A Foundation for Kilorobotic Exploration

Surya P. N. Singh and Scott M. Thayer
[spns, sthayer]@ri.cmu.edu
Robotics Institute, Carnegie Mellon University

***Abstract* – A new concept for cooperatively coordinating large-scale autonomous robot exploration teams having populations in the thousands is introduced along with four mapping scenarios. Inspired by constructs found within the human immune system, the Immunology-derived Distributed Autonomous Robotics Architecture (IDARA) was developed so that routine actions are refined and followed by specific and mediated responses based on each unit's utility and capability to timely address the system's perceived need(s). This method improves on initial developments in this emerging area by including often overlooked interactions resulting from the innate immune system to yield a stronger first-order, all purpose exploration mechanism. Using the IDARA architecture as a foundation, this work develops methods for flexible, kilorobotic exploration in dynamic environments. As characterized via a computer simulation with robot populations of up to 1,500 robots, IDARA-based exploration proved to be an efficient, robust, and compact method for large-scale multirobot control that to combine reflexive and deliberative control methods in an opportunistic fashion.**

***Index Terms*: kilorobotics, distributed robots, exploration and mapping, artificial immune systems**

I INTRODUCTION

Exploration is an integral feature of many robotic applications ranging from planetary exploration, hazardous environment assessment, urban warfare, to even domestic servitude. The use of robots for these exploration tasks minimizes human exposure to harm and automates banal operations. However, in hostile or dangerous environments the use of robotic platforms may become a necessity.

In this paper, we consider kilorobotics – robot colonies with large populations (in the thousands) – for exploration of uncertain and potentially dangerous environments that are complicated by variable, dynamic changes. 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 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 [1]. 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. Equally remarkable is the immune system's dynamic nature, which allows it to respond to dynamically 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 [2,3].

Using the IDARA architecture and immunity-based methods, this paper outlines the development of methods for exploration based on the human immune system. 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. They also give the system more degrees of freedom so that it may adapt to a wide range of variations in the environment.

However, kilorobotics comes with a caveat – to be effective they need an efficient and intelligent method for control, coordination, and communication. Without this, parallel resources are misallocated and may be counterproductive. Moreover, kilorobotics is not amenable to classic multi-robot control strategies as they can not be reasonably scaled to this domain. For example, one of the principal characteristics of traditional centralized coordination architectures is that they are highly communications dependent and exponential in complexity [4]. Newer behavioral artificial intelligence (AI) based methods are an increasingly active area of research in this area; however, many of its macroscopic, "bottom-up" (i.e., unified system level) approaches do not have the planning and strategy necessary for operations in complex environments [5]. Therefore, in order to fully reap the potential of kilorobotics, an architecture is needed that can dynamically optimize their function of simple, yet specific, plastic agents, especially in changing, unpredictable environments.

Traditionally the use of immunity-based approaches and artificial immune system models has been as a decentralized behavior arbitration mechanism for behavior-based AI. Using the human immune system as a basis, the IDARA-based methods uses a more extensive model of the immune system to not only arbitrate behaviors, but to coordinate the interaction of heterogeneous groups of robots/agents such that the unique talents of any individual are fully exploited [6-8]. In particular, IDARA's modeling of the general, first-order response of the immune system allows these robots to interact in new environments before they have an opportunity to fully learn or acquire information about these environments. The importance/necessity of this approach is evident by the analogue – when one travels to a "foreign" location their immune system may require time to fully adjust to the environment, but is still capable of providing basic defenses.

The paper investigates the development of this facet of the immunology-based algorithm. As detailed in the next section, the human immune system has evolved into a network of specialized interconnected systems that range from general immune cells to antigen specific lymphocytes. Together these systems perform various levels of immune response and functionality in efficient manner. The IDARA architecture described in this paper uses this facet of the immune system as a control and coordination mechanism for discovery with respect to directed interior exploration. This, in turn, allows IDARA to be capable of responding dynamically and efficiently without detailed data about the area to be explored.

II IMMUNE SYSTEM OVERVIEW & RELATED WORK

On the surface, the human immune system has a clear and basic role: the monitoring and preservation of the identity of the body. The operations of this diffuse system (it is scant more than 1-2% of a person's body weight) are individually simple, but combine to construct a rich and complex web of interaction and coordination that, while not optimal, display exceptional levels of robustness and flexibility, especially with regards to unknown situations and conditions.

1. Innate and Acquired Immunity

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: a general response mechanism that is not directed at any specific disease organism/pathogen (i.e., innate immunity) and a specific, antibody 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). This behavior can be viewed as a tradeoff between response time and specificity/effectiveness and is illustrated in Figure 1.

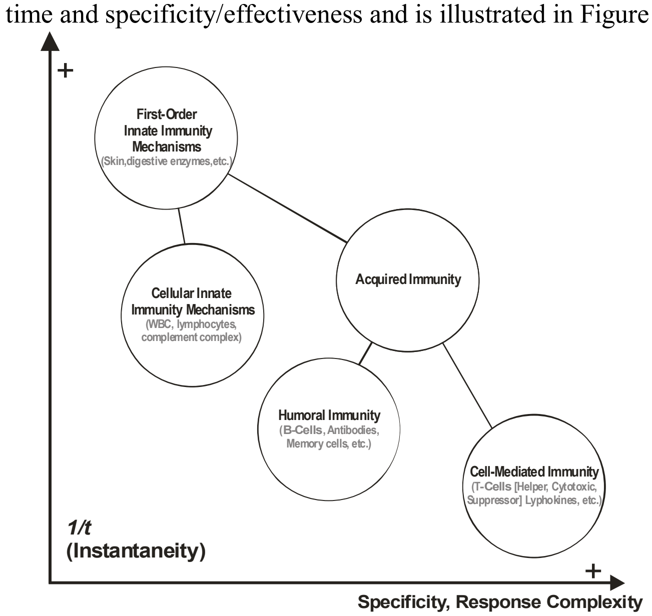

Figure 1: Process Diagram for Immune System Responses (Response becomes more specific and advanced with time)

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 decent as its primary recruitment method (see also Figure 2). This mechanism primarily operates by permitting self/non-self discrimination and by activating certain general kill mechanisms [3].

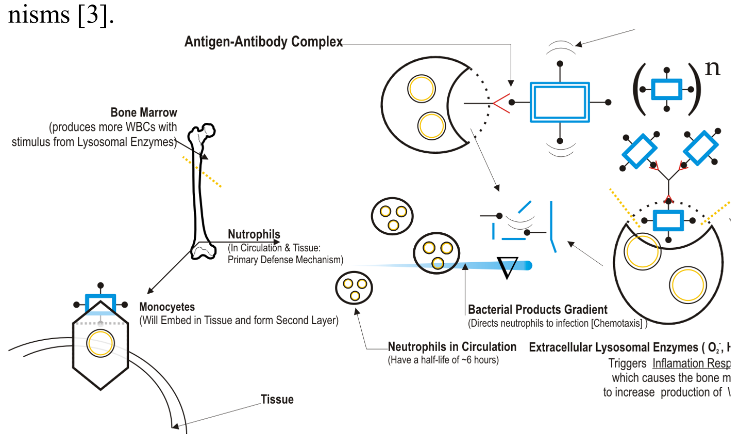

Figure 2: Model of Human Innate Immunity

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 proving B-cell assistance and orchestration). Both are initiated by antigens and signaled by antibodies (i.e., Y-shaped molecules which match key proteins based upon their encoded specificity; there are some 10 million in the immune system) [3].

As detailed in the next section these characteristics have spawned a new and growing research area in the development of algorithms and theories based upon this result.

2. Artificial Immunity Theory

Initial modeling of the human immune system was begun almost thirty years ago with the hope of applying "classical" systems and control theory to the immune system [9].Recently there has been a reversal of roles, now control and multi-robot theory is looking to immunology to gather insight on new methods of control. This has resulted in the relatively new area of artificial immune theory and simple artificial immune system (SAIS) research for a variety of research applications.

The use of immunity-based control in the form of SAIS algorithms is a developing area of research in AI and robotics. Often implemented through a probabilistic approach based on Jerne's Idiotopic Network Hypothesis whereby acquired immunity is used as a model for a new, intelligent problem solving technique. However, theses techniques are based on a very simplistic model of the acquired immune system and do not model more the more advanced learning and communications aspects of the immune system [2,3]. While principally being researched in software-agent coordination applications, the SAIS model research suggests 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 [9].

Most of the activity in this new and diverse field has centered on the modeling and use of the acquired immune system as a mechanism for mediating behaviors in behavior-based AI systems. Hofmeyr and Forrest have described the distributivity and robustness of SAIS [1]. Segal and Bar-Or have described how simple immunology models can be used to optimize effector performance and how the immune system can be seen as a distributed system [10]. Hunt has developed a sophisticated machine learning algorithm set (JISYS) that utilizes SAIS principles to perform a variety of "fuzzy" tasks (e.g., task classification, refinement, network generation, and interrelations) [9-11]. Finally, [12] discusses how to augment and enhance these architectures by employing statistical techniques and a recursive mechanism that varies generalization with refinement similar to the ladder shown in Figure 1.

3. Present Distributed Autonomous Mobile Robotics Approaches

Since the introduction of nouvelle (non-symbolic) AI algorithms over a decade ago, there has been an ever-growing interest in the development of multi-robotic systems. While many of these systems have been in the software-agent domain, many of the concepts under consideration can be successfully applied to the physical domain in the form of a nonlinear control and planning algorithms.

Many of popular multi-robot control systems available for object recovery and detection are based on centralized control and operations. For example, Albus and Stentz both base their results on the centralized, hierarchical approach to control a multi-robot system [13, 15]. While relatively easy to implement, the application and scaling of these systems has often been limited by the large computational and communications burden associated with their (centralized) operation [14].

A second approach is to use a highly distributed robot system that communicates via a series of peer-to-peer or implicit communications systems that are often based on the use of biologically inspired behavior-based control mechanisms. These approaches have been applied in various domains, but can be complicated to scale to larger, more complicated domains as many behavior-based approaches do not provide a convenient method for integration of learning throughout the whole system nor and applying machine learning algorithms (e.g., to filter large levels of sensor noise). Further examples of the principal research efforts in this field are outlined in Matarić's survey paper on distributed robotics [1,14].

Recently, hybrid approaches have been developed to combine the qualities of deliberative, centralized methods and, behavioral architectures. While these approaches resolve many of the problems associated with these two architectures, hybrid architectures have the disadvantage of increased system complexity, which limits how scaleable this architecture is to large heterogeneous colonies [9]. In addition, several more specific architectures exist for use in complex task domains. For example, Dias and Sentz's macroeconomic approach to mobile robot control resulted in a dynamic robot system that can simply and successfully execute tasks in dangerous environments [14]. In addition, Feddema has applied statistical methods and graph-theoretic approaches to coordinating hundreds to thousands of cooperative robotic agents [16].

A second approach that can be applied for coordinating robot colonies in dynamic environments with a relatively slow rate of change is to repeatedly apply a method with guaranteed coverage, such as [17], at a sufficiently high frequency. The most significant problem with this approach is that it essentially entails that the frequency of any environmental change is less than the bandwidth (or "refresh frequency") of the method being considered. Also, most of these methods operate in a "pseudo-steady" mode and make extensive use of the steady-state assumption. Thus, it is possible that transient effects of a dynamic environment could present unforeseen stability difficulties to this method.

III THE IDARA ARCHITECTURE

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). .

One of the principal advances of the proposed immunological control model over traditional SAIS approaches is the consideration of the entire response and not just mechanisms based on cell-mediated object recognition [3]. 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 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. Furthermore, it does so in a diverse manner so that unknown events and dynamic variations can be investigated efficiently. The IDARA architecture uses a multi-tiered response ladder to yield rapid, reactionary responses followed by deliberative responses that are focused and specific. No longer does an agent's design need to be constrained by traditional instability and recovery criteria, since the failure of an individual (disposable) agent is not detrimental to the entire system and may actually be beneficial to the overall action. Via this structure (as illustrated in Figure 3), 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.

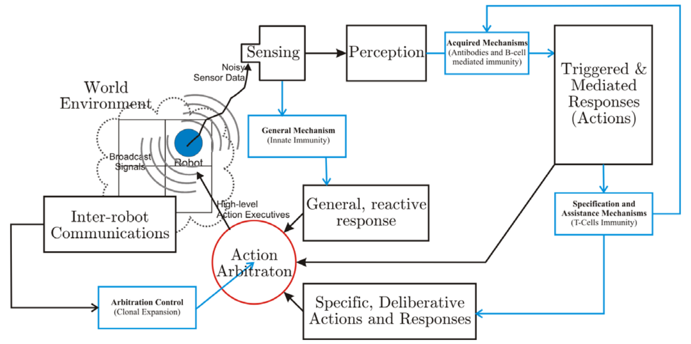

Figure 3: IDARA Software Architecture (Immunology analogous are shown in gray and (for comparison) the typical execution paths of reactive and deliberative/planner-based architectures are shown as dashed lines)

While particular details vary with the implementation and arbitration mechanism, a general description of the standard hierarchy used by IDARA is given as follows:

Unknown Response – The response used in the event when there is no information being perceived (such as during startup/bootstrapping, near completion, or in poor environment where antigens to be discovered are "rare" and sparsely located), the objective is to act in a manner that tries to change this situation in a fashion that does not disrupt the operation of the entire system. For example, this might be to randomly selected an action.

Default (Routine) Response – In the default case where there is some minimum level of information (i.e., sensor inputs are above some sensitivity threshold), the method should apply a general response that is often effective, but may not be optimal or fully exploiting all the information available.

Triggered Response – When the system is able to make a deeper inference or obtains more detailed information, such as a gradient signal communicating a previously located goal, a mediated response is initiated.

Deliberative Reponses – This level of response is the most intensive and complex and includes "machine learning" aspects of the immune system (i.e., the ability to use "memory B-cells" and recall patterns associated with certain cases of a task). In addition, this method would use the internal state and deliberative reasoning to try to develop a more optimal response. This approach allows the architecture a means of exploring around local-optima that would "trap" the triggered response level.

IDARA's multi-facetted response mechanism is a central feature of the architecture and the basis of the architecture's dynamic and scalable response mechanisms. Furthermore, it is 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. In other words, the response mechanisms and actions employed are no longer constrained by the low-level mappings of immune theory to robotic operations. However, this approach does not conflict with traditional SAIS design approaches as they can be implemented by varying the response blocks and placing the onus of control on them. An additional benefit of this design approach is that the IDARA architecture can be expanded to include features and mechanisms from other multi-robot coordination and control architectures, which should accelerate implementation and developmental process.

Similar to several hybrid architectures, IDARA uses a heuristically driven arbitration module to combine action directives being advanced by various levels of the architecture [13]. By using a vector-based approach in combination with the system's default random exploration routines, the arbitration mechanism calculates the best resultant action. The result of this procedure is a form of "directed randomness," in which the architecture varies and tunes the general nature or diction of its response from random exploration to specifically guided paths and actions.

The IDARA architecture has a variety of features compared to other algorithms and AI methods for the coordination of teams of robots (see Table 1). When applied directly at a lower level to a population of robots, this architecture will yield a mobile, robust, and adaptive control method. Such a method will combine the functions and critical mass of simple robots to solve complex tasks. This provides numerous advantages. First, the system will be more robust as failure in one component will have a minimal impact on the entire network. Second, the system will be more economically viable as simple, standard components could be used, as the individual failure modes no longer critically affect the device.

Characteristic	*Coordination Mechanisms*		
	IDARA	Microeconomic Cost Optimization [10]	SAIS Algorithms [12,15]
Massively Scalable	Yes	Yes	Some
Distributed	Yes	Yes	Some
Communications	Light	Medium	Medium
System-wide Approach	Yes	Yes	Yes
Adaptability (i.e., operates outside range)	High	Medium	Little
Learning	Yes	Some	Yes
Behavior AI based	No	No	Yes
Specificity	Yes	Some	Yes
A priori information needed	Can be utilized, but not needed	Some – Cost func. need to be defined	Yes
General/Instant Response	Yes	Some	Some
Fault-Tolerant	Yes	Yes	No
Optimal Solution	No	No	No

Table 1: Comparison of IDARA to Other Algorithms

While the IDARA architecture has a number of strengths, especially in the coordination and control of large robot colonies, it is not perfect. One weakness is that agents initially base interaction on Brownian motion until an antigen is found locally and then use local gradient optimization to follow the signals from initial interactions. This, however, predicates that there is an initial interaction between the two effectors. Thus, this architecture needs an inherent "critical mass" and may not operate well in small teams. Further, gradient techniques are only locally optimal. Thus, in order to obtain a highly (and perhaps globally) optimal solution, IDARA needs to be somewhat random in its initial motion so that it will be fairly well distributed. Finally, IDARA does not place a strong value on an individual unit and therefore can be highly unit sacrificial.

IV IDARA Applied to Kilorobotic Exploration

The primary goal of this method is the coordination of a large team of robots in the exploration of an environment in a manner that incorporates appropriate guidance (e.g., “explore the corners”) from a high-level controller (e.g., human or software agent). A second goal of the method was to any handle dynamic variations.

By focusing on the general and triggered response mechanisms for recruitment and coordinating actions and modeling interaction responses primarily on innate immunity the design of the exploration method was simplified. This also serves to emphasize the reactive attributes of the IDARA architecture (as compared to the architectures deliberative mechanisms).

The robots were modeled with only a classifying proximity sensor in order to detect if it is adjacent to a obstacle or robot and a radio beacon in order to send alert signals. These beacons signals, analogous to histamines in the immune system, were modeled with a decay function proportional to distance traveled and the elapsed time (i.e., simulation iterations) since discovery. Using these beacon signals each robot generated a signal gradient map that was used to evaluate its motion along each axis (dimension) under consideration. Finally, the simulation accounts for Type I and Type II sensor noise up to 10%.

Individual robots do not build the map. The mapping was performed by leveraging the IDARA coordination method. In particular, the “histamines” were received by and mapping was performed by a separate, coordinated set of fixed (and well localized) mapping stations. By sharing information via a network, these mapping stations are able to build a common map at a centralized server that is post-processed and delivered to the user. The map is constructed by triangulation of communications from the robots (i.e., “histamines”) and with any additional information that maybe coded in the received signal (e.g., object type or the robot’s estimate of its position).

The mapping process is shown in Figure 4. Operator inputs are used to vary the nature of the antigen map, which is done by placing antigens in a pattern that reflects the parameters supplied. For example, if no information is given then map would be uniformly distributed; by comparison, if the system is told to explore the center of an environment then majority of antigens are distributed in the center. Once this has been calculated, it is downloaded or set in the robots (e.g., via an initializing data-transfer broadcast) before they are deployed.

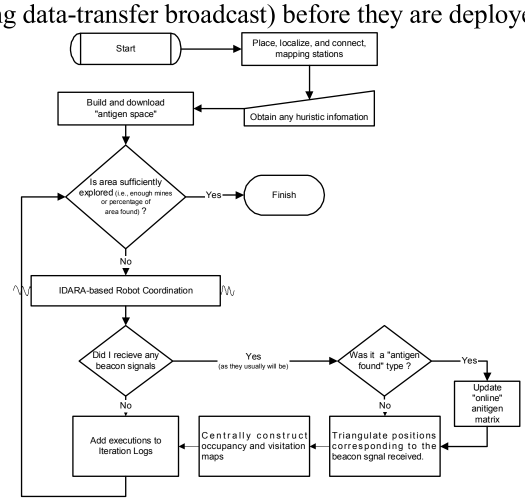

Figure 4: Robot-level Mapping Process

Unlike the antigens, the robots are distributed in a Gaussian pattern around some central point or locations. This was done because it is envisioned that the deployment of these robots would be done by someone locally distributing /bootstrapping the robots (e.g., after downloading the “antigen map”). As one would expect, and as confirmed experimentally, this pattern is less efficient than a random distribution of the robot population. Lastly, dynamic effects can be modeled by having the robot vary the antigen locations on-line.

To simplify the operation, the robot colonies consisted of identical robots whose only sensing was a proximity sensor (with a default misdetection rate of 10%). While the IDARA architecture does provide means for adjusting coordination and response based on heterogeneous populations of robots, the general case was studied as non-specific robots (i.e., teams lacking specialization or prior memory/knowledge) is a less efficient case. A few assumptions about the environment were made to further focus the exploration process. The first is that random trials and sensor noise are Gaussian and thus traditional statistical analysis is valid. Second, it was that the robots were allowed to operate and add costs as necessary (i.e., there is no limit to the cost function). Finally, it was assumed networking issues were not complicating factors. This simplified the initialization processes of the implementation; however, the algorithm does not depend on this assumption.

V Experimental Design and Results

To experimentally validate this architecture and its hypothesized interactions a series of directed exploration experiments were devised. These experiments were implemented, executed, and analyzed using MATLAB. Given a relative area and some general operator constraint, the simulator environment generates the field and distributes antigens within it. The simulation uses a sensor module, which looks at the “solution field” to see if it is adjacent to an antigen.

The goal in this case is to incorporate dynamically varying exploration based upon rough initial guidance from an operator (e.g., move west). This was implemented by focusing on the general and triggered response mechanisms for recruitment and coordinating actions, the system’s complexity is proportional the level of interactions within the humoral (acquired) immune system. This, in turn, does restrict the systems ability to learn and adapt to changing, hostile environments. Generally in exploration and mapping, the obstacles (and their) signatures do not change over the robot’s investigation. Modeling interaction responses primarily on the first-order, innate immunity system instead of the full immunological model not only to simplifies the design (and computation), but also serves to emphasize the reactive attributes of the IDARA architecture (as compared to its pattern recognition mechanisms).

1. Simulator Details

Through the model outlined above, a robot could be considered to mimic a macrophage and the cumulative response of the system to mimic first-order (i.e., innate) human immunity. Before running a variety of simulations to characterize the performance and nature of the IDARA-based exploration, a variety of quick experiments were conducted to establish good default values for the various control parameters that affect the simulation. Unless specially varied or otherwise mentioned the default values for the control parameters are the values tabulated in Table 2.

Variable Name	Default Value/Formula	Execution Order (Empirically derived)
Environment Size	30 × 20 m (650×650 pixel)	O(N2)
Antigen Density	0.60 %	O(N)
Robot Count	$\left(1-e^{-1}\right)P\left(fail\right)+e^{-1}$	O(N+P(fail)×N)
Antigen Motion Pattern	Random	N/A
P(fail)	e-1	N/A
Sensor Noise	10%	N/A

Table 2: Default Values for IDARA Simulation Parameters (Values used by simulator unless parameter being tested)

The simulation tested five types of distribution by skewing the antigen generation method as described earlier. Each experiment was repeated to verify the results. The following types were tested and their mean results shown:

- Center – A Gaussian at the center with: $\sigma_x = \sigma_y = 50$
- Perimeter – Antigens spread towards the edges
- Random – A random distribution in both x, y
- Side - A Gaussian centered along the first column (y=0)
- Uniform – Antigen locations spaced uniformly,

2. *Simulator Results*

Several sets of experimental runs have been performed in order to validate this approach and characterize the performance of the IDARA architecture. As detailed in the previous section, these experiments were conducted using the MATLAB-based multi-robot simulator package (that was developed as part of the IDARA architecture).

Figure 5 is a iteration-lapse sequence that shows the performance and coverage strategy for the IDARA-based exploration system. Its preference for structure recovery or coverage is determined by arbitration rules and the placement and density of antigens throughout the space to be explored, both of which are operator-controlled parameters.

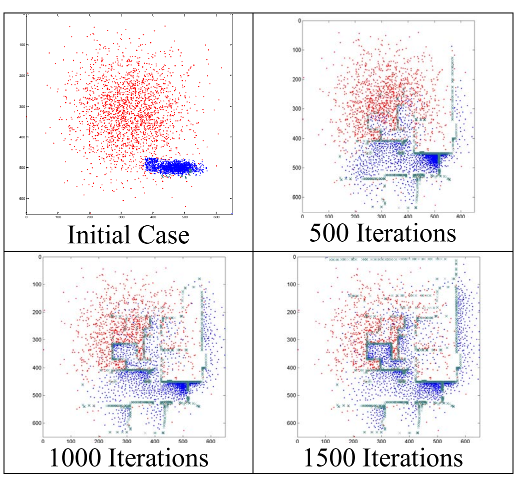

Figure 5: Robot time progression over 1500 Iterations

A key advantage of the IDARA-based method is that user preference can be used to direct the search. This “directed randomness” characteristic/feature targets the exploration process, but still ensures that there is enough variation that unexpected areas or features are not missed.

The visibility plots (see also Figure 6) show the number of times a space was visited by a robot in the colony during the exploration simulation (lighter colors signifying frequent visitation). In addition to showing potential redundancies in the search strategy, it also indicates the general efficiency of the exploration as a centralized method would be able to minimize repetitious visitations. In particular, the visitation maps are from simulations with initially 1500 robots 2500 antigens (i.e., local goals). These plots also show that the directed antigen distributions were efficiently able to guide, the method. Thus, the robots have a mechanism for dynamically changing their mapping behavior without any change to the architecture.

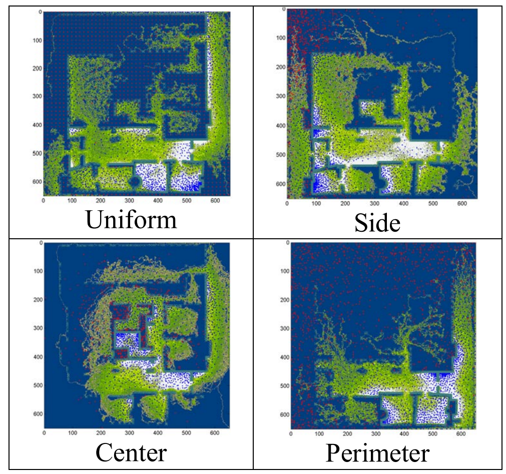

Figure 6: Comparative Visitation Maps

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. Table 3 shows that system was most efficient for disperse distributions.

Method	Iterations (#)	Cells Visited (#)	Energy (steps)	Efficiency (%)
Center	2500	161854	2845543	5.69
Perimeter	1184	126062	1452324	8.68
Random	528	195629	623920	31.35
Side	2561	184474	3122032	5.91
Uniform	692	127888	836436	15.3

Table 3: Summary of Energy Consumption for the Five Types Tested

VI CONCLUSIONS

We have developed a novel architecture for distributed multi-robot coordination and control of large populations of heterogeneous robots in exploration and mapping. This paper discusses the human immune system, its interactions/general control, and the important analogues it presents for robots and automation with large heterogeneous populations. In general, the results of the simulation were as hypothesized and show that the IDARA architecture was able to efficiently coordinate kilorobotic colonies.

The development of an exploration method was based on primarily modeling the interaction and character of innate aspects of the human immune system. This resulted in an architecture that can respond quickly, has a mechanism for learning, and can coordinate a team of robots effectively. In addition to the four cases presented, the simulation found IDARA-based exploration to be a useful mechanism for coordinating the large populations under investigation.

The IDARA architecture is principally characterized by the concepts of an increasingly specific response ladder and arbitration with "directed randomness." Together these will lead to the development of robust, highly effective, flexible model that can respond effectively to unknown situations, are highly efficient, can adapt/learn as new challenges arise, and will be efficient enough so that they can be implemented on hardware platforms with limited computational (and memory) resources (e.g., micro-robots and cellular robotics).

VII FUTURE WORK

It is the goal of the IDARA team to more fully simulate and implement this method for multi-robot control and to use insights gained from these simulations to refine this method of control. It is envisioned that future simulations will more fully integrate more advanced aspects of the immune system (e.g., B-cell learning and T-cell direction). In addition, variations to current techniques will also be investigated (e.g., non-gradient decent based optimization and use of a command history to better suggest subsequent actions).

VIII ACKNOWLEDGEMENTS

The authors acknowledge the assistance of the Field Robotics Center and the Cognitive Colonies project and all its team members. Further, the authors thank Dr. Wesley Thayer (M.D., Ph.D.) for his guidance and expertise. This research and paper were made possible, in part, through support from an Army Research Office sponsored National Defense Science and Engineering Graduate (NDSEG) program fellowship. The authors greatly acknowledge the generosity and leadership of this organization.

IX REFERENCES

1. S. Hofmeyr and S. Forrest, "Immunity by Design: An Artificial Immune System," *Proceedings of the Genetic and Evolutionary* Computation Conference (GECCO), pp. 1289-1296, 1999.
2. N. K. Jerne. "The Immune System." *Scientific American*, volume 259, pp. 52 - 60, July 1973
3. A. Guyton and J. Hall. Textbook of Medical Physiology. W.B. Saunders Co., Ninth edition, 1996.
4. R. A. Brooks. "Elephants Don't Play Chess." In Robotics and Autonomous Systems, volume 6, pp. 3-15, 1990.
5. M. J. Matarić, "Issues and Approaches in the Design of Collective Autonomous Agents," *Robotics and Autonomous Systems*, Vol. 16, pp. 321-331, 1995
6. A. 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.
7. A. Ishiguro, T. Kondo, Y. Watanabe, Y. Shirai, and Y. Uchikawa. "Emergent Construction of Artificial Immune Networks For autonomous Mobile Robots." In IEEE International Conference on Systems, Man, and Cybernetics 1997, volume 2, pp. 1222 - 1228, 1997.
8. A. Ishiguro, Y. Watanabe, T. Kondo, and Y. Uchikawa. "Dynamic Behavior Arbitration of Autonomous Mobile Robots Using Immune Networks." In Proceedings of IEEE International Conference on Evolutionary Computation, pp. 722-727, 1995.
9. 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.
10. L. A. Segel and R. L. Bar-Or. "Immunology Viewed as the Study of an Autonomous Decentralized System." *Artificial Immune Systems and Their Application*, Ed: D. Dasgupta, pp 65-88, 1999.
11. J. Hunt, J. Timmis, et al. "Jisys: The Development of an Artificial Immune System for Real World Applications." *Artificial Immune Systems and Their Application*, Ed: Dasgupta, pp. 157-186, 1999.
12. M. Neal, J.Hunt, and J. Timmis, "Augmenting an artificial immune network," *Proc. of Int. Conf. Systems and Man and Cybernetics*, 1998, 3821-3826.
13. R.C. Arkin, and T. Balch, "AuRA: Principles and Practice in Review", *Journal of Experimental & Theoretical Artificial Intelligence*, Vol. 9, No. 2/3, pp.175-188, 1997.
14. A. T. Stentz and M. B. Dias. *A Free Market Architecture For Coordinating Multiple Robots*. Technical Report CMU-RI-TR-99-42, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, December 1999.
15. A. Stentz and B. L. Brumitt, "Dynamic Mission Planning for Multiple Mobile Robots", *Proceedings of the International Conference on Robotics and Automation*, No. 3, pp. 2396-2401, 1996.
16. D.A. Schoenwald, J. T. Feddema, and F.J. Oppel, "Decentralized Control of a Collective of Autonomous Robotic Vehicles," *Proceedings of the 2001 American Control Conference*, Vol. 3, pp. 2087-2092, 2001.
17. H. Choset, D. Latimer IV, S. Srinivasa, A. Hurst, V, Lee-Shue, Jr., Towards sensor based coverage with robot teams. In Proceedings of the International Conference on Robotics and Automation, IEEE, 2002.
18. 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.