Development of a Virtual Manufacturing Framework: From End-User Performance Requirements to Robot Competitions

Stephen Balakirsky[*a], Raj Madhavan[a,b], Christopher Scrapper[a]
[a] National Institute of Standards and Technology,100 Bureau Drive, Gaithersburg, MD, USA 20899
[b] Oak Ridge National Laboratory, Oak Ridge, TN, USA 37831

ABSTRACT

This paper presents the motivation behind the new joint NIST/IEEE Virtual Manufacturing Automation Competition (VMAC). This competition strives to take the Automated Guided Vehicle (AGV) user community driven requirements and turn them into a low-entry-barrier competition. The objectives, scoring, performance metrics, and operation of the competition are explained. In addition, the entry-barrier lowering infrastructure that is provided to competitors is presented.

Keywords: Robot Competition, Simulation, Manufacturing, Automated Guided Vehicle, Performance Metrics

1. INTRODUCTION

Automated Guided Vehicles (AGVs) represent an integral component of today's manufacturing processes. They are widely used on factory floors for intra-factory transport of goods between conveyors and assembly sections, parts and frame movements, and truck-trailer loading/unloading. Automating these systems to operate in unstructured environments presents an exciting area of current research in robotics and automation. Research is currently being performed in areas such as path planning [5,6], multi-vehicle traffic management [8], and navigation in constricted spaces [9]. Unfortunately, the traditional entry barrier into this research area is quite high. Researchers need an extensive physical environment, robotic hardware, and knowledge in research areas ranging from mobility and mapping to behavior generation and scheduling. An accepted approach to lowering this entry barrier is through the use of simulation systems and open source software.

The systems described in this paper may be viewed as intelligent embodied agents. These agents require an environment to operate in, an embodiment that allows them to affect and move in the environment, and intelligence that allows them to execute useful behaviors that have the desired outcome in the environment. There are many aspects of the development of these agents in which simulations can play a useful role. If correctly implemented, simulation can be an effective first step in the development and deployment of new algorithms. Simulation environments enable researchers to focus on the algorithm development without having to worry about hardware aspects of the robots such as maintenance, availability, and operating space. Simulation provides extensive testing opportunities without the risk of harm to personnel or equipment. Major components of the robotic architecture (for example, advanced sensors) that may be too expensive for an institution to purchase can be simulated and enable the developers to focus on algorithm development. If performed correctly, the simulation-developed algorithms and technology may be seamlessly transitioned to real hardware for further test and validation.

The author's of this paper have received a grant under the *IEEE Robotics and Automation Society's New Initiatives Competition* in order to expose faculty and staff of regional universities to automation and manufacturing problems and to create a new competition based on improving AGV performance. The PIs have extensive experience managing the *de facto* standard open source systems for Virtual Urban Search and Rescue (USAR) development (www.sourceforge.net/projects/usarsim), and will base the Virtual Manufacturing Automation Competition (VMAC) on lessons learned from this effort. It is our belief that competitions are an effective means of stimulating interest and participation among students by providing exciting technological problems to tackle. The areas of factory automation and robotic simulation were new to some or all of the participants, therefore a workshop was held a NIST to bring the participants up-to-speed on the available infrastructure. This infrastructure includes the simulation framework, an

[*] email: stephen@nist.gov

implementation of the 4D/RCS Reference Model Architecture [1], and community forums for discussions and problem solving.

The remainder of this paper is organized as follows: Section 2 provides details on the RoboCup Competition that is the basis for the VMAC competition, Section 3 details specific information on the new competition, and Section 4 discusses additional details on the supplied infrastructure. Finally, Section 5 provides conclusions and areas of future work.

(a) Example of the cubicle area from the indoor environment.

(b) Example of an accident scene from the outdoor environment.

Figure 1: Example of competition worlds from the RoboCup 07 Virtual Robot Competition.

2. ROBOCUP COMPETITIONS

RoboCup is an annual international competition that brings together robotics researchers for competitions ranging from robotic soccer, to dance, to Urban Search and Rescue [7]. Last year's (2007) competitions were held in Atlanta GA and featured over 300 human teams from 33 countries and robots that included humanoid, wheeled, tracked, and flying varieties. The widely varied competitions highlighted competencies in domains that included hardware design, sensor processing, and planning and strategy algorithms. In addition to the competitions that featured real robots, several virtual competitions took place in the areas of robotic soccer and USAR. Figure 1 depicts samples of the virtual worlds that were used in the Virtual Robot Competition (VRC) of the USAR Simulation League.

The VRC is the basis for the design of VMAC. The VRC, which in 2007 saw its second annual running, is designed to foster collaboration and competition between research groups conducting research in the diverse areas of human-computer interfaces, map building, the formation of multi-agent communication networks, complex terrain navigation, and victim search and identification strategies. The competition was run over 7 days and consisted of two preliminary pass/fail rounds followed by three main competition rounds, 2 semi-final rounds, and 2 final rounds. The competition featured 8 teams from 5 countries.

The preliminary rounds of the competition were designed to verify that teams met a minimum set of competencies that were necessary to be competitive in the overall competition. Teams needed to control their robots through the use of a provided communications server (a software package that mimics the non-line-of-sight nature of a real disaster location), generate maps and find victims, and provide the judges with maps and victim locations that were in the proper format. All of the teams passed and moved onto the actual competition.

The competition rounds consisted of extensive indoor or outdoor terrain. The goal of the competition was to find as many victims while clearing as much area as possible before the batteries of the robot expired. Robots were given a battery life of approximately 20 minutes. In order to support a wide variety of research interests and lower the competition entry barriers by assuring that teams did not need to be experts in all fields, *a priori* data was provided on the difficulty of terrain traversal, of communicating with the base station, and of finding victims.

The competition is specifically designed to foster collaboration and learning amongst the participants. The duration of the competition allows teams to apply lessons learned to future rounds, and all of the final code becomes open source (posted at www.robocuprescue.org/wiki/). This allows the following year's teams should be a least as good as the best team from the current year. In addition, the top teams joined together to publish a joint journal article [2] that provides specific details on the various algorithms.

3. VIRTUAL FACTORY AUTOMATED TECHNOLOGY COMPETITION

The successful competition model in the area of Urban Search and Rescue prompted the National Institute of Standards and Technology (NIST) in cooperation with the IEEE Robotics and Automation Society decided to begin a new competition area in factory automation. To begin with, this competition will address the area of AGVs operating in unstructured environments. Under this effort, faculty members and their interested students from six universities in the Greater Washington Area (Washington D.C., Northern Virginia, and Baltimore) were introduced to this time-critical research area through the creation of a factory automation regional competition and tutorial. The first running of this competition will take place in April of 2008 at NIST. The top two teams will then be invited to participate in a competition demonstration at the International Conference on Robotics and Automation (ICRA) that takes place in May 2008 in Pasadena, CA.

3.1 Competition Rules

The actual rules of the competition are quite simple. There is a single fixed package loading station and a number of fixed unloading stations scattered around a maze-like factory floor. Each team must control multiple AGVs to load packages from the loading station, navigate through the factory floor, and deliver the packages to an unloading station that is specified in a radio-frequency identification (RFID) tag. This tag is mounted on the package. Points are awarded for successful package pickup and delivery. Points are subtracted for collisions, dropped packages, and incorrect deliveries.

The initial environment is constructed in such a way that the width of a route from the pick-up to drop-off point is proportional to the length of the route (with wider routes being longer). The areas around the drop-off points may also be cluttered, requiring some form of traffic management to allow for multi-vehicle operation. This initial configuration is specified in an *a priori* vector data file that is provided to the teams shortly before the competition. As points are accumulated during the competition, the environment becomes increasingly dynamic and difficult. The changes that will occur are not known beforehand by the competing team, and will take forms including:

- "Spills" that will block passageways. The AGVs will be forced to take narrower and narrower passageways.
- Passageways that are blocked in the *a priori* map may open.
- Dynamic objects may be added to the environment (i.e. people and human-operated mobile platforms). These objects may intermediately block access to certain unloading points.

In addition to the overall competition, "best-in-class" awards will be available for teams that perform the best on specific performance metrics. This includes such areas as docking the AGV to a loading/unloading table in a confined space, producing an updated map of an area without the use of a ground referencing system, and the vehicle that is able to negotiate the passageway with the least clearance.

3.2 Competition Objectives

While AGVs represent an important part of today's manufacturing processes, there are many areas that represent opportunities for improving the performance and range of applications of these systems. According to Bishop Consulting's report [4] on AGV Industry Next-Generation Technology Priorities, "In the eyes of the system vendors, the most prominent technology development area is in moving from today's AGVs, which require highly structured environments and reference markers installed throughout the plant, to operating in less structured or unstructured environments. In fact, the site preparation required to site and install these reference markers are a significant portion of the system cost… In implementing AGVs for unstructured environments, another need is for sophisticated simulation capability… no capability currently exists to simulate AGVs for unstructured environments". The idea behind the VMAC competition is to foster research and collaboration in this important area and to advance the state-of-the-art for AGV systems. In order to accomplish this, the VMAC competition is designed to contain aspects of AGV systems operating in unstructured, dynamic environments while striving to serve the three different communities of test developers, researchers, and industry.

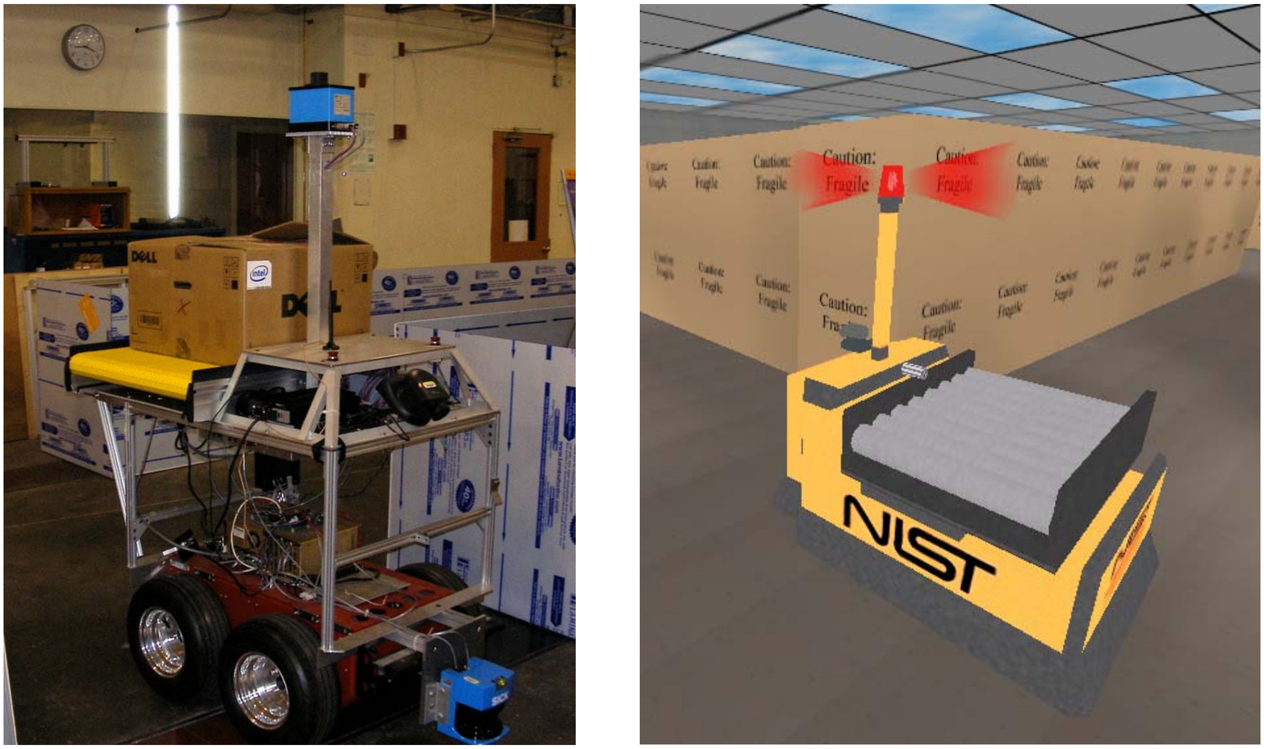

Figure 2: A performance metric that tests required AGV side clearance during turns and its virtual representation. This metric provides an adjustable width passageway. The robot is repeatedly driven through the metric with the width of the metric being reduced after each run.

As the test developer, NIST must solve the problem that there is currently no standard definition of what a "less structured" or "unstructured" environment consists of. In addition, there is no technique to measure the performance of these systems in such environments. Through discussions with both current and potential end users of AGV systems, we are attempting to define these terms and to distill new areas where next generation AGV platforms could prove useful. The discussions begin by understanding where users would like to apply automation, and why current systems are not up to the task. Once technology challenge areas have been identified, performance metrics may be developed that measure the performance of AGVs attempting these specific tasks. These performance metrics may then be built into the competition environment where they may both be evaluated for effectiveness and used as challenges for the research teams. An example of one such challenge is that many older, preexisting manufacturing sites were not designed with AGVs in mind. These factories may have very narrow passageways and limited areas where turning and docking maneuvers may be performed. A metric that was derived from this challenge is shown in Figure 2 (robots running the metric may be found in the video proceedings). The left-hand side of this figure shows a physical performance metric that tests the ability of an AGV to navigate in an adjustable width passageway with a turn. The passageway width may be varied to determine the minimum clearance that is required for successful navigation. In the virtual world, various width passageways are provided along with *a priori* information. This allows the teams to plan for which routes and thus which passageway widths will be attempted.

The researchers participating in this competition are all performing research into embodied intelligent agents. These agents require an embodiment (e.g. a robot), an environment (e.g. the world that the robot interacts with), and intelligence (e.g. some form of control system that causes the robot to respond to the environment). The competition structure is striving to provide relevant, interesting problems that will engage their teams while providing a means of objectively evaluating their performance. While they are looking for interesting and challenging problems, they are not looking to invest too many resources in physical infrastructure or research areas that are peripheral to their main thrust. In the RoboCup VRC, this was solved by having the Unified System for Automation and Robot Simulation (USARSim, described in Section 4.1) provide the embodiment and environment for the robots and the teams provide the intelligence. This freed the researchers from needing to maintain robot labs and complex test environments.

For VMAC, even without these infrastructural needs the individual teams would still need expertise in control issues ranging from fine motor control, to path planning, to multi-agent scheduling. Several of the VMAC teams had little experience in one or more of these areas. Furthermore, the level of performance that would be required to be competitive was beyond reach for this year's competition. For these reasons, additional support was offered. This support came in the form of Mobility Open Architecture and Tools (MOAST), a public domain control system that was able to perform the basic operations necessary to compete in the competition (described in Section 4.2). Our industrial partners are interested in novel ways to apply automation to their processes and in techniques that will gauge how successful this automation will be. The competition provides them with an infusion of new ideas and algorithms as well as a means to judge the maturity of the technology.

3.3 Scoring Metrics

Scoring may be decomposed into the areas of the formal competition and the "best-in-class" activities. For the actual competition, the scoring is based solely on the number of boxes successfully picked-up and delivered, and the number of collisions, dropped boxes, and incorrect deliveries. The teams have expressed that for this first competition, they prefer that the environment remain static and that they have perfect ground-truth and *a priori* information. These restrictions will be relaxed for future competitions.

Figure 3: Example of a Unit Loader AGV approaching a loading station for docking.

The best-in-class competitions have their own unique scoring metrics. Teams will be able to participate in additional competition rounds that are designed to showcase individual technologies. The first of these areas is in docking with a loading station. A simulated Unit Loader performing a docking maneuver is shown in Figure 3. For this test, the team's vehicle will sequence through a constellation of starting locations, where each location specifies a position and orientation. The number of successful docking maneuvers will be recorded along with the time that it takes for each docking procedure. If multiple teams are successful, the amount of free area for maneuvering will be reduced and the test will be performed again. This will continue until all teams are unsuccessful in docking.

The next best-in-class competition is in the area of map building without the assistance of ground truth. This technology is useful both for creating an initial map of the facility and for keeping the map up-to-date during normal operation. Since this year's competition will not include a dynamic environment, the mapping best-in-class competition will concentrate on the creation of an *a priori* map that may be used to set-up the AGVs and that will be used during normal operation. For this competition, the boundaries of the machine accessible area will be provided and the teams will place their robots in an "explore" mode. The deliverables that will be judged are a vector and/or raster map of the facility.

The judging of the team-generated map products will leverage the RoboCup VRC map performance metrics. These metrics judged the map products based on feature quality, multi-vehicle fusion, attribution, grouping, and utility. All of the delivered maps are required to be delivered as geo-registered images with specific color mappings or vector files. This allows the judges to directly compare competitor's maps to ground truth using geographic information services

(GIS) software. For VMAC, the areas of feature quality (metric accuracy), multi-vehicle fusion, and skeleton quality will be evaluated.

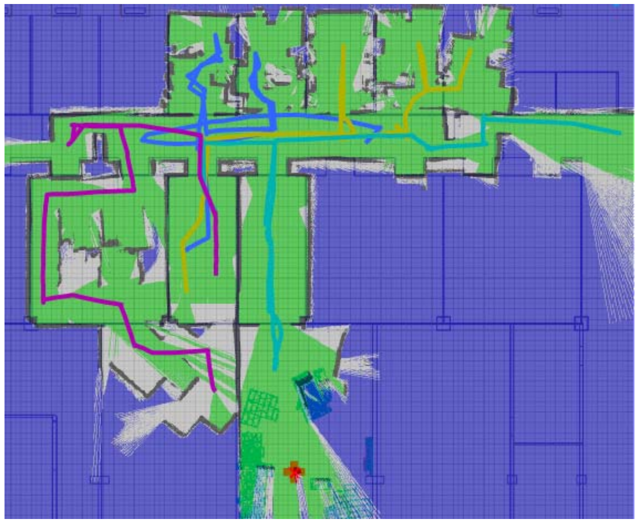

Figure 4: Competitors map overlaid on ground truth from RoboCup 2007. Annotations on RoboCup map include area explored (gray), area cleared (green), victims located (red cross), and robot paths (multi-colored lines).

- Feature quality (metric accuracy) – For the RoboCup VRC, feature quality was evaluated subjectively. As shown in Figure 4, geo-registered maps were overlaid on ground truth and were examined for the number of discrete errors. For example, on some maps it was obvious that a single error led to a piece of the map being rotated. False obstacle reports (a single wall being reported in multiple locations) and scaling issues were also noted. The maps were ranked from best to worst and then assigned points based on their ranking. In the case of factory automation, it is critical that a map be metrically accurate. Therefore, for VMAC a correlation will be performed between ground truth and the delivered maps. Points will be awarded on a per square meter basis for matching ground truth and subtracted on a per square meter basis for errors. RoboCup VRC is examining moving towards this criterion as well for the 2008 competition. The ability to automatically apply metrics will greatly simplify the competition scoring.
- Multi-vehicle fusion – As in the RoboCup competition, teams will only be permitted to turn in a single map file. Those teams that included the output from multiple robots in that single map will be awarded bonus points. The reason for this is that utilizing multiple vehicles speeds the map building process.
- Skeleton quality – A map skeleton reduces a complex map into a set of connected locations. The purpose of generating these maps for the manufacturing environment is to allow the automatic determination of AGV routes. For skeleton quality, several start and end locations will be provided and the teams must generate routes based on their collected maps. In addition, minimal route widths will be specified that the teams must conform to.

The final best-of-class competition area will be navigating through narrow passageways. The object of this competition is to determine how well teams are able to perform fine motion control and localization. *A priori* data will be provided on a series of passageways. The further into the series that one precedes, the narrower they become. This test will be conducted with and without ground truth. When ground truth is provided, the test will measure the team's abilities to control and plan for the motion path of their platform. When ground truth is not provided, a measure of the decrease in system performance will be obtained. This decrease should be due to localization and mapping errors and will be used to judge how the systems will perform in narrow unstructured environments. The distance traveled into the passageway, the number of collisions, and the time of traversal will all be utilized in the scoring criteria.

4. SUPPLIED INFRASTRUCTURE

In order to provide a uniform low-cost, no-maintenance platform to all participants, it was decided that the competition would be run in simulation. The highest achieving teams would then be given the opportunity to directly port (no changes in software) their code to a NIST supplied physical robot shown in Figure 2. The use of a simulation for the physical hardware also provides the benefit of having a common software interface at the robot hardware level. This helps to enable software sharing and compatibility amongst the various teams. The underlying simulation framework that was chosen for the competition is USARSim. This is the same system used in the RoboCup VRC.

4.1 USARSim

The current version of USARSim [3] is based on the UnrealEngine2[1] game engine that was released by Epic Games as part of Unreal Tournament 2004. This engine may be inexpensively obtained by purchasing the Unreal Tournament 2004 game (currently available as part of the Unreal Anthology). The USARSim extensions may then be freely downloaded from sourceforge.net/projects/usarsim. The engine handles most of the basic mechanics of simulation and includes modules for handling input, output (3D rendering, 2D drawing, and sound), networking, physics and dynamics. USARSim uses these features to provide controllable camera views and the ability to control multiple robots. In addition to the simulation, a sophisticated graphical development environment and a variety of specialized tools are provided with the purchase of Unreal Tournament.

The USARSim framework builds on this game engine and consists of:

- standards that dictate how agent/game engine interaction is to occur,
- modifications to the game engine that permit this interaction,
- an Application Programmer's Interface (API) that defines how to utilize these modifications to control an embodied agent in the environment ,
- 3-D immersive test environments,
- models of several commercial and laboratory robots and effectors,
- models of commonly used robotic sensors

USARSim does not provide a robot controller. However, several open source controllers may be freely downloaded. These include the community-developed MOAST controller (sourceforge.net/projects/moast), the player middleware (sourceforge.net/projects/playerstage), and any of the winning controllers from previous year's RoboCup competitions.

During the workshop, details were provided on the 4D/RCS reference model architecture and the MOAST controller that provides a sample implementation of the architecture. It was agreed that NIST would provide a baseline system that was capable of performing the competition task, and that the Universities would pick specific areas to improve upon.

4.2 MOAST

MOAST is a framework that provides a baseline infrastructure for the development, testing, and analysis of autonomous systems that is guided by three principles:

- Create a multi-agent simulation environment and tool set that enables developers to focus their efforts on their area of expertise without having to have the knowledge or resources to develop an entire control system.
- Create a baseline control system which can be used for the performance evaluation of the new algorithms and subsystems.
- Create a mechanism that provides a smooth gradient to migrate a system from a purely virtual world to an entirely real implementation.

MOAST has the 4D/RCS architecture at its core and consists of the additional components of control modules, interface specs, tools, and data sets. MOAST is fully integrated with the USARSim simulation system and may communicate with real hardware through the Player interface.

---

[1] Certain commercial software and tools are identified in this paper in order to explain our research. Such identification does not imply recommendation or endorsement by the National Institute of Standards and Technology, nor does it imply that the software tools identified are necessarily the best available for the purpose.

MOAST provides an implementation of primitive echelon through the section echelon of the 4D/RCS reference model architecture. This implementation is not designed to be complete, but rather is designed to provide examples of control strategies and a starting point for further research. The framework provides methods of mobility control for vehicles including Ackerman steered, skid steered, omni-drive, helicopter-type flying machines, and water craft. Control modalities increase in complexity as one moves up the hierarchy and include velocity/steering angle control, waypoint following, package delivery in a factory environment, and an exploration behavior.

All of the control modules are designed to be self-contained and fully accessible through well-defined interfaces. This allows a developer to create a module that conforms to the specifications and replace any MOAST provided system. The idea is to allow a researcher to utilize all of the architecture except for the modules that are in the area of their expertise. These modules would be replaced with their own research code.

In order to simplify this replacement, several debug and diagnostic tools are also provided. These allow for unit testing of any module by providing a mechanism to send any command, status, and data into a module that is under test. In this way, a module may be fully and repeatably tested. Once the system is debugged in simulation, the standardized interfaces allow the user to slowly move systems from simulation to actual hardware. For example, planning algorithms may be allowed to control the actual vehicle while sensing may be left in simulation.

5. CONCLUSIONS AND FUTURE WORK

While the first competition will not occur until April, planning for next steps is already underway. The next physical metrics that will be constructed will be a docking metric and a navigation solution metric. In addition, as part of the Performance Metrics for Intelligent Systems (PerMIS) conference that takes place in August 2008, a mapping camp will be held to evaluate and explain different navigation solutions. In addition, plans are underway for a special session at PerMIS to highlight accomplishments that have been made by the various Universities. Finally, planning is also underway for next year's competition. New teams are currently being sought to participate.

REFERENCES

[1] Albus, J. S., "RCS: A Reference Model Architecture for Intelligent Control," *Computer*, Vol. 25, No. 5, 1992, pp. 56-59.

[2] Balakirsky, S., Carpin, S., Kleiner, A., Lewis, M., Visser, A., Wang.J., and Ziparo, V., "Towards Heterogeneous Robot Teams for Disaster Mitigation: Results and Performance Metrics from RoboCup Rescue," *Journal of Field Robotics*, Vol. 24, No. 11-12, 2007, pp. 943-967.

[3] Balakirsky, S., Scrapper, C., Carpin, S., and Lewis, M., "USARSim: Providing a Framework for Multi-robot Performance Evaluation," *Proceedings of the Performance Metrics for Intelligent Systems (PerMIS'06) Workshop*, Vol. NIST Special Publication 1062, NIST, Gaithersburg, MD, USA, 2006.

[4] Bishop Consulting, "AGV Industry Next-Generation Technology," Contract 2005-08-08-01 (KT Consulting), Apr. 2006.

[5] Bostel, A. J., Gan, W. W., Sagar, V. K., and See, C. H., "Generation of optimal routes in a neural network based AGV controller," *Second International Conference on Intelligent Systems Engineering*, 1994, pp. 165-170.

[6] Chao-Huang Wei and Shang-Ping Chen, "VR-based teleautonomous system for AGV path guidance," *7th International Conference on Control, Automation, Robotics, and Vision (ICARCV 2002)*, Vol. 3, 2002, pp. 1262-1267.

[7] Lima, P. and Custódio, L., "RoboCup 2004 Overview," *RoboCup 2004: Robot Soccer World Cup VIII*, edited by D. Nardi, M. Riedmiller, and J. Santos-Victor Lecture Notes in Artificial Intelligence (LNAI), Springer, 2004, pp. 1-17.

[8] Naiqi Wu and MengChu Zhou, "AGV routing for conflict resolution in AGV systems," *IEEE International Conference on Robotics and Automation (ICRA '03)*, Vol. 1, 2003, pp. 1428-1433.

[9] Shih-Jie Chang, "Autonomous fuzzy parking control of a car-like mobile robot," *Systems, Man and Cybernetics, Part A, IEEE Transactions on*, Vol. 33, No. 4, 2003, pp. 451-465.