-commandants or even attempt to neutralize within a specified standoff area around the non-combatants.

During Phase I of the competition, the UGVs were required to enter the competition field through a designated entry point. During this portion of the competition the UGVs did not have a UAV feed and did not encounter mobile OOIs. The UGVs were required to map the area in its entirety and neutralize all static OOI. While completing this phase the UGVs encountered a maze made of felt covered boards, barrels used as position markers in assorted colors, parked cars, and corridors of chain link fence covered with a black fabric. One of the challenges of the maze involved the use of a laser range finder (LRF) for obstacle detection. Different materials reflect the laser beams differently. For example, darker objects may absorb more of the laser radiation, and, in the case of fabrics, the laser beams may travel all the way through the fabric and not give an accurate representation of where a fabric boundary is located.

The Phase II environment was more complex; the UGVs encountered mobile OOI (both combatants and non-combatants) as well as similar obstacles to those encountered in Phase I. In addition, the robots encounter a large sand pit which they could travel through or circumvent. A UAV feed provided additional information on where the mobile OOIs were located. Mobile OOIs moved in set patterns during the entire Phase II operation.

In Phase III, all of the complexities of previous challenges were encountered. In addition, that phase included a sniper capable of disabling robots and locations where mobile OOI (both non-combatants and combatants) may share a portion of the same path (paths would cross or walk side-by-side). This required timing and precision as the combatant could only be neutralized when alone. The number of mobile OOI encountered in Phase III was also greatly increased compared to Phases I and II.

DESIGN PHILOSOPHY

Our long term goal for the technology developments resulting from our competing in MAGIC 2010 was to develop relevant technology of long-term value to the war-fighter for control and coordination of multiple small UGVs. Our choices of platform, sensing, and computational engine had to provide a clear pathway to deployment. Decisions taken throughout the design process were based on providing a deployable low cost autonomous mission module. Cost was a large parameter in our sensor, localization and radio selections. We could have “bought our way out of the problem” in many circumstances by purchasing more LADAR sensors, higher-quality navigation components, or advanced computing hardware, but those decisions would have yielded a system that would be prohibitively expensive for broad acceptance for either military procurement or in the first responder sector.

RISK REDUCTION

Our overall approach was based on an up-front risk reduction process. That process flowed to both software and hardware requirements during the initial design phase. High-risk elements were identified and considered as key factors when performing trade studies and choosing among various approaches to solving the MAGIC problem.

Hardware risks were often mitigated using as much proven COTS sensors and components as possible. When custom hardware designs were required, they were fabricated early in the process to allow time for the “gotchas” that often follow with custom hardware. As an example: our custom navigation electronics solution running on the Talon™ platform is currently at revision level 8. Some revisions were based on required redesign efforts, and others were based on additional requirements discovered as the design process evolved. By identifying that component as a high-risk element early in the process and tackling the design early, each of those intermediate iterations were much lower risk, making for a better product in the end.

Software risks were also addressed by working on the hard problems first. By the time of our down-select site visit in the summer of 2010, we had already coded and tested substantial portions of the identified high-risk elements (coordination planners, real-time control and navigation of a tracked system, operations in comms-denied environments). Lower-risk items were moved to later in the schedule, since they did not have as much potential of changing the overall system architecture.

	Decision	Compromise	Advantage
	Use of tracked platform	Harder to build map, harder to control, harder navigation solution	Only pathway to deployment
	Single relatively inexpensive LADAR	Less pixels per second	Cheaper, more reliability with less parts
	Custom navigation solution	Development costs	Low cost accurate dual purpose solution: navigation/timing
	360 deg field of view video	Processing time	Improved situational awareness
	Higher speeds without stopping for mapping	Harder to map on the move	The Army will not buy sitting ducks.
	3D maps	Development of clever gimbal for LADAR	Capable of dealing with inclines, and 3D obstacles
	Shorten Mast	More difficult negative obstacle detection	Portability and weight advantages
	No obvious fiducials	Harder to detect each vehicle	If we put obvious fiducials, the enemy may see them too

Figure 2: System Design for Deployability Decisions

SYSTEM DESIGN FOR DEPLOYABILITY

For the competition we used a cost/benefit analysis for deployability to determine the system components for our platforms. For each examined component this method allowed the team to choose the features that would be most advantageous to the overall system. Because our team goal was not solely to win the competition, but to develop and demonstrate a deployable platform, our decisions were heavily motivated by factors that would optimize the system into one that would be of use to the warfighter on operationally relevant terrain. One example is the

single relatively inexpensive LADAR the team used for the competition. While this LADAR returns less pixels per second, its advantages were that it was a more cost-effective solution, it had less parts, and was more reliable than alternatives. The team decided to use a tracked platform for the competition even though it would make it more difficult to build the map, control, and develop a navigation solution. The tracked platform provided the only pathway to deployment of the system as similarly sized unmanned systems used by the majority of the military are tracked (Talon, Dragon Runner, etc.). A custom navigation solution was chosen for the platform because it provided a low cost, accurate dual purpose solution that provided both navigation and timing, even though it meant that the team had to spend the development costs. To ensure improved situational awareness, a 360 degree field of view video was chosen, even though it would result in additional processing time. The team chose for the platform to move at higher speeds without stopping while mapping, even though it is harder to map while moving; however this provided the benefit of having a platform that does not make as easy a target when used in a military environment. To provide the advantage of being capable of dealing with inclines and 3D obstacles, 3D maps were selected to be used for the competition which meant that the team had to develop a gimbal for the LADAR. A shorter mast made negative obstacle detection more difficult, but provided the team with both increased portability and weight advantages. No obvious fiducials were placed on the vehicle which made it harder to differentiate the vehicles. Again, for reason of adoptability for operational missions, we felt obvious fiducials would also provide an enemy an easy method for identifying and targeting particular robots, therefore, the small advantage that might be gained by an identification system was not incorporated. See Figure 2 for a Design for Deployability Decisions chart.

A RELEVANT PLATFORM

Selection of a platform meant weighing the advantages against the requirements of the competition and deployability. The total vehicle weight had to be under 80 pounds. At the minimum, 5 robots were required (but in practical terms, more would be needed). The total budget for Phase I was only \$100,000. The first consideration was to look at the vehicles the military was actually using. The National Institute of Standards and Technology (NIST) was establishing standards for testing small robots for first responders and the military, so we looked to their results for guidance.

The Intelligent Systems Division at NIST tests small robotic platforms utilizing a calibrated ASTM qualified course. Vehicles that are currently in-theatre or are being considered for deployment are tested in this “do or die” facility. Tested vehicles in the MAGIC weight category included: The Dragon Runner, Souryu IV, G2bot, Element, UMRS 2009, Kenaf, Quince, Matilda 1, Matilda 2, Packbot 510, Mutech-R4, Helios, KOHGA, Talon, TeleMax, Caliber MK3, Andros HD-1J and Andros Mini.

These vehicles had one common denominator: all were tracked. The US Department of Defense (DoD) understands that wheeled vehicles in this weight class are simply not relevant for current operations. This confirmed our decision to utilize a tracked vehicle, even though wheeled vehicles would simplify the navigation and mapping aspects of this competition. The Talon™ is a sturdy, rugged, capable vehicle, used in EOD/IED missions, so we carefully assessed the factors of its heavier weight vs. function. Based on its proven performance, we selected QinetiQ-NA's (QNA) Talon™ platform, however, we want to emphasize that the resulting solution can

be adjusted/adapted to smaller platforms, like the Dragon Runner. Figure 3 shows five of the RASR platforms during testing. The ability to obtain QinetiQ-NA as a RASR Team partner would also fit our decision model for deployability.




Figure 3: Eight RASR-Bots map the interior of a large building

SENSING

Sensors are often the culprits of cost escalation for autonomous systems. To keep the cost down, we utilized a single COTS LADAR sensor. Utilizing multiple LADARs would have made the problem simpler, and reduced the software requirements at the cost of reducing the deployability of the system. We chose to take the harder path and make up for it in software. Although this approach added risk, these risks were identified and mitigated early-on in our design process for an overall reduction in our risk assessment process.

NAVIGATION SYSTEM

IMUs (inertial measurement units) are another example of where we could have “bought our way out of the problem.” In urban warfare, it is common to encounter GPS-denied and multipath situations where GPS (even subscription versions) is not sufficient to provide the localization accuracy needed to accurately map the scenario.

There are a variety of COTS navigation solutions in the US \$50-\$100k range that would provide inertial backing to support the accuracy requirements. It is our opinion, that the overall cost of an autonomous system that has a component in this cost range would not be appealing for military or other applications. Therefore, we developed a navigation unit in the US \$10k range and tailored it to the mission constraints. The overall localization performance of this system is between .1% to .5% of distance traveled depending on the terrain condition and the vehicle calibration. While GPS was not used for MAGIC 2010, our navigation unit is capable of using GPS, when it is available.

DATA RADIOS

In current tele-operated systems, where a constant communication link is required, the DoD customers usually opt for radios that are not at the top of the price range. This is evident by the radio selection that is standard for both the Talon (QNA) and Packbot (iRobot), arguably the most popular platforms currently in theater. Therefore, we opted for radios similar to those in the DoD price range. Although not the most advanced radios were used, we compensated by

concentrating on the autonomy and software smarts to deal with comms losses (which are bound to happen -- no matter how expensive the radio). Comms bandwidth for our system is on the order of 20 kb of data per robot depending on the current operating mode.

COMPUTING PLATFORM

The MAGIC competition requires advanced multi-asset coordination and planning systems in order to accomplish the mission. There is a great temptation to “throw computers at the problem.” Our RASR team resisted, opting for designing algorithms that execute within a high-efficiency software architecture framework.

A single commercial computing platform, a Mac Mini, was selected. Although not rugged for field deployment, several ruggedized platforms have similar computational power. We implemented a mixture of real-time control algorithms with high- and mid-level planners that work as a unit to use as little computing power as needed. In addition, to lower cost and complexity, another advantage of minimizing computing hardware is to lessen the burden on the battery system, allowing for missions of up to 4 hours without recharging/changing the batteries.

RELEVANT AUTONOMOUS MOBILITY AND COORDINATION SOFTWARE

Since much of our design criteria heavily weights the overall cost and deployability of the team, the autonomous mobility system must be robust enough to cope with this decision. In particular it must be able to cope with the treaded platform, the midrange sensors capabilities and inexpensive IMU components. Moreover, the vehicle cooperation infrastructure must provide intelligent behavior while the vehicle is out of comms. We spent a significant portion of our budget and time designing and implementing software for the hard constraints that the MAGIC 2010 rules imposed, as well as our own constraints from what our team partners at General Dynamics Robotic Systems (GDRS) and QinetiQ-NA had experienced from deployed systems.

OVERALL SYSTEM DESCRIPTION

Hardware and software solutions are explained in detail in the following sections.

HARDWARE

The RASR Team is composed of eight platforms (6 sensor UGVs and 2 disruptor UGVs), and 2 Operator Control Units (OCUs). Each robot had a sensor pod providing 360 by 90 degree LADAR coverage, 360 by 90 degree camera coverage, an INU, two data radios (data and E-STOP), and a main computer. Figure 4 displays a flow chart of the hardware design for the RASR Team platforms.




Figure 4: Hardware Design

SYSTEMS ARCHITECTURE

The systems architecture is composed of a distributed/hierarchical control architecture (Lacaze A. , 2002) and a matching Human Machine Interface (Figure 5). The control architecture is composed of the Coordination Layer, the Autonomous Mobility Layer and the Vehicle Platform Layer (Balakirsky, 2002), (Coombs, 2000). Each layer in the control hierarchy contains Sensing, Modeling and Planning (S,M,P) modules (Balakirsky, 2000). The Coordination Layer maintains the overall situational awareness and exposure measures. This layer is distributed among the robots and the base station. At the OCU, a corresponding Coordination Human Machine Interface (CHMI) provides the operators with coordination oversight. The Autonomous Mobility Layer performs local path planning and OOI neutralization (Albus, 1996). A separate copy resides on each robot. At the OCU, a corresponding Autonomy HMI (AHMI) and a Neutralization HMI (NHMI) assign coarse scheduling tasks to provide oversight for these operations. The Vehicle Platform Layer provides low level control functions including path following, communication infrastructure and e-stop functions (Albus, J.S., et al., 2002). At the OCU, the Platform HMI (PHMI) provides the operator with oversight of platform issues. The different functionalities of the HMI have been integrated into a single interface.




Figure 5: The overall system provides autonomous mobility coordination and communications loss resiliency

GROUND VEHICLE COMPONENT & SYSTEMS

Based on current military relevant deployed platforms, the robot was a tracked, skid-steer vehicle. The navigation is more difficult than on a wheeled platform (that can more heavily rely on wheel encoders). Figure 6 shows the component parts of the RASR-bot.




Figure 6: RASR-Bot detail

Terrain sensors added to the base platform include a COTS LADAR configured in an innovative pattern and three fish-eye cameras. The cameras are used for Object Of Interest (OOI) detection and tracking, visual odometry, and teleoperation (when required). A custom-made INU (with GPS) supplies the core navigation solution. Two different radios provide a data link and a remote E-stop capability. A Core II duo processing board hosts the control system (Core II duo is fully utilized by the system). A custom built power distribution system allows hot swapping of batteries.

UVS Autonomy and Coordination Strategy

INTRODUCTION

MAGIC 2010 introduces a number of real world planning challenges:

- Coordinating groups of vehicles is a highly dimensional search problem that cannot be fully expanded using realistic computational capabilities.
- Uncertainties of radio communication makes centralized approaches vulnerable to outages.
- Algorithmic cost spaces that include the opposing requirements of searching and neutralizing, create unbalanced search heuristics.
- Autonomous mobility in a small platform has stringent weight, power, and size constraints.
- A hybrid set of system capabilities from both a mission standpoint (sensor UGVs vs. disruptor UGVs), and from a computational standpoint (UGVs vs. OCUs).
- The need to perform localization and mapping indoors, without GPS.

The RASR approach unravels these challenges to provide a computationally feasible coordination and autonomous system that is highly resilient to communications and GPS outages.

OVERALL PLANNING AND COORDINATION SYSTEM

The planning and coordination system is hierarchically organized providing a distributed coordination layer and a group of specialized planners to solve the mapping and neutralization problems. Figure 7 shows the overall system diagram.




Figure 7: The control system is organized as a semi-distributed hierarchical architecture to minimize complexity while providing resiliency to communication outages.

The system is organized so the elements at the top of the hierarchy are coarse, with slower planning cycles, while at the bottom of the hierarchy, the cycles are fast, and the resolution is high within a reduced scope. At the top of the hierarchy, the system has a coordination layer that plans the synchronized motion of the UGV group. It performs task allocation and rough scheduling of the group. This layer resides on each UGV as well as on the Operator Control Units. Each module in the layer is composed of a Coordination Planner (CP) which interacts with the Global Autonomous mobility Model (GAM) and the Global Mission Model (GMM). The OCU also contains a Situational Awareness Model (SAM). The Coordination Layer uses radio communications to maintain database coherency and to propagate plans.

Each vehicle has an Autonomous Mobility Layer (AM). AM is composed of a local version of a layered map and exposure database, the Local Autonomous mobility Model (LAM) and the Local Mission Model (LMM). The planner at this level solves the problem of single vehicle navigation and local coordination in the case of Neutralization. This layer receives coarse plans from the Coordination Layer, and provides plans that minimize local exposure and optimize mobility constraints for the UGV. These plans are sent to the Vehicle Platform Layer where the task is to transform these plans into actuator commands. The Vehicle Platform Layer maintains the navigation solution and provides the E-stop Controller (EC).

REPRESENTATION

The system provides three concurrent representations:

The Autonomous Mobility World Model provides traversability information that is used by the different levels to calculate the costs of the plans from a mobility standpoint.

Mission Specific World Model provides information about the Object(s) of Interest (OOIs), and their motion prediction. It is represented in the form of a probability density function of each OOI being present at a location at a particular moment in time. The system will also maintain the exposure to OOIs (mobile or static) based on the mapped areas. A probability of detonation is computed using both layers.

Situational Awareness World Model is designed for operator consumption. This representation will allow the operator to understand the environment, and intervene if necessary.

COORDINATION LAYER

The coordination layer resides on each UGV and on each OCU. Our overall philosophy embeds coordination capabilities on each robot in the architecture. The communications between robots is kept at a minimum by only propagating bounds of the solutions found in the nodes called “contracts.” When communications connect the UGVs and the OCUs; the coordination layer benefits from the larger number of computational units. In those cases, the larger number crunching capabilities of some nodes, like the OCUs, will provide search bounds to the rest of the robot team. When communications are poor and UGVs are isolated, they still can coordinate in their local communication neighborhood. It has been shown that this system is guaranteed to outperform an auctioning coordination strategy. The Robotic Research MPAC library (MPAC is software and system developed for autonomy of small unmanned surface vehicles) provides the search engine in the Coordination Layer Planner.

GENERALIZED FORMULATION OF AREA SEARCH

The search area is decomposed into a number of smaller areas called “countries”. Each country has a point, called the “capital”. From the capital, the robot can see every other point in the country. This depends on the range of the robot sensors and the line of sight around known obstacles. The creation of countries and capitals is described in the following section.

This generalized formulation allows other types of missions to be performed in addition to search-only missions. Additionally, the ability to plan for multiple types of vehicles is made possible by abstracting the tasks into the starting conditions, ending conditions, and

resources required for completion. Through this basic formulation, a wide variety of pertinent missions can be managed on-the-fly by a group of UGVs.

Bektas provides a good overview of some of the applications of the multiple traveling salesman problem as well as explaining the approaches that may be taken (exact solution, heuristics, or transformations) (Bektas, 2006). Our system includes the logic required for dealing with comms loss. This logic makes performance comparisons difficult as performance then becomes highly dependent on radio models and the morphology of the site being used. Ardekani, et al. (Ardekani, Arthanari, & Ehrgott, 2010) utilize performance metrics for a single traveling salesman problem (TSP) (Ali & Kennington, 1986), but we assume that similar advantages and disadvantages will transfer to the multiple TSP (Gavish & Srikanth, 1986). As is typical with this highly dimensional problem and with the added complication of the heterogeneous robot set it is not a simple task to compare the search algorithms as they rely on problem dependent heuristics to generate the bounds. These performances are irrelevant to a certain degree in real world scenarios since they are dwarfed by delays caused by communication and the morphology of the site.

K-MEANS LINE OF SIGHT: KML

A new algorithm was designed by the team to compute the countries and capitals. The idea behind KML is to find the smallest number of points (capitals) from which all the tiles in the desired search space can be viewed taking under consideration the line of sight. Once this is accomplished, the problem of coverage becomes a travelling salesman *without* having to do LOS checks or coverage propagation throughout the search. This means we have gains on the order of 20% to 30% over that of a horizon based approach for the specific type of terrain. The space of search collapses by several (hundreds in most cases) orders of magnitude. Figure 8 depicts the KML concept graphically. Some of the characteristics of KML:

- In the family of K-means, but guarantees LOS
- Minimizes number of points with similar convergence and warranties as K-means
- Minimizes the sums of squares. This is very important because it means that it computes paths that minimize distances to areas to be surveyed
- Efficient implementation



Figure 8: KML automatically subdivides the total area into smaller areas called “countries”, each with a “capital”. From the capital the robot can see the entire country, given the known obstacles.

ANYTIME AND MEMORY-CONSTRAINED SEARCH ALGORITHMS

All algorithms are designed to be “anytime algorithms.” These algorithms can find solutions quickly (fast response) and will continue to improve solutions when given more time (ongoing improvement until the optimum solution is found). This approach also allows the algorithms to quickly respond and replan as new information is gathered, e.g. a new blockage is discovered. In addition, these algorithms have been designed, implemented, and tested to work within the memory available (from a single megabyte to multiple gigabytes). For sufficiently complex missions, the distributed algorithm will return a non-optimal plan immediately while continuing to improve the solution in the background.

Complementary to the anytime/memory aspects is the online optimality guarantees. By nature of the bounded search, a set of upper and lower bounds on the optimal solution are constantly being calculated. These bounds guarantee three important benefits: (1) the system is able to estimate how close the current solution is to the optimal solution as the search progresses, (2) the system is capable of proving it has found the optimal solution, and (3) the system is guaranteed to find the optimal solution when given sufficient computational resources.

MULTI-VEHICLE COORDINATION ALGORITHMS

In a distributed environment, the propagation of information, the transfer of tasks, and communications topologies have been addressed. Robotic Research's "MPAC" system provides a contract-based multi-vehicle coordination system which provides an efficient way to share information and exchange task responsibilities even under degraded communications.

PATH PLANNING AT THE AUTONOMOUS MOBILITY LAYER

The autonomous mobility layer is based on the High Maneuverability Planner (HMP). This kino-dynamic planner has been utilized by U.S. Army unmanned platforms for a variety of programs including: the Collaborative Technology Alliance (U.S. Army Research Laboratory's (ARL) Robotics Collaborative Technology Alliance (RCTA)), Safe Ops (US Army, TARDEC, David Kowachek, PM), and the Autonomous Navigation System (PM FCS). The implementation on MAGIC is the first time it has been applied to small robots maneuvering in tight quarters.

This module generates trajectories for each UGV avoiding obstacles while meeting the constraints of the plans created by the coordination layer (Figure 9). An instance of this module resides on each UGV. The path planner's input is a 3D representation of its vicinity in a relative coordinate frame (LAM). It outputs a trajectory to be followed by the Vehicle Platform Layer (VPL). The HMP combines all sensor information into a single representation of the environment that is then utilized to evaluate the cost of performing different actions. As such, the environmental representation includes morphological information, as well as slippage characterization. The resulting trajectories are sequences of vehicle state/time pairs that the VPL follows. Among other things, the state information includes the desired vehicle position and attitude. (Lacaze A. M., 1998)




Figure 9: Simulation of the HMP generating a trajectory through a staircase with rubble. The HMP can handle both holonomic and non-holonomic platforms.

Sensors, Processing & Mapping for UGVs

Although sensing is not meant to be the core problem of the MAGIC 2010 challenge, several of the sensing requirements are not trivial. The system requires fusing LADAR and color vision to create maps and to Objects of Interest. Robotic Research, Del Services, and GDRS have a long history of developing and testing sensing algorithms in robotic systems. Previous research allowed us to identify which areas needed work for the small UGVs. The sensing system fuses LADAR and color information to classify terrain, map the areas, and autonomously recognize and classify the static and dynamic OOIs (Figure 10).




Figure 10: Sensor processing fuses video and LADAR information to sense, track and predict OOIs

SENSOR PROCESSING

Static feature detection is performed by fusing morphological information provided by the LADAR together with color and texture information provided by the cameras. The most challenging aspect of the sensing requirements is the detection and prediction of movers (Lacaze et al., 2010). The team has shown this capability in many programs, for larger platforms (SafeOps and CTA) and smaller platforms (SBIR DARPA SB082-29 Multi-Sensor Detection and Tracking using Traversability Based Prediction, Dr. Robert Mandelbaum, PM).

MAPPING REQUIREMENTS

Three different models are maintained for different purposes: autonomy, mission tasks and Autonomous mobility Models (AM).

Two maps of the world are maintained for mobility purposes; The Global Autonomous mobility Model (GAM) and the Local Autonomous mobility Model (LAM). LAM is a vehicle centered grid based 2½ D map. Each tile in the map provides classification, elevation, density, and the presence of positive or negative obstacles.

Features in the LAM are used by the LADAR registration algorithms to merge all of the individual LAMs into a single coherent GAM. The GAM is displayed on the OCU.


MAP PROCESSING

The OCU has a model of the maps with the current best estimate of the vehicles navigation solution. As navigation information becomes available, including single robot loop closures and multi-robot loop closures, maps are regenerated from relative maps solutions. The resulting maps are then displayed on the OCU and used for generating the final map submission.

MISSION MODEL (MM)

Since part of the mission is to correctly identify and neutralize OOIs, the MM is tasked with maintaining and predicting the knowledge about the humans. Dynamic OOIs are detected using onboard sensors and through the metadata provided by the UAV. Each vehicle tracks and classifies humans in its field of view and stores them in the Local MM (LMM), the Global MM is then used to maintain coherency of classification as the non-combatants and referees move in and out of the field of view of the UGVs. The MM at both levels also has the task of performing dynamic OOI prediction. In order to perform this prediction, Robotic Research's *Terrain Aware Coordination Tool for Intelligent Control* (TACTIC) is utilized. TACTIC was designed by Robotic Research for an Army War College challenge organized and funded by ARL in 2004. RR won this challenge by utilizing the TACTIC toolset.

TACTIC approximates the results of a Monte Carlo simulation at a much reduced computational burden. Figure 11 shows the motion prediction of TACTIC. In this simulation, a dynamic OOI starts at the far right. Based on its history, we assume that it is headed to the yellow line to the far left. The map has a series of mobility obstacles. Green areas represent predicted high probability areas, while red areas provide low probability areas.






Figure 11: Mission module predicting the motion of a dynamic Object of Interest around a building

In this case, the OOI is likely to walk north or south of the obstacle; however, the OOI is not likely to walk through it. This prediction is a probability density function (PDF) in (x,y,t) . MM generates and maintains these PDFs. Based on the PDFs, and using the areas that have been previously explored, the MM provides information about the likelihood of finding a particular OOI in a particular location at a particular moment in time and, most importantly, the probability of entering in a death zone created by a dynamic OOI at a moment in time. By utilizing the GMM and by predicting the vehicle's own velocity, the planners at the coordination layer can plan to rendezvous with OOIs to initialize the neutralization procedures. At the autonomous mobility layer, this cost discourages entering rooms (without first clearing the entrance) as well as cutting corners around the buildings too sharply before first clearing the areas. It also provides guidance during the neutralization procedure so that a set of collaborating sensor UGVs are not cornered into situations that would force them to enter the kill zone of the dynamic OOI.

Operations in GPS-denied environments

The onboard navigation system is responsible for determining the pose of each robot, both the position (x, y, z) and orientation (roll, pitch, yaw). This system must work in buildings, near obstructions, and in the open.

Based on our team's experience developing and using navigation systems for both robots and people, we enhanced our existing adaptive Kalman filter for the UGVs with inputs from a 6 DOF MEMS IMU, wheel encoders, Differential GPS, visual odometry, and LADAR based map registration (Figure 12). The result is a system that performs well indoors and outside and especially ensures a smooth transition between GPS and non-GPS environments.




Figure 12: Pose estimation takes under consideration wheel encoders, MEMS IMU, GPS, LADAR odometry and visual odometry.

Human-machine interface (HMI)

REALTIME INTERACTION

The Human Machine Interface (HMI) monitors and controls the three different levels of the system controller: the Platform, Autonomous Mobility, and the Coordination levels. The software is designed to be operated from a touch screen or from a more standard mouse and keyboard combination. The layout and organization is based on previous HMI designs and borrows aspects from online team video game strategies (Figure 13).

The Coordination Layer HMI has been designed to aid the operator in viewing and modifying the coordination level plans. At this level the operator is not required to interact with individual trajectories of vehicles, but rather with coarse tasks and coarse scheduling decisions.

The Autonomous mobility HMI utilizes existing control methods: teleoperation, drive by waypoints, movable waypoints. RR and GDRS previously developed and integrated these methods under the VTI program (VTI, POC Jillyn Alban TARDEC). The Neutralization HMI provides two modalities. To neutralize a moving OOI, two video streams are shown to the operator for each UGV involved with the neutralization procedure. Results of dynamic OOI detection are framed so as to minimize operator burden.




Figure 13: The HMI has been designed to provide efficient control from a keyboard or from a touch panel. It provides maps, area coverage, coordination and vehicle status.

The Platform HMI provides status of the different mechanical aspects of the UGVs and allows the operator to reboot subcomponents. This status is hierarchical, and additional information can be gleaned simply by touching the appropriate status symbol. Additionally, with warning and error states, commonly recommended operator actions are presented in a natural, seamless fashion.




One goal of this competition was to reduce the operator-to-robot interaction time to allow a single operator to control a larger number of robots. We believe our system used the least amount of operator time of the MAGIC 2010 finalists by taking full advantage of our robust autonomous mobility and automated coordination aspects. This shorter interaction time should allow the system to be scaled to larger areas easily. In addition, map sizes are adjustable to fit the current mission and have been used with map sizes up to 2 km².

A VARIETY OF MAPPING RESULTS

For this competition the team chose to utilize a variety of mapping results available at different points during the competition. The OCU real-time map was used during the competition to provide feedback to the operators. The MAGIC 2010 post-processed map, full post-processed data map, 3D mapping capabilities map, and RR View are additional mapping results available after completing the competition. See Figure 14 for additional details on Team RASR's mapping results.



ROBOTIC RESEARCH, LLC
555 QUINCE ORCHARD ROAD, SUITE 300, GAITHERSBURG, MD 20878
PHONE 240-631-0008 FAX 240-631-0092
WWW.ROBOTICRESEARCH.COM

Map Type	Use	Features
OCU real-time map 	Show the operator approximate coverage and approximate sensed areas of all robots	•uses “feed forward” navigation solution and onboard LADAR registration •Real-time, low bandwidth, sufficient for coordination of assets, allows operator to add control measures for the vehicles which will not show in final map. •The system generates Geotiffs using this method in between phases
MAGIC post-processed map 	Autonomous generation of maps for MAGIC from registered “mini maps”	•Accurate maps that are geo-registered to an a-priori image •Takes about 1 hour to run for the showground •The system generates a Geotiff from this solution
Full post-processed data 	Most accurate post processed data regenerating all maps from raw sensor data	•Requires full access to raw LADAR •Eliminate most slippage •Requires download of data with a stick from the robots •Not being used for MAGIC because of bandwidth and time restrictions
3D mapping capabilities 	Post processed maps can be generated 2D or 3D. Useful for blast volume computations and for measuring doorways	•Built in measuring tools •Provides superb situational awareness •Has a texture mapping feature •Requires to download data with a stick from the robots (not allowed for MAGIC) •VIDEO



ROBOTIC RESEARCH, LLC
555 QUINCE ORCHARD ROAD, SUITE 300, GAITHERSBURG, MD 20878
PHONE 240-631-0008 FAX 240-631-0092
WWW.ROBOTICRESEARCH.COM




Figure 14: A Variety of Mapping Results

AFTER ACTION CAPABILITIES

An After Action Review (AAR) toolkit was developed to provide the operator with information that he/she was not able to capture in real-time. RR-Viewer displays a virtual camera image geolocated with the generated map. As the robot traverses an area, the onboard cameras record the surroundings. The result is an enhanced viewing of the mission. The operator can change the pan, tilt, and zoom of the image and the software will generate the desired view from the omni-directional images collected onboard. It also displays the navigation solution from the vehicle as well as an optional 2D or 3D LADAR map display (Figure 15). In addition, a slightly different version of RR-Viewer has been commercialized, renamed FLASHBACK™. The FLASHBACK™ version is for 2 cameras, a touch screen and offers an overhead view.



Figure 15. RR-Viewer allows the operator to pan, tilt, zoom and teleport to navigate the scene. 2D and 3D maps are available as well as image recall and playback from prior locations.

Experimental Results and Lessons Learned

Each one of the problems that we encountered during the competition has a clear counterpart on current and future robotic DoD operations. From the list of our lessons learned (Figure 16) it is possible to see the validity of the MAGIC 2010 competition. One problem our team encountered was when we lost power to the antenna rotator which meant that communications only provided coverage during Phase I. The lesson learned from this was to not underestimate the cost of hardening and testing all elements of a system. Another problem was the lack of GPS ephemeris data for our day of the competition. Software that expected to have fresh ephemeris data malfunctioned when this data was not available that day. This cascaded into problems localizing OOIs. The lesson learned from this was to read ephemeris data from the GPS broadcast rather than trusting other services to provide it as well as implementing better error handling when data is not available.

Any time there is a one-shot competition; there are inherent problems with rules and unexpected situations. Unfortunately, the scores based on the metrics used during the competition were unavailable to the team and cannot be reported. The testing environment/facility used for MAGIC 2010 was exceptional. It was well prepared and exceeded our expectations. The measuring systems used for the competition also performed flawlessly, enabling a much smoother competition. One area of improvement would be in the rules. At times we thought we understood a rule when it turned out the organizers intended a different meaning. Having only a single shot at the course in a restricted timed course also highlighted the ambiguity of the directions. Overall, it was a

commendable effort for the departments of defense of both the US and Australia to push a concept like this forward as a competition, especially as it included a lot of firsts, such as autonomous mobility in small platforms (only tackled in small areas previously [NIST, IGVC, etc.]) and this went beyond basic capabilities to include coordination mapping in GPS denied areas.

Problem	Culprit
Communications issues We lost power to the antenna rotator, and therefore communications only provided coverage for Phase I.	An Australian to US converter caused us huge communication problems in more than ½ of the course
No GPS ephemeris data for our day of competition	Software expecting to have fresh ephemeris data got confused when data was not available that day.

Figure 16: Experimental Results

INNOVATION IN THE PROPOSED APPROACH

The RASR Team development for the MAGIC 2010 competition focused on coordination and automating the coordination aspect of operating multiple platforms. The resulting key design innovations, previously discussed, include:

- Coordination of COTS military platforms
- Innovative sensing suite (camera and LADAR suite)
- Automated distributed coordination (KML + distributed planning) rather than a horizon based approach

Team breakdown

Robotic Research, LLC – TEAM LEAD. System integration, hardware and software design, navigation, video processing, autonomous mobility, multi-robot coordination, operator control station, testing, and configuration management.

Del Services - system integration, LADAR perception, autonomous mobility, and testing.

QinetiQ-NA (Parent company of Foster Miller and Applied Perception) - base platform, control and Symphony interface, shipping, Talon support in Australia.

General Dynamics Robotic Systems - enclosure design, part fabrication, general team support.

Cedar Creek Defense - communications.

Embry-Riddle Aeronautical University (ERAU) – hardware design trades, laser pointer mount design, system assembly, testing. Two ERAU interns worked at Robotic Research during the summer of 2010 and at the MAGIC 2010 competition.

Summary

The RASR entry to the MAGIC-2010 competition provides robust autonomous mobility for small UGVs as well as an innovative coordination strategy capable of dealing with communication losses. The team took on the challenge of using a militarily relevant platform and a minimum set of relatively inexpensive navigation and LADAR sensors because this is the most likely pathway to deployment for the near-term benefit of the soldier.

References

Albus, J. L. (1996). Planning in the Hierarchy of NIST-RCS for Manufacturing. *Proceedings of the International Conference on Intelligent Systems: A Semiotic Perspective*. Gaithersburg.

Albus, J.S., Huang, H.M., Messina, E., Murphy, K.N., Juberts, M., Lacaze, A., Balakirsky, S.B., Shneier, M.O., Hong, T.H., Scott, H.A., Proctor F.M., Shackelford, W., Michaloski, J.L., Wavering, A.J., et al. (2002). 4D/RCS Version 2.0: A Reference Model Architecture for Unmanned Vehicle Systems. *NISTIR*.

Ali, A., & Kennington, J. (1986). The asymmetric m-traveling salesman problem: a duality based branch-and-bound algorithm. *Discrete Applied Mathematics* , 13, 259-276.

Ardekani, L., Arthanari, T., & Ehrgott, M. (2010). Performance of the Branch and Bound Algorithm on the Multistage Insertion Formulation of the Traveling Salesman Problem. *Proceedings of the 45th Annual Conference of the ORSNZ*.

Balakirsky, S. L. (2002). Value-Driven Behavior Generation for an Autonomous Mobile Ground Robot. *Outside Publication*.

Balakirsky, S. L. (2000). World Modeling and Behavior Generation for Autonomous Ground Vehicles. *Proceedings of the 2000 Intelligent Vehicles Conference*. Dearborn.

Bektas, T. (2006). The multiple traveling salesman problem: an overview of formulations and solution procedures. *Omega* , 34, 209-219.

Coombs, D. M. (2000). Driving Autonomously Offroad up to 35km/h. *Proceedings of the 2000 Intelligent Vehicles Conference*. Dearborn.

Gavish, B., & Srikanth, K. (1986). An optimal solution method for large-scale multiple traveling salesman problems. *Operations Research* , 34 (5), 698-717.

Lacaze, A. (2002). Hierarchical Planning Algorithms. *Proceedings of the SPIE 16th Annual International Symposium on Aerospace/Defense Sensing, Simulation and Controls*. Orlando.

Lacaze, A. M. (1998). Path Planning for Autonomous Vehicles Driving Over Rough Terrain. *Proceedings of the ISIC/CIRA/ISAS '98 Conference*. Gaithersburg.

Lacaze, A., Murphy, K. (Robotic Research, LLC), and Del Giorno, M. (Del Services, LLC). (2010). Reconnaissance and Autonomy for Small Robots (RASR): MAGIC 2010 Challenge. *Land Warfare Conference*. Brisbane.

Lacaze, A., Murphy, K., Uzunovic, N., and Putney, J. (Robotic Research, LLC). (2010). Autonomous Mobility for Areas with Large Numbers of Pedestrians. *PerMIS*.

Lacaze, Alberto (Robotic Research, LLC). (2011). Reconnaissance and Autonomy for Small Robots (RASR): MAGIC 2010 Challenge. *TARDEC Quarterly Robotics Workshop*. Michigan.