

University of Pennsylvania

MAGIC 2010 Final Report

J. Butzke, K. Daniilidis*, V. Kumar, A. Kushleyev,
D. D. Lee†, M. Likhachev, C. Phillips, M. Phillips
GRASP Laboratory
University of Pennsylvania
Philadelphia, PA 19104

Abstract

In this report, we describe the technical approach and algorithms that have been used by the Univ. of Pennsylvania in the MAGIC 2010 competition. We have constructed and deployed a multi-vehicle robot team, consisting of intelligent sensor and disrupter UGVs, that can survey, map, recognize, and respond to threats in a dynamic urban environment with minimal human guidance. The custom hardware systems consist of robust and complementary sensors, integrated electronics, computation, and highly capable propulsion and actuation. The mapping, navigation, and planning software is organized hierarchically, allowing autonomous decisions to be made by the robots while enabling human operators to interact with the robot team in an efficient and strategic manner. The ground control station interfaces integrate information coming from the robots as well as metadata feeds to focus the operator attention and rapidly respond to emerging threats. These systems were developed and tested by the team to complete two phases of the MAGIC 2010 challenge in a safe and timely manner.

* Alternate POC: *kostas@cis.upenn.edu*

† Point of contact: *ddlee@seas.upenn.edu*

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 10 JAN 2011		2. REPORT TYPE FInal		3. DATES COVERED 12-01-2009 to 31-09-2010	
4. TITLE AND SUBTITLE University of Pennsylvania MAGIC 2010 Final Report				5a. CONTRACT NUMBER FA23861014018	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) J Butzke; K Daniilidis; V Kumar; A Kushleyev; DD Leey				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) GRASP Laboratory, University of Pennsylvania, Philadelphia, PA 19104, Philadelphia, PA 19104, PA 19104, PA 19104				8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AOARD, UNIT 45002, APO, AP, 96337-5002				10. SPONSOR/MONITOR'S ACRONYM(S) AOARD	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) AOARD-104018	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT In this report, we describe the technical approach and algorithms that have been used by the Univ.of Pennsylvania in the MAGIC 2010 competition. We have constructed and deployed a multivehicle robot team, consisting of intelligent sensor and disrupter UGVs, that can survey, map, recognize, and respond to threats in a dynamic urban environment with minimal human guidance. The custom hardware systems consist of robust and complementary sensors, integrated electronics, computation, and highly capable propulsion and actuation. The mapping, navigation, and planning software is organized hierarchically, allowing autonomous decisions to be made by the robots while enabling human operators to interact with the robot team in an efficient and strategic manner. The ground control station interfaces integrate information coming from the robots as well as metadata feeds to focus the operator attention and rapidly respond to emerging threats. These systems were developed and tested by the team to complete two phases of the MAGIC 2010 challenge in a safe and timely manner.					
15. SUBJECT TERMS MAGIC 2010, robot, cooperating robot, UGV, autonomy, autonomous system					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 20	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

1 Introduction

The goal of the 2010 Multi Autonomous Ground-robotic International Challenge (MAGIC 2010) is to successfully field a team of robots that can explore and map a large dynamic urban environment, as well as locate, classify and respond to threats. This challenge will significantly accelerate the development of autonomous and unnamed vehicle systems so that a large robotic team can operate effectively with limited guidance from human operators. In the following sections, we describe the technical approach of the UPenn team, consisting of students and faculty from the School of Engineering and Applied Science at the University of Pennsylvania, to address the concomitant problems in this challenge, resulting in a second-place finish in the competition.

1.1 Statement of the problem

This challenge is unique in that it encompasses a large number of unsolved problems across multiple spatial and temporal scales. From fine-grained perception and control problems at the individual robot level to high-level human-machine interfaces and multiagent coordination, there are many technical issues that need to be addressed. In particular, we see the following problems as especially critical to this challenge:

- Perception, localization and navigation by an individual UGV
- Mapping static and dynamic features from multiple sensor sources
- Planning efficient search and neutralization strategies across multiple robots
- Interfacing autonomous behaviors with strategic human decisions in coordinated responses to threats
- Integrating visual, mapping, and metadata feeds for a large robot team without operator cognitive overload

It should be noted that these problems overlap with the objectives of several recent large-scale military robotics initiatives, involving teams of hundreds of researchers and very large budgets. The UPenn team is addressing these problems with a team and budget several orders of magnitude smaller in scale.

1.2 Conceptual solution proposed

The wide scope of the problems in the MAGIC 2010 challenge necessitates a broad solution that encompasses a wide range of spatial and temporal scales. Our proposed solution uses a carefully constructed *hierarchical* decomposition of perceptual, planning, and control algorithms shown in Figure 1. This hierarchical solution allows for efficient high-level interaction from the human operators while simultaneously allowing the robots to operate in an autonomous manner.

Our solution uses a bottom-up hierarchical organization of sensing and mapping tasks, integrating low-level sensor readings at fast time scales from individual sensor robots into a global

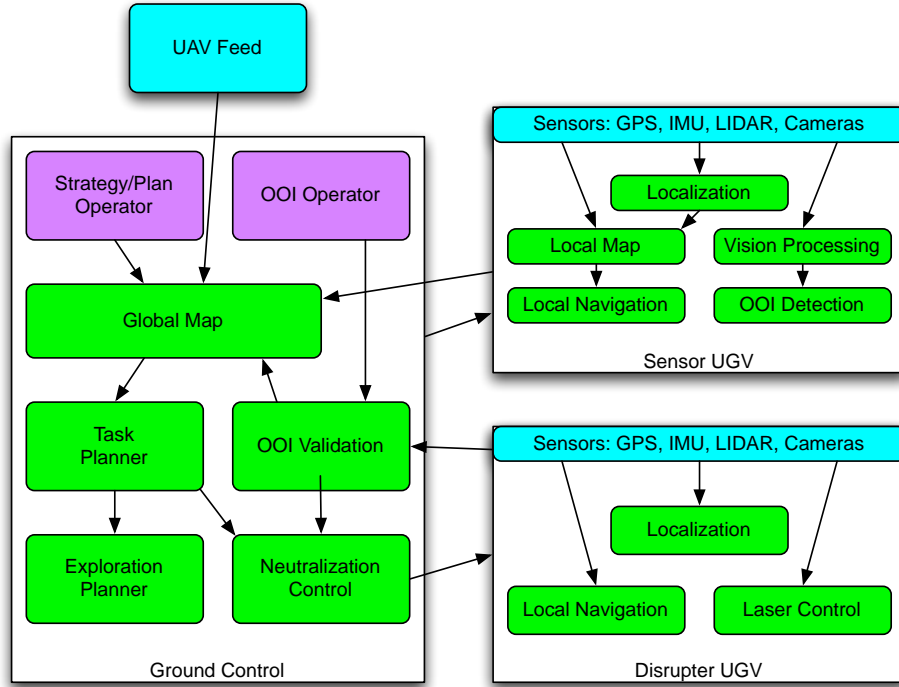


Figure 1: Our approach uses a hierarchical decomposition of perceptual, planning, and control tasks with high-level human operator commands integrated with low-level autonomous robot algorithms.

overhead view for the human operators. This high-level display of the world is then overlaid with metadata feeds, as well as validated object-of-interest (OOI) positions. A top-down hierarchical planning and strategic control decomposition allows the human operators to efficiently issue high-level task commands that get propagated down a chain of planners and controllers to ultimately navigate the individual robots to their desired locations and execute appropriate actions.

1.3 Overall systems architecture

Our overall system involves successfully integrating the various components shown in Figure 2. This is comprised of the following subsystems:

- **Sensor UGV:** Mobile UGVs with LIDAR and camera sensors, GPS, and IMU perform low level localization, mapping, and visual detection tasks. Each sensor UGV is capable of autonomous navigation to explore unknown terrain, as well as tracking mobile OOI's.
- **Disrupter UGV:** Highly mobile UGVs with sensors and a pan-tilt laser pointer are used to neutralize static objects of interest (OOI). These UGVs verify and validate potential OOI's with human operators.




Figure 2: The various subsystems that comprise our overall MAGIC 2010 system architecture.

- **Wireless communications system:** A long range wireless communications system over unlicensed bands provides redundant communication links between the UGVs and the ground control station. The system utilizes some of the sensor UGVs as relay repeaters to facilitate communications in difficult RF environments.
- **Ground control computers:** A cluster of multi-processor computers integrate all incoming sensor information, from the UGVs as well as the simulated UAV feed, into a global map with static and mobile OOI locations. This map is used by the human operators for strategic planning, and for display to the challenge judges.
- **Strategy/Plan control station:** A high-resolution graphical interface is provided to allow the strategy/plan operator to input areas of interest for exploration and mapping, and to allocate tasks to different robots.
- **OOI validation control station:** Potential OOI's identified by UGVs are displayed in both omnidirectional and high-resolution images for validation by the OOI operator.

The modules in each of these subsystems are organized hierarchically, and are connected to each other using fault-tolerant interprocess communications protocols. Common to the UGV platforms are modules that integrate sensory information at a local level for pose estimation and automatic object recognition. Sensor UGVs are specialized to build local maps of obstacles and static

OOI's and to quickly track mobile OOI's. This information is then forwarded to the ground control stations where human operators can make high-level command decisions to neutralize OOI's or avoid certain regions. Disrupter UGVs can then be tasked to navigate to validated OOI's to initiate the neutralization process.

The ground control station incorporates multi-core processors that integrate the various lower-level data streams. A hierarchical mapping module provides a global map view of the environment by aggregating and registering the local maps from the individual robots. An overlay with UAV feed information is then used to generate real-time updates of static object locations as well as dynamic object movements. The high level maps generated by these modules are available for real-time display to the human operators.

This information is utilized by the human operators to give high-level commands to a series of planning algorithms. An intuitive point and click interface allows for rapid human input to control robot task allocation as well as strategic planning. A series of exploration and task assignment planners then compute the optimal routes and actions for the robot team by balancing exploration, mapping, and threat neutralization objectives. Low-level commands are then relayed by the system back to the individual UGVs for execution.

2 Ground vehicle component and systems

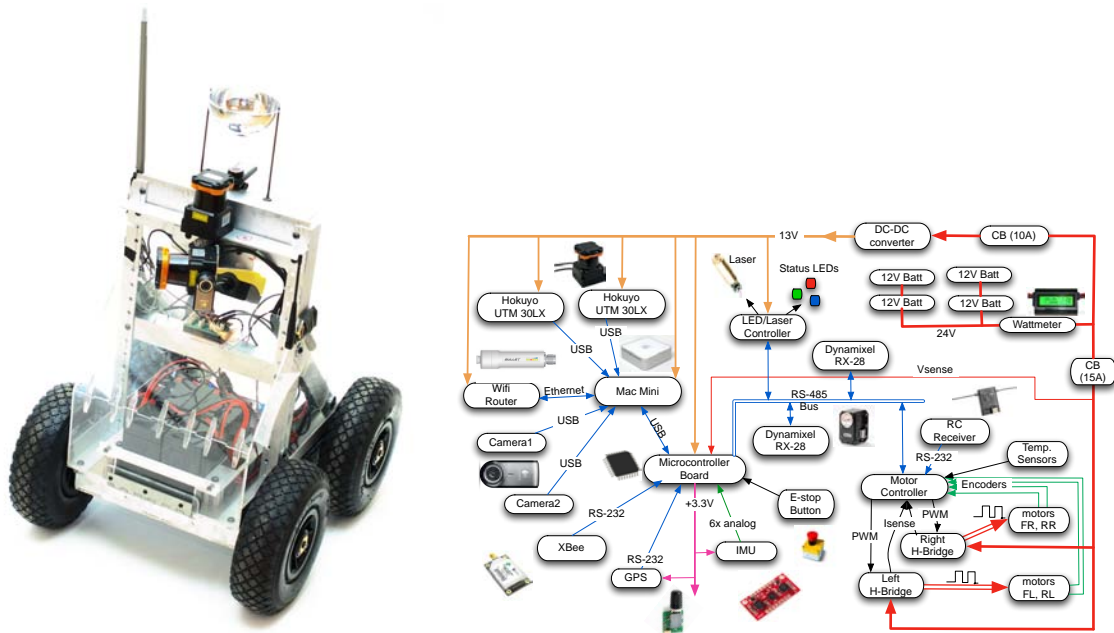


Figure 3: UGVs use an all-terrain aluminum robot base with complementary sets of LIDAR and camera sensors, GPS, IMU, battery-powered embedded computers and wireless connectivity.

Given the overall weight constraint of 40 kg on the UGVs, our robots are built on a lightweight, all-terrain robot vehicle base as shown in Figure 3. The vehicle base is constructed from welded

aircraft-grade aluminum, with a high current DC motor drive system connected to a set of four rugged 25 cm wheels, capable of traversing over 10 cm tall obstacles and speeds up to 2 m/s on flat terrain.

The power system uses two sets of rechargeable $2 \times 12\text{V}$ batteries that are configured so that they may be hotswapped in the field, eliminating the need to reboot the robot. The power is distributed via a series of switching DC-DC converters and customized control electronics to a large set of complementary sensors. The sensor suite for each robot consists of the following [1, 12]:

- Horizontal scanning Hokuyo LIDAR detector returning laser ranges from up to 30 m away
- Vertical scanning Hokuyo LIDAR detector on a panning servo motor returning ground laser ranges
- Omnidirectional catoptric color camera
- Panning frontal view color camera
- Hall-effect motor encoders for proprioception
- 6 degree of freedom strapdown inertial measurement unit integrating MEMS-based gyroscopes and accelerometers
- 50 channel helical GPS receiver

The disrupter UGVs have an additional pan-tilt degree of freedom that allows for aiming a green laser pointer to neutralize identified static objects of interest. Computational power onboard the robots is provided by an embedded Mac Mini Linux computer, with USB connections to the various sensors and microcontrollers. A 802.11g WiFi interface enables high bandwidth network connectivity to the ground control computers, while a Xbee radio link provides a redundant communications channel for emergency control and safety purposes.

3 UVS autonomy and coordination strategy

We constructed and fielded nine UGVs as shown in Figure 4. In order for only two human operators to effectively control this team of robots, there must be a great deal of autonomy built into the system to reduce the level of cognitive overload on the human operators[5, 8]. Here we illustrate how we use hierarchical planning techniques in an example scenario to control our robot team.

Figure 5 shows how a high level exploration planner coordinates the OOI search task among the different sensor UGVs. Here a large outdoor area has been partially mapped, and no static OOI's have been located on the ground. The planner automatically determines the most promising frontier areas to explore, and allocates locations for each individual UGV to navigate. This is done by optimizing the overall information gain in traversing states j

$$IG(j) = \sum_{k \in vis(j)} q_k \quad (1)$$



Figure 4: As many as nine UGVs need to be efficiently monitored and coordinated by two human operators.

where q_k is inversely related to the amount of knowledge about state k that is reachable from state j . These locations are then sent to a lower-level navigation planner onboard each robot to autonomously execute in real time.

Once a putative OOI has been located, a popup message is displayed to the OOI human operator asking to verify and validate the identity of the OOI. If the human operator ignores this message, the sensor UGVs will automatically continue with their search tasks. However, if the identity of the static OOI is confirmed, the location of the OOI is forwarded to the central strategic control computer for display on the global map.

At this point, the Strategy/Plan operator may choose to initiate a neutralization with a single click on the OOI location. The system automatically then cross-cues the nearest disrupter UGV to navigate towards the OOI location, simultaneously placing an avoidance region 3 meters around the OOI. This part of the procedure is also displayed in Figure 5 where the OOI location, avoidance region, and disrupter and sensor OOI's are highlighted.

Once the disrupter UGV has reached the OOI location at a safe distance away from the neutralization zone, the OOI operator confirms visual acquisition of the correct target using images streamed from the disrupter UGVs forward camera. The sensor UGVs cameras can also be used to increase situational awareness during the neutralization process. Once the OOI target has been centered, a command is relayed from the ground control station to the disrupter UGV to start neutralization by switching on its laser. When the neutralization process is complete, the OOI is marked as a neutralized OOI in the global map, and the UGVs automatically return to their previous tasks.




Figure 5: High level map and planners coordinate search, exploration and neutralization tasks among multiple UGVs.

On the other hand, when a dynamic OOI is discovered by a UGV or indicated in the UAV feed, two sensor UGVs are tasked to navigate to locations surrounding the path of the dynamic OOI. Once the UGVs have reached their desired locations, a visual confirmation is given by the OOI operator indicating the identity and location of the dynamic OOI. At this point, a command is relayed to the sensor UGVs to switch into an autonomous tracking mode. In this mode, the robot cameras are servoed to follow the moving OOI. Once visual lock has been achieved, the OOI operator checks the surrounding area for non-combatants and begins neutralization if the area is clear. When neutralization has finished, the sensor OOI's are switched back into their search and exploration mode.

4 Sensors, processing and mapping for UGVs

The onboard perceptual system processes information from the multiple cameras and LIDAR sensors on each UGV. There are two USB cameras mounted on each robot to provide a complementary set of visual images. An inexpensive hemispherical mirror is used with one of the cameras to provide a 360 degree field of view enabling rapid search for potential OOI's and regions of interest around the robot [4]. An example of an unwarped omnidirectional image taken by a UGV is shown in Figure 6. A second front facing camera is mounted on a panning servo that allows for closer inspection of any identified areas of interest in the omnidirectional view.

To assist in the labeling of static and mobile OOI's, each robot performs a red object detection routine on the omnidirectional and front facing camera images. This algorithm analyzes connected regions of high Cr values in the YCbCr image space, and returns bounding boxes of potential regions of interest along with corresponding scores for each region. These regions are presented to the OOI human operator superimposed on the images, with each region's bounding box high-




Figure 6: Omnidirectional image unwrapped to provide 360 degree field of view around each UGV.

lighted according to the saliency of the region's rank. Depths of potential objects are computed by correlating region size with information from the UGV LIDAR sensors.

The main localization and mapping module rely upon Hokuyo LIDAR sensors, configured to scan in horizontal and vertical planes. Each robot is equipped with two LIDAR's; one is fixed in a horizontal orientation, and the other scans vertically with the ability to rotate the scan plane with a servo motor. The horizontal sensor is fixed with respect to the robot frame, and its long range returns are used for 2D localization and mapping. The vertical sensor gathers dense information about the ground in front of the robot and helps determine the traversability of the surrounding terrain. As the sensor continuously scans the frontal area, dangerous three-dimensional obstacles such as low-lying structures and curbs can be robustly detected.

The range data acquired by these two sensors are integrated with measurements provided by the motor encoders and onboard inertial measurement unit. In this manner, the three-dimensional orientation of the robot can be tracked with very low latency. The odometry and inertial readings are combined with laser scan matching to localize the robot using a probabilistic Rao-Blackwellized particle filter to properly keep track of changes in orientation as well as translational motion [11]. This filter relies upon a factorized representation of the underlying pose:

$$P(x, y, \theta) = \sum_i \delta(\theta - \theta_i) N(x, y | \theta_i) \quad (2)$$

where the heading angle θ is sampled and the translational degrees of freedom (x, y) are conditionally independent given the heading.

Simultaneous localization and mapping (SLAM) techniques are employed to build a local map of the environment surrounding the robot. Readings from the two complementary LIDAR sensors are used to update a probabilistic grid map where both traversable ground and obstacles are




Figure 7: Using two complementary LIDAR sensors to map both traversable ground and obstacles simultaneously.

recorded. An example of such a map being built by a single UGV is shown in Figure 7. In this figure, traversable ground as determined by the vertical LIDAR is shown in blue, and obstacles such as walls and low-lying objects are shown in red.

Local navigation with this local map is accomplished by first constructing a basic 8-connected path using A*. This initial path is difficult to follow due to the sharp turns that are inherent in these simple planners. Our approach next uses a special lattice based planner with a special costmap [6, 7]. This planner uses a set of dynamically feasible “motion primitives” that allow greater flexibility than simple 8-connected grid moves. The cost map for this planner is derived from the local map and encourages UGV motions that maximize visibility of unknown areas. The final optimized path is then executed by a path following module on the UGV resulting in a smooth motion that is able to quickly avoid any collisions with nearby obstacles.

5 Operations in GPS-denied environments

Each UGV is equipped with a highly sensitive GPS receiver (50 channel D2523T module) which provides absolute position information at 1 Hz. However, because GPS is not available indoors and can be highly unreliable due to occlusions and multipath effects, we designed most of our systems to operate without relying upon these measurements.

Instead, our operations rely upon a globally consistent map built using a hierarchical map registration algorithm that does not use GPS. First, local maps are built on each UGV using only odometry, IMU, and LIDAR readings. These local maps are then sent incrementally to a higher level map registration module which merges the incoming local maps from each robot into a globally consistent map by maximizing the likelihood of the map cell occupancy statistics. Figure 8 shows an example of a globally consistent map computed by merging several UGV maps together in a large environment where GPS was not available [10].




Figure 8: Mapping a large indoor area from multiple sensor UGVs without GPS.

Only after this globally consistent map has been updated does our high level mapping module utilize GPS measurements. This final stage of processing registers the map to a global reference frame using only the very sporadic GPS readings when there are many satellites in view and the horizontal dilution of precision is very accurate. This is accomplished by finding the optimal transformation T^* between the global map’s coordinate frame and absolute UTM coordinates by optimizing the squared error between the good GPS readings and map coordinates:

$$T^* = \arg \min_T \sum_i (\vec{x}_{GPS} - T\vec{x}_{map})^2 \quad (3)$$

Thus, GPS readings are only used to determine a coordinate transformation between map coordinates and absolute UTM coordinates. Unreliable GPS readings will only skew this coordinate transformation, and will not affect the overall quality of the map used by the human operators for mission purposes.

6 Processing and fusion of metadata

In order to fully utilize the overhead imagery information, we first align and register it to our global map. This process is illustrated in Figure 9 where Google Earth imagery is shown overlaid on an indoor/outdoor global map constructed by the UGVs. First, image processing algorithms are applied to the overhead images to extract perceptually relevant features. The locations of these features are then registered with the current global map, and the probabilities within the corresponding cells in the map are updated accordingly. These maps can then be used for preplanning mission operations by highlighting areas of interest [3].




Figure 9: Incorporating and registering overhead imagery data with robot maps.

An additional ground control software module is responsible for connecting to the simulated UAV data feed via TCP, and converting transmitted OOI information to the global map module. These transmissions are logged and registered for display on the map, just as the data coming from the individual UGVs are recorded.

Thus, in this framework, incorporating metadata from a UAV is no more complex than fusing information from the different sensor UGVs. An overlay of the UAV feed can be switched on and off by the Strategy/Plan operator. This information can then be used to quickly determine desirable exploration or undesirable avoidance regions in the global map.

7 Situational awareness tools




Figure 10: Omnidirectional and forward camera views with automatically selected regions of interest displayed to the human OOI operator.

The key to having good situational awareness in a complex system is to have the system automatically identify potentially salient cues in the large incoming informational streams. In our system, we rely upon the omnidirectional images from the multiple UGVs to provide overall awareness of events occurring near the robots.

However, it is difficult for a single human operator to attend to all the images simultaneously. Instead, we use automatic image segmentation algorithms to identify salient areas of interest, such as large red objects that could be lethal OOI's. These areas of interest are then displayed to the OOI operator as illustrated in Figure 10, who can then take further action to focus in on these locations. In this manner, our system provides both a wide field of view of the overall situation in addition to focused attention on particular items of interest.

The individual UGVs also contain enough autonomy in their navigation modules to avoid dangerous situations without human control. If an obstacle or object comes too close, the robot will automatically stop to allow the human operators to assess the situation via the transmitted omnidirectional and frontal images. The human operators then can determine the overall threat level and can quickly issue high-level commands to particular UGVs to respond to the situation.

8 Human-machine interface



Figure 11: Strategy/Plan ground control display with intuitive human-machine interface controls.

We have constructed an easy-to-use interface to control and modify the strategic tasks of the individual UGVs in the team. Shown in Figure 11 is a user interface for the Strategy/Plan operator that is modelled after a real-time strategy (RTS) game. By pointing and clicking on various locations in the global map, the human operator can quickly retask and direct individual UGVs using very high-level commands to the underlying planners and controllers in our system.

These tasks and behaviors use simple MATLAB data formats and scripts to relay human operator commands to the software modules on the different robots. Through the continuously updated

real-time map display, the operator can quickly monitor the progress of the robot team and efficiently verify and possibly intercede during mission critical operations.

The architecture of our overall system has hooks for human operators to monitor the status of any incoming and outgoing messages between the various software modules. Human intervention is needed when there is large uncertainty in the information being passed between modules. For example, if there is uncertainty in the identity or location of an OOI, a human operator is notified to verify the putative identify and location before any further operation can continue. Thus, the key design approach in our system is to be “certain about uncertainty” regarding any operational aspect requiring human intervention.

9 Missions operations strategy




Figure 12: Defining mission-relevant areas of interest in the global map.

Our human-machine interface design allows the operators to quickly direct and guide operational policies by the UGV team. Figure 12 shows how the ground control station can be used to define particularly relevant areas to explore, and any potentially dangerous areas to avoid. To determine these areas, we plan to first identify areas of interest during a pre-operations stage where analysis of overhead imagery data is incorporated. These areas can then be preloaded into operational configuration files for use during particular phases.

Once the operational phase has begun, the global map and metafeed overlay can then be utilized to quickly modify operational policies for the UGVs by the Strategy/Plan operator. As the OOI operator provides visual situational awareness of the surrounding environment, directed response

actions by the UGVs are coordinated by the two human operators. We have had some opportunities to test these procedures during a short testing phase before the competition. Some of the team members served as a “Red Team,” setting up a mock course with hidden static and dynamic OOI’s. Evaluators then scored the performance of the robot team and human operators in order to determine deficiencies in system modules or human operational procedures. Although there has not been much development time for full testing, we hope that this limited testing procedure will help to increase the overall accuracy and efficiency of our team operations across all the various components.

10 Risk reduction strategy

Safety is the primary concern for both hardware and software design in robotic applications. We have previously demonstrated outstanding safety standards in our previous work, and was able to mitigate against any unforeseen risks in our system for MAGIC 2010.

10.1 EMI/RFI and electrical

Since only commercial off-the-shelf electronic components are used on our robots, we have verified EMI/RFI and electrical compliance in each of the procured subsystems. Rigorous stress testing, during which electric power consumption has been closely monitored has not revealed any potential electrical risks.

10.2 Vibration and physical

Mechanical design and assembly have been performed by students and researchers who have all had previous experience in building complex robotic systems. Structural elements and connectors have been subjected to testing in both laboratory and field settings. Any loose or vibrating parts have been duly noted, and the system has been continually upgraded to rectify any systematic problems.

10.3 Modeling and simulation

We have modeled our UGV robots using a three-dimensional physically realistic simulation environment. Shown in Figure 13 is the output of our real-time rendering engine showing the motions of various robots operating in a complex environment. The simulations are carried out using 3D models within the Open Dynamics Engine (ODE) open-source physics engine, running our software modules in MATLAB.

Unfortunately, due to the lack of development time, we have not been able to fully utilize insights from the simulation analysis. Most of our team’s effort have to date been focused on directly improving the the actual physical robots.



Figure 13: Real-time simulation and rendering of a team of our UGVs.

10.4 Safety, E-stop, freeze, and lost-link



Figure 14: Master fail-safe emergency stop operator console for the UGV team.

Communication with the vehicles will be handled on two levels. Mission critical information will be passed over a reliable long range low bandwidth RF link. This includes the ability to start, control, pause, resume and permanently disable the vehicle as well as to receive feedback such as the status of each platform. At any point in time, the operators have instant access to the vehicle controls, ready to override if the situation requires to do so.

Reliable communication is complemented by the software and hardware safety architecture originally developed for the Urban Challenge competition. Simple, redundant and tested, the sys-

tem uses heart-beat messages from the control station in order to verify vehicle status and is backed by a fail-safe hardware implementation. In an event of lost link, the system is programmed to perform a software pause, temporarily cut power to all actuation, or disable the UGV until the end of the mission phase. A simple to operate emergency control box for reliable operation of the robot team that will be provided to MAGIC 2010 organizers is shown in Figure 14.

10.5 Communications architecture

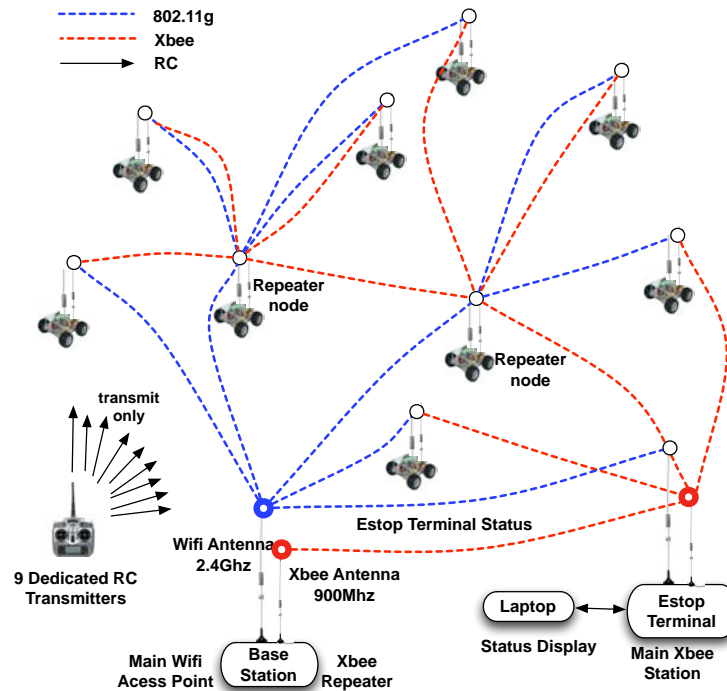


Figure 15: Communication links between the ground control station and UGVs.

Figure 15 shows our communication architecture. We use 2.4 GHz 802.11g WiFi connections for high bandwidth communications between the robots and the ground control station. High quality video streams and map updates are sent periodically from each robot to the GCS via UDP packets to avoid bottlenecks, whereas critical command parameters are transferred back to the UGVs via guaranteed interprocess communications TCP packets.

Additionally, we designed a completely independent, secondary radio link to control the UGVs in case the WiFi connections are severed. This second long-distance communications channel is based on a Xbee Xtend network. These radio modules are the key component of our fail-safe heartbeat and emergency stop system. This network uses frequency hopping at 915 MHz, so the connection is more reliable at long distance than standard WiFi frequencies. The Xbee network can also be configured using repeater nodes for relaying messages as well as for broadcast.

Finally, we can also use a more traditional Spektrum DX7 seven channel remote control operating at 2.4 GHz for manual operator control of the UGVs motor drive system if needed. With

these redundant ways of communicating with the UGVs, we seek to mitigate against any risk that a particular communications channel is blocked.

10.6 Spectrum plan and usage

These links will utilize the unlicensed 915 MHz and 2.4 GHz frequencies of the industrial, scientific and medical (ISM) bands of the radio frequency spectrum. RF power emissions are tuned to stay within US and Australian regulations, and the ground control stations will continuously monitor signal strengths. When necessary, wireless signals will be relayed using a subset of the sensor UGVs as repeaters using standard wireless distribution system (WDS) protocols. The routes and trajectories of these relay nodes will be planned to accordingly minimize interference among them.

If necessary, the wireless transmissions can be secured using either WEP or WPA encryption keys. A directional antenna at the ground control station will also be used to direct RF power emissions to the UGV team, and minimize interference with other surrounding equipment.

10.7 Test plan

As safety considerations are paramount, each fielded trial will begin with a preliminary check of the E-stop, freeze, and radio links. Only after verifying that low-level safety systems are operational will the higher level controllers and modules be started on the robots. Next, the higher level planners and controllers on the ground control station will initiate communication links with the robots, and only after verifying that all modules have been correctly initialized will the UGVs be allowed to begin their mission operations.

11 Summary

This document has presented the technical approach of the University of Pennsylvania team in the MAGIC 2010 challenge. We have designed and constructed a large team of UGVs with the appropriate set of computational, sensing, communications and actuation hardware. We use a hierarchical series of sensing, planning, and control modules to coordinate and direct the autonomous navigation and actions of the UGV team to achieve search, mapping, and neutralization mission objectives.

A complementary set of sensor readings is fused to produce static and dynamic maps of the surrounding environment. By properly representing uncertainty and filtering the inertial, odometry, visual and LIDAR measurements, each individual UGV constructs a local map without GPS. These local maps are then hierarchically merged at the ground control station with metadata from overhead imagery and the simulated UAV feed to construct a globally consistent map of the robot team activities. This global map is displayed to the strategy/plan human operator who can then issue appropriate high-level strategic commands to the UGV team. Visual imagery from the robots' omnidirectional and frontal view cameras are used by the second operator to verify the identities

and locations of OOI's, and provides the human operators with good situational awareness during mission critical events.

All efforts to reduce known risk factors in the overall system have been taken using proven safety systems. Although development time has been extremely short, we have taken every opportunity to test the various components of the system in increasingly complex environments. These systems were used to complete the first two phases of the MAGIC 2010 challenge in Adelaide, Australia, successfully locating and neutralizing eight separate objects of interest throughout the competition.

References

- [1] J. Bohren, et. al. (2008). Little Ben: The Ben Franklin Racing Teams entry in the 2007 DARPA Urban Challenge. *Journal of Field Robotics* 25, 598–614.
- [2] S. Chitta, P. Vernaza, R. Geykhman, D. D. Lee (2007). Proprioceptive localization for a quadrupedal robot on known terrain. *Proceedings of the IEEE International Conference on Robotics and Automation*.
- [3] D. Ferguson, T. Howard, and M. Likhachev (2008). Motion planning in urban environments. *Journal of Field Robotics* 25, 939–960.
- [4] C. Geyer and K. Daniilidis (2002). Omnidirectional Video. *The Visual Computer*.
- [5] M. A. Hsieh, et. al. (2007). Adaptive teams of autonomous aerial and ground robots for situational awareness. *Journal of Field Robotics* 24, 991–1014.
- [6] M. Likhachev, G. Gordon and S. Thrun (2003). ARA*: Anytime A* with provable bounds on sub-optimality. *Advances in Neural Information Processing Systems (NIPS) 16*. Cambridge, MA: MIT Press.
- [7] M. Likhachev and D. Ferguson (2008). Planning long dynamically-feasible maneuvers for autonomous vehicles. *Proceedings of Robotics: Science and Systems*. MIT Press.
- [8] N. Michael, J. Fink and V. Kumar (2008). Experimental testbed for large multi-robot teams: verification and validation. *IEEE Robotics and Automation Magazine* 15, 53–61.
- [9] N. Moshtagh, N. Michael, A. Jadbabaie, and K. Daniilidis (2008). Distributed, bearing-only control laws for circular formations of ground robots. *Proceedings of Robotics: Science and Systems*. MIT Press.
- [10] J.P. Tardif, Y. Pavlidis and K. Daniilidis (2008). Monocular visual odometry in urban environments using an omnidirectional camera. *Proceedings of the IEEE International Conference on Intelligent Robots and Systems*.
- [11] P. Vernaza and D. D. Lee (2006). Rao-Blackwellized particle filtering for 6-DOF estimation of attitude and position via GPS and inertial sensors. *Proceedings of the IEEE International Conference on Robotics and Automation*.
- [12] P. Vernaza, B. Taskar and D. D. Lee (2008). Online, self-supervised terrain classification via discriminatively trained submodular Markov random fields. *Proceedings of the IEEE International Conference on Robotics and Automation*, 2750–2757.