AN IDIOTYPIC NETWORK APPROACH TO TASK ALLOCATION IN THE MULTI-ROBOT DOMAIN

Use of an Artificial Immune System to Moderate the Greedy Solution

Amanda Whitbrook[1], Gabriel Gainham[2] and Wen-Hua Chen[2]

[1]*BAE Systems, Systems Engineering Innovation Centre (SEIC), Loughborough University*
Loughborough, Leicestershire, LE11 3TU, U.K.

[2]*Department of Aeronautical and Automative Engineering, Loughborough University*
Loughborough, Leicestershire, LE11 3TU, U.K.

Keywords: Multi-Robot Task Allocation (MRTA), Artificial Immune System (AIS), Idiotypic Network, Autonomous Mine Clearance.

Abstract: This paper presents and explains a set of equations for governing simultaneous task allocation in multi-robot systems and describes how they are used to construct a novel algorithm - the Idiotypic Task Allocation Algorithm (ITAA); the equations are based on Farmer's model of an idiotypic immune network but are adapted to include 2-dimensional stimulation and suppression and the use of affinity rather than concentration levels to select antibodies. This novel approach is taken to render the model suitable for simultaneous task allocation where robots must act individually; other idiotypic algorithms have only been applicable to problems where many robots are required to perform one task at a time using swarming behaviours. The paper describes the analogy between idiotypic network theory and the problem of task allocation and shows how the former can be used to increase the fitness of solutions to the latter, also discussing the types of Multi-Robot Task Allocation (MRTA) problem that might benefit from this approach. The results of applying ITTA to a number of simulated mine-clearance problems (with increasing numbers of robots and mines) are presented, and clear advantage over the greedy solution in both simple and more complex scenarios is demonstrated.

1 INTRODUCTION

There are many different types of Multi-Robot Task Allocation (MRTA) problem including varying combinations of single-task (ST) robots, multi-task (MT) robots, single-robot (SR) tasks, multi-robot (MR) tasks, instantaneous assignments (IA, with no planning for future allocations), time-extended assignments (TA, which allows for future allocation planning) and online assignment variations of IA (OA, where tasks are introduced one at a time). The interested reader is directed to Gerkey and Mataric (2004), which presents a comprehensive, architecture-independent taxonomy. In addition, robots may be heterogeneous in their capabilities and performance, and tasks may differ in complexity, difficulty and solution requirements. Whilst all types of MRTA problem may be solved by implementing a greedy algorithm, characterised by repeatedly taking the 'best' valid option (based on some measure of fitness) at a local level, optimization is not guaranteed. In addition, the Linear Programming (LP) approach, which guarantees optimality, cannot be applied to some of the more complex MRTA problem types including ST-SR-IA-OA, ST-SR-TA, ST-MR-IA, ST-MR-TA, MT-SR-IA, MT-SR-TA, MT-MR-IA and MT-MR-TA combinations, some of which are strongly NP-hard (Gerkey and Mataric (2004)). There is thus a need for heuristic approaches that are capable of providing fitter solutions than those offered by greedy algorithms, and much research effort has been directed towards the development of heuristic MRTA techniques. For example, auction-like allocation mechanisms are described in Guerreor and Oliver (2011), Nanjanath and Gini (2010) and Gerkey and Mataric (2002). There have also been a number of works published on market-based techniques (for example Dias et al. (2005)), coalition-formation methods (Shehory and Kraus

Whitbrook A., Gainham G. and Chen W..
AN IDIOTYPIC NETWORK APPROACH TO TASK ALLOCATION IN THE MULTI-ROBOT DOMAIN - Use of an Artificial Immune System to Moderate the Greedy Solution.
DOI: 10.5220/0003709000050014
In *Proceedings of the 4th International Conference on Agents and Artificial Intelligence* (ICAART-2012), pages 5-14
ISBN: 978-989-8425-96-6

(1998)), self-organisation (Fukuda et al. (1988)) and emergent systems (Liu et al. (2007) and Atay and Bayazit (2007)).

This paper presents a set of equations based on those developed by Farmer et al. (1986) that represent an idiotypic immune system approach (see Jerne (1974)) to solving the problem of task allocation in the multi-robot domain. The equations have been adapted to include 2-dimensional stimulation and suppression and the use of affinity rather than concentration levels to select antibodies. This is a novel approach that allows each robot to solve a separate task independently so that all tasks can be completed simultaneously.

The paper describes the analogy between task allocation and the idiotypic network theory of the immune system and shows how the equations can be applied to general ST-SR-IA problems with N robots looking for one task to complete and L tasks requiring one robot. It sets out how this approach differs from previous idiotypic implementations of MRTA and explains its advantages over them; in particular, other idiotypic algorithms have only been applicable to problems where many robots are required to perform one task at a time using swarming behaviours.

A set of experimental results on simulated problems of this type is presented, initially where $L = N$, N varies between 3 and 15, and robots are required to organise mine diffusion tasks in a way that minimizes travel costs. Some preliminary results for the case where $N \neq L$ are also briefly discussed. The results provide empirical evidence that the Idiotypic Task Allocation Algorithm (ITTA) described here is capable of outperforming the greedy approach such that mean fitness is significantly improved for these problem types.

2 BACKGROUND, PRIOR WORK AND MOTIVATION

2.1 Background

The purpose of the immune system is to identify and neutralize the molecules or cells that are dangerous to the body (antigens) without damaging healthy cells (Barra and Agliari (2007)). This is achieved through the interaction of many different types of immune cell, which each have specific roles. The main constituents of the adaptive immune system are B-lymphocytes (B-cells) and T-lymphocytes (T-cells), which have particular protein molecules on their surfaces called receptors. The receptors of B-cells can bind to antigens that 'match' them, allowing the B-cells to neutralize them.

The clonal selection theory of the immune system (Burnet (1958)) states that lymphocytes operate independently, and that once a match is established, B-cells proliferate (increase in concentration) by cloning and releasing free receptors known as antibodies. Binding takes place between a region of the antibody known as the paratope and a region of the antigen known as the epitope. In contrast, Jerne's idiotypic network theory of the immune system (Jerne (1974)) postulates that lymphocytes interact with each other so that the immune system functions as a global network of cells stimulated and suppressed by internal recognition and matching between themselves. This is because antibodies also serve as internal images of certain antigens and are thus themselves being detected and acted upon (Barra and Agliari (2007)), which keeps the concentrations of antibodies at appropriate levels. Antibody paratopes are thus not only matched to antigen epitopes but also to epitope regions on other antibodies, known as idiotopes. Figure 1 below shows the structure of an antibody and illustrates how antibody concentrations are suppressed by other antibodies that recognise their idiotope, and how concentrations are stimulated to increase when they recognise another antibody's idiotope.

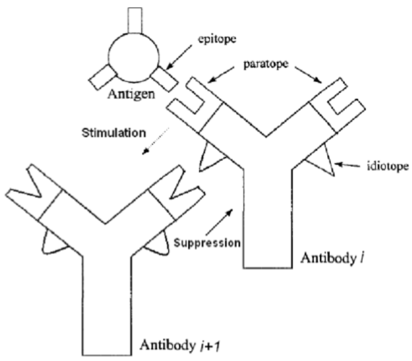

Figure 1: Antibody paratope and idiotope regions.

2.2 Prior Work

The dynamics of antibody and antigen concentrations are modelled computationally as differential equations in Farmer et al. (1986). This model is widely used for constructing Artificial Immune System (AIS) implementations of idiotypic networks, especially in navigational robotics, where the method has demonstrated flexible behaviour-mediation properties. However, in this field artificial idiotypic networks have largely been confined to single robot navigation problems, for example, Watanabe et al. (1998), Vargas et al. (2003), Luh and Liu (2004), and Whitbrook et al. (2007), where the individual behaviours of single robots are modelled as antibodies and environmental information is modelled as antigens.

On the other hand, the application of idiotypic principles to task allocation in the multi-robot domain is somewhat more scarce, especially utilization of the Farmer-based model, despite the fact that its decentralized yet cooperative and coordinated approach to problem solving lends itself very elegantly to such systems. Sathyanath and Sahin (2002) implement idiotypic mine detection but use a simplistic analogy rather than the Farmer model, i.e., idiotopes are not modelled and play no role in determining the stimulation and suppression levels of robots. Mitsumoto et al. (1995) implement swarm behaviour by using a clonal selection-based method rather than an idiotypic network; self-non-self discrimination is modelled and tactics between the robots are secreted and proliferated until swarming behaviours emerge. Dioubate et al. (2008) use a hybrid Farmer-based idiotypic network coupled with clonal selection and genetic evolution of lymphocytes to generate co-ordinated formation of robots behind obstacles. Lee and Sim (1997) use the Farmer model to develop idiotypic cooperative strategies leading to swarm behaviours; robots communicate their behaviours to each other on a local level and the behaviour (antibody) that shows the greatest stimulation is adopted by the whole group. Jun, Lee and Sim (1999) and Sun, Lee and Sim (2001) use an extended version of this model that includes additional T-cell control of concentrations to improve the adaptation capability. Razali et al. (2009, 2010) use the same model as Jun, Lee and Sim (1999), but also include memory enhancement to achieve shepherding behaviour for robot dogs managing robot sheep. Li et al. (2007) solve the same problem, also using Farmer's idiotypic model, but do not include T-cell control or enhanced memory.

2.3 Motivation

In all of the examples cited above either the Farmer model is not implemented or the goal is to adopt majority behaviour patterns rather than assign individual behaviours to individual tasks. The general Lee and Sim approach is thus suited to problems where a number of tasks that require many robots to solve them are completed in sequence (ST-MR-TE), but it is not applicable to the broader spectrum of problems including those that require instantaneous assignment (IA) of robots to different tasks. In essence, the Lee and Sim analogy is the same as for single robot navigation, i.e., behaviours are modelled as antibodies, and only one behaviour is adopted at a given time. If robots in the group are required to adopt different behaviours at a given time (as in IA problems), then a different model is clearly needed. Furthermore, there is a real requirement for IA-type assignment of heterogeneous tasks to heterogeneous robots within the military domain. In particular, it is envisaged that within the next twenty-five years autonomous military capabilities will undergo a major shift toward joint, multi-mission, collaborative operations between manned and unmanned vehicles (US Department of Defence (DOD) (2009)). For example, fleets of unmanned aerial vehicles (UAVs) and unmanned ground vehicles (UGVs) will be required to work together to accomplish reconnaissance, surveillance, mine detection and target-designation missions. Within such operations, successful task allocation and coordination of the many heterogeneous assets will be critical to mission success, but will also impose a great burden on central command and control as the number of assets increases. For this reason, an autonomous, decentralized, self-regulating coordination system in which the assets are able to allocate tasks independently of human control would be of great value to the military. In addition, if progressed through to use in theatre, a successful framework for decentralized coordination and control of heterogeneous, multi-agent, military systems would represent a significant step forward for autonomy. Indeed, the current US DOD Unmanned Air Systems Roadmap (2005) cites "distributed control" as the main criteria for achieving an autonomy level of 8 in the DOD scale (range = 1 to 10) compared with the remotely-operated systems that are typically in place at present; these are measured as between levels 1 and 3 on the same scale. This paper sets out, describes and tests the Idiotypic Task Allocation Algorithm (ITAA), which provides a potential solution to the problem of autonomous, decentralized, distributed task allocation for IA-type assignment of heterogeneous tasks to heterogeneous robots.

3 PROBLEM SPACE

A mine-clearance scenario has been selected as the test-bed for the ITAA as it has many properties that make it ideal. In particular, there is sufficient flexibility within the problem space to allow more simple variations to be implemented in the early stages of research, and to build and test more complex instances as the work progresses, for example beginning with ST-SR-IA experiments, equal numbers of identical tasks (mines to diffuse) and identical robots, incrementally building up to the inclusion of online assignments (OA), time-extended assignments (TE), unequal numbers of tasks and robots, heterogeneous tasks and robots, multi-task robots (MT), multi-robot tasks (MR), and real-time, real-world implementations that require additional features such as reactive obstacle avoidance modules.

In this paper, research begins with the problem of assigning a known number L of identical, un-diffused mines to a known number N of homogenous robots in simulation. Initially, it is assumed that:

1. the robots have equal capabilities and travel at the same, fixed speed;
2. the mines are equally accessible to all the robots, i.e., there are no obstacles to negotiate;
3. the level of difficulty in diffusing a mine is equal for all mines and constant throughout the operation;
4. the number of mines L does not change at any time during the operation;
5. the number of robots N does not change at any time during the operation;
6. the number of robots available is always equal to the number of mines needing diffusing, i.e., $N \equiv L$.
7. once assignment has taken place and the mines are diffused, all work is done.

Note that assumptions 1 to 3 allow the distance between robots and mines to be used as a measure of affinity between them. If this were not the case, then a more complex measure would be needed, i.e. one that also considers the ability of each robot to complete each individual task and the additional time that would be needed to negotiate (possibly moving) obstacles. This paper is chiefly concerned with validating the theory set out in Section 4 so use of the most simplistic case in the first instance allows the essential theory of the ITTA model to be tested independently of any real-world noise. The results of a preliminary investigation into cases where $N \neq L$ is also briefly discussed here but more complex experiments will be conducted in the future in order to establish whether the method stands up to the problems associated with real-world implementation.

4 THE IDIOTYPIC TASK ALLOCATION ALGORITHM (ITAA)

In the model presented here, antibodies are analogous to possible robot-mine pairs, and the affinity U of each antibody to the current antigen (physical positioning of all robots and mines) is the distance d between the robot and mine in the antibody pair. This is the antibody pre-affinity, before stimulation and suppression from other antibodies are taken into consideration. The post-affinity after stimulation and suppression is denoted as T. Writing this more formally, U_{ij} is the pre-affinity, T_{ij} is the post-affinity and d_{ij} is the distance between robot i, $i = 1, \ldots, N$ and mine j, $j = 1, \ldots, L$. The pre-affinity is thus given by

$$U_{ij} = d_{ij}. \tag{1}$$

After pre-affinities have been calculated, the initial allocation of robots to mines is achieved by executing a simple and intuitive greedy algorithm where the antibody with the smallest affinity is repeatedly selected as an allocation and then all pairs that contain that robot and mine are eliminated from future allocations until exactly one robot is allocated to exactly one mine. This greedy algorithm is also known as the Sequential Best-Pair Algorithm (SBPA, see Oliver and Guerrero (2011)). Let γj represent the index of the robot allocated to mine j based on pre-affinities. Under the ITTA model, this is the *antigenic* robot to mine j, one of a set of L antigenic robots (as exactly one robot is antigenic to each of the L mines). Similarly, let $i\gamma$ represent the index of the mine allocated to robot i based on pre-affinities. This is the *antigenic* mine to robot i, one of a set of N antigenic mines (as exactly one mine is antigenic to each of the N robots). After the antigenic robots and mines are known, the post-affinity is calculated using

$$T_{ij} = U_{ij} + \sum_{k=1}^{N} V_{kj} + \sum_{k=1}^{L} W_{ik} - X_{ij} - Y_{ij}, \tag{2}$$

where:

- V corresponds to suppression from the antibodies that represent robots competing for the same mine (they may suppress the antigenic robot if they have a higher fitness (lower affinity) or are close in fitness to it).

- X corresponds to stimulation of antibodies that represent robots competing for the same mine (the antigenic robot may stimulate other robots only if the other robots have a higher fitness (lower affinity) than it).
- W corresponds to suppression from the antibodies that represent mines competing for use of the same robot (they may suppress the antigenic mine if they have a higher fitness (lower affinity) or are close in fitness to it).
- Y corresponds to stimulation of antibodies that represent mines competing for use of the same robot (the antigenic mines may stimulate other mines only if the other mines have a higher fitness (lower affinity) than it).

Equation (2) is similar to the original Farmer equation but differs in two important respects. First, concentrations of antibodies are not modelled, only affinities, and second, there are two stimulation terms and two suppression terms (rather than one of each as in the original). This reflects the 2-dimensional nature of the model used here, i.e., stimulation and suppression are considered between robots and also between mines. To illustrate, if the affinities between mine-robot pairs were set out as a matrix, for example with each row representing a unique mine and each column representing a unique robot, then stimulation and suppression are measured both across the columns in the *x*-direction and down the rows in the *y*-direction, see Figure 2, which shows an example of 2-dimensional stimulation and suppression for the 3-robot, 3-mine case. Note also that stimulation terms are subtracted from the pre-affinity and suppression terms are added to it. This is because, in this case, the affinity is based on the distance the robot has to travel, and thus, a reduction is seen as an improvement.

In this model the total inter-robot suppression on antibody γj is given by the sum of the suppressions V imposed by antibodies kj (j = 1 to L, k = 1 to N) where

$$V_{kj} = \frac{U_{\gamma j}}{p_1 |U_{kj} - U_{\gamma j}|} \forall (k \neq \gamma \wedge |U_{kj} - U_{\gamma j}| < \zeta). \quad (3)$$

The inter-robot stimulation X on antibody ij is imposed by antibody γj (i = 1 to N, j = 1 to L) and is a single term given by

$$X_{ij} = \frac{(U_{\gamma j} - U_{ij})}{p_1} \forall (i \neq \gamma \wedge U_{ij} < U_{\gamma j}). \quad (4)$$

The total inter-mine suppression on antibody $i\gamma$ is given by the sum of the suppressions W imposed by antibodies ik (i = 1 to N, k = 1 to L) where

$$W_{ik} = \frac{U_{i\gamma}}{p_1 |U_{ik} - U_{i\gamma}|} \forall (k \neq \gamma \wedge |U_{ik} - U_{i\gamma}| < \zeta). \quad (5)$$

The inter-mine stimulation Y on antibody ij is imposed by antibody $i\gamma$ (i = 1 to N, j = 1 to L) and is a single term given by

$$Y_{ij} = \frac{(U_{i\gamma} - U_{ij})}{p_1} \forall (j \neq \gamma \wedge U_{ij} < U_{i\gamma}). \quad (6)$$

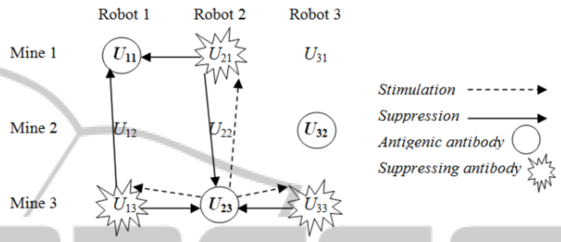

Figure 2: An example of 2-dimensional stimulation and suppression for a 3-robot, 3-mine case.

In the above equations, p_1 is a scaling constant that determines the overall level of stimulation and suppression and ζ is a constant that governs how closely antibodies have to match in affinity to become stimulated. After post-affinities have been calculated the SBPA is implemented again to allocate the new antigenic antibodies. The post-affinities then become the new pre-affinities and stimulation and suppression are calculated again. The algorithm proceeds in this way until some stopping criteria is met. The final, overall, theoretical fitness F of the task-allocation solution is determined as

$$F = \frac{10{,}000}{\sum_{j=1}^{L} d_{\gamma j}}, \quad (7)$$

where $d_{\gamma j}$ is the distance between a final antigenic robot and its allocated mine. Note that a different measure of fitness, for example use of time taken t to get to the mine (instead of d in the above equation) should be used when attempting to demonstrate the practical advantages of the ITTA, rather than the theoretical. However, in the experiments described here these measures are equivalent because of assumptions 1 to 3.

The ITAA, as described above, is original in its 2-dimensional approach to stimulation and suppression, its focus on affinities rather than concentrations of antibodies, its novel suppression and stimulation models, and its algorithmic implementation, which results in the assignment of a unique task to each robot, rather than the global adoption of majority behaviours as in previous idiotypic research within the multi-robot domain. Note that 1-dimensional models were trialled but failed to guarantee converge to a solution. In addition, a 2-dimensional model is a more accurate reflection of an idiotypic system, where interactions occur between all agents.

5 EXPERIMENTAL DETAILS

The ITTA was transcribed into MATLAB code and was programmed to store the current fittest solution after each iteration. The algorithm was stopped after a maximum of 15 iterations had elapsed and the best solution was accepted. In all cases, the initial positions of the N robots and mines were generated randomly on a square grid 30m by 30m in area, and baseline comparisons were made for each problem using the greedy (SBPA) algorithm (the solution after the first iteration). Initially, the ITTA was applied to 10,000 different mine diffusion problems for N between 3 and 10 in order to determine suitable values for parameters ζ and p_1, i.e., the above was repeated varying the parameter p_1 between 10 and 1,000 (values of 10, 50, 100, 150, 250, 500, 750 and 1,000 were trialled), and varying the parameter ζ between 0.5m and 4m in steps of 0.5m. Once suitable values were found, the ITTA was applied to a further 10,000 mine diffusion problems for N ranging between 3 and 15, in order to assess its performance against the baseline.

6 RESULTS

6.1 Parameter Selection

In all initial test cases p_1 values of 10 and 50 proved superior in performance to the others, with 10 tending to work better for smaller numbers of robots (3 to 7) and 50 tending to work better for larger numbers (8 to 10). Figures 3a and 3b show how the mean % improvement in fitness varies with p_1. Figure 3a summaries the results for the different ζ values and Figure 3b does the same for the numbers of robots N. Figure 3b also shows that mean % improvement in fitness tends to increase steadily with the number of robots; this is discussed more fully in Section 6.2.

The ζ value was more robust, demonstrating much less variation in performance than p_1. This can be seen in Figure 3a. Figures 4a and 4b also summarise the preliminary results for ζ; the charts show how mean % fitness improvement varies with ζ, with Figure 4a showing the results for each value of N and Figure 4b showing the results for each value of p_1.

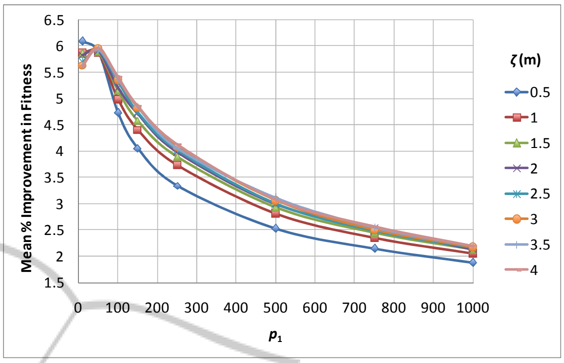

Figure 3a: Variation of mean % improvement in fitness with p_1 for different ζ values.

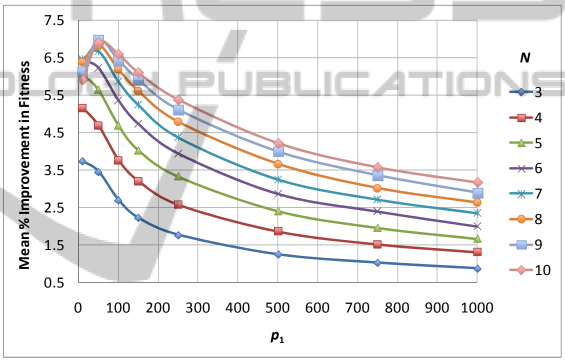

Figure 3b: Variation of mean % improvement in fitness with p_1 for different values of N.

Figure 4a also illustrates a clear trend for increase in mean improvement in fitness as the numbers of robots increases (see Section 6.2). In general, there is a slight improvement as ζ rises, but the differences are much less pronounced than for p_1. Figure 4b highlights the poorer performance when higher values of p_1 are used. It shows that values of either 10 or 50 are preferable and that a p_1 value of 50 has an almost constant performance across the ζ spectrum, whereas a p_1 value of 10 tends to work better for lower values of ζ, between about 0.5 and 1.5.

As it showed a consistent performance for all ζ and worked well with higher numbers of robots, a p_1 value of 50 was chosen for use in the performance assessment, where N would rise to 15. Initially, ζ values of 3.0m and 4.0m were selected, based on the preliminary results, but the best overall performance was obtained when ζ was set to 0.5m and p_1 was set to 50.

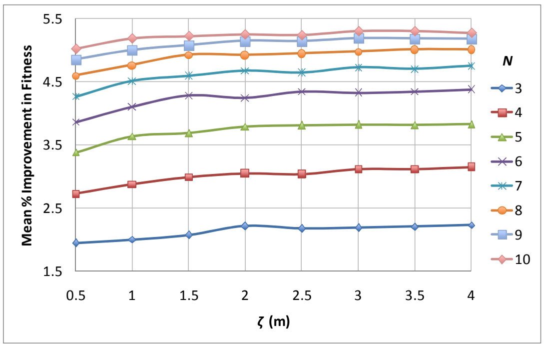

Figure 4a: Variation of mean % improvement in fitness with ζ for different values of N.

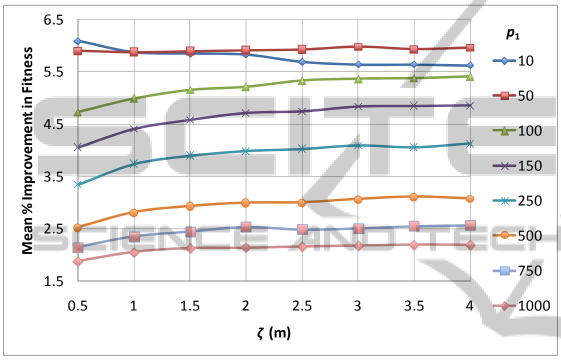

Figure 4b: Variation of mean % improvement in fitness with ζ for different values of p_1.

6.2 Performance Assessment

Table 1 summarises the performance of the ITTA (using the assigned parameters, p_1= 50, ζ = 0.5) compared with the baseline greedy SBPA algorithm. For all N there is an increase in the mean fitness of the solution, which ranges from about 3.5% (N = 3) to 7.2% (N = 9 and N = 11). Moreover, paired t-tests conducted on the mean fitness values show that ITTA fitness is significantly higher (at the 95% level) than the baseline for all values of N, for all results and also for the sub-set of improved cases. Figure 5 shows that the mean improvement in fitness starts off by increasing almost linearly with N but gradually reaches a plateau at about N = 9. This may be explained as follows; the likelihood that the initial, greedy solution may already be the optimal one is higher for smaller N, and so there is less room for improvement. This explanation is validated by examination of the % of improved cases, which also increases steadily with N up until reaching a plateau at about N = 10, see Figure 6. In addition, when the improved cases are examined in isolation, as expected, the % improvement is greater for all N, but this is more pronounced for smaller N, for example,

Table 1: Performance summary for ITTA and baseline.

N	Mean Fitness				Mean % Imp.		Mean Iterations		Max % Imp.	% Imp. Cases
	All Results		Improved Cases		All	Imp.	All	Imp.		
	SBPA	ITTA	SBPA	ITTA						
3	303	314	255	285	3.43	11.67	2.30	6.61	75.91	34.85
4	238	249	212	233	4.49	9.86	3.52	6.88	72.11	51.13
5	202	213	186	203	5.41	9.21	4.51	7.07	79.96	63.80
6	176	186	165	180	6.08	8.90	5.26	7.26	61.60	72.45
7	157	168	150	163	6.54	8.80	5.73	7.36	64.85	77.90
8	143	153	138	150	6.80	8.65	5.96	7.31	69.13	81.52
9	132	141	128	139	7.17	8.73	6.29	7.46	53.69	84.27
10	122	131	119	130	7.14	8.56	6.24	7.30	52.19	85.42
11	114	123	112	122	7.17	8.52	6.26	7.28	57.26	85.97
12	108	116	106	115	6.99	8.33	6.01	7.02	45.99	85.61
13	102	109	100	108	6.86	8.21	5.83	6.83	44.91	85.26
14	97	104	95	103	6.70	8.00	5.66	6.64	43.83	85.30
15	92	98	91	98	6.57	7.87	5.53	6.51	50.41	85.00

the difference is 8.25% for N = 3, but only about 1.30% at plateau values of N, see Figure 5.

For all N the maximum % improvement in solution is considerably higher than the mean, for example, for N = 3 it is about 76% and for N = 12 it is about 46%. This variable tends to oscillate locally, but has a general downward trend with increased N, see Figure 6. For all results, the mean number of iterations ranges from about 2.0 for N = 3 to about 6.5 for plateau values of N, see Figure 7, which is intuitive given the earlier explanation for plateau behaviour. For improved cases only this variable is much more consistent, tending to about 7.0 iterations.

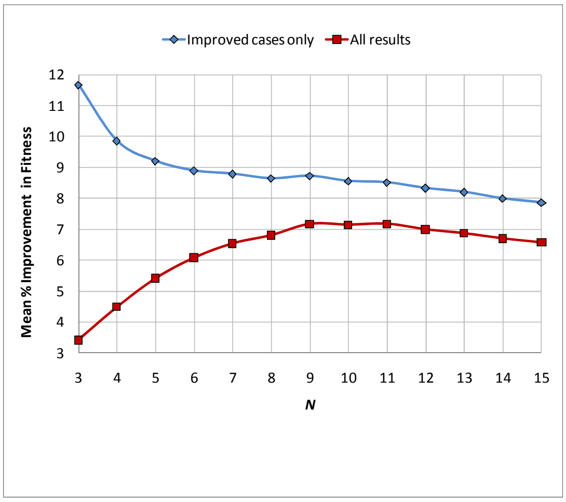

Figure 5: Variation of mean % improvement in fitness with N.

The above results suggest that the ITTA is able to make significant improvements over greedy strategies, and that, as the number of robots increases, an improvement in the solution is more likely. For the particular parameters used here there is approximately an 85% chance of generating a

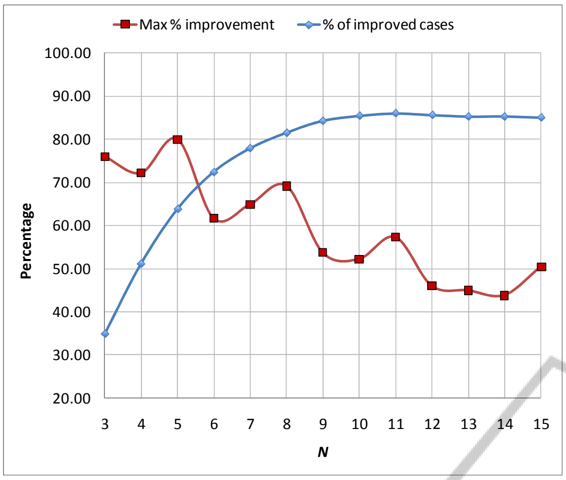

Figure 6: Variation of max % improvement in fitness and % of improved cases with *N*.

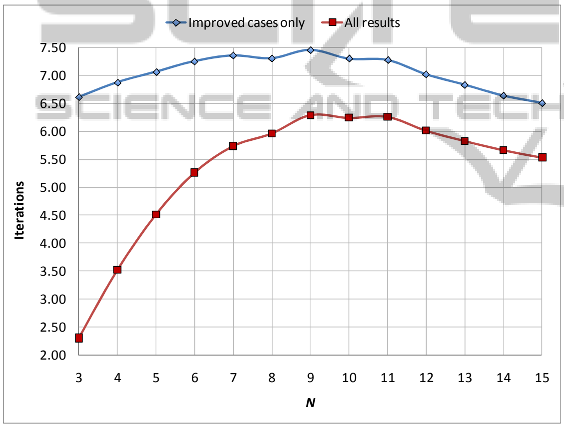

Figure 7: Variation of mean iterations with *N*.

better solution for N greater than 8. For N greater than 6, a mean increase in fitness of about 7% is expected, although individual increases of up to about 70% are possible. The ITTA has also proved to be a fast algorithm as the mean number of iterations for convergence is always below eight.

In addition, a further set of experiments that varied p_1 between 50 and 150 across the suppression and stimulation equations (3), (4), (5) and (6) has also been conducted for $N = 4$. The use of $p_1 = 50$ in all equations except (3) (which used 90 instead) increased the overall performance by a further 0.7%.

These results demonstrate the potential of the ITTA method and show that it is a good candidate for further investigation involving more complex problems (as described fully in Section 3), real-world implementations and more rigorous parameter tuning. Preliminary investigations have already shown that the method is easily adapted to cases where $L > N$ and $N > L$. Where there are more robots than mines ($N > L$) robots are simply marked as redundant when the SBPA part of the algorithm does not allocate them to a mine. In addition, ITTA consistently outperforms SBPA and, as N increases for a fixed number of L, performance improves. Conversely, when $L > N$ absolute performance drops, with the algorithm having to run repeatedly as robots change position, but ITTA still performs better than the greedy algorithm. Thus, relatively speaking, there is no noticeable drop in ITTA's performance when $L > N$.

7 FUTURE WORK

Future work will aim to develop an optimum stopping criteria and to test the algorithm in more complex scenarios, where online and time-extended assignments are required, there are heterogeneous tasks and robots, multi-task robots and multi-robot tasks. Real-world implementations that require additional features such as reactive obstacle avoidance modules will also be carried out in an outdoor environment. Further work also needs to be done to compare performance of the ITTA with state-of-the-art task allocation methods (for example market-based approaches) and linear optimization techniques such as Mixed Integer Linear Programming (MILP).

Note that in real-life implementations the algorithm would need to run independently on each robot in order to constitute a truly decentralized system. The robots would also need to communicate reliably in order to transmit their locations to one another, and there would need to be a level of assurance that each robot was receiving all the available information and compiling the same solution to the problem. Maintaining and sharing an accurate intelligence picture within an ad-hoc network has been the subject of a research program within BAE Systems Advanced Technology Centre (ATC), and the outputs have already produced a prototype data sharing framework. Future work will thus aim to integrate the ideas presented in this paper with the outputs of the data sharing programme in order to demonstrate decentralized multi-robot task-allocation in a real-world environment. In addition, the ATC has also been working on Consensus-Based Bundle Algorithms (CBBA) and Max-Sum task allocation mechanisms (Mathews et al. (2010)), so work will be undertaken to assess the feasibility of integrating the ITTA approach with those methods (see also Stranders et al. (2009)).

8 CONCLUSIONS

This paper has described an idiotypic AIS algorithm (ITAA) for solving task allocation problems in the multi-robot domain. The algorithm is novel since other idiotypic approaches have only been applicable to problems where many robots are required to perform one task at a time using swarming behaviours; in contrast ITTA is suited to problems that require members of a multi-robot team to act individually so that different tasks can be solved simultaneously. The algorithm is also original in its implementation of the Farmer equation, which ignores concentrations of antibodies and uses novel, 2-dimensional models for stimulation and suppression of the antibody affinities.

A series of initial tests have been carried out on the algorithm using simulated mine diffusion problems in MATLAB. These tests have helped to establish suitable parameter values for the stimulation and suppression terms and have provided statistical evidence that the ITTA is capable of out-performing the greedy Sequential Best-Pair Assignment (SBPA) algorithm in about 85% of cases for numbers of robots N exceeding 8. For smaller N the likelihood of outperforming the greedy solution rises almost linearly as N increases. The ITTA has also shown fast convergence to a solution; for N of 8 and above the mean number of iterations for arrival at the best solution is about 5, i.e., the solution can be produced almost instantaneously.

REFERENCES

Atay, N., Bayazit, B., 2007. Emergent task allocation for mobile robots. In *Proceedings of the Robotics: Science and Systems Conference (RSS'07), Atlanta, GA, USA*.

Barra, A., Agliari, E., 2007. Stochastic dynamics for idiotypic immune networks. *Physica A 389: pp. 5903-5911.*

Burnet, F.M., 1959. *The clonal selection theory of acquired immunity*, Cambridge University Press, Cambridge, U.K.

Dias M. B., Stentz, A., 2002. Opportunistic optimization for market-based multi-robot control. In *Proceedings of the IEEE/RJS International Conference on Intelligent Robots and Systems (IROS), Lausanne, Switzerland: pp. 2714-2720*

Dioubate, M., Guanzheng, T., Toure-Mohamed, L., 2008. An artificial immune system based multi-agent model and its application to robot cooperation problem. In *Proceedings of the 7th World Congress on Intelligent Control and Automation, Chongging, China: pp. 3033-3039*.

Farmer, J. D., Packard, N. H., Perelson, A. S., 1986. The immune system, adaptation, and machine learning. *Physica, D, 2(1–3): pp. 187–204.*

Fukuda T., Nakagawa, S., Kawauchi, Y., Buss, M., 1988. Self-organizing robots based on cell structures - CEBOT. In *Proceedings of the IEEE/ RSJ International Conference on Intelligent Robots and Systems (IROS), Victoria, British Columbia, Canada: pp. 145-150.*

Gerkey, B. P., Mataric, M. J., 2002. Sold! Auction methods for multi-robot coordination, *IEEE Transactions on Robotics and Automation 18(5): pp. 758-768.*

Gerkey, B. P., Mataric, M. J., 2004. A formal analysis and taxonomy of task allocation in multi-robot systems. *International Journal of Robotics Research 23(9): pp. 939-954.*

Jerne, N. K., 1974. Towards a network theory of the immune system. *Ann. Immunol. (Inst Pasteur), 125C(1/2): pp. 373–389.*

Jun, J-H., Lee, D-W., Sim, K-B.,1999. Realization of cooperative strategies and swarm behavior in distributed autonomous robotic systems using artificial immune system. In *Proceedings of the 1999 IEEE International Conference of Man and Cybernetics 6: pp. 614-619.* IEEE Press, New York.

Lee, D-W., Sim, K-B., 1997. Artificial Immune Network-based Cooperative Control in Collective Autonomous Mobile Robots. In *Proceedings of the IEEE International Workshop on Robot and Human Communication: pp. 58-63*.

Li, J., Xu, H., Wang, S., Bai, L., 2007. An immunology-based cooperation approach for autonomous robots. In *Proceedings of the 2007 International Conference on Intelligent Systems and Knowledge Engineering (ISKE 2007), 4*.

Liu, W., Winfield, A. F. T., Sa1, J., Chen, J., Dou, L., 2010. Towards energy optimization: Emergent task allocation in a swarm of foraging robots. *Adaptive Behavior 15 (3): pp. 289-305*

Luh G. C., Liu, W.W., 2004. Reactive immune network based mobile robot navigation. In *Proceedings of the 3rd International Conference on Artificial Immune Systems (ICARIS) 3239: pp. 119–132.*

Mathews, G., Waldock, A., Paraskevaides, M., 2010. Toward a decentralised sensor management system for target acquisition and track. In *Proceedings of the 5th SEAS DTC Technical Conference: pp. C2.*

Mitsumoto, N., Fukuda, T., Shimojima, K., Ogawa, A., 1995. Micro autonomous robotic system and biologically inspired immune swarm strategy as a multi agent robotic system. In *Proceedings of the 1995 IEEE International Conference on Robotics and Automation: pp. 2187-2192.*

Nanjanath, M., Gini, M., 2010. Repeated auctions for robust task execution by a robot team. *Journal Robotics and Autonomous Systems 58(7): pp. 900-909,* North-Holland Publishing Co. Amsterdam, The Netherlands.

Oliver, G., Guerrero, J., 2011. Auction and swarm multi-robot task allocation algorithms in real time scenarios. In *Multi-Robot Systems, Trends and Development, Toshiyuki Yasuda (Ed.), ISBN :978-953-307-425-2: pp. 437-456,* InTech.

Razali, S., Meng, Q., Yang, S-H., 2009. Multi-robot cooperation using immune network with memory. In *Proceedings of the 2009 IEEE International Conference on Control and Automation, Christchurch, New Zealand: pp. 145-150.*

Razali, S., Meng, Q., Yang, S-H., 2010. A refined immune systems inspired model for multi-robot shepherding. In *Proceedings of the 2010 Second World Congress on Nature and Biologically Inspired Computing, Kitakyushu, Fukouka, Japan: pp. 473-478.*

Sathyanath, S., Sahin, F., 2002. AISIMAM—An AIS based intelligent multi-agent model and its application to a mine detection problem. In *Proceedings of the 1st International Conference on Artificial Immune Systems (ICARIS).*

Shehory, O., Kraus, S., 1998. Methods for task allocation via agent coalition formation, *Artificial Intelligence 101(2): pp. 165-200.*

Stranders, R., Farinelli, A., Rogers, A., Jennings, N. R., 2009. Decentralised coordination of mobile sensors using the max-sum algorithm. In *Proceedings of the of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI-09): pp. 299-304.*

Sun, S-J, Lee, D-W., Sim, K-B., 2001. Artificial immune-based swarm behaviors of distributed autonomous robotic systems. In *Proceedings of the 2001 IEEE International Conference on Robotics and Automation (ICRA 2001), 4: pp. 3993-3998.*

US Department of Defence, 2005. *Unmanned Aircraft Systems Roadmap 2005-2030*.

US Department of Defence, 2009. *Unmanned Systems Integrated Roadmap FY2009-2034*.

Vargas, P. A., de Castro, L. N., Michelan, R., 2003. An immune learning classifier network for autonomous navigation. In *Proceedings of the 2nd International Conference on Artificial Immune Systems (ICARIS) 2787: pp. 69–80.*

Watanabe, Y., Ishiguro, A., Shirai, Y., Uchikawa, Y., 1998. Emergent construction of behavior arbitration mechanism based on the immune system. In *Proceedings of IEEE ICEC: pp. 481–486.*

Whitbrook, A. M., Aickelin, U., Garibaldi, J. M., 2007. Idiotypic Immune Networks in Mobile-Robot Control. In *IEEE Transactions on Systems, Man, And Cybernetics—Part B: Cybernetics, 37(6): pp. 1581-1598.*