

Kilorobot Search and Rescue Using an Immunologically Inspired Approach

Surya P. N. Singh¹ and Scott M. Thayer¹

¹ [\[spns.sthayer@ri.cmu.edu\]](mailto:spns.sthayer@ri.cmu.edu)

Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213

Abstract. This paper presents a new concept and simulated results for cooperatively coordinating autonomous robot teams via the Immunology-derived Distributed Autonomous Robotics Architecture (IDARA) to perform autonomous search and rescue operations. Primarily designed for the coordination and control of large-scale, kilorobot colonies, this architecture uses the unique stochastic learning and response mechanisms of the immune system as a basis to yield a more astute and adaptive response so that actions are varied from being reactionary to deliberative as indicated by environmental conditions and the architecture's perceived capabilities to address them. The IDARA architecture exhibits the guided stochastic search characteristics similar to those found in the human immune system. This characteristic was exploited to develop a series of methods for performing terrain search of dynamic environments. These methods were then evaluated in a variety of domains via computer simulations with robot colonies consisting of up to 1,500 robots. These experiments show that the IDARA architecture and framework provides a simple and robust method that is computationally efficient and does not degrade when coordinating and distributing large colonies of robots in either the terrain exploration and mapping or search and rescue problem domains. By providing new levels of scalability in noisy environments IDARA enables the full potential of micro-scale robotic for intelligent exploration, mapping, and search and rescue operations in a manner not afforded by traditional methods.

1 Introduction

Search is an integral aspect of nearly every robotic application ranging from planetary exploration, hazardous environment assessment, urban warfare, to even domestic applications. The use of robots for these exploration tasks minimizes human exposure to harm and automates tedious operations. In certain classes of hostile or dangerous environments the use of robotic platforms may be a necessity.

In this paper, we consider kilorobotics – large-scale, heterogeneous multi-robot teams having populations in the thousands – for exploration of uncertain and potentially dangerous environments that are complicated by variable, dynamic changes [1]. To fully serve the needs of an operator or higher-level layers of an

automation system, these robot colonies need a coordination method that distributes exploration tasks and allocates resources such that not only is the environment fully characterized, but that this is achieved in manner that takes into account any priors. In particular, this paper looks at developing an integrated solution to the terrain search and mapping problem. Although occasionally considered two different problem domains, these two tasks are intricately linked as search can improve its performance based upon any available *a priori* information from available map(s) and/or heuristic data and the map building operations can be handled by the careful acquisition of data by the search operation. Thus, through deft coordination of very-large colonies of robots the performance of these operations can be improved through improved distribution and information sharing.

In nature, we observe several cases where large populations work cooperatively in a cohesive and productive manner to achieve complex goals in a far more efficient manner than may be accomplished individually. Many of these groups of robots or agents consist of large populations that coordinate and cooperate on tasks as needed in the presence of substantial complexity resulting from a variety of factors including environmental uncertainty, noisy inputs, adversarial agents, and external threats. One prime example of this type of system in nature is the human immune system. The immune system is a remarkable example of a highly scalable distributed control and coordination system [2]. In nature, we observe that the human immune system is able to control and coordinate a massively scaled distributed object environment in a measured, decisive, dynamic, and seamless manner to deter bacterial or viral threats. For example, the immune system in an adult male coordinates over a trillion lymphocyte cells, which together utilize about 10^{20} (100 quintillion) antibody molecules. The immune system also responds dynamically to changing macroscopic and microscopic conditions. As an example, in the time it takes to make a cup of coffee the immune system produces 8 million new lymphocytes and releases nearly a billion antibodies. In other words, the immune system acts like a protective force that continually monitors the bioenvironment and, depending upon a perceived threat to the body, activates the necessary multi-agent control systems and responses [3,4].

Just as the nervous system can serve as a powerful construct for building deterministic intelligent systems (e.g., neural net classifiers), the immune system serves as a powerful basis for the design of robot/software architectures that respond and perform learning via a stochastic processes [5]. Although its fundamental goal is pathogen/non-pathogen selection and response, the immune system model gives insights to several methods for autonomous multi-robot colony operations based upon the native exploration methods found within the human immune system [2]. By using this as a basis, kilorobotics will be able to more fully exploit the comparative advantages inherent in autonomous multi-robot systems, namely: parallel execution, redundant operations, increased reliability, and robustness to point failures. That is, the more robots which are available will give the search method extra degrees of freedom and thus result in a greater potential of finding the object under consideration be it a victim, UXO, goal-point, or terrain feature.

2 Previous Work

Search and Rescue using kilorobotic teams is an active area of research. Developments from several areas of robotics research such as exploration and mapping, immunology, multi-robot theory/coordination, and microrobotics) can be combined to provide significant guidance on the design characteristics of the immunology-based multi-robotic architecture needed for this key robotic domain [7,8].

The coordination of very-large populations of robots/agents cannot be performed using traditional coordination strategies nor even by attempting to scale current (often behavioral-AI based) models of multi-robot control. That is, the sheer order and nature of kilorobotic robot populations calls for a completely different level of coordination than traditional multi-robot architectures are designed to support. The performance and operation of such a colony would be primarily determined by the joint interactions of the system and not operations/control of the individual robots/systems comprising the colony [9]. In other words, for this scale of multi-robot coordination, the architecture should focus at a higher level of abstraction, resulting in minimal importance being placed on the individual.

Multi-Robot Search and Rescue Architecture Developments

Several approaches have been developed to enable multi-robotic search. While these approaches vary in scope and development from basic architectures to complex systems, they are based on the same tenant: that communities of agents working cooperatively towards a common goal will do so more effectively and efficiently than if the same agents worked independently [10,11].

There are two general approaches to exploration (and mapping): topological and metric approaches. Topological approaches combine a series of interconnected landmarks that have been augmented with distance information and/or probability (data confidence) information to yield a final search pattern (and map). By comparison, metric approaches are a simpler representation and essentially view the world as an occupancy grid, which may or not be modified using additional data (i.e., probability/data confidence) [12].

Several algorithms and architectures have been proposed for collaborative search and rescue with a team of mobile robots, such as methods for detecting and removing (i.e., rescuing) victims [11, 13-15]. These works also show that from a more abstract robot architecture perspective, search and rescue can be considered a special case of the robot exploration problem; however, with the caveat that it should be performed using a generally directed, but stochastic, search. That is, given a general direction (such as towards certain key goals/targets in the environment) the architecture needs to be stochastic with respect to the exact motions of the robots so that there is significant variation to the exploration path/area [14].

Previous work in this field also suggests that for multi-robotic applications to be effective there needs to be an efficient and intelligent method for control, coordination, and communication that is compatible with the hardware platforms on which it will be executed. As a general rule, the millibots and micro-bots on which this architecture will likely be executed will have very limited memory and computational resources (e.g., 8-bit MicroPIC) [6,16].

Immune System Overview

To appreciate the operation and interactions within artificial immune systems, it helps to have a general understanding of the immune system on which the IDARA metaphor is based. The human immune system works on two levels both with the general goal of pathogen control: a general response mechanism that is not directed at any specific pathogen (i.e., innate immunity) and a specific, anti-body mediated response that encompasses many of the pattern recognition and situational memory aspects that are a core aspect of the human immune system (i.e., acquired immunity). Figure 1 illustrates the specificity ladder between response and effectiveness that governs the immune system.

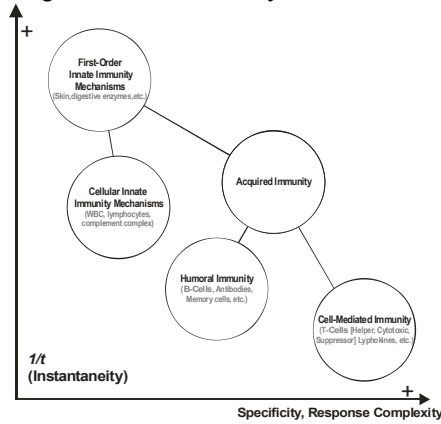


Figure 1: Cascading Response Model for Immune System Responses (Response becomes more specific and advanced with time)

Innate immunity is the natural and omnipresent resistance to a variety of pathogens. Its purpose is to act as the first-order, general defense mechanism. These innate mechanisms then couple with principal members of the acquired immune system to form a rapid, yet targeted, response that uses gradient descent as its primary recruitment method. This mechanism operates by permitting self/non-self discrimination and by activating certain general kill mechanisms [3].

In contrast to the innate system, acquired immunity is about specific responses to specific and known threats. Specific higher-level responses provide life-long critical immunity (e.g., a person with normal immunity can survive up to 100,000 times the dose/exposure of a pathogen that would be lethal without having acquired immunity). There are two types of acquired immunity: humoral (i.e., B-cells and antibody control/regulation) and cell-mediated (i.e., T-cells providing B-cell assistance and orchestration). Both are initiated by antigens and signaled by antibodies (i.e., Y-shaped molecules that match key proteins based upon their encoded specificity; there are some 10 million in the immune system) [4]. Together these operations give a recruitment mechanism (clonal expansion) where the recognition of a pathogen sets off a chain-reaction that generates a large population of antibody producing cells specific to the recognized antigen. In addition to scaling the response, this mechanism acts to provide a stochastic form of learning [5].

Simple Artificial Immune System (SAIS) Models

As detailed in the previous section, the immune system is a remarkable example of a highly scalable distributed control and coordination mechanism. Its unique and powerful characteristics have spawned the development of variety of algorithms from negative selection to anomaly detection [2,5]. As elaborated in [8], the main thrust of this research is not to mimic the immune system's operation, but rather to use its operation as a model for the construction of methods that coordinate large numbers of largely independent agents. This has resulted in the relatively new field of simple artificial immune systems (SAIS). Often implemented through a probabilistic approach based on Jerne's Idiomatic Network Hypothesis these techniques use a simplistic model of the acquired immune system as a new, intelligent problem solving technique [3,5]. While principally being researched in software-maintenance applications, SAIS methods suggest that adoption of a control architecture based on the immune's systems compound architecture will result in a powerful, yet dynamic, multi-robot control and coordination schemes [17]. For example, this technique has been successfully used to as a mechanism for mediating behaviors in behavior-based AI systems and to perform a variety of "fuzzy" tasks (e.g., task classification, network generation, and interrelations) [18,19].

3 IDARA Architecture Design

IDARA's central tenet is that immunology is a promising approach to the command and control of unprecedented numbers of robots. By focusing on the solution of general macroscopic guidance and coordination issues, rather than specific individual command and control, IDARA has lead to the development of a self-optimizing and dynamic robotic control architecture. While the current research has emphasized the use of these algorithms towards the development and demonstration of a first-order distributed robotics system, it is envisioned that the intelligence and robustness inherent to IDARA can be extended to other robot domains (e.g., to aid in task planning and allocation).

The core of the IDARA architecture is derived from the cascading level of response model for immune system operation, which illustrates the body's tradeoff between response time and effectiveness (see also Figure 1) and clonal expansion concepts that govern the operation and performance of the immune system. The use of the immune system and its stochastic learning mechanism (sometimes referred to as Hebbian learning) were chosen because they provide coordination in a robust, diverse, and non-deterministic manner [5].

One of the principal advantages of IDARA's control model over traditional SAIS approaches is the consideration of the entire response and not just mechanisms based on cell-mediated object recognition [17]. This consideration allows the system to respond quickly via a directed, but general, method and then focus its response in time as it proceeds through various levels of response. Finally, this model (unlike many SAIS approaches) can include interactions not easily linked to immune cell actions. Using the aforementioned model as a basis, the IDARA architecture was made by basing the fundamental immune functions of the immune system as modules in the software architecture.

The structure of the IDARA architecture is shown in Figure 2 and combines the immune system's multi-tiered response ladder to yield rapid, reactionary responses followed by deliberative responses that are focused and specific. In addition, the architecture uses the clonal expansion characteristic as a feedback mechanism that governs the arbitration process, which would allow a heterogeneous population of robots to tailor and learn responses to triggered events. Via this structure, the IDARA architecture combines the power of classic deliberative, thorough planning architectures, with the relative simplicity and rapid response of reactionary architectures in a unified framework. No longer does an agent's design need to be constrained by traditional instability and recovery criteria, since the failure of an individual agent is not detrimental to the entire system and may actually be beneficial to the overall action.

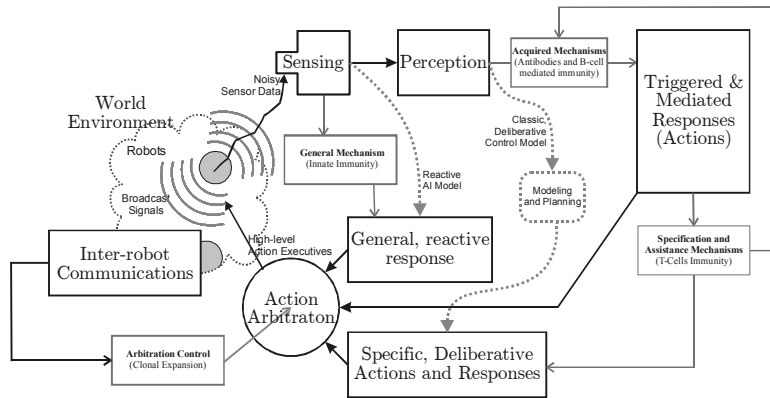


Figure 2: IDARA Software Architecture (Immunology analogous sections are shown in gray and (for comparison) the typical execution paths of reactive and deliberative/planner-based architectures are shown as dashed lines)

IDARA's multi-faceted response mechanism is a central feature of the architecture and the basis of the architecture's dynamic and scalable response mechanisms. Furthermore, it also represents a significant difference between IDARA and traditional SAIS algorithms; in that, IDARA maps different aspects and features of the immune system to various modules and tiers of the response ladder and not as to actions or certain robots. By placing the analogous operations at a high level of abstraction, the IDARA architecture becomes more flexible and easier to implement. The particular details of the modules in the IDARA architecture will vary the scope of the problem being considered; however, it is imagined that the step in responses from *General/Reactive* to *Triggered* to *Deliberative* will be a function of the level and quality of sensor data and/or inter-robot communications.

Inter-robot communications in IDARA are assumed to be a beacon broadcast to a triggered "pathogen" (e.g., the victim in a search and rescue case or a broken robot). The use of beacon broadcasts instead of point-to-point communications not only simplifies the communications model and makes it more compatible with the probable platforms, but also serves as a means of coordinating groups of robots.

Similar to several other architectures, IDARA uses an arbitration module to combine action directives being advanced by various levels of the architecture [7]. However, this operation and control of the module is refined over time as the system attempts to customize its response to beacon signals and other stimuli so that it may respond to the situation at hand and potentially learn a particular stimulus/response pattern.

Finally, the modular aspect of the IDARA architecture makes it compatible with the varying computational resources available on the variety of robotic platforms available as higher, more resource intensive forms of processing associated with deliberative methods can be processed after an initial action has taken place.

4 Search & Rescue Method Design and Simulation

With the general IDARA architecture in place, we examine its use in the design of an immunology-based search and rescue method, which we will then characterize via a computer simulation. The use of this architecture is particularly well suited for the development of a search and rescue method as they share many traits. For example, the highly unstructured environments typical of search and rescue operations benefit from a stochastic search that is capable of self-direction. The specificity ladder of IDARA is also relevant in search and rescue operations, as the method needs to be able to change its coordination strategy as it hones in on the object of the search operation.

A key control mechanism in the IDARA architecture is the placement of antigens in the environment as they not only serve as goal points to partially direct the search, but also as indicators of key environmental features. In the search and rescue architecture developed, the initial antigens were associated with the object of the search. Thus, as the robots (initially moving in an undirected manner) sensed the goal object they were able to broadcast an “antigen find,” which would spark the architecture’s colonial response recruitment mechanism. Furthermore, in order to exit the environment the robots would simply reverse the process by associating a second type of “antigen” with the external environment and repeating the search process to find the most appropriate exit, which may not be the point of entry.

Search and Rescue Simulation

The IDARA architecture was evaluated and tested via the development of a multi-robot search and rescue method for urban and terrain environments. The method was assessed via a series of computer simulations in a manner that would demonstrate the operating characteristics and performance of the architecture and its coordination of robot colonies having up to 1,500 robots.

The search and rescue methods developed were based primarily on the architecture as outlined in the previous section. Each robot was assumed to independently assess sensor inputs from a differentiating bump sensor (i.e., the robot was able to determine if there was an object or robot in an adjacent cell) and beacon signal that it received. From this information, the method calculated future actions as a motion vector according to the architecture’s three levels of operation (i.e., General, Triggered, and Deliberative). From these the arbitration mechanism essen-

tially used a weighted vector summation to determine the resulting motion, where the weights were originally equal, but could be modified by the arbitration control mechanism. If an object of interest (i.e., an antigen) was identified the method was issue a broadcast signal indicating this, which, in turn, was used to recruit neighboring robots to come search the area of interest.

The implementation of the simulation was studied using a MATLAB-based simulation with core libraries run as compiled C++ code integrated via the *Mex* interface in a manner similar to the simulator used to evaluate IDARA as a mobile minefield clearing architecture [8]. This simulation started with an operator indicating the goal points or areas of interest. These were subsequently tagged as a source of antigens in the environment. Using IDARA the robots then proceeded to coordinate in a manner that would identify and control the antigen source (i.e., the goal). Rescue was modeled as a reverse search where the robot would essentially looked for the simplest way to approximately return to its initial location.

In addition, the simulator included a Gaussian noise generator that modeled “practical” problems such as sensor noise, erroneous signal transmissions, robot failures, and signal decay.

The first step in was to setup the model environment for the simulator to search. A terrain rubble scan through a mountain to a cave was selected (see also Figure 3). The robots were initially deployed in a uniform distribution on the surface of the hill and the system was instructed to go from the outside to a goal near the center this structure.




Figure 3: Mountain cross-section and rubble figure that was subsequently studied.

In general, the simulator was implemented as detailed above. However, a couple of features were modified to aid simplify the computation. First, the robots estimated the distance to broadcasting robots using the ATRIA (Advanced Triangle Inequality Algorithm) approximate distance and k-nearest neighbor algorithm. Second, 8-way navigation was performed by applying the algorithm twice – once to determine the unit-step action (i.e., forward one step, backward one step, or no action) along the x -direction and the y -direction. It was assumed that the mechanism was able to accelerate in order to accommodate any diagonal commands and traverse the extra distance.

Finally, the simulator used varying initial distributions of the antigen density to show the effects of this control parameter on tweaking, but not fully controlling, the exploration behavior of the IDARA-based search and rescue method.

5 Experimental Results

The IDARA-based search and method was experimentally evaluated to validate this architecture and to characterize its “directed randomness” behavior. As detailed in the previous section, these experiments were conducted using a MATLAB-based multi-robot simulator package that was developed as part of the IDARA architecture. The results showed that this is a promising technique for coordinating a large colony of robots in highly unstructured environments.

The set of experiments evaluated the nature of the search method. As seen in the visitation map (Figure 4), IDARA was able to coordinate robots to proceed towards and then explore the goal area (marked by a star). Figure 5 is an iteration-lapse sequence that shows the progression of search as the robots proceeded from their initial position (on the surface of this mountain) to the goal in the center.




Figure 4: Visitation Map (areas visited more often are shown in white)




Figure 5: IDARA progression as seen at the 20, 60, 100, and 160 iteration points

For completely unknown environments a random or uniform exploration strategy provides the most efficient method for exploration. However, when priors are available (and can be encoded in the distribution of the antigens) the “directed randomness” of the IDARA method satisfies users goals while maintaining global exploration at the cost of reduced efficiency. For example, in the case of the mountain rubble scan the search was 37% efficient; that is, over a third of the motions made by the robots resulted in new information being collected.

6 Conclusions

We have developed a novel architecture for distributed multi-robot coordination and control for large populations of heterogeneous robots. Using the important human immune system analogues of a specificity response ladder and clonal expansion as a guide, the IDARA coordination architecture was developed along with a method for kilorobotic search and rescue. In general, the results of the simulation were as hypothesized and show that the IDARA methods were able to efficiently coordinate 1,500 robots in a complex task domain, such as the directed-stochastic search needed for performing search and rescue operations. Furthermore, IDARA's ability to perform searches in noisy, non-uniform environments was used in reverse to perform rescue operations after finding the goal (such as a victim in distress) because an efficient method of returning back to the start area in a dynamic environment may not necessarily be to back traverse the path taken.

The IDARA system builds upon immunology models and other related concepts and in the end results in a directed, but flexible, system that mimics that nature of the immune system's control structure. In conclusion, the IDARA method will allow kilorobotics to be able to more fully exploit the comparative advantages inherent in autonomous multi-robot systems, namely: parallel execution, redundant operations, increased reliability, and robustness to noise.

Finally, future versions of IDARA will incorporate three-dimensional terrain searches. This type of architecture will play a significant role as the manufacture and assembly of kilorobotic colonies is realized via recent advances in micro-robotics and MEMS (Micro-ElectroMechanical Systems) platforms mature allowing the mass-production and deployment of arrays of simple, miniscule robots.

7 References

1. D. Gage, private communications on the term "kilorobot," June 2001.
2. S. Forrest and S. Hofmeyr, "Immunity by Design: An Artificial Immune System," Proceedings of the Genetic and Evolutionary Computation Conference (GECCO), 1999, pp. 1289-1296.
3. N. K. Jerne. "The Immune System." Scientific American, volume 259, July 1973, pp. 52-60
4. Guyton and J. Hall. Textbook of Medical Physiology. W.B. Saunders Co., Ninth edition, 1996
5. D. Dasgupta, "An Overview of Artificial Immune Systems and Their Applications," Artificial Immune Systems and Their Applications, D. Dasgupta, Ed., Springer, 1999, pp. 3-21.
6. G. Whitesides and J. C. Love, "The Art of Building Small," Scientific American, volume 270, September 2001, pp. 39-47.
7. R.C. Arkin, and T. Balch, "AuRA: Principles and Practice in Review", Journal of Experimental & Theoretical Artificial Intelligence, Vol. 9, No. 2/3, 1997, pp.175-188.
8. S. Singh and S. 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.
9. R. A. Brooks, "New Approaches to Robotics," Science, Vol. 253, September 1991, pp. 1227-1232.
10. M. J. Mataric, "Issues and Approaches in the Design of Collective Autonomous Agents," Robotics and Autonomous Systems, Vol. 16, 1995, pp. 321-331.

11. R. Murphy, J. Casper, M. Micire, and J. Hyams, "Mixed-initiative Control of Multiple Heterogeneous Robots for USAR," *IEEE Transactions on Robotics and Automation*, 2000
12. G. Dedoglu and G. Sukhatme, "Landmark-based Matching Algorithm for Cooperative Mapping by Autonomous Robots," *Distributed Autonomous Robotic System 4*, Eds: L. Parker, G. Bekey, and J. Barhen, 2000.
13. Kobayashi and K. Nakamura, "Rescue robot for fire hazards," in *Proceedings of the 1983 International Conference on Advanced Robotics*, 1983, pp. 91-98.
14. J.G. Blicth, "Artificial intelligence technologies for robot assisted urban search and rescue," *Expert Systems with Application*, vol. 11, no. 2, 1996, pp. 109-124.
15. 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
16. G. Caprari, P. Balmer, R. Piguet, R. Siegwart, "The autonomous micro robot 'Alice': a platform for scientific and commercial applications," *Proceedings of the 1998 International Symposium on Micromechatronics and Human Science*, 1998, pp. 231-235.
17. 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.
18. 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.
19. 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.
20. C. Merkwirth, U. Parlitz, W. Lauterborn, *TSTOOL User Manual*, January 2001, pp. 13-18.