

# Immunology Directed Methods for Distributed Robotics: A Novel, Immunity-Based Architecture for Robust Control & Coordination

Surya P. N. Singh and Scott M. Thayer

Robotics Institute, Carnegie Mellon University  
5000 Forbes Avenue, Pittsburgh, PA 15213-3815

## ABSTRACT

**This paper presents a novel algorithmic architecture for the coordination and control of large scale distributed robot teams derived from the constructs found within the human immune system. Using this as a guide, the Immunology-derived Distributed Autonomous Robotics Architecture (IDARA) distributes tasks 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 area by including often overlooked interactions of the innate immune system resulting in a stronger first-order, general response mechanism. This allows for rapid reactions in dynamic environments, especially those lacking significant *a priori* information. As characterized via computer simulation of a “self-healing” mobile minefield having up to 7,500 mines and 2,750 robots, IDARA provides an efficient, communications light, and scalable architecture that yields significant operation and performance improvements for large-scale multi-robot coordination and control.**

**Keywords:** distributed, robots, self-healing, mobile minefield, demining, artificial immune systems

## 1 INTRODUCTION

While previous efforts have centered around the use of immunity-based approaches as a decentralized behavior arbitration mechanism for behavior-based AI, IDARA 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 [1-3]. In particular, IDARA’s modeling of the general, first-order response of the immune system allows this robot to interact in new environments before it has had 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.

Multi-robot theory generally states that a group of distributed robots can concurrently perform a class of tasks in parallel and thus be more effective than a single agent. However, this comes with a caveat – to be effective it needs an efficient and intelligent method for control, coordination, and communication. Without this, parallel resources are misallocated and may be counterproductive.

One of the principal characteristics of traditional centralized coordination architectures is that these systems use control schemes that are exponential in complexity and highly communications dependent when directing large numbers of agents [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]. Thus, these approaches are not amenable with very-large scale distributed robotics applications, where traditional control and coordination methodologies can quickly overwhelm the system and saturate the communications pipeline. In order to fully reap the potential of large-scale distributed robotic systems, an architecture is needed that can dynamically optimize their function of simple, yet specific, plastic agents, especially in changing, unpredictable environments.

The immune system is a remarkable example of a highly scalable distributed control and coordination system. 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 to defeat the threat. Thus providing the necessary protection that is essential for survival with minimal harm/impact to an individual [6].

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 functional model of the immune system as a control and coordination mechanism. This, in turn, allows IDARA to be capable of responding dynamically and efficiently without detailed information about the environment.

## 2 HUMAN IMMUNE SYSTEM OVERVIEW

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.

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 level 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 compares these responses and illustrates the body's tradeoff between response time and effectiveness.

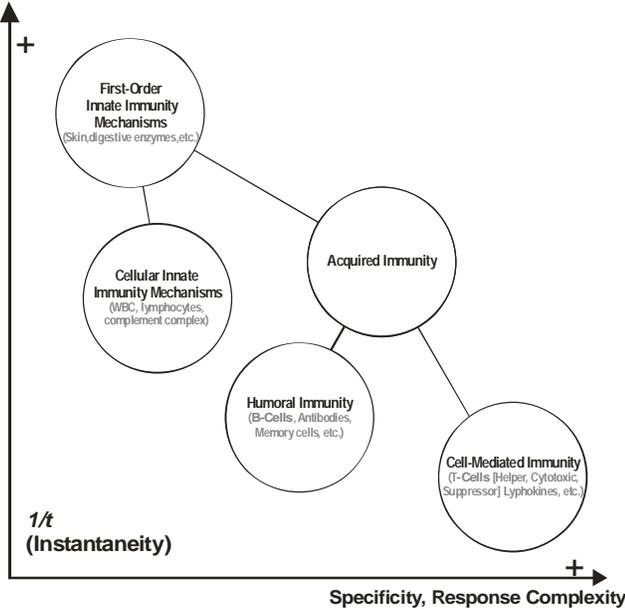


Figure 1: Process Diagram for Immune System Responses (Response becomes more specific and advanced with time)

## 2.1 Innate Immunity

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. This mechanism primarily operates by permitting self/non-self discrimination and by activating certain general kill mechanisms. The innate immune includes the following:

1. Tissue Macrophages—Within minutes after inflammation phagocytic activity (i.e., engulfing of bacteria and other invaders) begins. This first line of defense is fairly primitive in and can be thwarted by a variety of pathogens. Further, it does not have a mechanism for specific responses (e.g., some bacteria have adapted to thrive in macrophages).
2. Destruction of ingested pathogens by digestive enzymes and acids in the stomach.
3. Resistance by the skin to invasion from organisms
4. Chemicals in the circulation that attach to and help destroy foreign antigens (e.g., basic polypeptides, complement complex proteins)
5. Natural killer lymphocytes.

Innate immunity mechanisms 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. Figure 2 illustrates first-order responses by the immune system to a bacterial infection for which some immunity has been established (i.e., basic anti-body mediated “learning” has already occurred) [7].

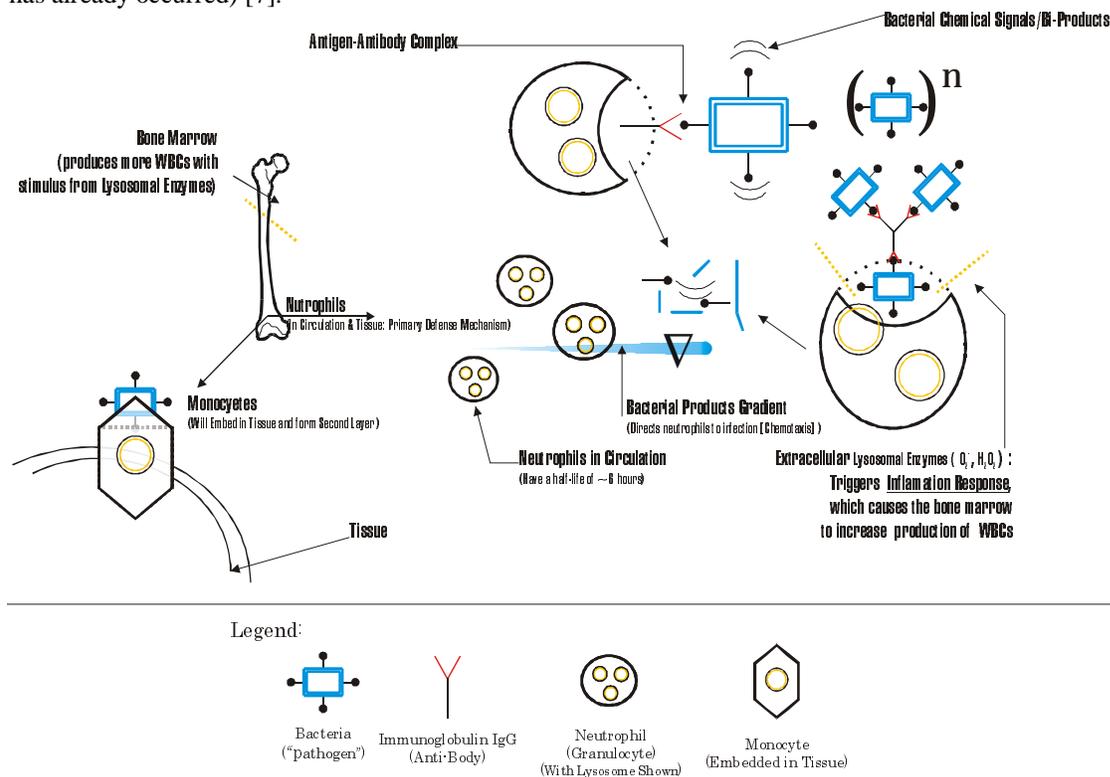


Figure 2: Model of Human Innate Immunity

## 2.2 Acquired 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 are signaled by antibodies (i.e., Y-shaped molecules which match certain proteins based upon their encoded specificity; there are some 10 million present in the immune system) [14].

## 2.3 Salient Characteristics for Mobile Robots

The immune system exhibits a host of useful characteristics, which define its response and are essential to operation. Salient features/characteristics of the immune system for multi-robot applications include:

- a) **Scalability** – In order to provide the extreme levels of protection that are essential for survival, its operation literally depends on the coordination of trillions of cells and molecules that together act to continually monitor the bioenvironment and, depending upon the perceived threat to the body, activates the necessary control systems to defeat this threat.
- b) **Pattern Recognition** – All types/levels of the immune system use a chemical markers as a signature in order to determine the antigen and its nature. Further, pattern recognition is used for learned “self-tolerance,” which arrests autoimmune cells during cell maturity, thus preventing damage to oneself.
- c) **Learning** – A feature resulting from the acquired immune system and its pattern recognition layer, the immune system exhibits remarkable and diverse learning and memory characteristics. This aspect alone has been the focus of SAIS and machine learning research.
- d) **Kill Ladder** – The immune system has a partially redundant, yet specific and exact sequence for typing and resisting almost all types of antigens.
- e) **Specificity** – As documented extensively in the medical literature, the immune system has an incredibly elegant response mechanism that can very specifically respond to a host of antigens.

## 3 PREVIOUS WORK & BACKGROUND

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 non-linear control and planning algorithms.

### 3.1 Distributed Mobile Robotics

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 [9,11]. 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.

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 and a summary of the principal research efforts in this field are outlined in Mataric’s 1995 survey paper on distributed robotics [5,10].

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 Stentz’s macroeconomic approach to mobile robot control resulted in a dynamic robot system that can simply and successfully execute tasks in dangerous environments [10]. In addition, Feddema has applied statistical methods and graph-theoretic approaches to coordinating hundreds to thousands of cooperative robotic agents [12].

### 3.2 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. 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 SAIS (simple artificial immune system) research and a variety of related 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 Idiotoxic Network Hypothesis whereby acquired immunity is used as a model for a new, intelligent problem solving technique. However, these 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 [6, 13]. 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 scheme.

Most of the activity on 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 [1-3]. 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 [7]. Hunt has developed a sophisticated machine learning algorithm set (JISYS) that utilizes AIS principles to perform a variety of “fuzzy” tasks (e.g., task classification, refinement, network generation, and interrelations) [8].

#### 4 THE IDARA PRINCIPLE

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 an uncomplicated 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 [14]. 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.

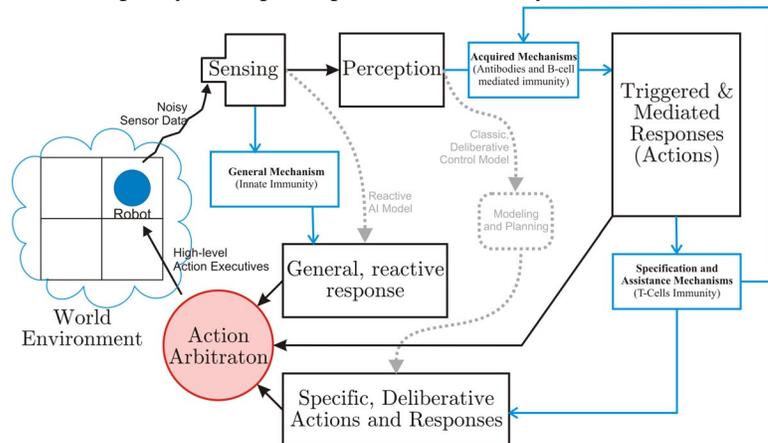


Figure 3: IDARA Software Architecture (Immunology analogous are shown in gray and (for comparison) the typical execution paths of reactive and deliberative/ planner-based architectures are shown as dashed lines)

Similar to several hybrid architectures, IDARA uses a heuristically driven arbitration module to combine action directives being advanced by various levels of the architecture [9]. By using a vector-based approach in combination with the system’s default random exploration routines, the arbitration mechanism calculates the best resultant action. Thus resulting in a form of “directed randomness” which allows the architecture to rapidly vary the nature 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 impact the device.

<i>Characteristic</i>	<i>Coordination Mechanisms</i>		
	<b>IDARA</b>	<b>Microeconomic Cost Optimization</b> [10]	<b>SAIS Algorithms</b> [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
<i>A priori</i> information needed	Can be utilized, but not necessary for random action	Some – Cost functions need to be defined	Yes
General/Instant Response	Yes	Some	Some
Fault-Tolerant	Yes	Yes	No
Optimal Solution Guaranteed	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.

## 5 EXPERIMENTAL DESIGN AND RESULTS

To experimentally validate this architecture and its hypothesized interactions a series of object searching and characterization experiments were devised. These experiments were conducted using a MATLAB implemented simulation of the case problem, the robot architecture, and the environment.

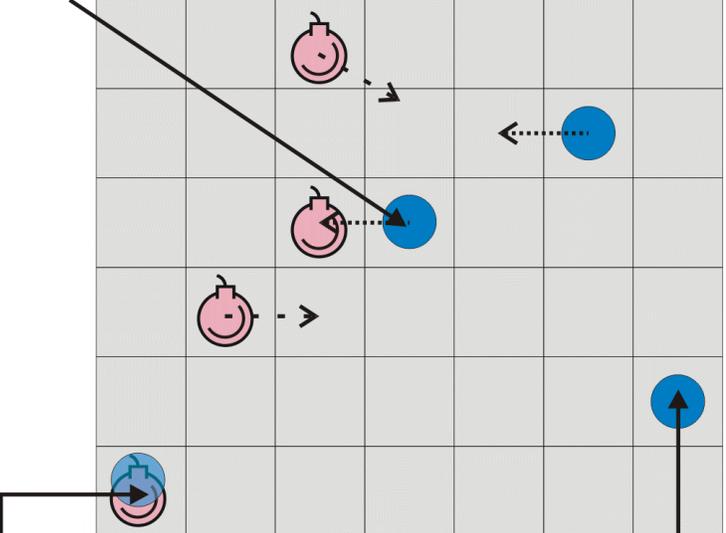
### 5.1 Case Problem: Mobile Minefield Clearing

The mine clearing problem is a generic object recovery problem in which some form of “proximity” sensor is used to locate objects until their density falls below some threshold. Thus, it follows, that the objective of any mine-clearing algorithm is to couple sensory inputs to subsequent actions to find (and presumably remove) the objects while minimizing damage to the members (robots) of the system.

The version of this “classic” problem tested relaxes the static mine assumption. In other words, it is assumed that the mines can move either randomly or specifically in response to any clearing efforts. Relaxing this assumption not only makes the mine-clearing problem dynamic (and, thus, not practically solvable using a traditional methods such as linear programming or constraint optimization) but also thwarts and defeats traditional “raster scan” approaches as the mine motion will act specifically to backfill any cleared areas thus negating any clearing efforts [15].

Thus, a multi-robot approach to this problem needs to vary its response and distribute resources to not only clear the mines but check the cleared space to prevent/minimize backfill. The IDARA architecture, as previously described, was applied to this domain in order to characterize its response in a dynamic environment. Furthermore, the simulation added both Type I and II sensor error/noise and probabilities for robot failure (“death”) based on motion and mine interaction (Figure 4).

Proximity Sensing: A proximity is used to sensor to detect what it is adjacent to. If it is a mine, an “attraction flag” is broadcast and the robot will attempt to remove it. If it is a robot, a “repulsion flag” is broadcast.



Robot-Mine Interaction: When a robot is in the same space as a mine it will attempt to remove it. There is a probability that it will fail when doing so (i.e.,  $P(\text{death})$ ).

Robot Motion: Robot uses a gradient of attraction flags to select an 8-way move or stay still. If it moves, there is a probability that it will fail when doing so (i.e.,  $P(\text{failure})$ ).

Legend	
● Robot	.....> Robot Motion
💣 Mine (UXO)	- - -> Mine Motion

Figure 4: Sample Minefield (Illustration of a subset of the larger grids tested)

**5.2 Simulation Details**

The simulation was implemented, executed, and analyzed using MATLAB. Given a relative area and the mine density (i.e., the percentage of the area that has been mined), the simulator environment generates a minefield and randomly distributes mines within it. This “solution” minefield is then hidden from the rest of the simulation and only access by a sensor routine in the simulator which checks to see if it is adjacent to a mine. As the simulation progresses, the mines are moved in a random or space-filling manner (to simulate “tumbleweed” and breach-healing mines respectively) [16].

Unlike the mines, the robots in the simulation were *not* distributed randomly, but rather in a Gaussian pattern around some central point. This was done because it is envisioned that the deployment of these robots would be done by someone locally distributing/bootstrapping the robots. As one would expect, and as confirmed experimentally, this pattern is less efficient than a random robot distribution. To simplify the simulation, the robot colonies consisted of identical robots whose only sensing was a mine/robot 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 simplified case was studied as general robots (i.e., lack of specialization) is a less efficient case.

In order to focus the simulation, a few assumptions were made. These assumptions are given as follows:

1. **Normality** – That the random trials are statistically normal and this traditional statistical analysis is valid. This is a fair assumption given the large sample size for each parameter set and a more extensive analysis on some randomly selected simulation runs appears to confirm this (i.e., the system has a central tendency about the mean, etc.).
2. **Gaussian Sensor Noise** – The noise created by the sensors is Gaussian and has a central tendency with a zero mean. Furthermore, noisy measurements are assumed to be independent and randomly occurring. These assumptions simplified the noise generation aspects of the simulator’s sensor model.
3. **Rectangular minefield** – Assuming a grid-world allowed for easier computation and representation.
4. **Uniformly Randomly Dispersion** – Mines are initially randomly dispersed within the minefield.
5. **Infinite budget** – The robots are allowed to operate and add costs as necessary (i.e., there is no limit to the cost function).
6. **An estimate of initial mine density is known (or calculated)** – Simplifies implementation and provides an exit criteria; that is, a lack of interaction threshold associated with a desired end density. The algorithm does not depend on this assumption and it can be removed by adding a mechanism to calculate the probability that the field is acceptably cleared.

An IDARA-based method was developed (see Figure 5) and tested on three experimental cases: increasing area, increasing probability of robot failure, and increasing levels of sensor noise. The robots were modeled with only a classifying proximity sensor in order to detect if it is adjacent to a mine or robot and a radio beacon/receiver in order to send/receive alert signals. These beacon 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.

Random scanning was used as the control mechanism for these experiments as it represents a classical solution to the mobile mine clearing problem. Raster (or Boustrophedon) scanning was not used because the backfilling nature of mobile minefields is specifically designed to thwart simple raster scanning [17]. The control runs were executed via a similar method.

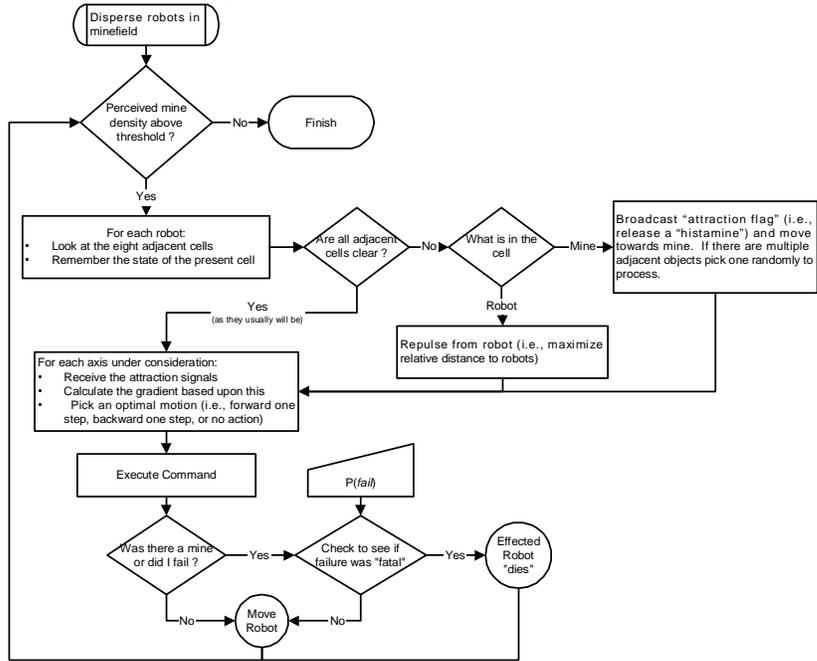


Figure 5: Flowchart for the IDARA-based Mine Clearing Method

### 5.3 Simulation Results

Through the model outlined above, the 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 mine-clearing method, 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)
Minefield Size	350 grids	$O(N^2)$
Mine Density	1 %	$O(N)$
Robot Count	$(1 - e^{-1})P(fail) + e^{-1}$	$O(N + P(fail) \times N)$
Robot:Mine Velocity	2	N/A
Backfill Pattern	Breach healing	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 was being varied in an experiment)

The first simulation consisted of varying dimension (from  $250^2$  to  $500^2$ ) and the mine density of the field (from 0.5% to 3%). Using the aforementioned metaphor as a guide, the IDARA simulation was implemented a simple energy/cost function that tallied the cumulative distance traveled. This was one of the metrics used to analyze the performance of the IDARA mine-clearing method. The average energy used (in grid units of unit distance traveled) by each of the robots until the estimated mine density was 2% of the initial density is shown in Figure 6.

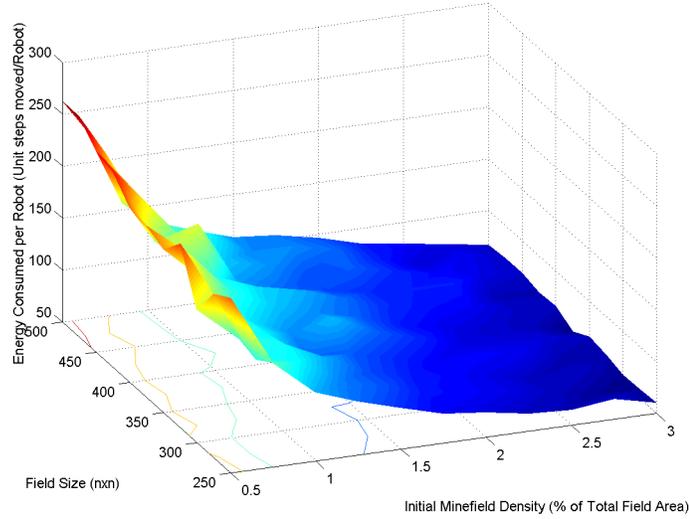


Figure 6: IDARA Mine-Clearing Method Relative Energy Expended  
(Area varied from 250x250 to 500x500 & mine density varied from 0.5% to 3%)

This result indicates that the density of mines is a primary variable governing the operation and performance of the IDARA method. Further analysis of the increase in motion (energy) for low initial minefield densities showed that this was a result of the sparse nature of these fields which resulted in insufficient sensory inputs to drive the model and thus undirected, random motions until sufficient interactions were present.

The second set of experiments studied the effect of varying the robot-mine failure probability (i.e., the chance that a robot will fail when it interacts with a mine). This was calculated by estimating the mine recovery percentage for both the random and IDARA-based method after they were allowed to run for a fixed number of iterations. The percent difference between the two samples was then calculated. As shown in Figure 7, the IDARA method was negatively impacted as failure probability was varied from 25% to 75%; however, even with a 75% failure rate, its mean recovery rate was about 15% better than the random control for the dense minefield and over 80% than the random control better the sparse minefield. The large gain in performance for sparse minefields is due to the characteristically weak performance of the control method for sparse cases.

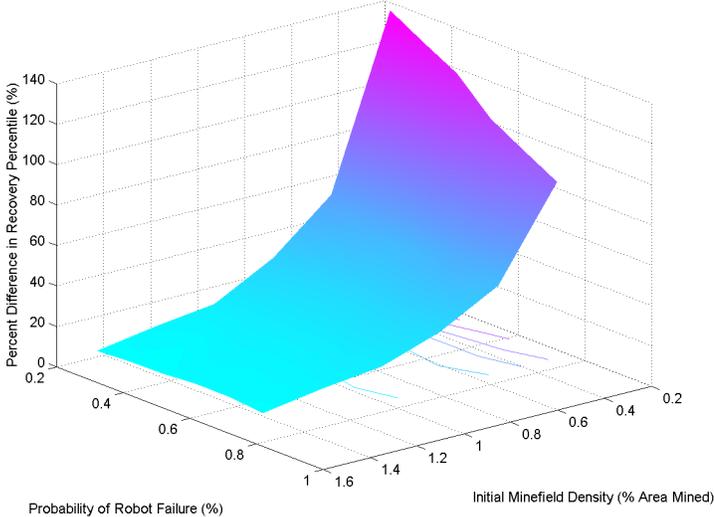


Figure 7: IDARA Percent Difference Over Random Search for Cases of Increasing Likelihood of Robot Failure on Mine Interaction

The third set of experiments (similar to the second) involved varying the sensor noise from none to 25%. This was done in order to characterize the operation and performance of the IDARA method in various noisy environments. In general, the performance of the IDARA method was reduced by increased noise. However, in sparse conditions the random variation improved the initially discovery of mines and thus the recruitment process. The random control was negligibly affected because its actions are totally random and essentially independent of the sensor values. As shown in Figure 8, even with significant amounts of noise the relative mean performance of the IDARA ranged from 8% to 136% better than the control.

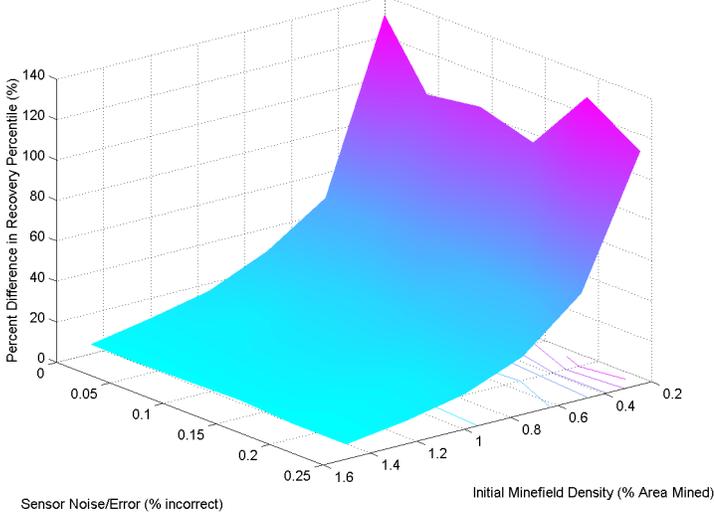


Figure 8: IDARA Percent Difference over Random Search For Cases With Increasing Sensor Noise

## 6 CONCLUSIONS

We have developed a novel architecture for distributed multi-robot coordination and control of large populations of heterogeneous robots. 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 process the mobile minefield more efficiently than traditional techniques. Comparison with the control case shows that this was true even for the case with small, densely populated fields, that typically were not as expected. It was theorized that the raster scan would be preferential to IDARA under these conditions due to the increased processing of IDARA. The simulation shows, however, that this is not the case and the IDARA technique is still beneficial even in this extreme case

The development of this architecture was based on modeling the interaction and character of both the innate and acquired aspects of the human immune system. This results in an architecture that can respond quickly, has a mechanism for learning, and can coordinate a team of robots effectively. The minefield simulation analyzed over 2,300 cases (each with an average population of approximately 735 robots and 1,250 mines). Analysis of the simulations found IDARA to be more effective than its control for all of the cases analyzed (i.e., IDARA had a mean mine recovery percentage of 98.6% compared to the control's 75% mean rate). This is not to imply that the architecture is guaranteed to outperform the control, but rather that under the scope of variables studied (i.e., minefield size from 250 to 500 square units, probability of failure due to robot-mine interaction from 25% to 75%, and added sensor noise from 0% to 25% of sensor measurements taken) the probability of this occurrence is fairly small.

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 (unlike classical neural-networks and similar AI control methods) 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).

## 7 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). The IDARA architecture will be expanded and further developed for other cases that will clearly benefit from very large-scale multi-robot technology, such as tactical field coordination and exploration and mapping.

## 8 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 and Mr. David Steck for his assistance with computer issues. 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.

## 9 REFERENCES

1. 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.
2. 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.
3. 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.
4. R. A. Brooks. "Elephants Don't Play Chess." In Robotics and Autonomous Systems, volume 6, pp. 3-15, 1990.
5. M. J. Mataric, "Issues and Approaches in the Design of Collective Autonomous Agents," *Robotics and Autonomous Systems*, Vol. 16, pp. 321-331, 1995.
6. N. K. Jerne. "The Immune System." *Scientific American*, volume 259, pp. 52 - 60, July 1973.
7. 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.
8. 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: D. Dasgupta, pp. 157 - 186, 1999.
9. 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.
10. 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.
11. A. Stentz and B. L. Brumitt, "Dynamic Mission Planning for Multiple Mobile Robots", *Proceedings of the IEEE International Conference on Robotics and Automation*, No. 3, pp. 2396-2401, 1996.
12. D.A. Schoenwald, J. T. Feddema, and F.J. Opperl, "Decentralized Control of a Collective of Autonomous Robotic Vehicles," *Proceedings of the 2001 American Control Conference*, Vol. 3, pp. 2087-2092, 2001.
13. D. Dasgupta and N. Attouh-Okine. Immunity-based systems: A survey. In IEEE International Conference on Systems, Man, and Cybernetics, volume 1, pp. 369-374, 1997.
14. A. Guyton and J. Hall. *Textbook of Medical Physiology*. W.B. Saunders Co., Ninth edition, 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. DARPA ATO Self-Healing Minefield Project, website: <http://www.darpa.mil/ato/programs/APLA/main.html>
17. H. Choset, "Coverage of Known Spaces: The Boustrophedon Cellular Decomposition," *Autonomous Robots*, Vol. 9, pp. 247-253, 2000.