Four Generations of Robotic Mapping and Exploration in Extreme Environments

Scott M. Thayer

Robotics Institute
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213

sthayer@ri.cmu.edu

Abstract

An overview of four generations of robotic mapping and exploration systems for operations in extreme environments is presented. Theoretical and empirical insights gained during research, experimentation, and field expeditions with map reconstruction, navigation and object recognition are illustrated in two distinct classes of environments: nuclear response and military operations. In particular, the issues associated with the deployment of monolithic vs. distributed agents are examined with respect to each environmental classification. Experimental and simulated results detailing mapping operations with robot populations ranging from 1 to 1000 agents are discussed with respect to mission objectives and the environment type.

1. Introduction

Extreme environments have been characterized by their ability to inflict damage on exposed agents that will ultimately be manifested in component, sub-system or system level failures [1]. Two types of environments that differ in the character of the presented hazards and the types of robotic systems required to successfully executing mapping and exploration missions therein are examined. Table 1 provides a look at the dominant characteristics of each environment

The first type of environment is dominated by incremental damage through exposure (radiological, chemical, etc.) that accumulates in susceptible components. This damage results in correlated reductions in component performance until utility is exhausted or outright failure occurs. Post containment nuclear response environments can be characterized in this manner. This type of environment is usually structured or semi-structured with a prior information regarding the concentration and/or distribution of hazards. In addition, established procedures for operations within environments of this type are well defined for humans and can be extended to machines with minimal effort.

Table 1: Characterization of Dominant Failure Modes in Monolithic and Distributed Ops.

	Nuclear	**Military**
Failure Mode	Gradual	Sudden
Radiation Exp.	Yes	No
Chemical Exp.	Yes	No
Structured	Yes	No
Elemental Exp.	Minimal	Yes
Prior Knowledge	Yes	No
Adversarial	No	Yes

The second type of environment is dominated by sudden, catastrophic damage that results most often in system-level failures. Damage during battlefield operations is usually incurred by previously functional systems and most often realizes catastrophic losses to the robotic system in terms of its ability to continue effective mission participation. In addition, this environment usually contains an intelligent adversary that may actively seek conflict. Prior knowledge is usually limited and engagement scenarios are highly unstructured and dynamic.

Distinct patterns of failure have driven systems for nuclear and military response down different respective design paths. Post containment nuclear response robots can leverage priors in terms of exposure, facility, and operations knowledge during system design. Custom designs enable monolithic robots to perform during a life expectancy that is governed by the validity of the priors that shaped robot design and expected exposure profiles. These highly specialized and complex robots successfully operate in lethal environments primarily through optimization and contingency planning during design.

Military scenarios, no less lethal, are uncertain and adversarial in nature and require a level of operational redundancy that necessitates using larger number of less complex agents in order to achieve robustness and to insure that loss of individuals has minimal impact on the mission execution. Success in military operations is achieved through the coordinated action of teams of flexible agents.

The discussion begins with the early work done at Carnegie Mellon with mapping and exploration systems for nuclear response robots

2. Nuclear Mapping and Exploration: a monolithic approach

Nuclear accidents often pose radiological, chemical, and fire hazards to those first responders, usually firefighters and HAZMAT teams, whose mission focuses on control and containment. After containment, long-term activities demand cleanup and environmental remediation. Our autonomous systems focus on post-containment operations, where the highest risk to robotic systems is from accumulated radiation and possibly chemical exposure.

Two systems for mapping and exploration class operations within nuclear environments are described. Artisan, a laser-based object recognition system, was designed to address the problem of decontamination and decommissioning of retired U.S. nuclear facilities. The mapping system from the Pioneer robot was designed to provide three-dimensional information from the interior of the failed Unit 4 reactor at the Chornobyl Nuclear power plant [2]. These scenarios are characterized by mid to high levels of radiation exposure,

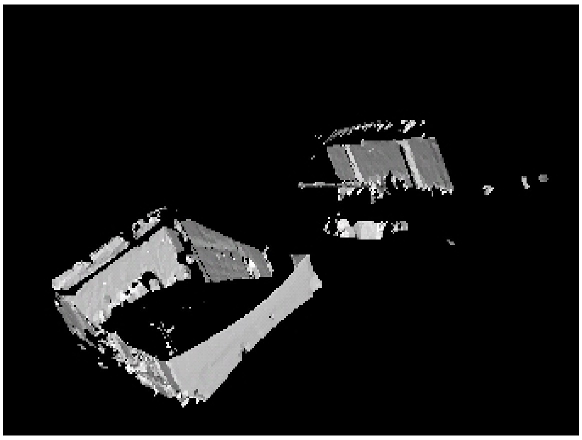

Figure 1: An example range image from the Z&F Scene Modeler**.**

minimal elemental exposure, and the absence of an adversarial element. Artisan was designed to work in highly structured environments with good prior knowledge, such as floor plans, radiation maps, etc. Pioneer was designed to operate in accident regions where any prior knowledge is unreliable and the need to update and discover new information is critical.

2.1 *Artisan: Laser-Based Mapping*

Artisan was constructed at Carnegie Mellon University to perform “As-built” analysis, or inverse CAD, of U.S. Department of Energy Nuclear facilities. An example of a sub-sampled surface mesh [3] constructed from a single viewpoint range data acquisition on the Z&F Scene Modeler is displayed in Figure 1. Artisan’s design centered on two features that were deemed critical for the application class: (1) the accurate and timely reconstruction of facility-level range data sets (minimize exposure), and (2) the identification and localization of key objects of interest within these datasets (sorting contaminated components).

The ability to reconstruction facility level data sets synthesized from individual range images acquired from arbitrary locations within a facility can be decomposed into two key requirements: (1) “matching” and (2) “merging” of co-registered surface meshes into an integrated representation. Together, matching and merging allow a robotic system to co-register overlapping 3-D data sets and to synthesize from that co-registered set a single, integrated uniformly sampled data set suitable for object recognition.

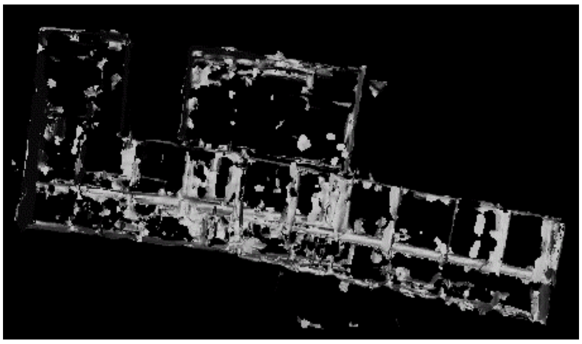

Figure 2: An integrated surface mesh composed of the integration of 32 individual scans

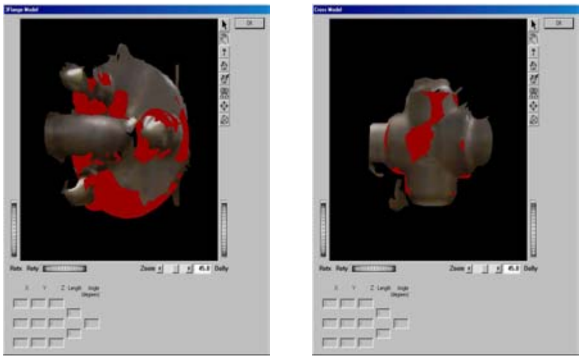

Figure 3: Spin-Image based recognition of standard U.S. industrial components. (a) flange (b) tee

Matching is defined as the process of estimating the relative transformation between overlapping sensor acquisitions. In Artisan, matching is based on Hebert and Johnson's notion of constructing a "spin-image" for a subset of the oriented points that comprise a surface mesh [4]. Once a characteristic set of spin images are defined for each mesh from the partially overlapping range images, local shape similarity measures, obtained by pair wise cross correlation between individual spin-images, are used to provide a set of candidate points matches. Points with similar spin images are considered to have similar local shape. Once correspondences are found between the data sets, the transformation aligning the datasets may be estimated. Global matching, or registration, can be accomplished through successive pair-wise registration of datasets until a common coordinate frame is realized.

Mesh merging is the construction of seamless, textured surfaces from an arbitrary number of co-registered surfaces meshes. It requires both a robust and efficient mechanism for accumulating support evidence in an incremental fashion, and robust operators that extract uniformly sampled surfaces from them. Artisan uses a modified implementation of Johnson and Kang's surface synthesis technology, based on 3-D occupancy grids [5]. Their method represents surfaces using a probabilistic model that encapsulates a sensor error model (stereo or laser) and an artificial point spread function. Each point from a surface mesh, along with the surface normal for that point, is distributed into the occupancy grid using the appropriate models. This evidence, both the surface likelihood estimate as well as the surface normal estimate is accumulated for each point from each input surface mesh. Surface extraction is performed using robust ridge detection on the surface probabilities in the voxel space that forms an implicit surface using the likelihood gradient and the surface normal. Polygonization of the implicit surface is performed using the Marching Cubes algorithm [6]. Figure 2 displays an integrated surface mesh from the "car barn" facility at Carnegie Mellon.

The final requirement for Artisan, recognition and localization of key objects, is a very similar problem to that of mesh matching. As such, a more general form of the spin image-matching algorithm is used [9]. This algorithm can efficiently find multiple known objects within a given scene mesh, storing their position and orientation for future use. The spin image is designed to be resistant to clutter and occlusion--if part of an object in a scene is hidden, recognition is still possible. Using Principal Components Analysis (PCA), the spin images may be compressed, allowing recognition of objects from large libraries in a space- and time-efficient manner [8]. Figure 3 contains a pair of sample object recognitions performed using the Minolta Vivid 700 desktop laser scanner.

Artisan is capable of object recognition and the construction of 3-D photo-realistic models from range data. Taken together, these capabilities serve important roles in the arena of semiautomatic teleoperation within structured nuclear facilities. The production of abstract, fully 3-D models via object recognition is essential to the rapid, accurate, and forceful interaction with a nuclear environment that minimizes exposure and automates object recognition tasks.

2.2 *Pioneer: Stereo based Mapping*

The Pioneer mobile robot (see Figure 4) was designed to enter Chernobyl's sarcophagus and deploy devices to measure radiation, temperature, and humidity; acquire core samples of concrete structures for subsequent engineering analysis; and make photo-realistic three-dimensional (3D) maps of the interior structure [7].

Figure 4: The Pioneer robot

Pioneer utilized a trinocular stereo based mapping head and artificial lighting designed to provide a reliable and effective means to create range maps from indoor scenes with horizontal and vertical edges. Additionally, the mapping system was required to perform the following:

- Co-register auxiliary sensory information with the 3-D surface maps.
- Measure length, surface area, and volume within 5 cm. in each dimension.
- Integrate color texture information with the 3-D surface maps to provide a photo-realistic rendering of the target facilities.

The hardware of the mapping sensor was designed to withstand an environment, which included radiation, invasive decontamination procedures, and the vibration and shock resulting from vehicle motion over rough terrain. The gamma radiation field will result in a 10^6 R of total accumulated dose over the planned mission with dose rates up to 3.5 kR/hr. Average dose rates in Unit 4's reactor room range from 500 - 1000 R/Hr giving the robot a mean lifetime of 1000 hrs. Electronic components on the vehicle where radiation hardened, shielded, or considered sacrificial. For unshielded components, the induced effects of radiation were negligible for the expected dose rates, or transient so that a given mapping session is not interrupted by equipment failure or poor performance of a component. Components deemed sacrificial, e.g. the cameras, where inexpensive enough to be replaced frequently and performed within acceptable limits during periods of maximum exposure. The system was designed to allow any sacrificial components to be easily replaced under field conditions. The mapping camera system (see Figure 5) is removable as a unit so that needless

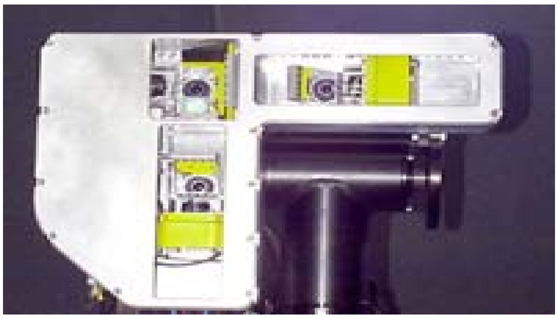

Figure 5: The trinocular stereo mapping head of the Pioneer robot. Cameras are wrapped in 1.27cm of lead.

exposure of the cameras is avoided when the vehicle performs non-mapping operations.

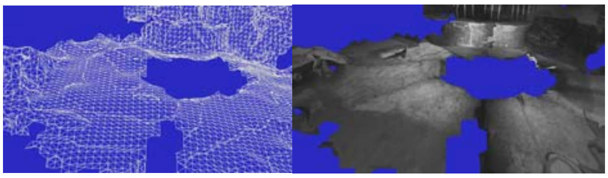

Figure 6: (left) An integrated wire frame mesh of a section of the RedZone mockup, (b) mesh with texture applied

Pioneer mapping software leveraged the mesh matching and integration technology from Artisan to create photo-realistic reconstructions of facilities from stereo range images. Figures 6 and 7 contain sample models from a test facility in Pittsburgh[1]. Visual discontinues are a result of artificial lighting residing on the Pioneer system.

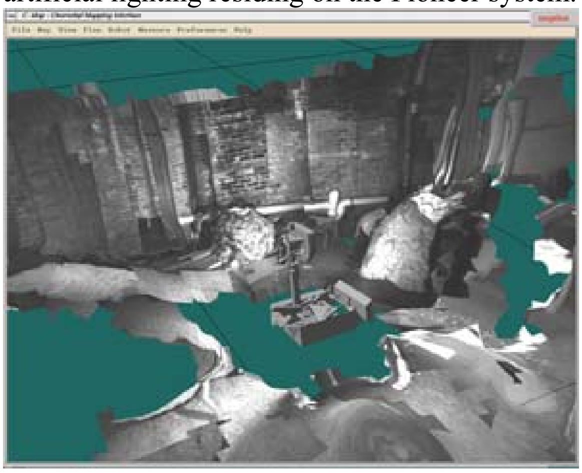

Figure 7: A 3-D reconstruction of the RedZone "mock-up" using 24 textured stereo images

[1] RedZone Robotics, Inc.

Pioneer's mapping unit is capable of the construction of 3-D photo-realistic models from stereo range data in unstructured nuclear environments with lethal radiation fields. Since much about the environment to be explored was known a priori, robot design was optimized around maximum exposure levels and a predetermined set of requirements.

Two systems for post-containment nuclear operations were described for structured and unstructured environments – Artisan and Pioneer, respectively. In the next section, two distributed systems examine the issues in mapping with teams of robots in military scenarios.

3. Military Mapping and Exploration: a distributed approach

Two systems for exploration and mapping in military scenarios are presented. Both approaches focus on the production of reliable, accurate, and fault-tolerant operations by exploiting the redundancy inherent in a team of mobile robots – a robot colony. The first system models multi-agent dynamics in terms of economic activity; thus, the motivation for robot effort is directly coupled to the rewards of task-execution [16]. The second focuses on a model derived from mammalian immunology [12] to simulate control of thousands of robots engaged in a special class of exploration and mapping problems - search and rescue. It uses models of antigen/antibody interaction to motivate massively parallel interaction with the world.

3.1 *Mapping with Market Economics*

A distributed approach enables mapping and exploration techniques that employ teams of robots that scout remote sites, maintain operational tempos, and successfully execute tasks, including

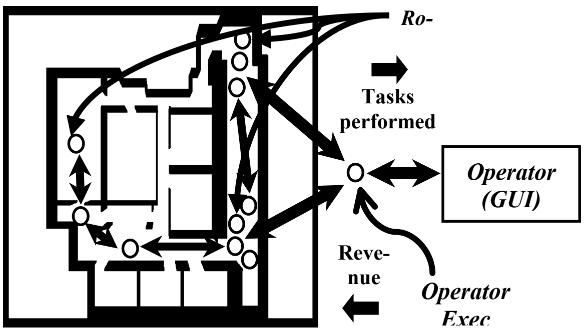

Figure 8: Organization structure for a colony of robots engaged in a mapping and exploration task

Figure 9: A ten-robot mapping and exploration colony

the construction of site maps, despite multiple system-level failures. Using an economic model of agent interaction based on market architectures first described by Dias and Stentz [16], a robotic colony was synthesized where task execution does not directly depend on the status of particular individuals within the colony.

This architecture was verified by tasking a team of mobile robots with distributed mapping [1]. The mobile robots, deployed from starting formation in an unknown world, were assigned the task of mapping the space. The robot colony conceptual structure is illustrated if figure 8 with the actual robots shown in ring formation in Figure 9.

This problem is equivalent to the distributed traveling salesman problem, where the goal points are the cities to visit. Each agent was able to incrementally construct a map of the world, which enabled it to reason about the cost associated with visiting each goal. The costs were the lengths of the shortest paths between goals in an eight-connected grid. The interface between the human operator and the team of robots was a software agent, the *operator executive* (*exec*), which conveyed the operator's commands to the members of the team, managed the team revenue, monitored the team cost, and carried out the initial goal assignments. Once the *exec* had completed the initial assignments, the robots negotiated among themselves to subcontract goal assignments. Only single-goal deals were considered, and the robots continued to negotiate among themselves until no new, mutually profitable deals were possible. Thus, initial negotiations ceased once the system settled into a local minimum of the global cost. As territory is discovered, new goals are injected into the economy and re-bidding occurs within the colony. Figure 10 displays a composite sonar map from five robots engaged in conference room exploration.

ted resp
and men
ntal com
also an
as gene
idual ac
s to grou

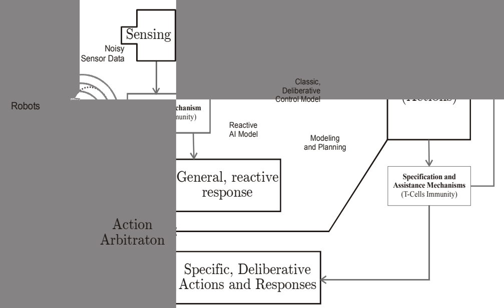

Figure 12: The IDARA software architecture showing potential response paths within Artificial Immune Systems and their mapping onto reactive and deliberate response in robot control theory.

IDARA's ability to command and control large number of robots in the exploration and mapping task is enabled by its focus on general "macroscopic" coordination laws rather than command and control of individual components or systems. Figure 12 displays how IDARA attributes immune system features to high-level control modules that can provide tactical mixes of reactive and deliberate responses that drive kilorobotic colonies to explore a dynamic and uncertain world.

Immunology-based exploration and mapping is applicable in scenarios that are highly unstructured and dynamic and the catastrophic failure of the individual agent is commonplace. The fundamental control mechanism IDARA uses to control the tempo and direction of exploration is the placement of antigens that serve to globally direct the search. In addition, terrain obstacles, other robots, and goal points are modeled as antigens that illicit an appropriate innate or acquired response from individual robots within the kilorobotic colony. As robots move through the world and encounter antigens, IDARA's response mechanisms [12] trigger individual and coordinated action.

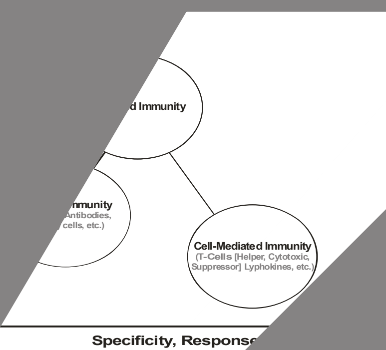

omplexity vs. responsi
Systems (AIS).

describes the rel
xity and the res A's
nation scheme une sys
responds simila eral-resp
ate immunity by a spec

[2] "Kilorob Kilorobot"
las Gag

Figure 13: Mountain cross-section and simulated rubble field

Initial tests addresses the coordination of up to fifteen hundred robots engaged in a search and rescue mission [13] a super class of the exploration and mapping problem, of a simulated mountainous rubble pile (see Figure 13). Each robot was assumed to have a discriminating proximity sensor, i.e. distinguish another robot from obstacle), a local communication mode, and basic mobility. Robot deployment onto the rubble pile was simulated with a uniform distribution of agents onto the rubble surface, a process known as seeding.

Once the colony is seeded, robot operators guide the colony through the injection of virtual antigens into the robot population that bias the search vectors to incorporate any priors that might be available. For example, if scanning data from ground-penetrating radar indicates the presence of a cavity at a particular depth, virtual antigens provide a statistical motivation that favors search in the desired regions.

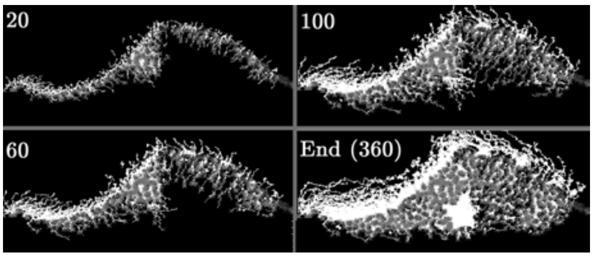

Figure 14: Search progression snapshots through time

Innate immune mechanisms direct the search (see Figure 14) by selecting reactive behaviors that are stimulated by sensing the proximity of obstacle, goal, or other robots. In the absence of sensory stimulation a random exploration strategy is selected. Adding a "stochastic element" to exploration is absolutely critical as it is necessary drive search progress in the presence of uncertain or conflicting information and competing goals.

The visitation map in Figure 15 encodes the relative frequency of visits by robots to a particular location within the rubble grid. As such, it implies a measure of efficiency with respect to the energy expended by the colony during the search operation that is necessary to realize its goal. In this example, 37% of the movements resulted in new information being acquired. Additionally, despite the operators biasing of the rubble search rubble space, IDARA's "directed randomness", or stochastic search, completely covered (explored) the simulated rubble pile. This feature is critical in the exploration of noisy, non-uniform, and dangerous environments. IDARA is distinguished as a multi-agent control architecture in that it manifests coordinated parallel execution, robustness, and reliability in presence of noise, uncertainty, and fragility of the individual search agents.

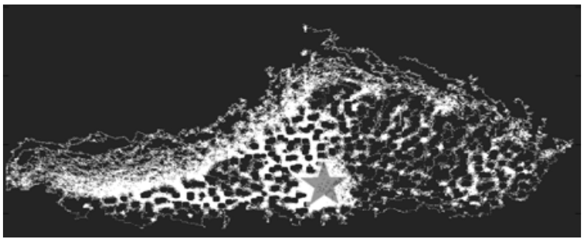

Figure 15: Visitation map (frequently visited areas shown in white)

4. Summary

Four generations of mapping and exploration systems for nuclear and military response were presented. Artisan and Pioneer were targeted at the production of component and structural analyses of nuclear facilities using laser and stereo range data, respectively. They used monolithic robots that incorporated priors regarding expected exposure levels, operational procedures, and facility knowledge directly into design constraints. Successful operations in these lethal environments depended heavily on specialization and contingency planning during robot design and sound operational procedure during field experimentation and deployment.

In military scenarios, expected failure modes necessitate a different paradigm – a transformation to simpler robots with sufficient redundancy (in terms of numbers) together with an effective, nimble coordination strategy. Coordination methods based on market economics and immunological principles where discussed. An economic paradigm to drive exploration and goal exchange in multi-robot systems allows robots to respond to situational and environmental dynamics with ease. This is analogous to the use of military tactics by humans to respond to uncertainty and change on the battlefield. Immunological principles allow an unprecedented economy of scale that enables the deployment of kilorobotic colonies to coordinate in the execution of complex tasks. Further, the adaptability of an Artificial Immune System's innate and acquired response mechanisms drives a probabilistic mix of reactive and deliberate response mechanisms in exploration tasks. This situational mixing of deliberation and reaction enables tactical response to emerge from the interaction of thousands of robotic agents.

References

[1] S. M. Thayer, B. Digney, M.B. Dias, A. Stentz, B. Nabbe, and M. Hebert, **"**Distributed Robotic Mapping of Extreme Environments**"** ", *Proceedings of SPIE: Mobile Robots XV and Telemanipulator and Telepresence Technologies VII*, Vol. 4195, November, 2000

[2] A. R. Sich, The Chornobyl Accident Revisited: Source Term Analysis and Reconstruction of Events During the Active Phase, Dept of Nuclear Engr., Massachusetts Institute of Technology Ph.D. Dissertation, 1994.

[3] J. R. Shewchuk. "Engineering a 2D Quality Mesh Generator and Delaunay Triangulator", First Workshop on Applied Computational Geometry, Philadelphia, PA. pp. 124-133, ACM, May 1996.

[4] A. Johnson and M. Hebert. "Surface registration by matching oriented points." Proc. Int'l Conf. on 3-D Digital Imaging and Modeling (3DIM '97), Ottawa, Ontario, May 12-15, 1997.

[5] A. E. Johnson and S. B. Kang, "Registration and Integration of Textured 3-D Data", International Conference on Recent Advances in 3-D Digital Imaging and Modeling, Ottawa, Ontario, May 12-15, 1997.

[6] W. Lorensen and H. Cline. "Marching Cubes: A high-resolution 3-D surface construction algorithm." Computer Graphics (SIGGRAPH '87), pp. 163-169, 1987.

[7] Maimone, M., L. Matthies, J. Osborn, E. Rollins, J. Teza and S. Thayer, "A photorealistic 3-D mapping system for extreme nuclear environments: Chornobyl," Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, October 13-17, 1998.

[8] A. Johnson and M. Hebert, "Efficient Multiple Model Recognition in Cluttered 3-D scenes." Proc. IEEE Conference on Computer Vision and Pattern Recognition. Santa Barbara, June 1998.

[9] A. Johnson and M. Hebert. "Recognizing objects by matching oriented points." Carnegie Mellon Robotics Institute Technical Report CMU-RI-TR-96-04, 1996.

[10] T.T. Blackmon, S. Thayer, J. Teza, V. Broz, J. Osborn, M. Hebert, and et al, "Virtual Reality Mapping System for Chernobyl Accident Site Assessment," *Proceedings of the SPIE*, Vol. 3644, February, 1999, pp. 338-345.

[11] S. Singh and S.M. Thayer**, "**Immunology Directed Methods for Distributed Robotics: A Novel, Immunity-Based Architecture for Robust Control & Coordination**"**, *Proceedings of SPIE: Mobile Robots XVI*, Vol. 4573, November 2001.

[12] D. Dasgupta and N. Attoh-Okine. Immunity-based systems: A survey. In IEEE International Conference on Systems, Man, and Cybernetics, volume 1, pp. 369-374, 1997.

[13] R. Murphy, C. Lisetti, L. Irish, R. Tardif, and A. Gage, "Emotion-Based Control of Cooperating Heterogeneous Mobile Robots," IEEE Transactions on Robotics and Automation: special issue on Multi-robots Systems, 2001

[14] Ishiguro, T. Kondo, Y. Shirai, and Y. Uchikawa. "Immunoid: An Architecture for Behavior Arbitration Based on the Immune Networks." In Proceedings of the 1996 IEEE/RSJ International Conference on Intelligent Robots and Systems, volume 3, pp. 1730 - 1738, 1996.

[15] Lee, D., Jun, H., et al. "Artificial Immune System for Realization of Cooperative Strategies and Group Behavior in Collective Autonomous Mobile Robots." In Proceedings of the Fourth International Symposium on Artificial Life and Robotics (AROB). Oita, Japan, January 1999.

[16] M. B. Dias and A. X. Stentz, "A Free Market Architecture for Coordinating Multiple Robots" CMU-RI-TR-99-42, Carnegie Mellon University, December 1999.