Proceedings of the 2008 Winter Simulation Conference
S. J. Mason, R. R. Hill, L. Mönch, O. Rose, T. Jefferson, J. W. Fowler eds.

AUTOMATED RED TEAMING: AN OBJECTIVE-BASED DATA FARMING APPROACH FOR RED TEAMING

C.L. Chua

DSO National Laboratories
20 Science Park Drive
Singapore 118230

CPT W.C. Sim

Singapore Armed Forces
Operations Research Office
Ministry of Defence, Singapore
Singapore 669645

C.S. Choo

DSO National Laboratories
20 Science Park Drive
Singapore 118230

Victor Tay

Defence Science & Technology Agency
71 Science Park Drive
Singapore 118253

ABSTRACT

In this paper, we describe an objective-based Data Farming approach for red teaming called Automated Red Teaming (ART). The main idea is to develop an ART framework using Evolutionary Algorithms (EAs), Parallel Computing and Simulation, and apply it to uncover exploitable gaps in military operational concepts, complementing the Manual Red Teaming (MRT) effort. The capability of the ART framework was evaluated vis-à-vis MRT using two maritime security scenarios addressed at the International Data Farming Workshops (IDFWs) 14 and 15. The evaluation showed that, in general, results from ART were better than those obtained from MRT, some of which were non-intuitive and surprising solutions.

1 INTRODUCTION

Red teaming is a technique commonly used in the military Operational Analysis (OA) community to uncover system vulnerabilities or to find exploitable gaps in operational concepts, with the overall goal of reducing surprises, improving and ensuring the robustness of the Blue operational concepts (Upton and McDonald 2003; Upton et al 2004). It is currently a manually intensive technique that typically brings together experts relevant to the system under consideration and who are then charged with identifying the systems weaknesses. However, the vulnerability assessments made are usually "bounded" by the knowledge of these Subject Matter Experts (SMEs).

Under the sponsorship from the Singapore Defence Science and Technology Agency (DSTA), DSO National Laboratories (DSO) has developed an ART framework that leverages on the advancing technologies of EAs, Parallel Computing and Simulation. A detailed description of the ART framework was presented by Choo et al (2007), and its capability evaluated using a military scenario in urban operations, with encouraging results. The ART framework was further evaluated vis-à-vis MRT using two maritime security scenarios at IDFW 14 and IDFW 15.

2 OBJECTIVE

The objective of this paper is to share the findings of the evaluation done at the two workshops. It concludes with a description on the follow-on work and also some suggestions on potential applications.

3 OBJECTIVE-BASED DATA FARMING

Data Farming is a process made possible by a convolution of advancements in Agent Based Models, computing power and the ability to organize, analyze and visualize data (Horne and Meyer 2004, Horne et al. 2005). The objective of Data Farming is to generate and observe a large number of possible outcomes for the studied scenario and to obtain insights as to what factors drive the occurrence of each outcome. ART can be considered a specialized variant of Data Farming. Instead of exploring the entire parameter space to generate a response surface, a point (or objective) on the response surface is first identified, followed by a search for the parameters that result in this point. In the context of red teaming, this point corresponds to the scenario where the Red defeats the Blue, and the associated Red parameters set the conditions for this to happen.

A class of algorithms known as Evolutionary Algorithm (EA) uses an iterative process, inspired by biological mechanisms of evolution, to optimize a certain fitness value that can be considered to be the objective function. The growth of the population in an EA is determined through a set of operators such as

recombination, mutation, and selection. EA has been identified as the search algorithm for the objective-based Data Farming.

In the next two sections, an overview of the architecture of the ART framework and an Agent Based Simulation (ABS) called MANA are presented.

4 ARCHITECTURE DESIGN OF THE ART FRAMEWORK

The architecture of the ART framework consists of the following key components (see Figure 1):

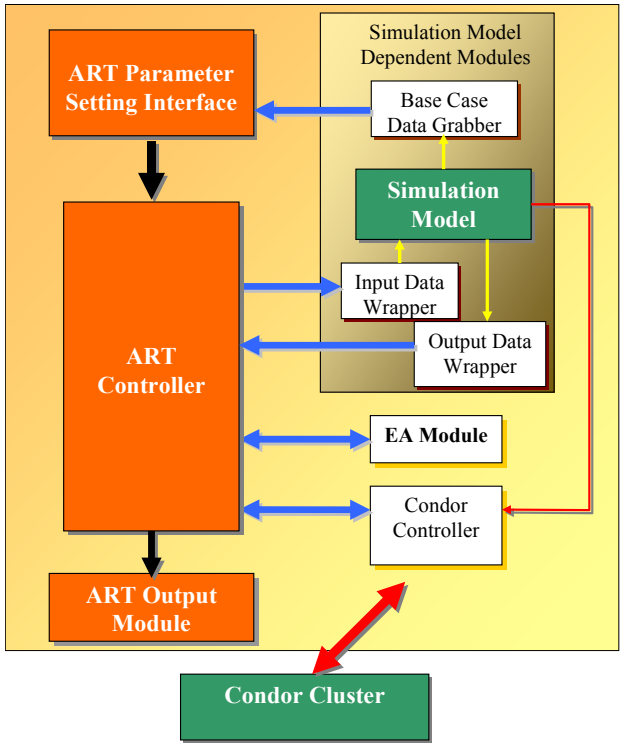

Figure 1: Architecture Design of the ART Framework

The architecture was designed to be modular and flexible enough to incorporate new simulation models (non man-in-the-loop type) and EAs. Starting clockwise from the top left:

- The ART parameters setting interface allows the initial selection of the parameters that are to be varied.
- The simulation model dependent modules add a layer of data flow to and fro the ART framework and simulation models. Data flowing into the simulation models would be the parameters to be executed and the data outflow will be the results from the simulation runs. These data are translated with specially created wrappers from the simulation format to the ART framework data structures.
- The EA module houses the EA library in which the user can choose. The EA module currently contains the Strength Pareto Evolution Algorithm Version 2 (SPEA2), an Evolutionary Multi-Objective (EMO) optimization algorithm. The library is also expected to expand with the addition of other algorithms. The role of the EA module would be to prepare the parameters for the individual simulation, analyze the results, and distill the desired red teaming objectives. The role of the execution is handled by the Condor controller.
- The Condor controller will submit the run of each individual simulation to the Condor cluster. The current Condor cluster in DSO comprises 48 compute nodes which can run in parallel 48 simulations. It will monitor the completion of the individual runs and flag to the ART controller for further processing.
- The ART output module will provide feedback on the whole process, updating the user on the selected parameters and the run results.
- Finally, the ART controller is the heart of the framework providing co-ordination on the whole process.

Detailed explanations to the various modules, including the choice of EA are given in Choo et al (2007).

5 MANA

MANA (Map Aware Non-uniform Automata) is an ABS developed by the Operations Analysis group at Defence Technology Agency (DTA) in New Zealand and is one of the suite of models supporting Data Farming in the International Data Farming Workshop (Lauren 2002). MANA has been used in a number of studies involving land combat, civil violence management, and maritime surveillance.

The strengths of MANA include the user friendliness of the interface, the relative ease in creating a scenario from scratch, and fast execution from being run on pre-compiled executable. Additionally, MANA supports the experimentation with intangibles like behaviors which add complexities to the models. These intangibles include proxies for aggression, leadership, and determination. The tool as a whole allows analyst to investigate warfare as a complex adaptive system and to observe the emergent behavior.

6 MARITIME SECURITY

With shipping at the heart of the global economy, maritime security is required to ensure freedom of the seas and to facilitate freedom of navigation and commerce. Two key

aspects of maritime security are protection of Key Installation (KIN) and anchorage against threats from terrorists and criminals.

IDFW is a series of workshops, led by the Naval Postgraduate School (NPS), where teams use advanced experiment designs, fast running simulations and parallel computing to address challenging military and homeland security related questions. At IDFW 14 in Monterey (Sim et al 2007), a Singapore-led team explored how ART could be applied to ensure robustness of plans for protection of KINs against intrusion of fast boats. The same team continued the effort at IDFW 15 in Singapore (Wong et al 2007), studying a scenario on anchorage protection.

7 PROTECTION OF KEY INSTALLATION

7.1 Scenario

At IDFW 14, the team looked at a scenario where Blue force was conducting coastal patrols to guard against threats on KINs, which were Coastal Surveillance Radar (CSR) equipped with minimum level of self-protection. Red force would attempt to penetrate the Blue defence and inflict damages using various approaches. Any damages to the coastline could be seen as a severe psychological blow to the Blue defence force. It was assumed that the Area of Operation (AO) was far away from the main shipping traffic, and hence neutral shipping was not considered. Similarly, the effects of weather and sea state were not considered. The scenario (see Figure 2) was modeled in MANA. MANA is one of the models that have been incorporated into the ART framework.

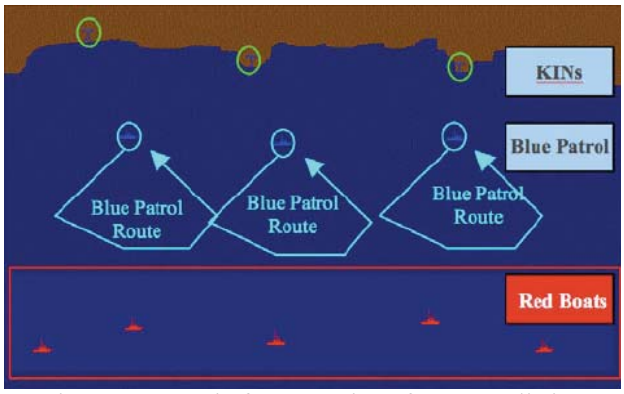

Figure 2: Scenario for Protection of Key Installation

7.2 Blue Force

The Blue force consisted of three KINs and three Patrol Vessels (PVs). The KIN's and PV's operational characteristics are distilled for modelling in MANA. Each KIN was protected by General Purpose Machine Guns (GPMGs). Each PV conducted normal patrol at 15 knots and gave chase at a maximum speed of 25 knots. The PVs were assumed to be capable of neutralizing the Red boats by closing in within 0.5 nm and maintaining this distance for 1 min. The dynamics of the close water combat was not modelled. In addition, the PVs would be activated to investigate detections made by the CSRs so as to achieve target identification and neutralization. The Blue force was assumed to have perfect communication. A summary of the key inputs used for the KIN and PV are given in Table 1 and Table 2 respectively.

Table 1: Key Inputs for Blue KIN

CSR Detection Range (nm)	5
Weapon Range (km)	2
Weapon Single Shot Probability of Hit	0.1

Table 2: Key Inputs for Blue PV

PV Speed [Patrol] (knots)	15
PV Speed [Chase] (knots)	25
PV Detection Range (nm)	3
PV Identification (ID) Range (nm)	1

7.3 Red Force

Five Red boats were modelled as small fishing boats with a maximum speed of 25 knots and loaded with explosives. These boats had a short visual detection and identification range of 1 nm. The five Red small boats would act independently without communicating with each other. By limiting the Red's detection/identification and communications ability, we are modelling small autonomous low technology units which are the likely kind of threats maritime security units are facing during peace time. Table 3 summarizes the key inputs used for the small boat.

Table 3: Key Inputs for Red Boats

Maximum Speed (knots)	25
Detection/ID Range (nm)	1

7.4 Measures Of Effectiveness

There were two Measures Of Effectiveness (MOEs):

- Mean Red Mission Success, defined as the number of successful Red attacks on KINs/Coastline. Red mission was considered successful when at least one boat managed to penetrate the Blue defence.
- Mean Red Attrition.

7.5 Manual Red Teaming

The team at IDFW 14 was first tasked to work on the Red attack plan against a fixed Blue patrol plan through MRT. Two plans were developed, based on the dual objectives of maximizing the mean Red mission success and minimizing mean Red attrition.

The first plan was based on "Flanking" where the Red small boats penetrated through two flanks and in doing so stretched the Blue resources, i.e. the Blue PVs (see Figure 3). For this tactic, Red force was able to achieve a mean mission success of 100% and a mean attrition of 0.85.

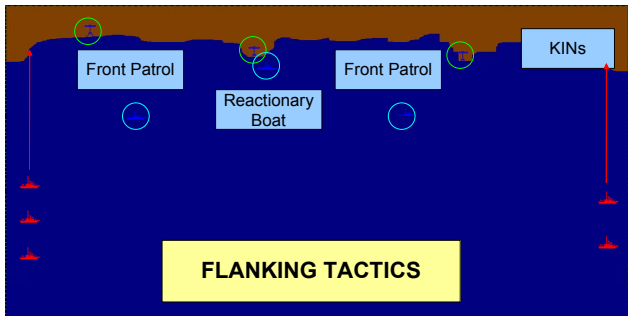

Figure 3: Flanking Tactics developed through MRT

The second plan was based on "Saturation". In this tactic, three Red small boats penetrated through the centre, and tried to saturate the Blue PVs, leaving some gaps for the remaining two small boats to sneak through by the sides (see Figure 4). The mean mission success and mean attrition for Red were 100% and 3.05 respectively.

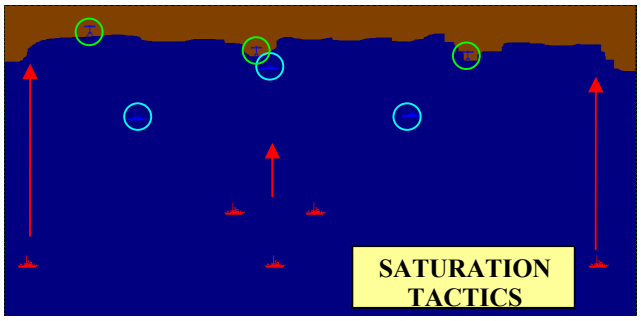

Figure 4: Saturation Tactics developed through MRT

7.6 Automated Red Teaming

Next, the team subjected the Blue patrol plan to ART. Besides evolving the Red penetration plan, the team also looked at how Red behaviours - aggression, cohesiveness and determination, could improve the plan. Aggression towards Blue PV, unit cohesiveness, and determination in moving towards objective (KIN/Coast) were chosen to represent the behaviours of the Red force.

The intangible parameters were presented and effected in MANA through a range of values. A negative value for aggression meant that the Red boats feared the Blue PV, a position value would imply the opposite, and zero would mean indifference. Similarly, a negative value for cohesiveness would imply the Red boats tended to spread out, positive means preferring to cluster, and zero being indifference to each other's presence. For determination, a positive value implied a pre-disposition towards the final objective, zero was indifference, and negative being avoid reaching the final objective.

It was interesting to note that the ART framework produced a decoy tactic that surprised the team, as shown in Figure 5. In this tactic, one of the Red small boats (the one in the centre and acted as a decoy), was deployed to lure the Blue PV on the left towards the right side to create an opening for the other two small boats to charge towards their objectives. The decoy's sweeping movement caused enough distraction to leave the left side of the map exposed for the two small boats to succeed in their mission. The mean Red mission success and the mean Red attrition achieved were 100% and 1.89 respectively. Thus, the effectiveness of ART generated plan was somewhere between the "Saturation Tactics" and the "Flanking Tactics".

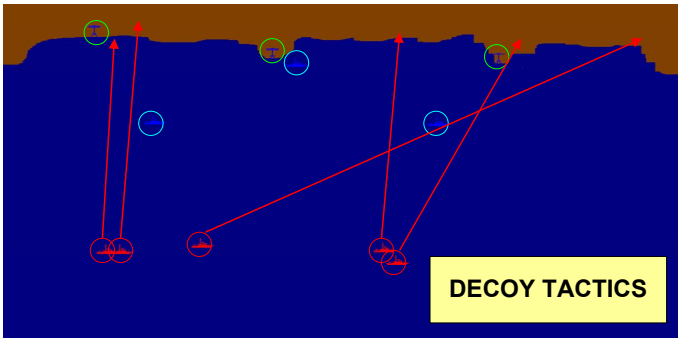

Figure 5: Decoy Tactics developed through ART

It was also observed that in general, the Red small boats were very focus in charging towards the KINs (highly determined) while trying to avoid the Blue PVs (negative values for Red aggression). However, Red small boats were more cohesive in the decoy tactics. The findings of the ART runs, as compared to the MRT effort, are summarized in Table 4.

Table 4: Summary of Results for Protection of KINs

	Flanking Tactics (MRT)	Saturation Tactics (MRT)	Decoy Tactics (ART)
Red Aggression	-60	-60	-83
Red Cohesiveness	-100	-100	8
Red Determination	60	60	53
Red Mission Success	100%	100%	100%
Red Attrition	0.85	3.05	1.89

8 ANCHORAGE PROTECTION

8.1 Scenario

Following the investigation done at IDFW 14, the same team continued the evaluation of ART framework, looking at anchorage protection. In this scenario, Blue force conducted patrols to guard against threats on an anchorage, with ten commercial ships anchored in the protected area. Red force would attempt to penetrate the Blue defence and inflict damages on the anchored vessels, using various approaches. Any damages to the commercial shipping will deal a severe psychological blow to the Blue defence force. As in the previous scenario, the AO was assumed to be away from the main shipping traffic. The anchorage covered an area of 30 nm by 10 nm. The AO was designed to be 100 nm by 50 nm so as to allow the Red force greater depth in their movement. The scenario was modelled in MANA, as shown in Figure 6.

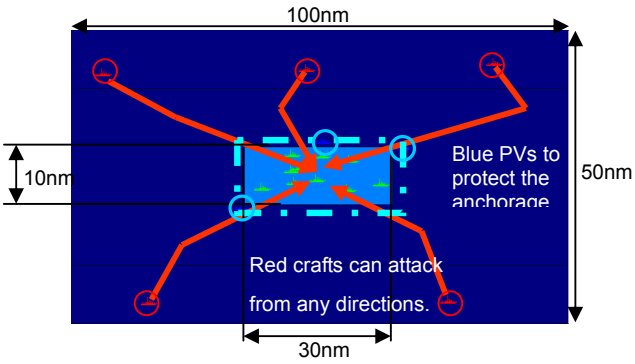

Figure 6: Scenario for Anchorage Protection

8.2 Blue Force

The Blue force consisted of three PVs, each conducting normal patrol at 8 knots and pursuing threats at a maximum speed of 16 knots. The PVs were assumed to be capable of neutralizing the Red boats by closing in within 2 nm. Once again, the dynamics of the close water combat was not considered, and the PVs were assumed to have perfect communication. Table 5 provides the key inputs for the PVs.

Table 5: Key Inputs for Blue PV

PV Speed [Patrol] (knots)	8
PV Speed [Chase] (knots)	16
PV Detection Range (nm)	6
PV Identification (ID) Range (nm)	2

8.3 Red Force

There were five small boats with a maximum speed of 16 knots and loaded with explosives. Each boat had a detection and identification range of 2 nm, and would act independently. Table 6 shows the key inputs for the small boats.

Table 6: Key Inputs for Red Boats

Maximum Speed (knots)	16
Detection/ID Range (nm)	2

8.4 Measures Of Effectiveness

There were three Measures Of Effectiveness (MOEs):

- Mean Red Mission Success, defined as the number of successful Red attacks on Neutral Commercial Shipping. Red mission was considered successful when at least one boat managed to penetrate the Blue defence.
- Mean Red Attrition.
- Mean Neutral Shipping Destroyed.

8.5 Manual Red Teaming

The team at IDFW 15 started to develop the Red attack plan against a fixed Blue patrol plan through MRT. The restriction was that three Red small boats had to start from the northern edge of the AO, and the other two small boats from the southern edge of the AO. The starting point was also limited to a 100 nm by 10 nm area off the starting edge so as to avoid placing the start points too close to the anchorage.

The team decided to fully utilize the Red force numerical advantage and launch a simultaneous attack on the anchorage area to saturate the Blue PVs (see Figure 7). This was developed based on maximizing the mean Red mission success, minimizing mean Red attrition, and maximizing mean Neutral shipping destroyed.

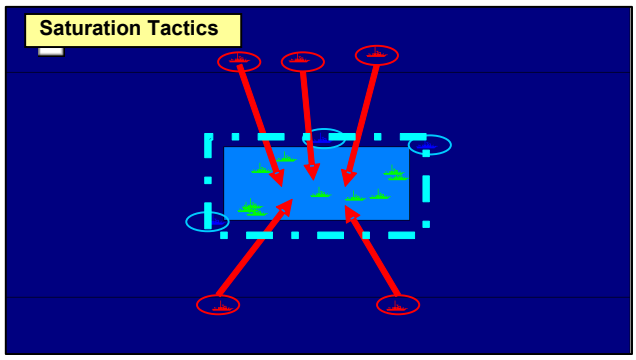

Figure 7: Saturation Tactics developed through MRT

The results obtained were 100% for mean Red mission success, 1.96 for mean Red attrition and 3.05 for Neutral shipping destroyed.

8.6 Automated Red Teaming

ART was applied next. The same intangible parameters were chosen to be evolved - aggression towards Blue PV, unit cohesiveness, and determination in moving towards objective (anchorage) representing the behaviors of the Red force.

Similar to the plan developed by MRT, ART generated a simultaneous red attack tactic towards the centre of the anchorage area with re-attack flexibilities. This would cater for cases where the anchored vessels were dispersed nearer to the anchorage edges. This is a refinement over the MRT tactic which objective was to reach the anchorage. The ART results pushed it further with the Red boats traversing the anchorage looking for dispersed vessels. The ART has given insights that it is not enough to just stop the Red boats from reaching the anchorage but it is also important to prevent the leaked boats from maneuvering with in the anchorage. Insights like these can help the planners to refine the security tactics. The ART tactic for the red team is as shown in Figure 8.

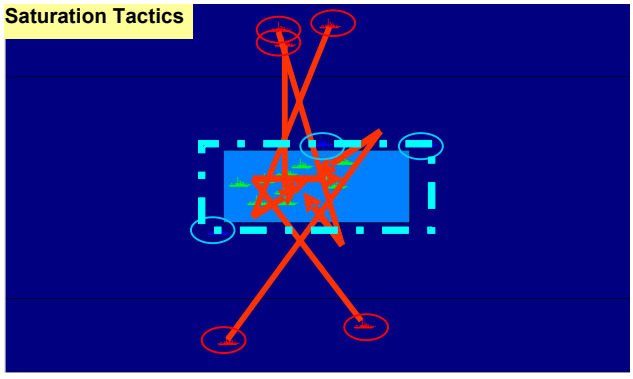

Figure 8: Saturation Tactics developed through ART

The ART generated tactic was able to perform better than the MRT plan, achieving a mean Red mission success of 100% with a lower mean Red attrition of 0.48, and a higher mean Neutral shipping destroyed of 4.52. It was also observed that the Red small boats had only a mild fear of the Blue PVs (-4 as compared to -60 for the MRT case), and they were more cohesive (-16 as compared to -100). The findings of the ART runs, as compared to the MRT effort, are summarized in Table 8.

Table 8: Summary of Results for Anchorage Protection

	Saturation Tactics (MRT)	Saturation Tactics (ART)
Red Aggression	-60	-4
Red Cohesiveness	-100	-16
Red Determination	60	45
Red Mission Success	100%	100%
Red Attrition	1.96	0.48
Neutral Attrition	3.05	4.52

9 CONCLUSION

The evaluation of the ART framework vis-à-vis MRT at IDFW 14 and IDFW 15 showed that ART could complement the manual efforts by introducing unique and surprising solutions (as in the Decoy tactic at IDFW 14) or providing refinement to the manual plans (as in the Saturation tactic at IDFW 15), which might otherwise, be overlooked.

However, it is important to stress that the objective of ART is not to replace MRT, but to complement the MRT effort. There is still a need to involve analysts to make sense of the ART results, at least for the current state of ART.

10 FUTURE WORK

The ART framework is still in its early stage of development. Besides the plan to incorporate new EAs and new models into the framework, further tests are required to benchmark its performance and test its robustness.

There are also some potential spin-offs from this project. It is natural to extend ART to fulfill the concept of Blue Teaming vs Red Teaming, i.e. Automated Co-Evolution (ACE) where the both sides evolve and adapt against changing tactics. The R&D work on ACE is currently on-going. Another potential application is in the calibration of model, e.g. what values should be assigned to the parameters that will result in certain desired outcomes.

ACKNOWLEDGEMENTS

We would like to thank the following organizations that helped to make this R&D work possible:

- Directorate of Research & Development in the Defence Science & Technology Agency (DRD-DSTA), Singapore, for funding this research.
- Defence Technology Agency, New Zealand Defence Force, for sharing the Agent Based Model, MANA.

REFERENCES

Choo, C. S., C. L. Chua, and V. Tay. 2007. Automated Red Teaming: A Proposed Framework for Military Application. In *Proceedings of the 9th Annual Conference on Genetic and Evolutionary Computation*, 1936–1942. ACM.

Horne, G. E., and T. E. Meyer. 2004. Data Farming: Discovering Surprise. In *Proceedings of the 2004 Winter Simulation Conference*, eds. R. G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, 807–813. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Horne, G. E., A. J. Forsythm and S. C. Upton. 2005. Marine Corps Applications of Data Farming. In *Proceedings of the 2005 Winter Simulation Conference*, eds. M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Joines,1077–1081. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Lauren, M. K., and R. T. Stephen. 2002. Map Aware Non-uniform Automata (MANA) a New Zealand Approach to Scenario Modelling. *Journal of Battlefield Technology* 5(1):27–31.

McIntosh, G., Galligan, D. P., Anderson, M. A., and M. K. Lauren. 2006. Recent Developments in MANA Agent-based Model. In *Scythe* 1(1):38–39.

Sim, W. C., Choo, C. S., Ng, E. C., Martinez-Tiburcio, F., Toledo-Ramirez, E. R., and K. Lin. 2007. Applying Automated Red Teaming in a Maritime Scenario. In *Scythe* 1(2):26–29.

Upton, S. C., and M. J. McDonald. 2003. Automated Red Teaming Using Evolutionary Algorithms. *WG31 – Computing Advances in Military OR*.

Upton, S. C., Johnson, S. K., and M. J. McDonald. 2004. Breaking Blue: Automated Red Teaming Using Evolvable Simulations. In *Proceedings of Genetic and Evolutionary Computation Conference 2004*. Available via < http://www.cs.bham.ac.uk/~wbl/biblio/gecco2004/prof99.html> [accessed August 19 2008]

Wong, A., Sim, W. C., Chua, C. L., Lim, Y. L., Chin, S. C., Teo, C., Lampe, T., Hingston, P., and B. Abbott. 2007. Applying Automated Red Teaming in a Maritime Security Scenario. In *Scythe* 1(3): 3–5.

AUTHOR BIOGRAPHIES

CHING LIAN CHUA is currently a Member of Technical Staff (MTS) in Operations Research Laboratory (ORL), DSO National Laboratories, Singapore. He received his B.E. in Mechanical and Production Engineering with specialization in Mechatronics from Nanyang Technological University (NTU), Singapore 2002, and his M.S. in High Performance Computation for Engineered Systems from Singapore-MIT Alliance (SMA), National University of Singapore (NUS) in 2003. His research interests lie in operations research, simulations and algorithms.

CHWEE SENG CHOO is currently a Principal Member of Technical Staff in Operations Research Laboratory (ORL), DSO National Laboratories, Singapore. He received his B.S in Physics from the National University of Singapore (NUS) in 1992, and his M.S. in Operations Research from Stanford University in 1997. His areas of interests include Combat Modelling, Simulation and Analysis, Experimental Designs, Evolutionary Computation and Data Farming.

CPT WEE CHUNG SIM is currently a military analyst in SAF Ops Research Office, Singapore. He has keen interest in the area of Agent Based Simulations, Data Farming and Statistical Analysis. He holds a Bachelor of Engineering (2nd Class Upper Honours) in Electrical and Electronics Engineering from Nanyang Technological University, in 2003.

VICTOR TAY is a Principal Engineer in the Defence Science and Technology Agency (DSTA), Singapore, and currently, he holds the appointment as Programme Manager in the Directorate of Research and Development (DRD) in DSTA, responsible for the long-term technology master planning of the Modelling & Simulation portfolio. He received his M.S. in Interactive Simulation from the University of Central Florida in 1999.