UNCLASSIFIED

Behaviour in Simulated Combat

Adaptation and Response to Complex Systems Factors

Dr. Kevin Sprague
Land Forces Operational Research Team

Dr. Peter Dobias
Land Forces Operational Research Team

DRDC CORA TM 2008-044
November 2008

Defence R&D Canada
Centre for Operational Research and Analysis

Land Forces Operational Research Team

National Defence Défense nationale

Canada

Behaviour in Simulated Combat

Adaptation and Response to Complex Systems Factors

Dr. Kevin B. Sprague
Land Forces Operational Research Team

Dr. Peter Dobias
Land Forces Operational Team

DRDC CORA

Technical Memorandum
DRDC CORA TM 2008-044

November 2008

Principal Author

Original signed by

Dr. Kevin B. Sprague

Defence Scientist

Approved by

Original signed by

Dr. D.S. Haslip

Section Head, Land and Operational Commands

Approved for release by

Original signed by

Dale Reding

Chief Scientist DRDC CORA

Defence R&D Canada – Centre for Operational Research and Analysis (CORA)

Abstract

Tactical combat has been demonstrated to exhibit properties of complex adaptive systems (CAS). In some cases, CAS dynamics can lead to large-scale (possibly catastrophic) impacts that are not in line with expectations under traditional thinking. In this paper, recognizing and exercising some degree of influence over combat CAS dynamics is investigated. In particular, we discuss approaches to selectively “drive” a conflict towards more favourable regions of the available phase space. The issue is addressed at the fundamental level of entity behaviour and interaction. Two key factors for consideration towards such a goal are 1) optimization of combatant behaviour and 2) awareness of and response to complexity within the system. A set of complex systems measures of effectiveness (CMOEs) appropriate to combat, drawn from various complex systems factors available in literature, are proposed and investigated. These measures provide a window into the dynamical progression of a combat CAS, while behaviour modifications offer the means to adapt to its changing conditions. The interplay between the two factors comprises the underlying theme of this paper. Candidate CMOEs discussed include: the fractal dimension, Shannon entropy (two forms: Carvalho-Rodrigues and spatial entropy), the Hurst coefficient, the self-similarity parameter, and symmetropy. Simulations are used to illustrate how a CAS mindset and adaptive behaviour can be leveraged to improve (simulated) mission results. CMOEs are first analyzed in repeated simulations of a simplistic combat scenario to help develop optimal behaviours. The derived behaviours are later combined with real-time awareness of a CMOE to assist in decision-making during a simulated mission of comparable nature to the original combat scenario. It was found that, overall, this new degree of freedom improved the quality of the mission outcome. Also, several CMOEs were found to be of apparent use only in repeated simulations (as opposed to real-time) due to data requirements and high variances in individual runs.

Résumé

La preuve a été faite que le combat tactique possède certaines des propriétés des Systèmes adaptatifs complexes (SAC). Dans certains cas, la dynamique des SAC peut avoir des répercussions à grande échelle (potentiellement catastrophiques) qui ne correspondent pas à ce que pouvaient laisser prévoir les modes traditionnels de pensée. Ce document se penchera sur la reconnaissance et l'exercice d'un certain degré d'influence sur la dynamique des SAC de combat. En particulier, nous discuterons des approches permettant de « mener » sélectivement un conflit vers les régions les plus favorables de l'espace des phases. La question est abordée au niveau le plus fondamental du comportement et des interactions des entités. Pour atteindre ce but, deux facteurs-clés seront étudiés : 1) l'optimisation du comportement du combattant et 2) la reconnaissance de la complexité à l'intérieur du système et la réponse face à celle-ci. Un ensemble de mesures de l'efficacité des systèmes complexes (MESC) convenant au combat, puisés parmi de nombreux facteurs de systèmes complexes que l'on retrouve dans la documentation, est proposé et étudié. Ces mesures constituent une fenêtre d'observation dans la progression dynamique des SAC de combat. Des modifications du comportement constituent pour leur part un moyen d'adapter les SAC aux nouvelles conditions. Les actions réciproques entre ces deux facteurs constituent le thème sous-jacent de la présente étude. Les MESC dont il est question ici sont les suivantes : la dimension fractale, l'entropie de Shannon (deux formes : Carvalho-Rodrigues et l'entropie spatiale), le coefficient de Hurst, le paramètre d'autosimilarité et la symétropie. On a fait appel à des simulations pour illustrer de quelles façons une attitude favorable aux SAC et un comportement adaptatif peuvent avoir un effet de levier sur l'amélioration des résultats (simulés) d'une mission. Les MESC sont d'abord analysées dans le cadre de simulations répétées d'un scénario de combat simple afin d'aider le développement de comportements optimaux. Les comportements dérivés sont par la suite combinés avec la connaissance en temps réel des MESC afin d'aider le processus de décisions au cours d'une mission simulée de nature similaire à celle du scénario de combat original. On a découvert que, dans l'ensemble, ce nouveau degré de liberté améliore la qualité des résultats de la mission. De plus, on a également découvert que plusieurs MESC ne semblaient avoir une fonction apparente que dans un contexte de simulations répétées (par opposition à des simulations en temps réel) en raison des exigences relatives aux données et des variances élevées dans les séquences individuelles.

Executive summary

Behaviour in Simulated Combat: Adaptation and Response to Complex Systems Factors

Introduction: Traditional analysis of combat has focused on equation-based models and attrition-based measures that essentially ignore the spatial component of combat (e.g., Lanchester equations). Additionally, Normal (Gaussian) statistics are implicitly utilized for many parameters of interest. However, tactical combat has been demonstrated to exhibit properties of complex adaptive systems (CASs) with underlying fractal statistics for both spatial and non-spatial (e.g., radio traffic) components. This is true both for situations where opposing forces try to eliminate one another, and also for desirably non-lethal engagements such as hostile crowd management. In some cases, CAS dynamics can lead to large-scale (possibly catastrophic) impacts that are not in line with expectations under traditional thinking. In others, such events are unlikely. In all cases, there exists the potential to observe the evolution of processes within the combat system wherein the interacting components give rise to emergent behaviour and/or special patterns/symmetries (among other possibilities), having varying contributions to events on all time horizons.

Recognizing and exercising some degree of influence over CAS dynamics is thus an important aspect of command and control (C2) and should not be overlooked. In particular, it is desirable to seek out methods to selectively “drive” a conflict towards more favourable regions of the available phase space (i.e., the better portion of the set of all things that can happen). Two key factors for consideration towards such a goal are 1) the optimization of combatant behaviour and 2) awareness of and response to complexity within the system. Complexity is observed and interpreted with respect to a set of complex systems measures of effectiveness (CMOEs) that appear to be relevant for tactical combat. CMOEs provide a window into the dynamical progression of the system while behaviour modification offers the means to adapt to risks and exploit emerging patterns intrinsic to the particular CAS under study. In this work, several candidate CMOEs are investigated through closed combat simulations, including: the fractal dimension, Shannon entropy (two forms: Carvalho-Rodrigues and spatial entropy), the Hurst coefficient, the self-similarity parameter, and symmetropy. The simulations illustrate how a CAS mindset and adaptive behaviour can be leveraged to achieve better C2 and improve chances of mission success. CMOEs are analyzed in repeated simulations of a simple combat scenario to help develop optimal behaviours using a genetic algorithm (GA). The derived behaviours are later combined with a real-time awareness of a CMOE to assist in decision-making during a simulated mission of comparable nature to the original combat scenario.

Results: Results reflect the notion that knowledge of complexity coupled with behaviour optimization can be leveraged to improve (simulated) mission outcomes. Using averaged properties of multiple simulations, information from each of the CMOEs introduced was leveraged to improve upon mission success rates through

1) Understanding the overall progression of dynamical aspects of combat; and

2) Partitioning the scenario and developing optimal behaviour profiles for friendly force entities (BLUE) within each partition via a GA.

The optimized behaviour profiles were then used in real-time to improve upon the quality of results for a separate, but related, simulated mission. Use of the behaviour profile was triggered by recognizing, through limited situational awareness (SA) of a CMOE, a change in the spatial pattern of disorder within the opposing (RED) force. Note that the scenarios investigated constituted small confrontations and consequently the data sets used were sparse. As a result, few CMOEs were of practical value in the real-time scenario due to the required support data. Some CMOEs required hundreds or thousands of data points to be useful, which is reasonable for repeated simulations but not so within a real-time scenario of the magnitude investigated.

A tertiary product of the study is a set of lessons-learned and insights regarding the use of the MANA GA for evolving optimal combatant behaviour profiles in combat simulations:

- Multiple simulations should be performed for each set of behaviour variables explored by the GA so as to buffer against the effects of randomness in the outcomes of a conflict. The MANA 'multi-runs' option provides this capability (10+ were used in this study);

- The final solution provided by the GA should be heavily validated through repeated simulation. Furthermore, it is prudent to compare the performance of the final solution with other solutions that performed extremely well in previous runs in case a good solution was 'lost'. Furthermore, solutions using different GA operator settings (i.e., mutation rates, presence/absence of cross-over) should be explored and compared. Lastly, it may be instructive to compare the final GA results with validations of 'best guess' solutions formed by the practitioner before running the GA; and

- Testing for genetic drift (evolution in the absence of fitness pressure) and evolving the system using a high mutation rate may help to eliminate extraneous variables (genes), thus improving performance and simplifying the interpretation.

Significance: Planning for missions likely to encounter conflicts can be aided by employing optimal, GA-evolved behaviour profiles for agent combatants, allowing one to explore the impact of a large space of possible tactics in a relatively short time. Moreover, unanticipated patterns of beneficial behaviour might be discovered by the GA search.

A deeper understanding of the dynamical progression of a conflict is possible through analysis of CMOEs. Specifically, non-attrition based dynamics are more readily observed. Depending on the scenario, a real-time tactical advantage may be gained by exploiting a not-so-obvious complexity degree of freedom acquired through SA.

Future plans: The identification and utility of additional CMOEs will be investigated, as well as the incorporation of these and other CMOEs into the GA fitness function. Designing more robust (as opposed to strictly optimal) GA solutions will also be pursued to buffer against solution breakdowns with respect to variation within a set of prescribed performance tolerances. Better GA management will also be investigated to address tractability concerns in the general optimization problem for evolving tactics in a combat simulation system.

Dr. Kevin B. Sprague and Dr. Peter Dobias; DRDC CORA TM 2008-044; DRDC CORA; November 2008

Sommaire

Behaviour in Simulated Combat: Adaptation and Response to Complex Systems Factors

Introduction : L’analyse traditionnelle du combat porte sur des modèles basés sur des équations et des mesures reposant sur l’attrition qui essentiellement ignorent la composante spatiale du combat (p.ex., les équations de Lanchester). De plus, les méthodes statistiques normales (Gaussiennes) sont implicitement utilisées pour de nombreux paramètres d’intérêt. Toutefois, il a été prouvé que le combat tactique possède les propriétés des systèmes adaptatifs complexes (SAC) grâce aux statistiques fractales sous-jacentes pour les composantes spatiales et non spatiales (p.ex., le trafic radio). Cela est vrai pour les situations où des forces opposées tentent de s’éliminer mutuellement de même que dans le contexte d’affrontements préférablement non meurtriers, tels que la gestion de foules hostiles. Dans certains cas, la dynamique des SAC peut avoir des répercussions (potentiellement catastrophiques) à grande échelle qui ne correspondent pas à ce que pouvaient laisser prévoir les modes traditionnels de pensée. Dans d’autres cas, de tels événements sont improbables. Dans tous les cas, il est possible d’observer l’évolution des processus au sein même du système de combat, là où les éléments qui interagissent engendrent des comportements émergeants et/ou des modèles spéciaux/des symétries (entre autres possibilités), qui affectent les événements à des degrés divers sur tous les horizons temporels.

Reconnaître et exercer un certain degré d’influence sur la dynamique des SAC sont donc ainsi des aspects importants du commandement et du contrôle (C2) et ceux-ci ne devraient pas être négligés. En particulier, nous discuterons des approches permettant de « mener » sélectivement un conflit vers les régions les plus favorables de l’espace des phases (c.-à-d. la meilleure portion de l’ensemble de tout ce qui peut se produire). Pour atteindre ce but, deux facteurs-clés seront étudiés : 1) l’optimisation du comportement du combattant et 2) la reconnaissance de la complexité à l’intérieur du système et la réponse face à celle-ci. Cette complexité est observée et interprétée en fonction d’un ensemble de mesures de l’efficacité des systèmes complexes (MESC) qui semblent être pertinentes pour le contexte des combats tactiques. Les MESC constituent une fenêtre d’observation dans la progression dynamique du système. Les modifications du comportement offrent quant à elles la possibilité de s’adapter aux risques et d’exploiter les modèles émergeants intrinsèques aux SAC étudiés. Dans ce document, plusieurs modèles de MESC sont étudiés par l’entremise de simulations de combats rapprochés, notamment la dimension fractale, l’entropie de Shannon (deux formes : Carvalho-Rodrigues et l’entropie spatiale), le coefficient de Hurst, le paramètre d’autosimilarité et la symétropie. On a fait appel à des simulations pour illustrer de quelles façons une attitude favorable aux SAC et un comportement adaptatif peuvent avoir un effet de levier permettant de meilleures conditions de C2 et ainsi améliorer les chances de succès de la mission. Les MESC sont analysées dans le cadre de simulations répétées d’un scénario de combat simple afin de permettre le développement de comportements optimaux en faisant appel à un algorithme génétique (AG). Les comportements dérivés sont par la suite combinés à une reconnaissance en temps réel d’une MESC pouvant aider à la prise de décisions lors d’une mission simulée de nature similaire au scénario de combat original.

Résultats : Les résultats témoignent du fait que la notion de reconnaissance de la complexité d’une situation, associée à l’optimisation du comportement, peut servir de levier pour améliorer

les résultats (simulés) d’une mission. En utilisant les propriétés moyennées de simulations multiples, on a augmenté l’influence de l’information provenant de chacune des MESC adoptées afin d’améliorer les taux de réussite des missions grâce

1) À la compréhension de la progression globale des aspects dynamiques du combat; et

2) Au découpage du scénario et au développement des profils de comportement optimums pour les entités appartenant à la force amie (BLEUE) dans chaque partition au moyen de l’AG.

Les profils de comportement optimums ont ensuite été utilisés en temps réel pour améliorer la qualité des résultats pour une mission simulée distincte, mais apparentée. L’idée d’utiliser un profil de comportement provient de la reconnaissance ─ par l’intermédiaire d’une connaissance de la situation (CS) limitée d’une MESC ─ d’un changement de la configuration spatiale du désordre au sein de la force d’opposition (ROUGE). Veuillez noter que les scénarios étudiés faisaient appel à de petites confrontations et, en conséquence, les ensembles de données utilisés étaient rares. Il en découle que seul un petit nombre de MESC possédaient une valeur pratique dans le scénario en temps réel en raison des données justificatives nécessaires. Certaines MESC nécessitaient des centaines ou des milliers de points de données pour être utiles, ce qui est raisonnable dans le cadre de simulations répétées, contrairement au contexte d’un scénario en temps réel de l’ampleur de celui étudié.

L’étude a eu un résultat tertiaire : un ensemble de leçons retenues et de réflexions concernant l’utilisation de l’AG MANA pour l’élaboration des profils de comportement de combattants optimums dans les simulations de combat.

- Des simulations multiples devraient être utilisées pour chaque ensemble de variables de comportement explorés par l’AG, de manière à diminuer les effets du hasard dans les résultats d’un conflit. L’option « multi-runs » du MANA offre cette capacité (10+ ont été utilisés au cours de la présente étude);

- La dernière solution fournie par l’AG devrait être abondamment validée par des simulations répétées. De plus, il est prudent de comparer le rendement de la dernière solution avec celle d’autres solutions qui ont eu des résultats extrêmement intéressants lors d’autres essais, au cas où une bonne solution ait été « perdue ». En outre, les solutions faisant appel à différents paramètres de fonctionnement de l’AG (c.-à-d. les taux de mutation, la présence/l’absence de recoupements) devraient être explorées et comparées. Finalement, il pourrait être constructif de comparer les résultats finaux de l’AG avec la validation des meilleures estimations des solutions énoncées par le spécialiste avant de faire appel à l’AG; et

- Effectuer des tests de dérive génétique (évolution en absence de pression d’adéquation) et faire évoluer le système en utilisant un taux de mutation élevé pourraient aider à éliminer les variables extérieures (gènes), améliorant de ce fait le rendement et simplifiant l’interprétation.

Importance : La planification de missions susceptibles de faire face à des conflits peut être facilitée par l’utilisation de profils de comportement optimums élaborés à partir d’un AG pour les agents combattants, permettant ainsi à l'un d'entres eux d'explorer les répercussions de l'utilisation d'une vaste gamme de tactiques possibles en un lapse de temps relativement court. De plus, il est possible que des modèles imprévus de comportement bénéfique puissent être découverts grâce à la recherche faisant appel à l'AG.

Une meilleure compréhension de la progression dynamique d'un conflit est possible grâce à l'analyse des MESC. Plus précisément, les dynamiques qui ne sont pas fondées sur l'attrition peuvent être plus facilement observées. En fonction du scénario, un avantage tactique peut se trouver en temps réel en exploitant un niveau de liberté acquis par la CS et qui ne serait pas immédiatement évident autrement.

Perspectives : L'identification et l'utilité de MESC supplémentaires fera l'objet d'une étude, tout comme l'intégration de ces MESC et d'autres MESC dans la fonction d'adéquation de l'AG. La conception de solution AG plus solides (par opposition à strictement optimales) fera également l'objet de recherches supplémentaires afin de se protéger contre les découpages des solutions relativement aux variations à l'intérieur d'un ensemble de limites de tolérance fixées liées au rendement. Une meilleure gestion de l'AG sera aussi étudiée afin de répondre aux inquiétudes liées à l'exploitation face au problème général de l'optimisation permettant de mettre au point des tactiques dans un système de simulation de combat.

Sprague, K.B.; Dobias, P.; Équipe DRDC CORA TM 2008-044; R et D pour la défense Canada – CORARO; novembre 2008.

This page intentionally left blank.

Table of contents

List of figures

List of tables

Acknowledgements

The authors are grateful of the exceptional efforts put forth by reviewers of this material.

1 Introduction

1.1 Background

Traditional warfare made use of results based on the analysis of classical combat dynamics such as, for example, the Lanchester equations describing attrition rates (Lanchester 1914). Most such models are implicitly based on the assumption of Normal underlying statistical distributions for salient characteristics such as kill probabilities and consequently attrition. However, it can be shown that in many instances the dynamics of combat obey fractal, rather than normal, statistics (Lauren 2003, Dobias 2008a). Examples include the temporal distribution of casualties, the spatial distribution of forces, the distribution of radio traffic and the frequency of conflict intensity. In addition, combat dynamics have been observed to display signatures of self-organized criticality (SOC) when viewed as a complex adaptive system (CAS)[1] (Lauren 2001, Lauren *et al.* 2002, Ilachinski 2004, Dobias 2008b). This means that once the system reaches its critical point[2] (typically an attractor), a rapid transition can occur, possibly leading to catastrophic events having a tremendous impact on the outcome of the conflict under study. If present in a combat system, ignoring these factors can lead to a perilous misconception of the risks involved in an operation (Ilachinski 2004). Furthermore, many general systems exhibiting SOC also contain precursors to large-scale events. The idea that precursors might exist in a combat CAS is also of relevance. Thus recognizing and exercising some degree of influence over complex adaptive system dynamics, especially at or near critical points, is worthy of investigation in the context of C2 and general mission success[3]. In particular, it is of potential value to develop schemes capable of selectively "driving" a conflict to move towards more favourable regions of the available phase space (Ilachinski 2004). Two key factors for consideration towards such a goal are 1) the optimization of combatant behaviour and 2) awareness of and response to complexity within the system. In simulated combat, behaviour encompasses how entities respond to what they know about different situations and what their goals (or commands) are. Knowledge of complexity, on the other hand, is represented by set of complex systems factors having the potential to impact combat effectiveness with regard to attaining the prescribed goals. Such factors are referred to herein as complex systems measures of effectiveness (CMOEs). CMOEs offer just one more piece of information that an entity can use to arrive at a decision concerning how to act at any given time. They provide a window into the dynamical progression of the system, and may even offer some hints as to what to expect when further along. It is the interplay between behaviour and the CMOEs in combat simulations that comprises the underlying theme of this paper.

To understand the dynamics of a complex system, appropriate measures must be established. In this paper the CAS is described via a few indicators, each of which is normally associated with the degree of disorder from a particular, meaningful perspective; namely,

[1] A complex adaptive system is any dynamical system composed of many simple, typically nonlinearly interacting parts, wherein the parts are capable of adapting to a changing environment (Ilachinski 2004).
[2] A critical point is typically characterized by the absence of preferred scales. System characteristics typically exhibit power law relationships in time and space, which are consistent with underlying fractal statistics.
[3] Note that there are also risks associated with adopting CAS-dependent schemes as components of operational warfare. See, for instance, Scherrer (2003).

1) Carvalho-Rodrigues (CR) entropy;

2) Spatial entropy;

3) The fractal dimension;

4) Symmetropy;

5) The Hurst coefficient; and

6) The self-similarity parameter.

The first four of these measures are mainly local in scope time-wise and relate to, or are explicit forms of, entropy, whereas the last two address long-term correlations in the system including the degree of persistence or lack thereof. The famous *1/f* noise, a signature of SOC, relates mainly to item 6–the self similarity parameter. Some measures overlap in information content (most notably item 2 with 3, and also item 5 with 6 for the scenario investigated), thus a reduced set of measures may suffice in general. Most of these measures have been applied to combat dynamics previously and each measure is described, in turn, below (see Section 2).

Since complex systems generally straddle the boundary between order and disorder, and indeed it is this mix that contributes heavily to the fascinating dynamics, it makes sense to observe the temporal evolution of disorder in the system in relation to major dynamical events or possibly even *potential upcoming events* of significance. Achieving this task is not an obvious or trivial matter, since even whether disorder is increasing or decreasing in a CAS is itself a matter of perspective. For instance, in many complex systems, parameters that describe so-called macro-properties (or *emergent* properties) of the system suggest that the unpredictable, nonlinear interactions of system components may self-organize to such an extent that they generate a larger-scale sense of order (i.e., 'disorder' is decreasing). These emergent properties are not easily derivable by analyzing any single component (e.g., attributes of an individual fish do not directly lead one to imagine the shape and behaviour of a school of fish). Such 'macroscopic' order, however, often hinges on 'microscopic' disorder (e.g., Second Law of Thermodynamics–'disorder' is increasing)[4]. Thus it is important to identify exactly what is meant by disorder in a CAS, and furthermore to identify what perspectives are relevant for addressing the problem at hand. Note that in this work the nature of the mechanism behind self-organization and any resulting criticality is not directly measured, but rather is inferred from general observations of the system dynamics.

Characterizing, measuring and tracking the degree of self-organization in a CAS as it progresses to a criticality holds potential value for future analyses. The Hurst coefficient and self-similarity parameter may provide some insight into self-organization, since self-organization is typified by a response to some memory process within the system, and these measures capture long-term correlations (a form of 'memory'). Entropy has been criticized as being inappropriate as an all-inclusive measure of self-organization since it does not adequately describe the complexity resident in several physical and biological systems (Shalizi *et al.* 2004). Nevertheless, it is easily

[4] The second law of thermodynamics states that the thermodynamic entropy (a measure of disorder) of an isolated system not in equilibrium will tend to increase over time, approaching a maximum value at equilibrium.

seen that in many instances entropy does provide information concerning states of the system near or away from relevant order-disorder mixtures. Also, entropy measures are often easily computed which lends to their widespread use. Other measures not addressed in this study have been proposed to characterize self-organization or in some way measure complexity. They include the often intractable 'Kolmogorov complexity' (Kolmogorov 1965) (also referred to as 'algorithmic complexity'), 'statistical complexity' (Shalilzi *et al.* 2004), 'total information' and the related 'effective complexity' (Gell-Mann *et al.* 1996).

The arena employed here to explore complexity in combat is an agent-based distillation (ABD) called 'Map-Aware Non-Uniform Automata' (MANA) (Lauren *et al.* 2002), described below. The main benefits of using the MANA ABD are 1) complex dynamics have been observed in MANA combat scenarios (Lauren 2001, 2002, Dobias 2008b), 2) MANA is easy to set up and use, and 3) simulations produce abundant data for statistical analysis. On the downside, MANA is a laboratory–there is no guarantee that what is observed in simulations will bear any resemblance to what might unfold in a true-to-life scenario. All one can really hope for is to observe some potentially useful property that seems plausible enough to be of relevance outside of simulation, and then follow up under more realistic circumstances.

ABDs form a subset of the more general class of agent-based models (ABMs). Several leading ABMs in the combat regime are based on the philosophy of cellular automata (CA) (Ilachinski 1997, 2000, Lauren 2002). They contain entities (agents) that are controlled by decision-making algorithms rather than by an interactive player. The behaviour of the agents is not predetermined; each agent makes its own decisions based on built-in algorithms, pre-set personal preferences, and its situational awareness (SA). ABMs have been successfully utilized to model a variety of scenarios where emergent behaviour rather than specific technical properties were analyzed. Agent-based distillations are a highly abstracted subset of ABMs. They focus only on the most generic characteristics of an analyzed system while ignoring many detailed features. For instance, a tank might be modeled as a medium speed, armoured vehicle with a significant direct fire capability. In an ABD, the focus is NOT to rely on excessive detail with regard to rigorous physical correctness for every aspect of the model, but rather to capture the main aspects of the environment and behaviour while permitting a less-constrained exploration of the parameter space of possibilities. By abstracting the physical laws, one can focus more on general scenario-exploration without the burden of specifying all the realistic (often irrelevant) details to high accuracy, which can quickly become overly taxing given the payoff on time and effort invested. Also, the more highly specialized (or *deep set*) a model is to a given realistic scenario, the less adaptable it will be to other situations that may display similar dynamics but in a different context or environment. The simplicity of ABDs makes them particularly attractive for analysis and interpretation.

1.2 Aim

The aim of this manuscript is to discuss and illustrate how knowledge of complexity in combat can be characterized and how it might lead to tactical advantage within a few conceptually simple (simulated) combat situations exhibiting fractal properties. Insights gained through simulations can then be leveraged at a later time to aid in the evaluation of more realistic scenarios of a practical concern.

1.3 Scope

Optimizing agent behaviour in a simulated combat operation is addressed. The main engine of optimization is the MANA genetic algorithm (GA). In this study, CMOEs are introduced into the optimization process as an aid to understanding the dynamical properties of the combat system and also to produce an extended mix of conditions that an agent can learn from and adapt or respond to. The former is achieved via the analysis repeated simulations, and the latter via real-time response to a CMOE. The utility of CMOEs is judged with respect to how they contribute to the success and/or efficiency of the optimization process, both in an absolute sense and also in relation to reasonable expectations without knowledge of CMOEs. The interpretive value of responding to CMOEs in real-time is likewise assessed.

The scenario employed is a difficult, closed, small unit operation that was chosen rather arbitrarily. It was not known or expected *a priori* to exhibit any particular patterns with respect to CAS or SOC dynamics.

Traditionally, the enemy force (hereafter referred to as 'RED') has been assumed to be a regular (conventional) force. However, in the current security environment RED can range from conventional forces, to insurgent groups, to gangs and hostile crowds. Although a more conventional force is utilized in the simulations, other types are also discussed with regard to the representative CMOEs.

The basic scenario investigated plays out as follows: a friendly force (hereafter referred to as 'BLUE') of simulated, autonomous agents is pitted against a formidable RED force. Using the MANA GA, the BLUE force is first permitted to evolve optimal tactics against RED based on a conventional knowledge of the battlefield environment (e.g., detections of RED, benefits/drawbacks of clustering, and movement patterns under various circumstances). The simulations are then analyzed using the CMOEs, and this knowledge is later exploited to improve upon the first iteration of behaviour optimization. In this manner, it is shown how knowledge of complexity in repeated simulations can be used to better BLUE's chances of success. In a subsequent phase, BLUE agents are endowed with knowledge of complexity <u>during</u> a simulation run. BLUE agents utilize this new degree of freedom to gain further advantage over their opponent in real-time.

Before delving into simulations, the chosen CMOEs are described in detail and the framework for behaviour representation and development is explained. Furthermore, GAs are portrayed as a search tool designed to 'find' the optimal state of behaviour for BLUE agents in various circumstances given limited situational awareness.

2 Theory

2.1 Complex systems measures of effectiveness in combat

In traditional combat models and wargames the primary measure of effectiveness is often attrition–whether measured directly (number of killed, loss-exchange ratio[5], etc.) or indirectly (attrition-based definition of mission success). However, in some cases the focus on attrition actually ignores the complexity of combat (Dobias 2008b). Furthermore, given the quality of the force protection of modern militaries and the often asymmetric nature of warfare, standard attrition-based measures might be misleading and/or inappropriate for describing combat dynamics with potentially detrimental effects on the mission outcome (a good example of such a case is hostile crowd management).

In this section several potential CMOEs are described–one attrition-based and the others spatial- or vector-based. These measures are deemed by the authors to be appropriate for dynamical analysis of a wide range of combat systems when viewed as CASs. Nevertheless, it should be noted that the applicability of a given measure is situation-specific–none are universally relevant and a particular CMOE is not guaranteed to provide <u>any</u> relevant information whatsoever for a general situation of interest. A key determinant seems to be how disorder unfolds in the system. Therefore, the main focus of most measures at this juncture is on the progression and degree of disorder within the system from various perspectives. As a scenario plays out, order and disorder appear as variations in a stream of data dominated by the rise and fall of the chosen reference patterns. For CR-entropy the reference patterns correspond to phases of attrition, for the fractal dimension they represent self-similarity and clustering, spatial entropy relates to force concentration, and for symmetropy disorder is rated with respect to a spatial pattern symmetry basis. The Hurst coefficient and the self-similarity parameter rate vector deviations from what is essentially a random (long-term) walk. They each rate the degree of disorder with respect to a reference notion about a particular aspect of a CAS that seems to make sense to observe. It is the selection of the reference views that determines what is found and what is not. The list of measures compiled herein is not complete, and disorder is not the only concept of interest. For some other prospective CMOEs worthy of investigation see, for instance, Grassberger-Crutchfield-Young's statistical complexity (Shalizi et al. 2001, Shalizi 2004) and Gell-Mann *et al.*'s (1996) total information and effective complexity, to name a few.

After describing the individual CMOEs, later in this paper it is demonstrated how an appropriate subset of these measures can be applied to a specific (simulated) combat mission to gain a tactical advantage over an enemy. It is also demonstrated that some combinations of measures yield redundant information, and furthermore some are of greater utility when studying repeated simulations rather than a single, real-time scenario.

[5] Ratio between the RED and BLUE killed.

2.1.1 Carvalho-Rodrigues entropy

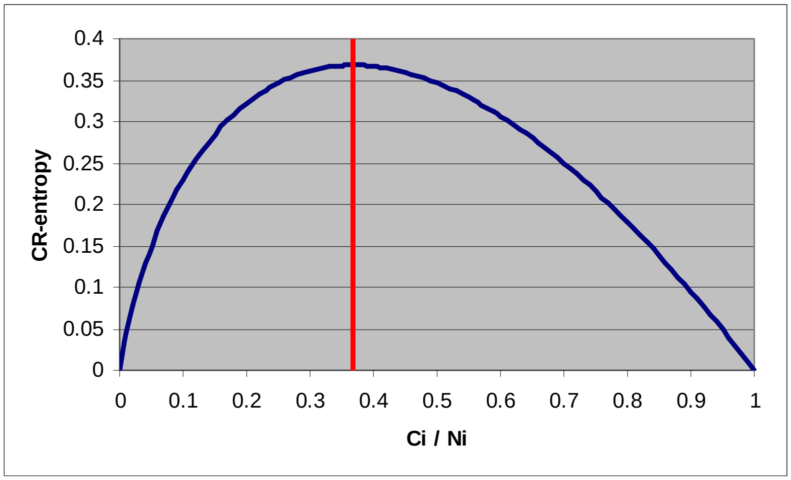

Figure 1: The CR-entropy function.

Several measures have been proposed to capture the complexity of combat. Possibly the oldest of them is *entropy*, of one form or another. As mentioned previously, entropy is a measure of disorder in a system from a particular perspective that varies depending on the application. Carvalho-Rodriguez (1989) (CR) proposed an attrition-based definition of combat entropy for the i^{th} force (i being RED or BLUE) of the form

$$S_i = \frac{C_i}{N_i} \ln \frac{N_i}{C_i} .$$

In the above definition C_i represents the number of casualties and N_i is the force strength of the i^{th} force at time t ($N_i = N_{i0} - C_i$, where N_{i0} is the initial strength of the i^{th} force). The overall combat entropy is then defined as $\Delta S = S_{Red} - S_{Blue}$. Early stages of attrition cause the combat entropy S_i to rise until reaching a maximum at $C_i/N_i \sim 0.37$ (see Figure 1)[6]. Up until this point, high CR-entropy translates to a more precarious position for the force in question. The maximum value qualitatively corresponds to a point of complete breakdown of combat capabilities, with the attrition reaching a positive feedback stage. Once the breakdown point has been reached, however, the interpretation of CR-entropy shifts to a somewhat opposite meaning–as attrition continues the entropy now decreases. The quotient C_i/N_i for two opposing forces can be used to define rough indicator stages of battle. The indicators provided below are based on dividing the range of C_i/N_i for each force into three regions: 1) $C_i/N_i < 0.37/2$ (=0.185) (less than half way to the disintegration point), 2) $0.185 \leq C_i/N_i < 0.37$ (more than half way to the disintegration point), and 3) $C_i/N_i \geq 0.37$ (disintegration).

[6] Note that in Figure 1 a closed combat system is assumed (no reinforcements or deserters, thus $0 \leq C_i/N_i \leq 1$).

- *Advantage BLUE*: There are three cases when BLUE has a clear advantage. The first two imply a moderate advantage and the final one suggests that the advantage is high:
    - Moderate:
        - $C_{\mathrm{BLUE}}/N_{\mathrm{BLUE}} < 0.185$; $0.185 \leq C_{\mathrm{RED}}/N_{\mathrm{RED}} < 0.37$
        - $0.185 \leq C_{\mathrm{BLUE}}/N_{\mathrm{BLUE}} < 0.37$; $C_{\mathrm{RED}}/N_{\mathrm{RED}} \geq 0.37$
    - High: $C_{\mathrm{BLUE}}/N_{\mathrm{BLUE}} < 0.185$, $C_{\mathrm{RED}}/N_{\mathrm{RED}} \geq 0.37$
- *Advantage RED*: Analogous to above:
    - Moderate:
        - $C_{\mathrm{RED}}/N_{\mathrm{RED}} < 0.185$; $0.185 \leq C_{\mathrm{BLUE}}/N_{\mathrm{BLUE}} < 0.37$
        - $0.185 \leq C_{\mathrm{RED}}/N_{\mathrm{RED}} < 0.37$; $C_{\mathrm{BLUE}}/N_{\mathrm{BLUE}} \geq 0.37$
    - High: $C_{\mathrm{RED}}/N_{\mathrm{RED}} < 0.185$, $C_{\mathrm{BLUE}}/N_{\mathrm{BLUE}} \geq 0.37$
- *Balanced*: Neither RED nor BLUE has a notable advantage. RED and BLUE CR-entropies are comparable.

Similarly, the difference ($C_{\mathrm{RED}}/N_{\mathrm{RED}} - C_{\mathrm{BLUE}}/N_{\mathrm{BLUE}}$) is also a relevant parameter to monitor if both forces have C_i/N_i on the same side of the maximum value or are proximal to it, especially considering that small differences in C_i/N_i near categorical boundaries should be interpreted as balanced rather than assigning an advantage to one side. Note that the definition of CR-entropy ignores the spatial dimension that is so important in modern manoeuvre warfare. Nevertheless, it is a useful quantity that contributes to spatiotemporal interpretations when combined with other measures. Furthermore, CR-entropy and/or straight attrition counting can provide valuable information about the underlying causes of spatial disorder. For example, in some cases they may help to resolve whether a burst of attrition is likely the result of a particular state of spatial disorder or if the spatial disorder arose out of attrition.

CR-entropy is a special case of a more general definition of entropy devised by Shannon (1949) in the field of Information Theory. The Shannon expression for entropy is

$$S = \sum_i p_i \ln \frac{1}{p_i}.$$

In the above expression, p_i denotes the probability of the i^{th} option and the summation is over all of the options considered in the model. Considered options may include, for example, the number of incapacitations (leading to CR-entropy), spatial distribution, or detections at certain ranges.

2.1.2 Spatial entropy

The computation and meaning of spatial entropy are akin to the fractal dimension (box counting dimension) (below). In fact, in Section 4 it is shown that the two are nearly indistinguishable for the scenario examined. Ilachinski (2004) suggested this specific form of Shannon entropy based on the spatial distribution of soldiers. A combat area of size *B* is split into a number of sub-blocks of size *b*. If, at any given moment, N_i out of *N* soldiers reside in the i^{th} sub-block, the probability of finding a particular soldier in that sub-block is $p_i(b) = N_i(b)/N$. Then Shannon entropy takes the form

$$S(b) = \frac{1}{2\ln(B/b)} \sum_{i=1}^{(B/b)^2} p_i(b) \ln\left(1/p_i(b)\right).$$

The factor $1/(2\ln(B/b))$ is introduced as a normalization coefficient. Unlike CR-entropy, spatial entropy characterizes combat dynamics independently of attrition. Therefore, it could be used to characterize the spatial dynamics of a conflict even in the absence of attrition (e.g., hostile crowd management).

For randomly distributed individuals $p_i = (b/B)^2$, and entropy $S = 1$. If all of the individuals are in a single sub-block, then $S = 0$. Thus, when individuals are tightly clustered relative to the chosen resolution (scale), entropy is close to 0. Conversely, if they are uniformly distributed over the entire battlefield, entropy is close to 1. In this fashion, entropy is capable of quantifying force cohesion and manoeuvres, and the temporal dependence of spatial entropy provides information about the overall combat dynamics.

2.1.3 Fractal dimension

Another option to describe the dynamics of a combat system is to use the fractal dimension as a measure of the spatial distribution of units (crowd, BLUE force, etc). Fractals are basically geometric objects whose 'parts' appear much like the 'whole', that is, they are self-similar at different length scales. The fractal dimension more-or-less measures the minimum number of variables needed to specify a given fractal pattern (Ilachinski 2004). For instance, a straight line can be described by a single variable, so its fractal dimension is one. A plane requires two variables, so its fractal dimension is two. In these two examples, the fractal dimension is equivalent to the intuitive *Euclidean dimension* of the objects. What is interesting is that certain geometric objects are best described by non-integer fractal dimensions.

Fractal dimensions cannot be computed exactly for real-world data sets and so they must be estimated. Claims based on fractal dimension estimates that are not supported by large amounts of data or corroborating evidence have to be considered as somewhat suspect. The most natural of many possible fractal dimensions to describe the spatial dynamics of combat is the box-counting (or capacity) dimension D_F[7]. Strictly speaking, D_F computed in this manner is more accurately

[7] Note that there exist various 'flavours' of fractal dimension with differing interpretations and methods of computation (e.g., Hausdorf dimension, Renyi dimensions, correlation dimension). The box-counting

termed the ‘box dimension’, however, it is so often used to measure the dimension of fractal sets that it is commonly referred to as the ‘fractal dimension’–the convention employed herein. Albeit somewhat rough, it expresses the relationship between the size of a box ε, and the minimum number $N(\varepsilon)$ of boxes needed to cover all of the subject units. Generally, the dependence is a power law:

$$N(\varepsilon) = (L / \varepsilon)^{D_F} .$$

In the expression above, L is the size of the battlefield. For units uniformly distributed over a two dimensional (2D) battlefield, effectively $D_F = 2$ using this technique (see Euclidean plane example, above). Taking the logarithm of both sides of the equation for $N(\varepsilon)$ given a sufficiently small ε, a formula for D_F is obtained:

$$D_F = \lim_{\varepsilon \to 0} \frac{\ln N(\varepsilon)}{\ln L / \varepsilon} .$$

Practically, ε just needs to be reasonably small compared to the battlefield size L. The battlefield is then divided into $(L/\varepsilon)^2$ squares, and all of the squares that contain at least one agent are counted. Then the ratio $\ln N(\varepsilon) / \ln(L / \varepsilon)$ is calculated. The fractal dimension is qualitatively very similar to spatial entropy, the main difference being that with spatial entropy the probability of finding a given agent in a particular square is considered (and therefore the number of agents within the square), whereas only the presence or absence of agents is considered when computing the fractal dimension. Note that as of version 4.0, calculation of D_F has been incorporated into MANA.

2.1.4 Symmetropy

A new quantity was proposed on the basis of Shannon entropy that measures the combined symmetry and entropy of a given pattern or shape. In this instance the measured quantity in question is the spatial distribution of units. This measure is called symmetropy (Nanjo *et al.* 2001). It captures not only the spatial distribution, but the symmetry of the distribution as well. The definition of symmetropy utilizes a two-dimensional Walsh transform as follows: The battlefield is divided into M x M cells where it is assumed that $M = 2^q$, q being a positive integer. The two-dimensional Walsh transform (Walsh 1910) is then applied,

$$a_{m,n} = \frac{1}{M^2} \sum_{i=0}^{M-1} \sum_{j=0}^{M-1} x_{i,j} W_{m,n}(i, j) ,$$

where $m, n = 0, 1, 2, \ldots, M - 1$, $x_{i,j}$ is the gray-scale value or intensity (e.g., “black” and “white” could have values “1” and “0” respectively) in the i^{th} row and the j^{th} column. $W_{m,n}$ is the two dimensional Walsh function:

---

technique was chosen since it has been applied in other studies of combat dynamics. See, for instance, Lauren (2001), Ilachinski (2003) and Moffat (2006).

$$W_{m,n}(i,j) = \prod_{k=0}^{q-1} (-1)^{\left(b_k(j)b'_{q-1-k}(m) + b_k(i)b'_{q-1-k}(n)\right)} .$$

In the above expression the function $b_k(i)$ denotes the k^{th} bit in the binary representation of i. For instance, when $q=3$, for a number $6 = (110)_2$ the values of b are $b_0(6) = 1$, $b_1(6) = 1$, and $b_2(6) = 0$. $b'_k(m)$ is a transformed function for the binary representation of the number m. The transformation is defined as

$$b'_0(m) = b_0(m),$$
$$b'_k(m) = \left(b_k(m) + b_{k-1}(m)\right) \bmod 2, \quad 0 < k < q$$

This transformation is necessary to obtain a proper ordering of the Walsh functions to allow for calculating projections into the four principal symmetries (vertical, horizontal, centro-symmetric or diagonal, and a double symmetry). The symmetries are as follows. If m is odd and n is even the $W_{m,n}$ measures horizontal symmetry; if m is even and n odd a $W_{m,n}$ has a vertical symmetry; if both are odd it is centro-symmetric, and finally if both are even, double symmetric (Nanjo *et al.* 2001) (see Figure 2). $W_{0,0}$ is the exception.

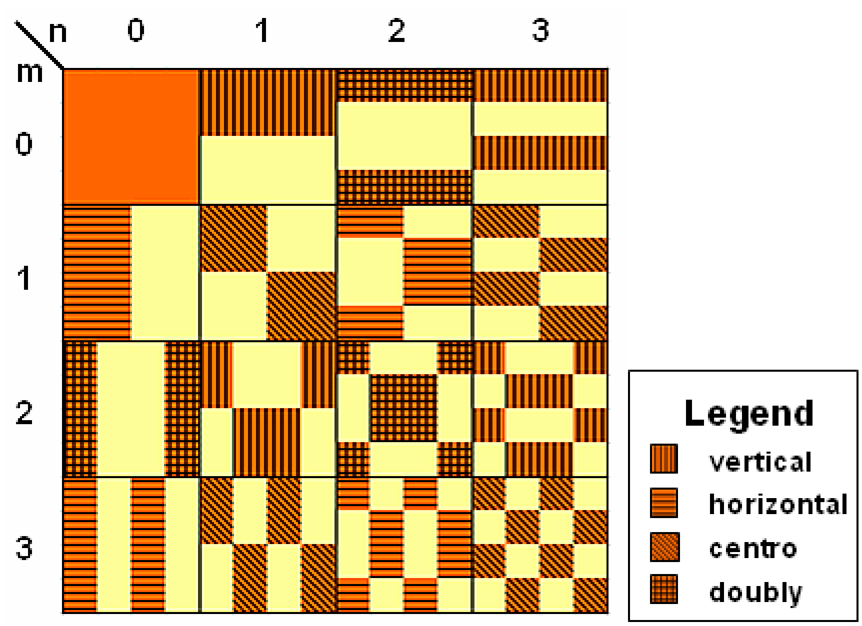

Figure 2: The Walsh Kernel when q=2.

The probability for each of the four types of symmetry (vertical, horizontal, central, and double symmetry) is then

$$P_k = \sum_{(m,n)^{Sk}} (a_{m,n})^2 \Bigg/ \left(\sum_{n=0}^{M-1} \sum_{m=0}^{M-1} (a_{m,n})^2 - (a_{0,0})^2 \right), \; k = 1,2,3,4$$

In the expression above, $(m,n)^{\mathrm{Sk}}$ denotes a sum over a particular symmetry (odd/even, even/odd, odd/odd, even/even). The probabilities satisfy the normalization condition

$$\sum_{k=1}^{4} P_k = 1 .$$

Then Shannon's formula $S = -(1/2)\sum_k P_k \log_2 P_k$ can be applied. The 1/2 factor serves to normalize the symmetropy so that the maximum value is 1. For a random pattern (randomly distributed black and white cells), the symmetropy is 1.0[8]. Moreover, if the value of a symmetry component P_k in the pattern is significantly higher than the others, the pattern is rich in the corresponding symmetry. When the values of the four components are nearly equal, the pattern is poor in symmetry (Nanjo *et al.* 2000, 2005). Thus, for a spatial pattern changing in time, an increase (decrease) in symmetry is captured as a decrease (increase) in *S*. Note that when evaluating symmetropy for a set of point objects (i.e., no area), as *q* gets large whether a point is in a light square or a dark square tends towards randomness for the higher, more refined, (*m,n*) patterns. This phenomenon occurs irrespective of the 'bulk' appearance of the pattern and tends to damp the signature of the symmetropy as seen at lower *q* values (see Figure 3 for a specific example). For this reason, it is suggested that either small values of *q* be used (2 or 3), or that objects be given an appropriate extent and the dependence on *q* tested.

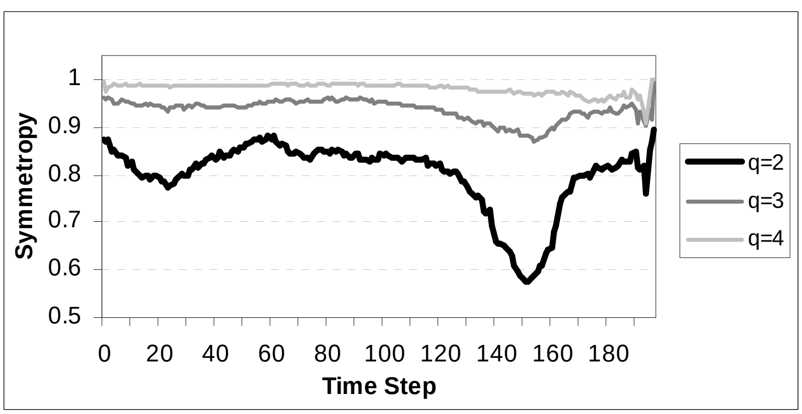

Figure 3: Symmetropy Variation with q for Point Data.

---

[8] Note that in Nanjo *et al.* (2001), the author does not use the normalization factor 0.5. Consequently, therein the maximum value for symmetropy (equal to the symmetropy of a random distribution) is 2.

2.1.5 Hurst coefficient

The interpretation of the Hurst coefficient bears a strong resemblance to that of the self-similarity parameter (below). The calculations for the two quantities also share similar features. In fact, in Section 4 it is shown that they are closely related for the scenario examined.

Temporal and spatial correlations in agent velocity (speed and direction) can be measured using the Hurst coefficient. Such characteristics could possibly provide additional insights into the system dynamics. The correlations are calculated independently for each velocity component. The Hurst coefficient H (also referred to as the *Hurst exponent* in some literature) for velocity is characterized by a scaling between the number of steps and the root mean square distance (RMSD) traveled by an agent. For random (Brownian) motion the relationship between the RMSD (L) and the number of steps (N) is $L = \lambda N^{1/2}$, λ being the length of a single step. The generalized expression relating the number of steps and the MSD via the Hurst coefficient is $L = \lambda N^{H}$. If the Hurst coefficient is $H = 0.5$, a random, Brownian motion is recovered. If $H > 0.5$, the motion is correlated, or in other words, *persistent*. As H approaches 1, the RMSD becomes directly proportional to the number of steps and $L = \lambda N$. This corresponds to intentional travel in a particular direction. If $H < 0.5$ the motion is anti-correlated, meaning the RMSD is less than the corresponding distance for the random walk. For the extreme case of $H = 0$ the RMSD is constant (e.g., circling around a fixed point).

The Hurst coefficient has been used to provide insight into the dynamics of crowds (Dobias 2008b). For a random group of people, such as pedestrians on a street in a downtown area, the speed and direction of individuals is uncorrelated ($H = 0.5$). On the other hand, for marching troops, or a parade, or a demonstrating crowd, the motion can be highly correlated. The Hurst coefficient for such systems would be greater than 0.5. A Hurst coefficient $H < 0.5$ suggests that the mean distance between any two individuals is more-or-less constant.

Various methods are available for computing the Hurst coefficient (Hurst 1951, Feder 1988, Jones 1996, Kaplan). Wavelet transform methods (Jones 1996) and the R/S method (Feder 1988) are frequently encountered in literature. The R/S method as per Kaplan is described briefly below.

At first the data series is divided into boxes of length n. Within each box, the data is (locally) integrated. The integration equation for a data series D_i of N points (within one box) is given by:

$$X(k) = \sum_{i=1}^{k} [D_i - D_{ave}^{box}] \quad ,$$

where k ranges from *1* to N.

Next, the range R is computed for each box as the difference between the minimum and maximum $X(k)$ values:

$$R = Max(\{X(k)\}_k) - Min(\{X(k)\}_k) \, .$$

Now a rescaled range R/S is computed for the box, where S is the standard deviation of the $X(k)$ series. Rescaled ranges are computed for each box of size n and then averaged, which we denote $R/S(n)$. This process is repeated for various box sizes n. Finally, the log-log plot of $R/S(n)$ vs. n is used to calculate a slope, which in turn provides the Hurst coefficient. Since the box size n is limited by the sample size, it is necessary to have a sufficiently large sample to obtain meaningful results. The large number of data points requires replicating a model a large number of times, or including large numbers of entities in the scenario (or both). Also, generally it is best to use values for box sizes uniformly distributed in logarithmic space. This allows for a better fit of log $R/S(n)$ as a function of log n.

2.1.6 Self-similarity parameter

The self-similarity parameter (SSP) α can be viewed as a measure of the 'roughness' of a time series (Peng *et al.* 1994, 1995). Furthermore, as mentioned previously it shares many interpretive properties of the Hurst coefficient (above). For velocity correlations, like the Hurst coefficient the SSP is calculated independently for each velocity component. One advantage of using the SSP over H (as computed above) is that it can be applied to a non-stationary time series. A time series is self-similar if the process $y(t)$ shares the same statistical properties as a properly rescaled process given by $a^{\alpha}\, y(t/a)$. The quantity α has the following interpretation (Goldberger *et al.*):

- $0 < \alpha < 0.5$: The series is anti-correlated. The interpretation is consistent with that of Hurst coefficient in this range;

- $\alpha = 0.5$: Like the Hurst coefficient, this corresponds to a random walk. The data series is uncorrelated (*white noise*);

- $0.5 < \alpha < 1$: Persistence is present in the long-term correlations. The interpretation is consistent with that of Hurst coefficient in this range;

- $\alpha = 1$: This corresponds to $1/f$ noise (or *pink noise*);

- $\alpha > 1$: Correlations exist, but they no longer follow a power law;

- $\alpha = 1.5$: This corresponds to *Brownian noise* (the integration of white noise).

The SSP also relates to self-affinity. A 2-D fractal is self-affine if it scales differently in the x and y directions.

A method called detrended fluctuation analysis (DFA) is commonly used to calculate the SSP (Peng *et al.* 1994, 1995). DFA was designed specifically to deal with non-stationarities (trends) in nonlinear data. For instance, variations in stock indices are composed of two parts. One is a small long-term increase; the other is the deviation from this trend. To analyze long term correlations in the deviations, the trend needs to be removed first. The DFA is based on a root mean square analysis of a random walk. The procedure is briefly summarized below.[9]

To begin with the entire data series is integrated and then divided into boxes of length n. The integration equation for a data series D_i of N points is given by:

[9] DFA has also been applied to computing the Hurst coefficient.

$$y(k) = \sum_{i=1}^{k} [D_i - D_{ave}] \quad ,$$

where k ranges from *1* to N and D_{ave} is the computed average of all N points.

Next, a least-squares fit is performed for each box. The linear fit represents the local trend in the analyzed variable for the box. For a given box size n, values $F(n)$ are computed as root mean squared deviations of the data series $y(k)$ from the local trend $y_n(k)$.

$$F(n) = \sqrt{\frac{1}{N} \sum_{k=1}^{N} [y(k) - y_n(k)]^2}$$

This process is repeated for various box sizes n. Finally, the log-log plot of deviation $F(n)$ vs. n is used to calculate a slope, which in turn provides the SSP. As with the Hurst coefficient, since the box size n is limited by the sample size, it is necessary to have a sufficiently large sample to obtain meaningful results. Also, generally it is best to use values for box sizes uniformly distributed in logarithmic space. This allows for a better fit of log $F(n)$ as a function of log n.

2.2 Application to conflicts

The fractal dimension and corresponding power-laws have been used to describe the statistical distribution of the intensities of wars (Roberts *et al.* 1998), warfare statistics (Richardson 1941) and attack casualties (Lauren 2001, Dobias 2008a), to name a few. In particular, when applied to the spatial pattern of force confrontations on a turbulent battlefield, the fractal dimension expresses how the forces engage each other by forming clusters, and to what extent a large cluster of combatants might itself be viewed as a collection of smaller clusters (i.e., self-similarity) and so on (Lauren 1999). Furthermore, the fractal dimension has been leveraged to explain spatial properties of the battlefront and characterize how dispersed a force is within the overall pattern formed (e.g., tightly grouped versus widely dispersed).

Spatial entropy was employed by Ilachinski (2004) to characterize the spatial distribution of soldiers on the battlefield, force concentration, and the degree of disorder.

CR-entropy was first used to address logistical concerns during military exercises (Carvalho-Rodrigues 1989). Dockery *et al*. (1993) use historical data to argue that CR-entropy is a useful predictor of the outcome of a battle during certain phases of combat.

The Hurst coefficient has been utilized to describe motion in hostile crowd management and, in particular, signal a phase transition between a group confrontational mindset and the inclination to disperse (Dobias 2008b). The SSP potentially could have been used in an analogous manner.

To our knowledge, the concept of symmetropy is a new quantity of consideration for combat dynamics[10]. It seems to hold promise for spatial pattern recognition under a degree of disorder, possibly extending to the identification or classification of forces or force organizational states based on limited SA. It also holds promise for identifying the overall state of a complex system. Examples from the geological sciences involving earthquakes and/or acoustic transitions leverage symmetropy values and corresponding symmetry projections to describe various dynamical aspects of the system in question. For a fault model with SOC (see below), fault patterns of critical states and sub-critical states[11] are distinguishable via symmetropy–sub-critical fault patterns show nearly constant symmetropy values whereas various values are taken on during critical states (Nanjo *et al.* 2005). Work with microfracturing in rock indicates that the process evolves under a constraint of increasing richness in double symmetry (a trend towards low symmetropy indicates that symmetry is building in the system) (Nanjo 2000). Since the general dynamics of complex systems are shared across multiple domains in Nature, it is not unreasonable to expect that symmetropy might exhibit meaningful variations in certain combat systems. It might also serve as a good starting point towards evaluating disorder with respect to other relevant patterns of combat (e.g., a known enemy formation).

2.3 SOC and precursors

The concept of self-organized criticality was introduced to explain the behaviour of systems with a slow storage and a rapid, avalanche-like release of energy, such as earthquakes, forest fires, and especially sand-piles. Sand-piles have become the quintessential prototype of SOC (Bak *et al.* 1988) and so a brief description is in order. As grains of sand are dropped onto the pile, the pile grows and the slope increases. The increasing slope causes some of the sand to roll down due to gravity. The grains of sand falling off the pile generally are not directly related to the grains added.

After a certain slope is achieved, the number of grains falling off is on average the same as the number of added grains. This stationary state is independent of the way the grains are added to the pile or the way the grains fall off. It is a characteristic property of the sand-pile. A sand-pile in this state is a special case of SOC. The pile evolves into this state independently of the driver (in this case, the mechanism of adding the grains). From the point of view of complexity, SOC represents an attractor for the sand-pile, which means that no matter what the initial state was, the system will organize itself in such a fashion that it leads to the attractor.

For a sand-pile at the point of criticality, a single added grain can trigger an avalanche of grains falling off, subsequently decreasing the slope. The dependence of the frequency of incidents on the number of grains falling off at each incident generally obeys a power law. In other words, the frequency of avalanches is higher for small avalanches than for large ones.

[10] The article by Dobias (2008b), in press at the time of writing, also applies the concept of symmetropy to combat. Therein, symmetropy is found to be quite effective in capturing system-wide changes in conflicts exhibiting SOC.
[11] A critical state of an SOC is one that exhibits a scale-free distribution of event sizes, whereas a sub-critical state is one that is not near such a criticality. As an example, SOC sand-pile models evolve through sub-critical states before reaching a critical steady state.

Large-scale near-critical events in many dynamical systems are sometimes preceded by smaller, more frequent events (precursors). In some systems, such precursors have been interpreted as the result of a weakening process with regard to a stress in the system, eventually giving way to a larger-scale release of the same stress. Examples from natural complex systems include tremors (foreshocks) that may precede large earthquakes (Narteau 2007) and pseudo-breakups or field-line resonances preceding the onset of a magnetospheric substorm (Samsom *et al.* 2003, Voronkov *et al.* 2004). Such precursors can facilitate early response to the possibility of a large-scale event in the near-time horizon, although similar small-magnitude events may also occur after a large-scale event (e.g., aftershocks). However, systems in a near-critical state are inherently unpredictable, and as such large-scale events might characteristically initiate without warning (e.g., chaotic seismicity (Narteau 2007)). Precursors can be thought of as probabilistic indicators *when they occur*, and interpretations of precursor characteristics can help to estimate the chance of observing, for instance, a single large-scale event or even a cluster of large-scale events in a short time frame (based on the fractal distribution of events in the system).

The existence of the precursors (pseudo-breakups) in a magnetospheric system was explained by the lack of sufficient free energy in the system (Dobias *et al.* 2006). This is consistent with a sub-critical system near a critical point (discharge event systems). Thus it is reasonable to expect that in some military combat systems, a major system change (phase transition) likewise will be signalled by precursors of a similar nature. Since the identification of precursors varies from one phenomenon to another, it seems reasonable to assume that, in general, the characterization of precursors in conflicts (if present) will depend upon the specific dynamics of the system under scrutiny. There are certainly obvious inferences to be made in many circumstances (e.g., the sudden detection of a concentrated enemy force may lead to a burst of casualties); however, less obvious features might prove equally predictive (e.g., a change in the spatial ordering of an opposing force might signal that current tactics are working or that a modification of tactics is now in order [sc. 'stage 2' or 'plan B']).

As an illustrative example of a possible precursor related to combat, Figure 4 shows the results for the spatial entropy for a crowd-confrontation scenario modeled using MANA. Entropy shows a slight increase and then a dip preceding the main increase due to crowd dispersal (time steps ~ *150-250*). This is consistent with the change in the system's state corresponding to a phase transition.

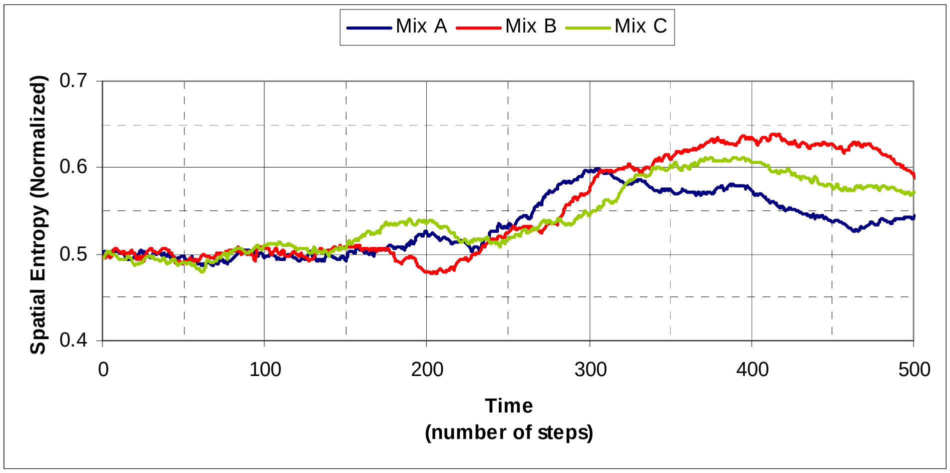

Figure 4: Average spatial entropy for three different BLUE configurations in a crowd confrontation scenario modeled with MANA.

2.4 Evolving behavioural *combat prowess* within an agent-based distillation

2.4.1 Motivation

With enough simulations, the likelihood of the various possible outcomes for a virtual battle becomes apparent. This fact alone can provide valuable insight into the dynamics of a real conflict and strategies needed to overcome difficult challenges. However, it is natural to ask, "What combatant behaviour maximizes an operation's chance of success?"

Given fixed forces, weaponry and equipment, the success probabilities of a particular side in a conflict modeled using an ABD can vary widely depending on the behavioural settings of agents belonging to friendly and enemy forces. Thus it is advantageous to examine exactly how these settings are impacting the probabilities. Fixing behaviour settings leads to situation-dependent patterns of movement and engagement. The collective pattern of behaviour that emerges for a group under fixed settings relates to their *combat prowess*, or in other words their level of superiority and skill on the battlefield. Behavioural combat prowess in this context amounts to manoeuvre tactics and involves different ways of traversing and responding to events in the simulation environment relating to obstacles, allies, opponents, neutrals, and preferred paths (e.g., towards a waypoint). Unless otherwise stated, hereafter the term 'behaviour' refers to this particular variety of combat prowess.

ABDs such as MANA (Lauren *et al.* 2002), 'Irreducible Semi-Autonomous Adaptive Combat' (ISAAC) (Ilachinski 1997), 'Enhanced ISAAC Neural Simulation Toolkit' (EINSTein) (Ilachinski 2000, 2003) and 'Warfare Intelligent System for the Dynamic Optimization of

Missions' (WISDOM) (Yang *et al.* 2004) provide a convenient environment for exploring such lines of interest for several reasons. A few are listed:

- They provide a means to represent a battle scenario, from a whole of system point of view, to a measurable (but not overly-burdensome) degree of realism[12];

- They provide controls to vary the behaviour of combatants through built-in agent parameters covering personality, weaponry and sensor capability;

- Conflict scenarios can be run multiple times and the average effectiveness of various sets of equipment, tactics and behaviours can be measured and compared; and

- In some, the behaviour of friendly and opposition forces can be arranged to automatically adapt, within user-specified ranges, according to an embedded GA, allowing one to explore a large parameter space of behavioural possibilities in the search for an optimal solution.

2.4.2 Genetic algorithms in MANA

In this section GAs in general are reviewed with specific reference to capabilities of the MANA GA.

The MANA GA has been employed by several authors to investigate the effects of behaviour on combat (Luscombe et al. 2005, McIntosh et al. 2006, Parunak et al. 2006). Simply stated, a GA is a method of searching a given parameter-space for the optimal solution of a fitness function. The mechanics are loosely based on the manner in which organisms have evolved as solutions to the problem 'How can a species live and procreate on planet Earth?' The fitness function measures how good a solution is with respect to the problem environment. GAs maintain a population of candidate solutions which evolve over successive generations in response to the fitness function. The fitness assignments permit a competitive ranking of solutions within the population.

Evolution proceeds by first coding 'solutions' (chromosomes) as a set of parameters (genes) and ranges that cover the entire solution space. Over successive generations, a new population of solutions is *bred* from the existing one. Breeding new solutions typically involves three processes or *operators*:

- Selection: solutions are measured for fitness and paired up according to some rule. The rule usually involves a degree of randomness and favours pairing fit solutions together (e.g.., fitness-proportional pairing), often to the exclusion of unfit pairs. These pairs become the 'parents';

- Crossover: 'child' solutions are generated by randomly combining genes of 'parents'. The children represent new, possibly unexplored solutions; and

[12] The "realism" of a scenario is an emergent property of the modeled system. It does not imply that the individual agents behave in the same manner as real soldiers would.

- Mutation: some genes may be altered via a random, generally small probability, change in value.

The idea is that high-fitness parents have the best chance of producing higher-still fitness children. In many implementations, the fittest individuals are carried over to the next generation unaltered to hedge against the destructive nature of crossover and mutation operators. The GA run terminates either when the desired level of fitness is attained or after a specified number of generations have been processed. The ‘solution’ is generally the parameter set within the chromosome of highest fitness in the final generation. The approach can break down in problems where independently good solutions combine in such a way that gains made are repeatedly lost or the optimal solution is sufficiently isolated in the fitness landscape so as to avoid being found. Note that a solution found by a GA is limited by the accessible degrees of freedom - it cannot evolve ‘outside of the box’. Thus the practitioner must be able to determine, minimally, the boundaries of the solution space in addition to how to distinguish a good solution from a poor one at the required granularity.

Since combat involves sources of randomness, the evolution of the population is somewhat complicated by the fact that, in the case of combat simulation, *the fitness function necessarily measures the outcome of a probabilistic chain of events*. It may ‘miss’ an optimal chromosome due to what essentially amounts to ‘bad luck’. For example, a high fitness solution to the problem can actually be discarded if it failed miserably to accomplish the operation set out in the simulation, despite the fact that the chance of failure may have been small. This effect can be buffered somewhat using the ‘Multi-run’ option in the MANA GA. However, doing so can greatly increase the computation time, even for modest settings (e.g., 10 multi-runs translates to 10 times the computational effort). Furthermore, chromosome evolution is still vulnerable to a string of bad luck, so a balance must be struck that depends on the particulars of the situation. Thus methods must include validation of the evolved solution (e.g., via simulations measuring the performance of a single chromosome). Furthermore, it is instructive to pay special attention to ‘spikes’ in the fitness function occurring throughout the various stages of evolution. One should ascertain whether or not such combinations of genes were just lucky or the result of a (possibly lost) highly effective solution. Note that such a practice, however prudent, goes against the above assertion that the highest fitness chromosome in the <u>final</u> generation is the solution. The term *final*, though, can be exploited since it is somewhat arbitrary, loosely conceptualized as the point at which one either is satisfied with the solution or has decided it is not worth pursuing further.

The use of event-driven changes of state (MANA *triggers*) with the GA provides increased flexibility for evolving agent behaviour. Using triggers, one can vary the response of agents to various stages within the conflict operation. For example, one set of behaviour parameters could apply (and evolve) when no opponents are within detector ranges and another set if opponents of a given type have been detected.

2.4.3 Monitoring and measuring GA performance

Monitoring the performance of a GA can direct a run towards faster convergence and avoid unproductive regions of the parameter space. Furthermore, it can be used to help define and refine the quantity and ranges of evolving parameters. The key measure in a GA is fitness. The distribution of fitness within a generation and how that distribution changes from one generation to the next provide indicators of algorithm performance. In many cases, it is also possible to estimate the fitness of the next generation and/or characterize the steady-state limit to a measured accuracy. Also, monitoring the individual progression of the evolving parameters (genes) can be of value, especially during the early phases of problem representation and structural scoping (i.e., defining a minimal list of parameters, anticipating architectural 'building-blocks', or recognizing 'genetic drift' (Rogers *et al.* 1999)[13]).

Although numerous 'involved' methods are available for evaluating and monitoring the performance of GAs (Bornholdt 1998, Goldberg 1989a/b, Holland 1975, Prügel-Bennet *et al.* 1994, 1997, Rogers *et al.* 2006), many potentially useful techniques require exceedingly more information than is readily available for analysis within the MANA GA environment (e.g., Markov chain analysis (Nix *et al.* 1992) requires knowledge of chromosome transition probabilities). Nevertheless, even a small subset of these methods can provide enough insight to infer important characteristics about the progression of GA runs. In the simulations that follow, GA progression is evaluated and monitored by following the genes of the fittest member of the population as the generations proceed. Additionally, mean population fitness is tracked and the effects of varying attribute settings of the genetic operators themselves are addressed, albeit briefly. Factors examined include: fitness criteria (MOEs), gene set, population size, number of repetitions, mutation rate/size and the use of trigger states.

A final point concerns the robustness of a GA solution. A highly optimized solution to a problem may hinge on minor features (or even errors) within the simulation environment, rather than the main drivers. Such a solution runs the risk of performing poorly when the simulation conditions are slightly altered. For example, an optimal tactical solution to a simulated surveillance mission might rely on the exact distribution of enemy forces using detection sensors with precisely defined ranges, such that if either are slightly modified the tactic fails. This is an example of a *brittle* solution. Thus one has to be careful not to 'over-evolve' the population. This typifies a general problem with optimization methods–if the solution is too specific it can't be used for anything else. To counteract this effect, randomization of minor effects and sensitivity testing within realistic ranges can aid the process of robust solution generation.

[13] In a GA, genetic drift refers to the convergence of a gene parameter over many generations to some value in the absence of environmental pressure to do so, or in simpler terms, convergence by pure chance.

3 Method

3.1 Outline

In brief, the method adopted for exploring some potential uses of complexity in combat simulations was as follows:

1. A simple combat scenario was devised consisting of a small BLUE force against a formidable, double-classed RED force. By double-classed we mean that RED is composed of two distinct squads with differing capabilities and tactics;

2. The BLUE force was permitted to evolve its tactics using the MANA GA, first without the use of triggers, to improve the probability of mission success;

3. The scenario was analyzed using the complexity measures of Section 2. These measures were qualitatively used to help describe the combat system dynamics and identify at what points trigger states that partition the scenario might be beneficial. The BLUE force was then permitted to evolve separately within the partitions. Note that the partition points chosen turn out to be intuitively obvious without the use of CMOEs. However, a less obvious, potential transition point was also revealed. Unfortunately, it could not be leveraged without extended simulation capabilities;

4. A new, but related scenario was then devised wherein BLUE agents were endowed with some knowledge of a CMOE relevant to their situation. This new degree of freedom was then exploited to gain further advantage against RED. Note that the CMOE employed did not provide direct knowledge of the global system dynamics, but rather the system dynamics as perceived locally through limited, range-dependent SA. The CMOE employed identified key changes in spatial disorder within RED's pattern of movement.

The basics of the scenario are described in more detail below.

3.2 Scenario

To reiterate, the purpose of the chosen scenario was two-fold—to demonstrate application of the GA in MANA and to explore how knowledge of complexity in combat can be utilized to achieve tactical advantage. To begin with, complexity was ignored and the optimal behaviour was found for a BLUE force pitted against a formidable RED force. Then CMOEs appropriate to the given circumstance were chosen and used to plan and execute a challenging (virtual) mission. Three simulations were conducted: *Sim I, II*, and *III*. In these simulations the focus was on the 'control' aspect of C2 through modified behaviour. The idea was that the influence of agent SA would drive the decision-making about what kind of behaviour to adopt during a given encounter (or situation) of interest.

The overall objective for BLUE was to embark from their starting position 'A' and reach a waypoint 'B' across a billiard table battlefield. In their path were two distinctive RED squads, differing from one another in patterns of movement, sensors and weapon capabilities. It was a tough scenario for BLUE to overcome—conditions were tailored to make their situation extremely difficult, with the hope that behaviour would emerge that would significantly improve chances of success.

In *Sim I*, the optimal BLUE force strategy was found without complexity SA. The scenario was treated as a single obstacle for BLUE to overcome—meaning that only one set of optimal behavioural parameters were sought to deal with the situation as a whole. In subsequent simulations (*Sim II* and *III*), the scenario was subdivided and each subdivision dealt with accordingly. In *Sim II*, the optimal BLUE force strategy for a given encounter was found utilizing MANA triggers (state changes altering behaviour) defined with the aid of CMOEs that monitored the overall system dynamics of *Sim I*. Both *Sim I* and *Sim II* involved multiple simulation runs to evolve combatant behaviour. The population size for the GA was set to 50 in both cases (i.e., 50 randomized, candidate solutions to the problem), and the fitness for each potential solution (chromosome) was determined by averaging over 10 runs using the built-in 'multi-run' feature. Between 100 and 200 generations were evolved, depending on the apparent convergence rate of the case at hand. Note that initial attempts to evolve behaviour with smaller population sizes and multi-run values did not produce stable results. Finally, in *Sim III* the feasibility of real-time response to the complex system dynamics was explored—awareness of a complexity measure was used by the BLUE force in a new, but similar, situation to define state changes *on-the-fly* between human-imposed and GA-evolved behaviour patterns.[14] The CMOE signalled changes in the (sparse) pattern of spatial disorder within RED force opponents detected by BLUE sensors, and this signal was used to switch an indirect fire capability (IDF) of BLUE off at (apparent) key moments.

[14] Note that in *Sim III* only one tactic was permitted per situation – in reality this would be undesirable since it would allow an opponent to learn the tactic and capitalize on it in future encounters. An extension would be to predefine a number of 'good' strategies for a given situation and then pick one unpredictably to execute.

4 Results

4.1 Sim I: Evolving combatant behaviour in a simple 'A to B' scenario

In this simulation, 6 BLUE soldiers depart from point A and make their way through a 12-member RED patrol of comparable (individual-wise) combat power and proceed to attack 6 RED site defenders having double the kill probability (0.2 versus 0.1) and slightly longer range sensors (25 versus 20 distance units). The weapons employed on both sides target individual agents rather than groups, meaning no area weapons. The RED patrol departs from point B at the beginning of the simulation on a heading towards waypoint A. The site defenders remain proximal to position B, moving randomly within a confined area. The setup is displayed in Figure 5. The measure of success (fitness) for the GA is defined as the number of BLUE combatants within the first cluster of agents to reach waypoint B under a time constraint of 500 steps[15]. Baseline MANA settings for this simulation can be found in Section A.1 of the Appendix.

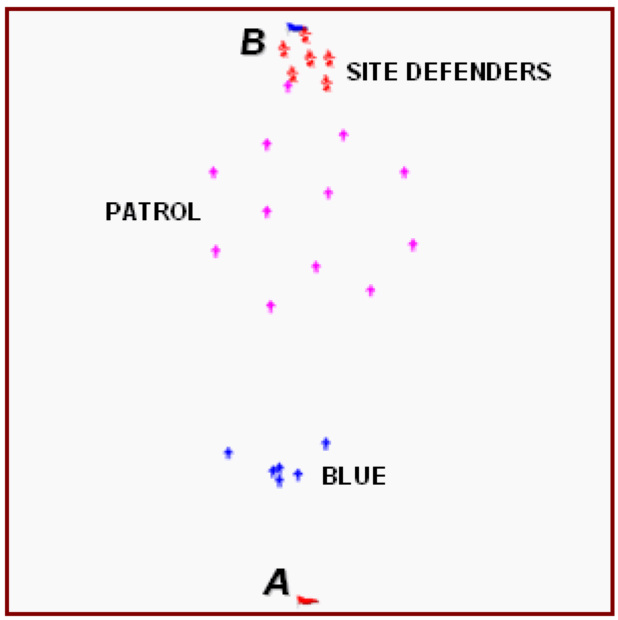

Figure 5: MANA Scenario for Sim I & Sim II.

A single set of personality traits was evolved for the BLUE team to optimize the situation described above to their favour. To keep the options open without overburdening the search algorithm, only a few key degrees of freedom (genes) were selected—neither a minimal set nor

[15] The simulation stop condition was triggered when BLUE reached their goal, and the MOE set in the MANA GA was the number of BLUE agents within a 25 unit radius of the goal, evaluated at the stopping time.

an overly large set (below). Note that positive values in the ranges correspond to attraction and negative values correspond to repulsion.

a) Attraction/repulsion to friends (psFriends, -100 to 100);

b) Attraction/repulsion to enemies detected personally (psEnemies, -100 to 100);

c) Attraction/repulsion to enemies detected by others [SA] (psOrgThreat3, -100 to 100);

d) Attraction to waypoint B (psNextFlag, 50 to 100).

Simulations indicated that the evolved behaviours had at most an 18% success rate (1000 repetitions, 1% standard error). The most common and significant beneficial trait was high attraction to friends (clustering). This allowed the BLUE force to concentrate firepower helping them get past the dispersed RED patrol. BLUE combatants remaining after this encounter proceeded with a less-than-fair chance to attack the RED defenders at waypoint B using the same strategy. Note that if BLUE somehow completely avoided the RED patrol, 100 simulations suggest only a 54% chance (5% standard error) of defeating the RED site defender squad using optimized tactics, which represents an approximate theoretical upper limit in this scenario.

From the perspective of the BLUE force, a typical successful mission roughly followed the timeline below:

1. First time step (1): BLUE departs from point A on a heading towards waypoint B;

2. 70-75 time steps: BLUE encounters the RED patrol;

3. 120-125 time steps: BLUE passes the RED patrol;

4. 135-145 time steps: BLUE encounters the RED site defenders;

5. 180-190 time steps: BLUE reaches waypoint B.

The process of arriving at relatively stable GA results within MANA required some exploration in and of itself. Main lessons learned from this simulation are summarized:

- The multi-runs option should be used when running a GA to buffer against the effects of randomness in the outcomes of a conflict (we used at least 10);

- The final solution provided by the GA should be heavily validated through repeated simulation. Furthermore, it is prudent to compare the performance of the solution with other solutions that performed extremely well in previous runs, and with solutions obtained using different GA settings. Lastly, it may be instructive to compare the results with ‘best guess’ solutions formed by the practitioner;

- Testing for genetic drift (see Section 2.4.3) and evolving the system using a high mutation rate may help to eliminate extraneous variables (genes), thus improving performance and simplifying the interpretation.

Three main lines of GA settings were pursued to find the optimal solution. The best performing solution, judged via validation of the highest performer in the population of the final generation, was achieved using crossover with a high mutation rate (referred to as the 'HM series'). This solution clearly outperformed a crossover-only prospect (dubbed the 'C series') and marginally surpassed cross-over with a moderate mutation rate (the 'CM series'). Success rates and settings for various GA options appear in Table 1: 1) the *HM series* is one of high mutation [rate: 50%, strength: 20%], 2) the *CM series* balances crossover and mutation [rate: 2%, strength: 20%], 3) the *C series* uses crossover only [mutation rate is set to zero], and 4) the *Default Settings* refers to a baseline, non-evolved 'solution'. Additional details are provided in the next section, where *Sim I* and *II* GA results are compared head-to-head in tabular form.

It is of interest to observe the evolution of solution fitness and gene convergence as the generations progress. Recall that the fitness of a solution during a GA run is based on 10 repetitions, whereas validation runs of a final solution are based on 1000 repetitions. Referring to Figure 6, it can be inferred that although the cross-over solution (C series) scored lowest of the optimal solutions in the validation runs, the C series population converged the fastest in mean fitness. By 100 generations, the C series population was completely homogeneous and since mutation was forbidden, further evolution was blocked. The CM series, on the other hand, overall displayed steady gains and losses. On average, the CM series performed at peak in the range of 140-180 generations, dipping slightly near the final generation (200). In contrast, the HM series performed rather consistently, on average, from the onset. Maximum population fitness shot upwards in the final generations, approximately matching the final stage fitness of the CM series. In the HM series, highly varying parameter values for genes other than psFriends (i.e., non-convergence), evaluated for the fittest member of the population within each generation, suggest that this trait is of primary importance. The phenomenon was not observed in the CM or C series.

Table 1: Validation of Sim I optimized solutions for BLUE behaviour.

Solution	Success Rate	psFriends	psEnemies	psOrgThreat3	psNextFlag
GA, HM Series	18%	100	60	20	70
GA, CM Series	17%	60	60	-20	68
GA, C Series	15%	71	14	84	77
Default Settings	3%	0	0	0	50

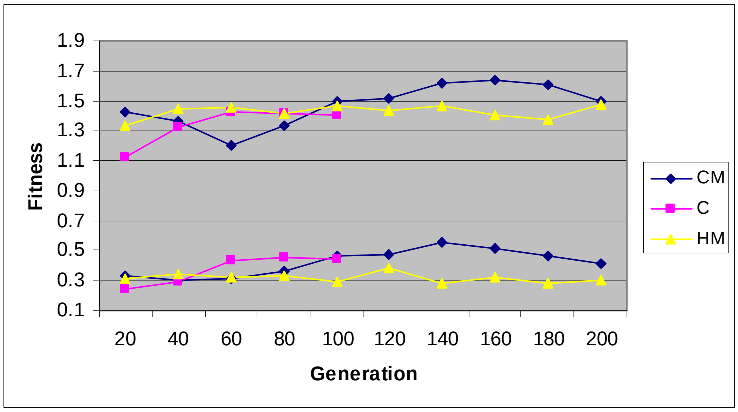

Figure 6: Maximum Fitness (upper) and Mean Fitness (lower) Evolution. Each node represents an average of the 20 preceding generations.

4.2 Sim II: Evolving combatant behaviour in stages

It is apparent that some benefit could be gained by partitioning *Sim I* into different stages. It is of interest to explore how the partitioning relates to the CMOEs of Section 2.1. To achieve as complete a picture as possible, complexity is viewed from various perspectives and scales. Note that in the figures that follow, RED forces are combined unless otherwise stated. To begin with, the fractal dimension is plotted for BLUE and RED forces at two different length scales: 1) the entire battlefield (200x200) and 2) minimal containment (see Figure 7). The latter restricts the evaluation space to a minimal, X and Y axes-oriented bounding box surrounding the squad of interest. The center of the box is the centroid of the squad and the box is always square. This box moves and resizes over time as the agents redistribute themselves spatially or are eliminated via attrition. In other words, it is the smallest axis-oriented square box that contains the subject agents at a given time step in plan view. The motivation for viewing the system from a local perspective has to do with the next section (*Sim III*) where the CAS is to be evaluated by agents based on limited (local) SA. As might be expected, the local perspective on the fractal dimension turns out not to be particularly insightful in this scenario. A similar conclusion is drawn for the spatial entropy. Symmetropy, on the other hand, does show some interesting behaviour at this scale. Long term correlations (recall the Hurst coefficient and SSP from Section 2.1) also display interesting behaviour, but are not evaluated for minimal containment since the spatial scale is not of concern for these measures.

The fractal dimension plots in Figure 7 are time step-averaged over many simulations--168 for the case of BLUE success and 832 for BLUE failure (1000 total simulations conducted). The same can be stated for all Figures comparing the two cases. At battlefield scale (Figure 7a), the RED force starts off tightly clustered (dot-like) as evidenced by the low fractal dimension. The dimension then increases as they spread out (near line-like (D_F=*1*)) and then decreases for two main reasons: 1) RED attrition, 2) RED patrols' arrival at their goal. Also at this scale, it can be seen that the BLUE forces' fractal dimension remains small (dot-like) as BLUE agents maintain close proximity to one another. The branching of the fractal dimension for the two different cases (BLUE success versus BLUE failure) can be seen reasonably clearly. In Figure 7b, the minimal containment results show the branching as more pronounced. For BLUE it occurs early on at ~ 70 time steps, whereas for RED ~ 150 time steps. BLUE no longer starts off with a small fractal dimension—the dimension now reflects the distribution of BLUE agents within the 'dot' as viewed from battlefield scale. Note that the graphs suggest that BLUE and RED are both more successful when they are able to maintain a higher fractal dimension (an explanation of exactly why follows, below).

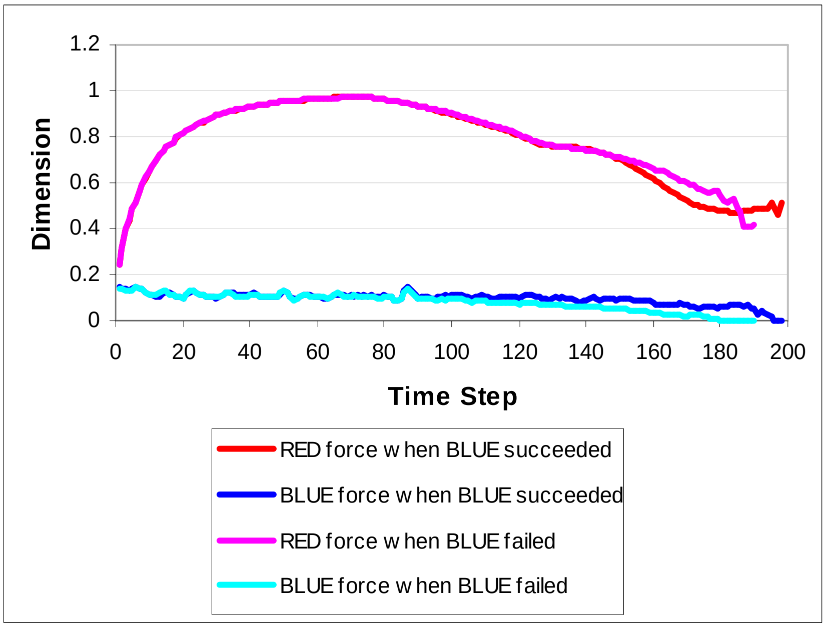

a)

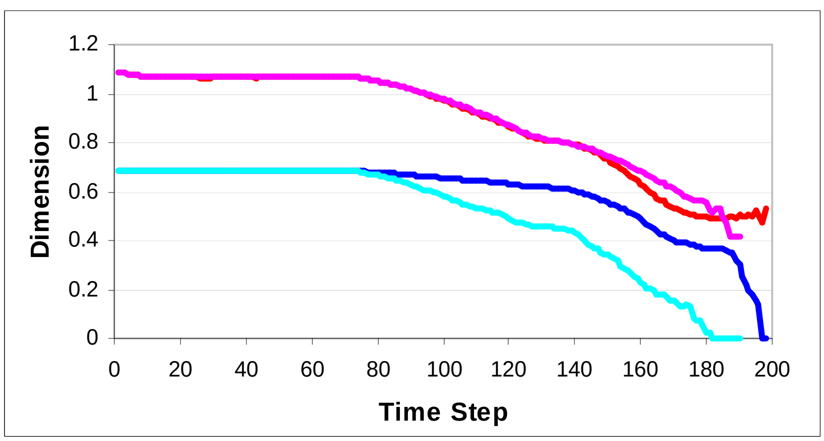

b)

Figure 7: Fractal Dimension for a) entire battlefield and b) minimal containment.

Figure 8a shows the spatial entropy for the fixed battlefield at 25 partitions per edge. The RED force spatial entropy plots are nearly identical in form to the corresponding fractal dimension plots of Figure 7a, and the BLUE force plots are similar, except that variations are more pronounced in the spatial entropy plot. Taking averages of 5-25 partitions, a process somewhat reminiscent of the fractal dimension computed via box counting, we see that Figures 8b and 8c are of nearly identical form to the fractal dimension plots of Figures 7a and 7b respectively.

Figures 7 and 8 suggest that something important (spatially?) is happening to BLUE around time step 70, and that something important happens to RED around time step 150. The underlying patterns in the spatial entropy and fractal dimension plots relate to a simple quantity—the mean (normalized) combat strength (CS) of the RED and BLUE forces evaluated at a given time step, defined as:

$$CS_i(t) = {(N_{0i} - C_i(t))}\big/{N_{0i}}, \; i = RED, BLUE\,.$$

N_{0i} is the original size of the i^{th} force, and $C_i(t)$ is the number of casualties at time step t. Figure 9 suggests that the 'important' events merely correspond to the timing of significant changes in the combat strength of the various forces (and attrition rate, of course); the fractal dimension and spatial entropy are more-or-less counting the surviving number of agents for each force following an encounter. This is especially evident in the minimal containment case (compare Figure 9 with 7b and 8c). So, although the fractal dimension and spatial entropy point out when the dynamics make critical transitions (branch), neither provides any insightful information beyond what is derivable from a simpler measure in this scenario. Likely contributing factors include the sparse number of combatants, the duration of the encounters (short 'bursts' of activity), and the evolved behaviour from *Sim I*, that is, the tendency for BLUE to tightly cluster. To shed some light on the complex nature of the dynamics, the remaining CMOEs from Section 2 are now explored.

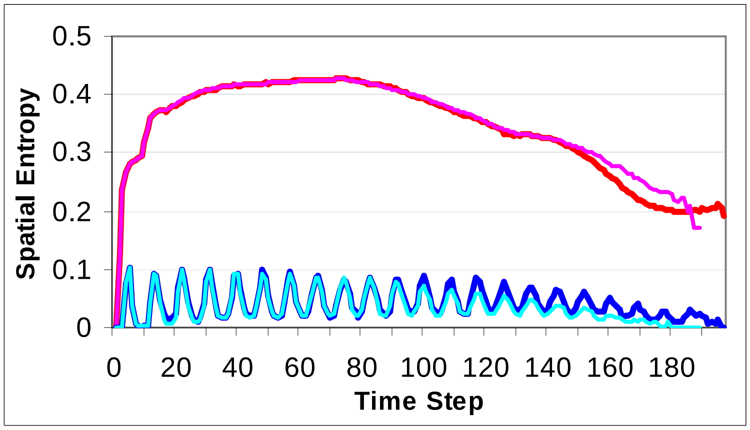

a)

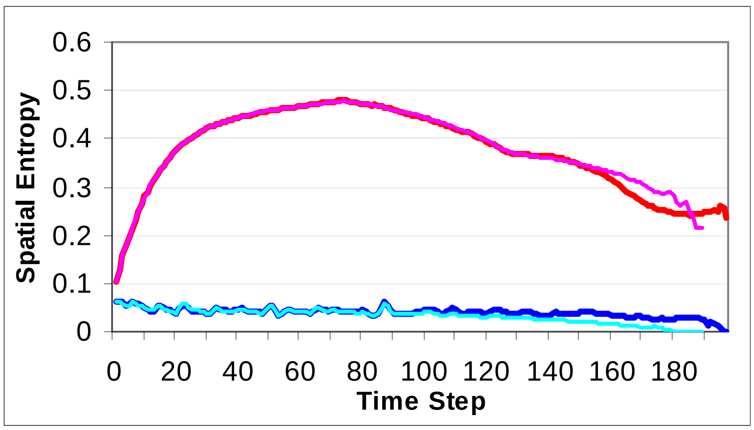

b)

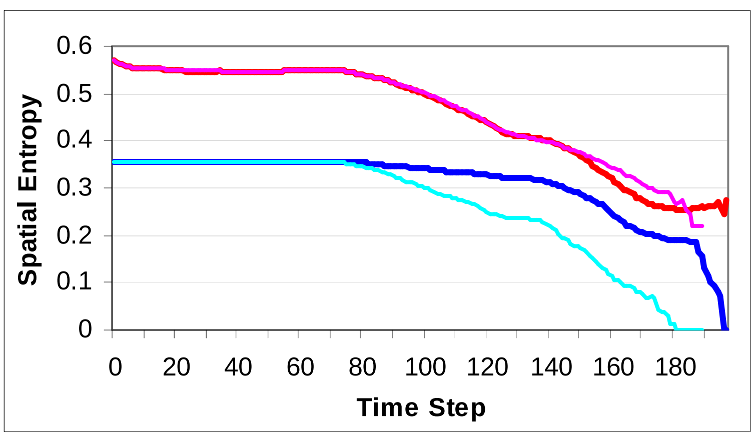

c)

Figure 8: Spatial Entropy for a) fixed battlefield (25 partitions), b) fixed battlefield (av. 5-25 partitions), and c) minimal containment (av. 5-25 partitions). (Legend of Figure 7 applies)

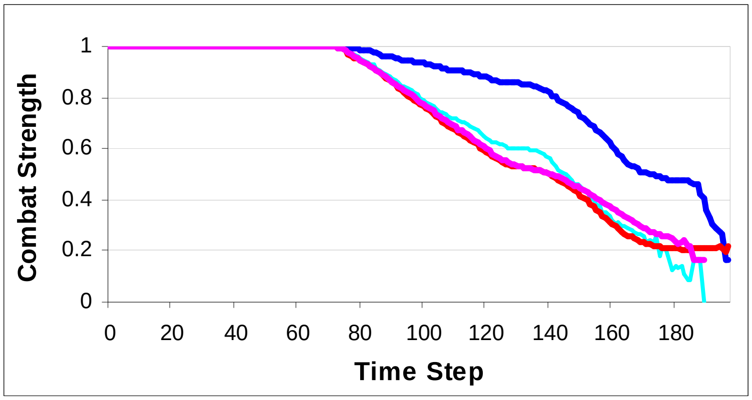

Figure 9: Combat Strength (Legend of Figure 7 applies)

Recall that symmetropy is a measure taking into account entropy with respect to a kernel of spatial symmetries. Symmetropy plots for the BLUE and RED forces are shown in Figure 10. All of the symmetropy data were computed using minimal containment.

In Figure 10a, the BLUE symmetropy paths diverge around time step 100 (in the midst of the battle with the RED patrol). At this point the successful branch maintains its course while the unsuccessful branch falls below it slightly. The branching may reflect the internal reorganization of the BLUE squad as it attempts to compensate for reduced numbers in light of the preferred formation tendencies (i.e., attraction towards friends, repulsion from enemies, etc.). Later, near time step 153 the successful curve makes a downturn, crossing the unsuccessful curve at ~160. Thus, for whatever reason, BLUE success favours steadiness with regard to the distribution of symmetries during this interval. Afterwards, the successful branch proceeds downwards while the unsuccessful branch shoots upwards. Note that the climb in BLUE symmetropy after time step ~160 (failure case) seems to be weighted by the timing of the elimination of the BLUE force[16]. On average, few members of the BLUE force remain in the contributing simulations by this time step, resulting in a decrease in the dominant symmetry (double symmetry, see Figure 11a). This decrease in double symmetry (*Pd* in the figure) is accompanied by a convergence of the other symmetries (*Pv*, *Ph*, and *Pc*) to a value near 0.25 (i.e., all four approximately equal indicating a lack of preferred symmetry). In the case of BLUE success, double symmetry is more consistently maintained (Figure 11b).

For RED symmetropy (Figure 10b), a small branching effect is evident near 130 time steps, while the main branching occurs near time step 150. The pattern of branching is similar to that of BLUE: when RED is more successful, its corresponding symmetropy curve first rises above the unsuccessful curve, and then falls underneath.

Figure 12 shows the symmetropy of the combined forces. The curves for BLUE success and failure branch near time step 110. The BLUE success branch rises above the failure branch and remains as such until the run terminates. The timing correlates well with the initial branching of

[16] Note that the rapid, upward shoot of the BLUE success curve in Figure 10a around time step 190 is an artefact of having very few simulations lasting that long—those that did had few remaining agents.

the BLUE symmetropy curve, suggesting once again that something pivotal happens in the vicinity 100-110 with regard to BLUE's spatial pattern, corresponding in this case either to a persistence of symmetry (success) or lack thereof (failure) (e.g., as eluded to above, perhaps by this time it can be determined whether BLUE had an overall good or bad encounter with the RED patrol). Next, long-term correlations are addressed.

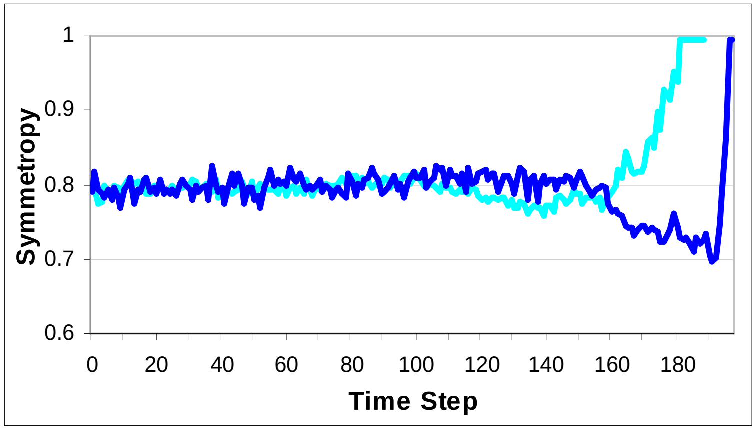

a)

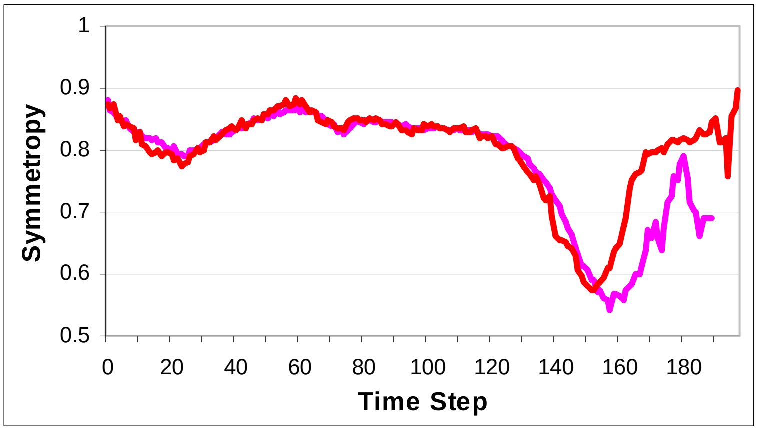

b)

Figure 10: Symmetropy (q=2) for a) BLUE force and b) RED force.

(Legend of Figure 7 applies)

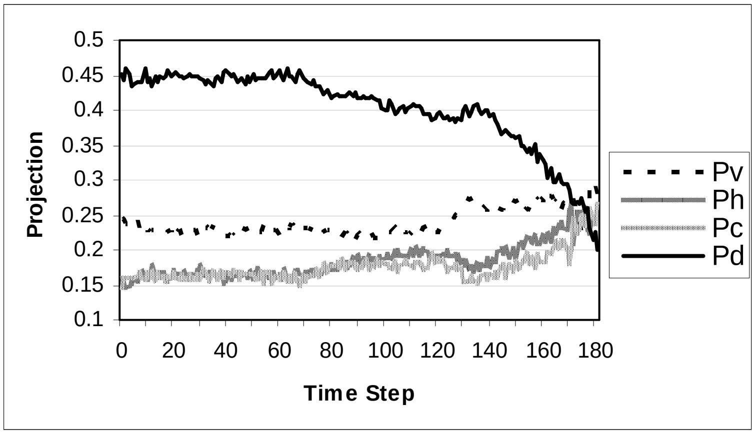

a)

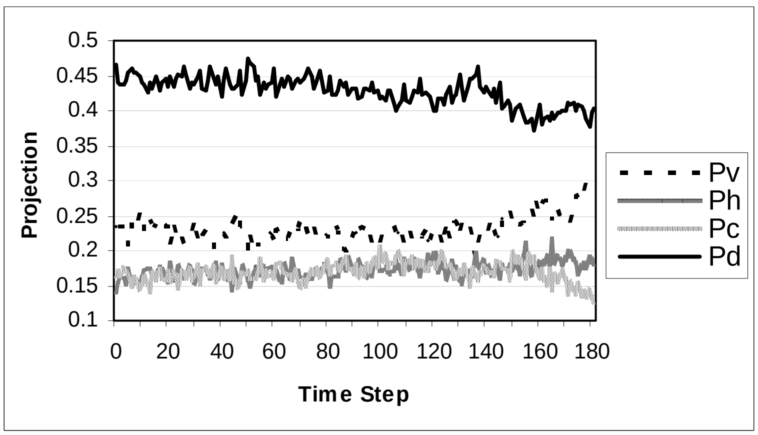

b)

Figure 11: BLUE Force symmetry projections: a) BLUE failure and b) BLUE success.

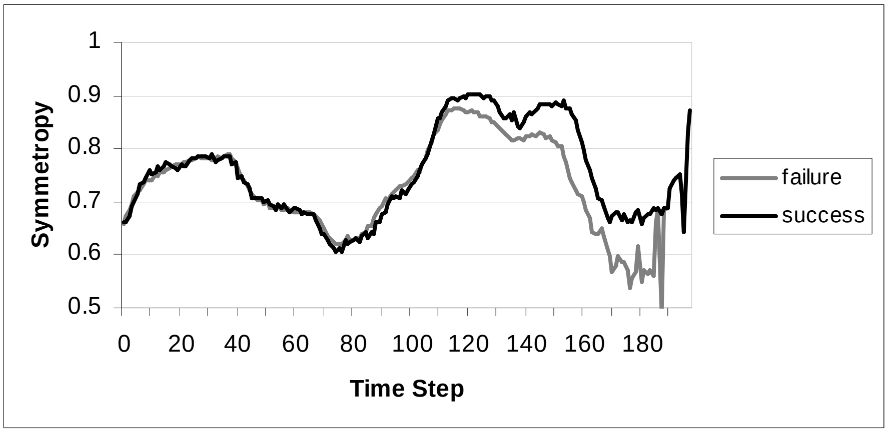

Figure 12: Combined forces symmetropy

Figure 13 shows the SSPs for BLUE force ‘X’ (lateral) and ‘Y’ (longitudinal–towards the goal) coordinates under the cases of BLUE success and BLUE failure. The coordinates represent components of the vector directions of agents at each time step, and the SSP measures correlations in movement at a given time step over a number of simulations. For ease of viewing, each point on the various SSP plots represents an average of the ten preceding data series points.

The first noteworthy feature is that the plots for SSP in the X and Y directions are quite different. This could be an indication of self-affinity (see Section 2.1.6). In the Y direction (Figure 13b), the crucial time step occurs around step 78–soon after the RED patrol is encountered. When BLUE is successful, they are able to maintain a higher level of persistence in motion towards the waypoint B. BLUE failure coincides with motion at this juncture tending toward randomness or even anti-correlation. Just prior to this branching point, persistence in motion rises for both curves. This could be due to BLUE having detecting RED. For most BLUE agents, initial knowledge of RED would be acquired through squad SA, giving them moments to react before the encounter (hence the sudden persistence). In successful runs, BLUE loses on average 0.91 ± 0.07 agents during this encounter and in unsuccessful runs 2.45 ± .05 agents, evaluated at time step 135 (standard errors are given). The loss of too many BLUE agents might account for the disruption in the persistent motion.

In the X direction (Figure 13a), time step 135 is where the branching occurs—again, immediately before an encounter with RED (this time, the site defender force). In opposition to Y, the better path for BLUE is one of anti-correlated motion. Together, the two seem to suggest that a higher degree of self-affinity might be beneficial (i.e., the ‘fractal’ scales differently in X and Y). Another possibility is that when BLUE fails they are simply forced towards non-self-affinity (i.e., they have no real choice in the matter).

Figure 14 shows the Hurst coefficient for the Y direction. Notice that it follows a nearly identical pattern to that of the corresponding self-similarity plot (Figure 13b). The same is true of the X direction (not shown).

Figure 15 shows the SSP for the RED site defender squad. This squad randomly moves about a small area proximal to waypoint B until encountering BLUE agents. It is interesting to note that X and Y correlations are nearly identical to one another, suggesting non-self-affinity. The branching of the curves in the time interval of 110-130 steps suggests a potential non-locality in the movement data. This feature is somewhat interesting since it does not specifically relate to attrition (i.e., the combat strength plot of Figure 9). Although the timing is roughly coincident with symmetropy bifurcation points for RED and BLUE forces (above), evidence suggests that this is coincidental since BLUE is well outside of sensor range of the RED site defenders until, on average, ~ 136 steps in both success and failure cases. Moreover, the earliest such detection time by RED site defenders (recall they have superior sensors) in all 1000 simulations occurs at time step 114, which is beyond the first drop in the SSP after the branching point (the SSP for time step 110 averages the values for time steps 101 to 110). Thus, BLUE is not aware of the RED site defenders and since there is no SA exchange between the two RED squads, the RED site defender squad is not aware of BLUE at this time. Thus the RED site defenders cannot be *reacting* to BLUE's close proximity. Rather, BLUE success seems to *select* a particular configuration as being the more favourable one. The initial movement pattern of the RED site defenders as BLUE is proximal and approaching must have an influence on the outcome of the encounter. This can be likened to catching the RED site defenders 'off-guard'. In an average sense, the drop in SSP could be interpreted as relating to a precursor as discussed in Section 2.3.

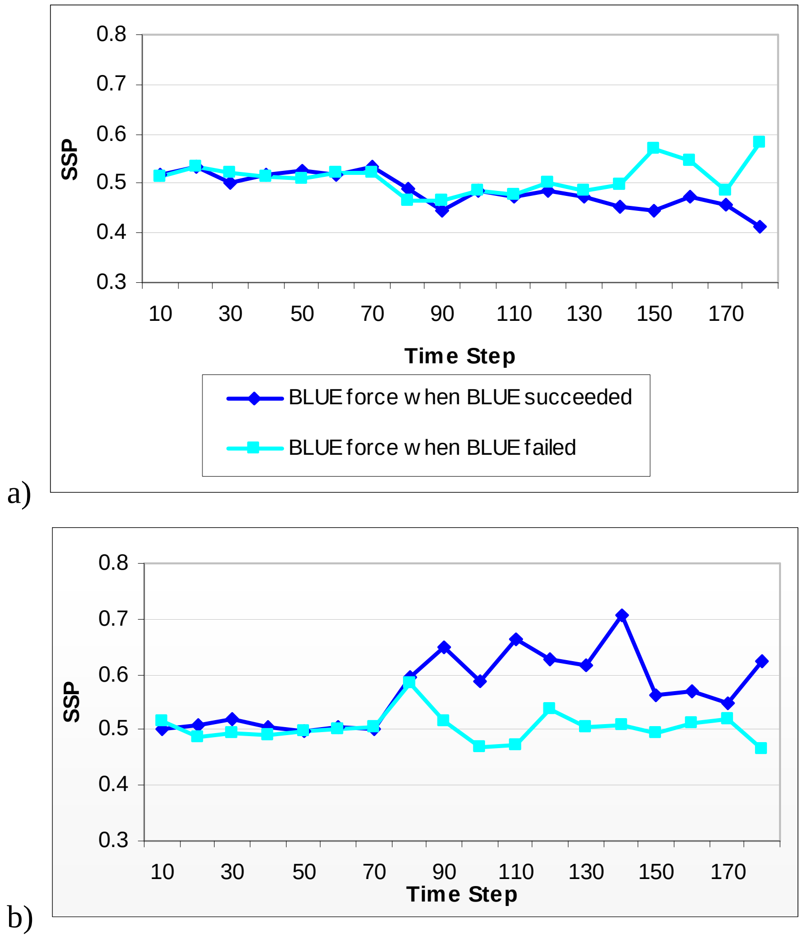

Figure 13: Self-similarity parameters for the BLUE force a) 'X' velocity component and b) 'Y' velocity component.

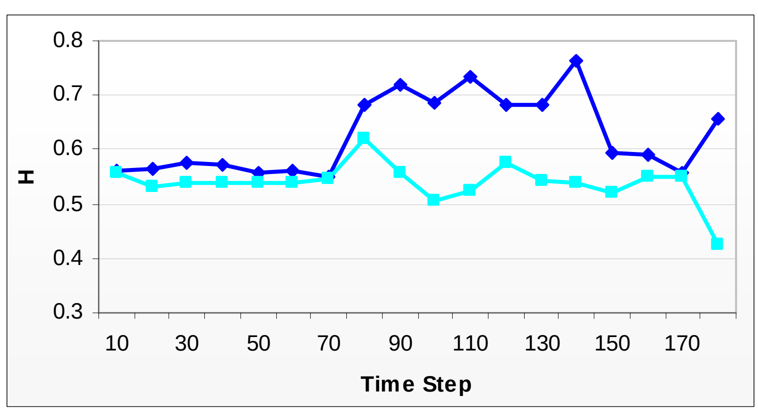

Figure 14: Hurst Coefficient for BLUE 'Y' velocity component.

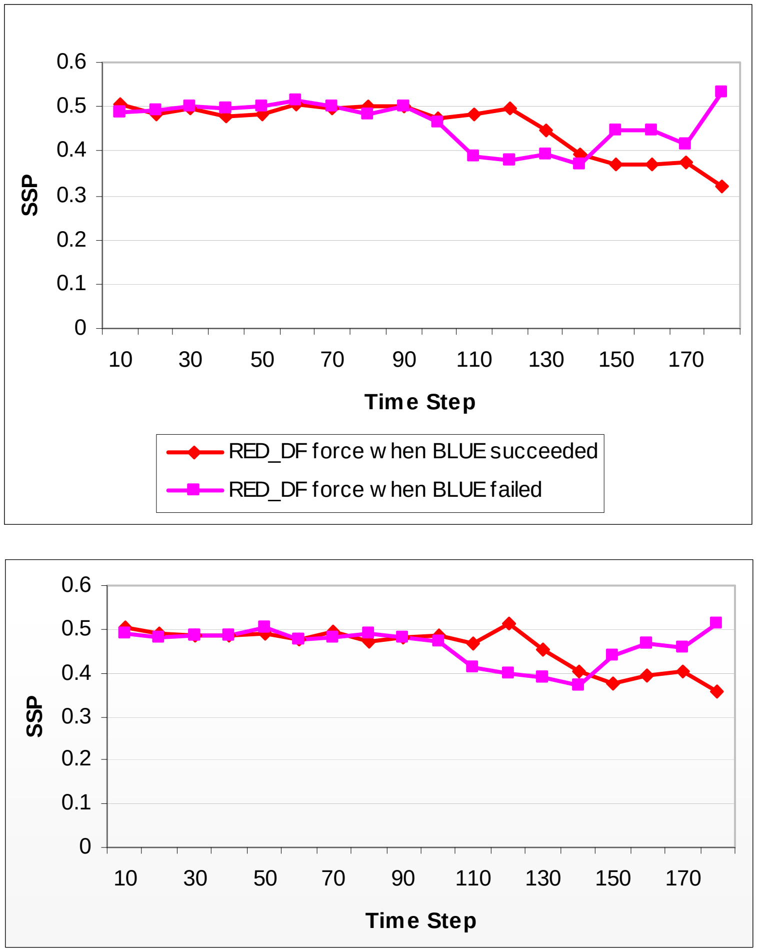

a)

b)

Figure 15: Self-Similarity Parameters for the 'RED site defender' Squad a) 'X' component and b) 'Y' component. (Legend of Figure 7 applies)

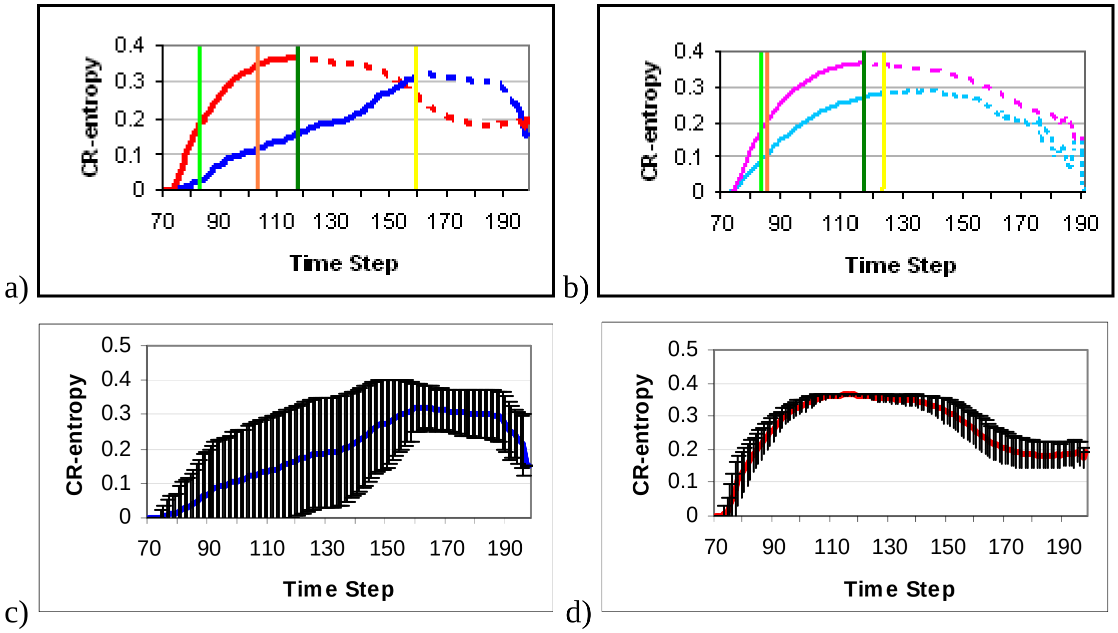

Figure 16: CR-entropies are shown for the two forces when a) BLUE succeeded and b) BLUE failed. The spread (standard deviation) is shown for c) BLUE and d) RED forces in the case where BLUE succeeded.(Legend of Figure 7 applies)

Finally, the CR-entropy is evaluated beginning at the first sign of attrition (around time step 70, see Figure 16). In many ways, the indications of CR-entropy at particular time steps seem to underlie the patterns observed in other measures. In Figures 16a and 16b, the BLUE and RED force curves are divided into intervals corresponding to CR-entropy states (recall Section 2.1.1). Prior to the light green marker (~ time step 82), attrition is low and the relative state is balanced. After the light green marker, and prior to the orange marker, BLUE has the advantage. Note that when BLUE is successful (Figure 16a), the advantage is maintained for around 20 time steps, whereas when BLUE fails (Figure 16b) the advantage is quickly lost. After the orange marker and before the dark green one, attrition is high but comparable, so the state is somewhat balanced. Surprisingly, this state ends at the same time step for both cases (~ time step 115), transitioning to one of BLUE advantage. This advantage is due to the disintegration of the RED force (the 0.37 threshold has been passed). When BLUE fails (Figure 16b), the advantage is once again short-lived and transition to the BLUE disintegration phase follows (~ time step 125), whereas under BLUE success (Figure 16a) the transition is delayed until approximately time step 160.

Note also the standard deviation (or *spread*) of the CR-entropy data as shown in Figures 16c and 16d for the case of BLUE success. BLUE is highly variable throughout, whereas RED is tightly controlled, reducing to almost nil near the onset of REDs' disintegration phase. The same can be stated for the BLUE failure case (not shown). The high variance in BLUE CR-entropy greatly blurs the above interpretation of the actual path followed by BLUE when successful or not, suggesting that in many simulations BLUE circumstances fell somewhere in the midst of the two paths.

It can be argued that the partitioning of the combat scenario is loosely conceivable as a response to a kind of phase transition inherent in the system dynamics. Indeed, the various CMOEs suggest that approximately 10-time step neighbourhoods around times 80, 110, and 150 represent critical and distinct dynamical events in the system. Since there are no available means to respond to the events near time step 110, we group the first two of these together and simply allow the BLUE agents' behaviour to be partitioned by encounter type (RED patrol or site defenders), rather than any particular timing. However, it would be interesting to determine if a third behaviour, defined between say time steps 100 and 130, holds any benefit if the BLUE agents are given extended SA. Specifically, a behaviour that enabled BLUE agents to monitor the site defender force and wait for the right moment to advance seems appropriate. The 'right moment' would have to be determined by analysing the RED site defenders movements in respect of BLUEs' advance vector, position, and the location of waypoint B. This avenue was not pursued.

We proceed by partitioning the scenario according to the intuitively obvious transition events (i.e., by encounter type) which are furthermore supported by the temporal dependence of the various CMOEs. Partitioning was implemented via MANA triggers. To define the triggers, the RED constituent forces were assigned different threat levels so that BLUE could respond to each one differently. This allowed for the evolution of two behaviour profiles appropriate for dealing with each encounter separately (note that detection, rather than hard-coding, of the transition point is a subject of the next section). For the sake of comparison, it is of interest to put forth initial *guesses* of what the optimal GA behaviour settings might be for the BLUE force. This allows one to compare solutions and determine the added value of employing a GA for this task, as opposed to strictly relying on the intuition of the practitioner. Both guesses that were made involve a fixed full attraction to squad members (+100) and moderate attraction to the waypoint B for the patrol encounter (+50). Settings not mentioned are defaulted as in *Sim I*.

<u>Guess 1</u>: Avoid contact with the patrol and then proceed directly to waypoint B

- When confronting RED patrol: Full repulsion to all enemies (-100)
- When confronting RED site defenders: Full attraction to waypoint (+100)

<u>Guess 2</u>: Punch through the patrol and redirect slightly away from site defenders.

- When confronting RED patrol: Default settings.
- When confronting RED site defenders:
- Strong attraction to waypoint (+75)
- Partial repulsion to all enemies (-50)

Guesses 1 and 2 yielded marginal gains for success rates (recall 18% from *Sim I*), given by 24% and 21% respectively (1000 runs – standard error reported as 1% by MANA).

Next, we employ the MANA GA to evolve optimal behavioural settings for the BLUE agents. Personality settings for BLUE while running the GA are provided in Section A.2 of the

Appendix. The GA settled on the following optimal settings, given a population size of 50 with 10 multi-runs per chromosome, mutation rate 2% and strength 20%:

> GA result: Avoid RED patrol contacts detected by other squad members (through SA), but proceed as normal when the detection is personal. Furthermore, rush RED site defenders detected by others, but run away from those detected personally.

a) When confronting RED patrol:

- Full repulsion to enemies detected by others (-100)
- Indifference to those detected personally (0).

b) When confronting RED site defenders:

- Full attraction to waypoint (+100)
- Full attraction to enemies detected by others (+100)
- Full repulsion to enemies detected personally (-100).

Validation revealed a success rate of 26% (1000 runs–standard error reported as 1% by MANA), improving over the solutions without triggers and the guesses with triggers. Although the attrition rate was not part of the fitness function, it is interesting to note that this solution displayed the lowest average casualties for BLUE and the highest for RED (see Table 1 for comparisons). In Table 2, the first three solutions are from *Sim II* and the remaining from *Sim I*. Note that only the success rates directly contributed to the fitness function (MOE). The RED force was the same in all instances. Error ranges shown are those reported by MANA.

The GA 2-trigger solution itself was somewhat surprising—characterized by major differences, even complete polarity, between personal versus squad (SA) detections of enemies. Full attraction to the waypoint when up against the RED site defenders was not surprising (see *Guess 1*). Success rates for various GA settings applied in *Sim I* (without triggers) are also provided in Table 2 for reference. For details concerning the HM, CM, and C series solutions, see Table 1.

Analysis of the gene evolution under a high mutation rate in *Sim II* did not reveal any definitive convergence patterns. On its own, this could indicate that either a rather delicate balance of parameters is necessary (i.e., mutation keeps destroying convergence) or that blind luck dominates (i.e., the settings don't really matter much). Relatively high success rates in the validation runs seem to confirm the former. Also, fitness maximums and population means are significantly higher here than those found in *Sim 1*, beginning early in the run. This suggests in and of itself that the two-trigger approach is superior to the single state approach of *Sim I*, as expected.

Table 2: A comparison of various solutions for BLUE behaviour.

Solution	Success Rate	BLUE Casualties	RED Casualties	Mean Time Steps	Validation Runs
GA, 2 triggers	26%	5.22 ± 0.05	10.78 ± 0.10	319 ± 2.0	1000
Guess 1, 2 triggers	24%	5.26 ± 0.05	10.45 ± 0.10	341.2 ± 2.2	1000
Guess 2, 2 triggers	21%	5.41 ± 0.04	10.64 ± 0.10	211.9 ± 1.1	1000
GA, HM Series	18%	5.46 ± 0.04	10.50 ± 0.10	184.6 ± 0.9	1000
GA, CM Series	17%	5.52 ± 0.04	10.18 ± 0.10	155.7 ± 0.7	1000
GA, C Series	15%	5.57 ± 0.04	10.18 ± 0.10	160.5 ± 0.8	1000
Default Settings	3%	5.94 ± 0.01	8.44 ± 0.11	132.2 ± 0.6	1000

4.3 Sim III: Real-time response to complexity factors for tactical advantage

This simulation illustrates how knowledge of combat complexity can be characterized in real-time and how it may lead to tactical advantage within a conceptually simple combat situation exhibiting fractal properties. Various C2 options were exercised by monitoring and responding to the temporal evolution of a CMOE defined above (Section 2).

In *Sim I & II* mission success was improved upon through use of the MANA GA capability. The behaviours so developed can be applied to larger simulations involving encounters with RED forces of a similar make-up with a reasonable chance of success under the right conditions of use. In the above simulations, the information about *which* element of RED was encountered was hard-coded into the trigger definitions, rather than inferred from RED's spatial dynamics or attrition entropy. Thus, the problem to address next is how to use real-time, localized CMOEs to quickly identify an encounter type (e.g., patrol or site defenders) so as to trigger the appropriate response (i.e., the appropriate behaviour profile). The ideal situation would be to find a 'precursor' to correctly identify the nature of the next encounter (see Section 2.3). This possibility is discussed below.

At first glance, the arguments used for partitioning *Sim II* do not seem to hold much practical value for real-time response. Upwards of one thousand simulations were needed to identify significant patterns in the CMOEs in relation to important events. In general, high variance in the value of the measures preclude their use as a basis for reliable forecasting in real-time for a single run involving small forces—the precognitive signatures sought are definitely not evident in the averaged results.

On the other hand, computing the fractal dimension and symmetropy of RED based on *limited range* detections by BLUE could conceivably produce distinguishing features for the different encounter types. This is akin to detecting a change in the pattern of spatial disorder within RED to signal a state change. However, it is important to consider that SA would be limited to a few detections before a course of action must be decided upon to qualify, intuitively, as a precursor event. Accordingly, given that sparse data are expected, coupled with the fact that the fractal dimension is more suited to characterizing data clustering, it seems inappropriate to rely on the fractal dimension of detections of RED by BLUE in this case. Furthermore, it was demonstrated in *Sim II* that the fractal dimension was more or less just counting the surviving RED agents, which is of no utility here. Analysis of the spatial entropy for RED and BLUE would certainly lead to a similar conclusion. Given an extended SA for BLUE, monitoring the SSP or Hurst coefficient for RED detections could possibly reveal the identity of the type of force about to be encountered given that movement patterns have been pre-established through simulation or otherwise, especially since one of the RED components tends towards stationarity (site defenders clustering around the BLUE waypoint). Nonetheless, as stated previously, computing the SSP or Hurst coefficient is data intensive and the real-time scenario is not likely to be capable of producing the required support data (several hundreds to several thousands of data points).

Therefore, since symmetropy alone is not overly constrained in the case of sparse data, it is the only measure investigated as a prospective CMOE for real-time determination of the encounter-type in this situation. Like the Hurst coefficient, it also requires at least a slight SA advantage to be particularly useful. A symmetropy signature would combine RED force spatial patterns with their degree of disorder. The signatures would have to be established before the operation through simulation or otherwise. The symmetropy patterns of detection preceding an encounter should provide a reasonably accurate cue about what to expect. The feasibility of real-time response to the complexity in the system based on recognizing symmetropy as a precursor is explored below[17]. Note that the term 'precursor' is greatly abused here – there is no real evidence to support that it has anything to do with a phase transition or SOC in this scenario. Nevertheless, the intuitive meaning fits well and so the mistreatment is overlooked.

To begin with, a new, but similar challenge for BLUE is designed (see Figure 17 for a visual representation, agent settings are in Appendix A.3). In this simulation we alter the above scenario somewhat, but not so much that we cannot draw upon the results of *Sim II*. In the new scenario, two 6-member BLUE patrols (A_1 and A_2) are 'searching' for waypoint B occupied by 10 RED site defenders. The heading is actually predefined (marked by grey lines in Figure 17), but the capability of the agents to recognize the waypoint for what it is relies on the proper interpretation of a CMOE. To get to the waypoint, the BLUE squads expect multiple encounters with a RED force similar in function to the RED patrol above (however, in this case, it is more of an occupying force than a patrol and is dispersed throughout the battlefield). RED patrol members are to be identified and eliminated by the indirect fire capability (IDF) available to BLUE. When near the waypoint, BLUE anticipates that they will face RED site defenders as defined above. IDF is <u>not</u> to be used at this stage–they must fight their way in (e.g., to protect against accidental targeting of civilians in a hostage situation), so the mechanism should not fire. To accommodate the IDF support, BLUE is given a slightly longer sensor range than RED (50 versus 40 units), and IDF is connected to the squad SA. Therefore, IDF has to quickly classify an encounter as a PATROL or a SITE based on the available CMOE data. When BLUE reads RED contacts, local SA information is passed to IDF, which determines if it should fire on RED or not. The determination is based on the encounter-type signature recognition from 'precursory' measures. These reference signatures are predefined using pre-existing contact data (e.g., as in *Sim II*).

---

[17] Note that only advantages of employing a CMOE are addressed here. Tradeoffs and possible adverse effects must also be assessed to achieve a full understanding of the benefits and vulnerabilities associated with employing such a technique.

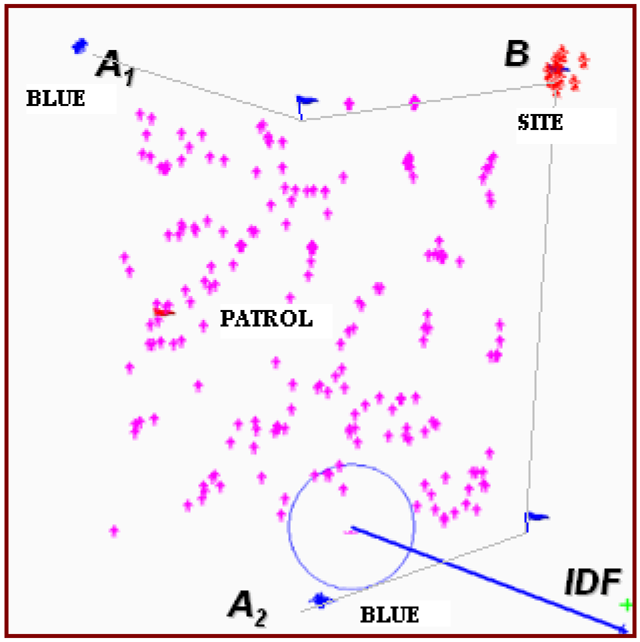

Figure 17: MANA Scenario for Sim III.

The first potential symmetropic signature investigated was the mean (or average) force symmetropy. Comparing mean symmetropies of the two RED force constituents (both fully and also partially based on detections) turns out to be a poor way to typify their differences. The mean values are close together for a given quantity of information and symmetropy exponent *q* (recall matrix is 2^q X 2^q), and the spread is high enough to blur any distinction[18]. The mean symmetropy data are presented in Table 3 (format is 'value, spread (standard deviation)'). Local symmetropies of 3 and 5 detections are shown in addition to full force symmetropies, averaged over numerous time steps. If the mean values had been significantly different, they could have been used to determine the encounter type (SITE or PATROL) and hence fix the decision whether or not to use the IDF support. Unfortunately, it is clear from Table 3 that real-time use of the computed mean symmetropy is of no value in this case.

There is, however, another option worth exploring. The detection data can be separated into distinct symmetropy *modes*. Symmetropy modes are defined here as dominant combinations of the basis symmetries that appear in the symmetropy data as recurring numbers[19]. These modes are a reflection of commonly encountered patterns in the symmetropy matrix. Rather than refer to the specific, unwieldy numeric mode values, they are simply labelled consecutively as Mode 1, Mode 2, etc. In Figure 18, the frequency of symmetropy modes is shown for the two RED encounter

[18] The standard error, computed as the standard deviation (spread) divided by the square root of the number of observations (N), was not shown in Table 3 since N varied considerably between measurements. The spread is less sensitive to N and so provides a better relative measure of uncertainty here. Standard errors were all below 0.04.

[19] Note that this does not account for different patterns that produce the same symmetropy value. A more concise method would be to track the individual symmetry projections as in *Sim II*, but there is no need for such precision within this demonstration.

types SITE and PATROL (the sample is 30 sets of 5 detections, each set of detections in a 30 time step or less time interval and is made by a single squad). The spike at Mode 6 is the sought-after signature. It accounts for 40% of all SITE detections, and only 13% of PATROL detections. Plus, the distribution of the PATROL symmetropy modes is far more uniform than that of the SITE modes. Mode 8 also adds to the signature, although it is weaker than Mode 6.

Table 3: Mean symmetropies of encounter types for Sim III mission reference.

RED Force Constituent	Force Strength	q	Symmetropy Of Entire Force	Symmetropy 3 Detections	Symmetropy 5 Detections
Site Defenders		2	0.81 , 0.12	0.86 , 0.05	0.80 , 0.11
“SITE”	10	3	0.95 , 0.04	0.92 , 0.07	0.94 , 0.08
		4	0.99 , 0.02	0.96 , 0.06	0.99 , 0.02
Main Unit		2	0.82 , 0.04	0.86 , 0.04	0.82 , 0.12
“PATROL”	200	3	0.93 , 0.03	0.93 , 0.08	0.89 , 0.12
		4	1.00 , 0.00	0.98 , 0.06	0.98 , 0.03

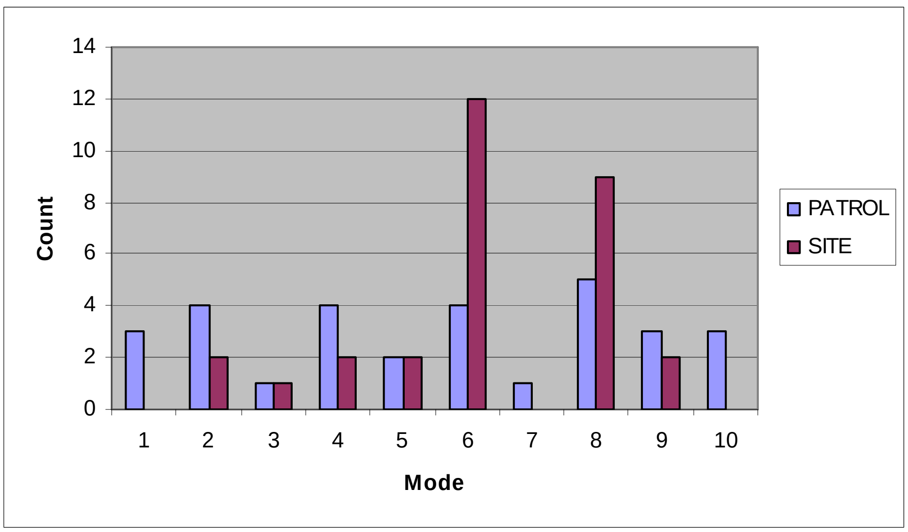

Figure 18: Symmetropy Modes of the RED Force (q=3).

Now it is time to capitalize on the distribution of symmetropy modes and devise a strategy for the *Sim III* mission. Ideally, the approach would be to simply use IDF support until the detections indicate a preponderance of Mode 6 and, to a lesser extent, Mode 8–above what would normally be expected when encountering a PATROL. Next, one could impose limits on the IDF support so as not to fire against those RED agents identified as belonging to the SITE force and let the BLUE assault team handle the encounter on their own. At this point BLUE agents (presumably) proximal to the RED site defenders would be designated to switch into the trigger state found in *Sim II* to be most successful against this group. A level of risk tolerance must be fixed before the simulation though, which essentially defines the cut-off between the expected Mode 6 (8) detections from a PATROL and unusually high Mode 6 (8) detections (indicating that the SITE has been found). The higher the cut-off, the more certain BLUE is that the target site defenders have been correctly identified.

Due to the sparseness of detection data in this scenario, a conservative approach was taken – temporarily turn off the IDF whenever Mode 6 appears and ignore all other detections (including Mode 8). Note that this slightly magnifies BLUE force exposure to the risk that the IDF might be turned off too soon. To minimize this exposure three steps were taken: 1) a minimum symmetropy bounding box length, equal to the BLUE detector range of *50*, was introduced – this helps to ensure that the symmetropy signature is spatially no smaller than the observed scale, 2) the initial start-up of the scenario was not processed (first 500 time steps = 1 time window, as defined below) and 3) the trigger state was given a lifetime of 200 time steps. Detection processing occurred within a running time window of width equal to 500 time steps, and each 'signature' set was composed of exactly *5* detections. The long time window was necessary to accumulate enough detection data to relate to the reference pattern and was gauged manually. The MOE for a run was defined as the number of RED site defenders alive at the time of the earliest correct Mode 6 discovery[20] (i.e., excluding false positives). The idea is that once the RED site defenders are positively identified, the role of the CMOE has been fulfilled. In a sense, the CMOE is an indirect measure of collateral damage assumed to occur if the IDF fires in the near vicinity of waypoint B. The overall mission is deemed 'completed' if BLUE reached the waypoint B.

The results of the simulations and baseline runs, 'Default' and 'Random', are displayed in Table 4. 'Default' runs are simulations (10) performed without utilizing the GA-evolved trigger states and without knowledge of the CMOE to signal a behavioural state change. Thus a single behaviour profile is employed by BLUE agents—the starting profile. 'Random' runs (10) employ the same set of trigger states and behaviour profiles as the CMOE-aware runs, except in this case switching to the state catered to dealing with RED site defenders is triggered at random, depending on when BLUE encounters any of ten randomly wandering neutral entities inserted into the simulation environment. The IDF is triggered on or off as appropriate depending on the current state of the agents. The maximum number of time steps for any run was set to 2000.

According to Table 4, on average 35% (47 of 133) of the BLUE force's opportunities to detect and classify the RED site defenders were successful. Statistically, the result is in line with expectations based on the histogram of symmetropy modes indicating roughly 40% (Figure 18). False positives occurred in 3.3% of cases (16 of 485), somewhat less than the expected 13% as judged from the histogram. The mission was completed in all *Sim III* runs for the Default reference case and the case using CMOEs and triggers. The mission was not completed in many of the Random case runs. The principal difference between the Default and CMOE-triggered simulations (numbered 1-10) lies in the value of the MOE judging mission success. The MOE was improved upon by over ten-fold compared to the Default case and over 3-fold compared to the Random case, on average. This translates to a significant increase in the quality of the end-result. Nevertheless, there is still much room for improvement via a more detailed analysis.

[20] When there were no SITE detections by the end of the run (e.g., simulation #6 in Table 4 in addition to the 'Default' and 'Random' cases), ten (10) subtract the number of RED site defenders killed by the IDF was substituted for the numerator of the MOE.

Table 4: Classification and mission success results.

Simulation Number	First SITE Id Time Step	Correct SITE Ids	False Positive SITE Ids	MOE	Mission Completed
1	667	2 of 2	0 of 88	5/10	YES
2	1249	4 of 16	0 of 0	7/10	YES
3	899	2 of 5	2 of 22	3/10	YES
4	1575	3 of 6	0 of 67	5/10	YES
5	1092	11 of 26	0 of 113	7/10	YES
6	NA	0 of 4	11 of 26	1/10	YES
7	1039	19 of 49	0 of 4	9/10	YES
8	1276	1 of 4	2 of 68	4/10	YES
9	744	1 of 9	1 of 94	5/10	YES
10	726	4 of 12	0 of 3	5/10	YES
SimIII Avgs	*1030 (52%)*	*35%*	*3.3%*	*5.1 /10 (51%)*	*100%*
Default	NA	NA	NA	0.4/10 (4%)	100%
Random	NA	NA	NA	1.6/10 (16%)	40%

Overall, in interpreting these results it is important to realize the mindset of the BLUE force. From BLUE's perspective, the CAS amounts to detections on an SA map corresponding to locations of allies and targets. At some point, the pattern of spatial disorder in that map changes to a known pattern. At this juncture, from BLUE's perspective, a transition is in order since it seems likely that the CAS has changed and a new set of dynamics is at work. BLUE then carries on in a new state of readiness to deal with the perceived threats in the most efficient way known to them.

The salient result of this simulation is that a CMOE was successful in improving mission success for a real-time combat scenario; this in spite of the fact that the forces were sparse and hence data were quite limited. The means through which CMOEs might contribute as a useful degree of freedom in a simulated conflict were not specifically known *a priori*; nevertheless, an opportunity was eventually uncovered. In other conflicts, these measures may contribute significantly to the acquisition of combat system knowledge, or only marginally over and above traditional measures. Lastly, measures not covered by the limited set of CMOEs used here may apply.

This page intentionally left blank.

5 Discussion

The simulations *Sim I & II* conducted in the previous section revealed several interesting features concerning the use of a GA to help define behaviour patterns in combat operations. *Sim III* provided a glimpse of the possible advantages of endowing agents with an awareness of complexity in the system, when the system is likened to a CAS.

A key benefit of employing a GA to find optimal behavioural patterns is the potential generation of new concepts that combine the available degrees of freedom (genes) in ways that the practitioner may not have considered otherwise. *Sim I* was useful in that it highlighted which genes contributed strongly to the fitness of an individual solution and which safely can be ignored (recall the HM series in *Sim I*), permitting efficient progression to *Sim II* where more variables were under consideration (due to the use of triggers to partition the search). In *Sim II*, global CMOEs displayed varying patterns and some showed distinctive features at time steps in the vicinity of key events, inducing a partitioning of the behavioural pattern search space into two distinct groups—one for each encounter type. A third prospective partitioning point that was not attrition related was revealed, but could not be utilized since it involved adopting a behaviour profile outside the scope of the selected parameters of interest. Nevertheless, it provided new insight into important system dynamics not easily obtained otherwise. Overall, the end result of the partitioning was a significant improvement in mission success rates. Regarding the GA used in *Sim II*, a somewhat surprising, unanticipated result was generated. That is, the opposing movement pattern that BLUE evolved for dealing with personal versus squad detection of enemies[21]. This is not the first time a surprising result was obtained using the MANA GA. In McIntosh (2006), RED agents evolved an unexpected, optimal behaviour in a combat scenario that allowed them to remain still, despite the fact that the option to remain still did not exist in any single gene.

Sim II also provided a test-bed for application of such SOC concepts as phase transition and critical point. Although the presence of a phase transition or SOC could not be demonstrated, the *idea* of approaching a possible ‘criticality’, adapting to a ‘phase transition’ and looking for ‘precursors’ fit well as an approach to framing a strategy for improving mission success. In fact, whether or not such features actually exist, as per complex systems doctrine, in a general combat CAS is not clear. Intuitively though, the mindset parallels a sensible and careful approach to optimal mission planning. Thus, despite some lack of rigor, it was demonstrated that the complexity indicators introduced in Section 2 can display recognizable and distinctive patterns in combat, and that these patterns can be leveraged to achieve better insight into the combat dynamics and possibly lead to a tactical advantage.

In *Sim III* it was shown that the real-time tracking and response to CMOEs can be of value in conflict scenarios. In the case considered, a CMOE was used in a precursory-like fashion, hinting at the nature of an imminent near-future change in the system dynamics. The precursors took the form of specific patterns of spatial disorder with respect to a set of predefined symmetries (Walsh functions) residing within the enemy force. Recognized via limited SA, the patterns were successfully resolved by the BLUE force and used to distinguish between RED encounter types

[21] In retrospect, perhaps this should not have been overly surprising or unanticipated, since it merely expresses a preference to attack as a group, rather than as an individual.

and to call off IDF support when appropriate. Furthermore, the use of precursors was combined with state changes and the partitioned, evolved behaviour found in *Sim II*. Mission completion rates when combining a CMOE with state changes were the same or better than the baseline cases, and furthermore showed an overall improvement in the quality of the end conditions in alignment with the main purpose of the mission (i.e., reduced use of IDF support against the RED site defenders, see MOE column in Table 4). With proper support, this result has potential application for automated recognition of, and early response to, an upcoming change or pivotal event (or perhaps criticality) in an observed conflict system; either as a warning or to highlight a budding opportunity regarding the possible onset of a large-scale event. It is of interest to determine if complexity measures can be used as such to detect precursors in a more subtle context and in real-time. Although the situation presented in *Sim III* is artificial, the methodology seems to show promise. Note that other methods, not directly linked to disorder or complexity, could have been devised to achieve a similar effect—there are many differences between the two kinds of encounters to capitalize on. Further study is required to establish whether some combination of complexity measures can provide unique capabilities relevant to C2 in the general case.

In any event, tracking and responding to CMOEs in combat simulations seems to have potential for enhancing awareness about the underlying complex system dynamics at work in a conflict. Such an awareness can also translate into an advantage, both in real-time and also through statistical analysis of repeated simulations.

6 Conclusions

Overall, GA-evolved behaviour profiles for agent combatants were found to significantly improve mission success probabilities within the simulated conflicts investigated. Moreover, unanticipated patterns of beneficial behaviour were discovered by the GA search.

Several CMOEs appropriate for a variety of conflict scenarios were described in this paper: CR-entropy, fractal dimension, spatial entropy, Shannon entropy, self-similarity parameter, Hurst coefficient and symmetropy. All but CR-entropy are based on the spatial dynamical properties of the system rather than on attrition, making them better suited to capturing certain aspects of the complexity of combat. It was shown that in some cases the information gained from distinct measures overlapped to a high degree. Specifically, there was no need to employ both spatial entropy and fractal dimension, and similarly no need to employ both SSP and the Hurst coefficient for the scenario investigated. Symmetropy and the long-term correlation quantities (SSP and the Hurst coefficient) were found to provide insights into the combat dynamics not easily obtained through other means, including aspects that were not directly related to attrition.

It was also suggested that precursors to large scale events (e.g., a wave of casualties) may exist in some combat systems as they do in natural complex systems such as earthquakes, and that CMOEs potentially could be used to help identify and capitalize on these precursors or a somewhat relaxed form of precursor (e.g., hint of an upcoming encounter).

The combat scenarios faced by the BLUE force in this paper presented a difficult challenge to overcome. Mission success rates and agent response capabilities were generally enhanced by adapting agent behaviour based on the knowledge of complexity in the system via CMOEs. The dynamics of the combat system were well described by the CMOE analysis. Factors that contributed to the improvements were 1) how to partition the system on the basis of various CMOE time series, and 2) the early determination of enemy type (or by extension, state) based on an entropy/symmetry measure (symmetropy). The scenarios investigated constituted small confrontations and consequently the data sets used were sparse. This prohibited the use of several CMOEs for use in real-time complexity tracking due to lack of data.

References

Bak, P., Tang, C., Wiesenfeld, K. (1988), Self-organized criticality, *Phys. Rev. B* 38 (1), 364-374.

Bornholdt, S. (1998), Genetic Algorithm Dynamics on a Rugged Landscape, *Phys. Rev. E* 57, 3853.

Dobias, P., J.A. Wanliss, J.C. Samson (2006), Non-linear stability of the near-Earth plasma sheet during substorms: 9 February 1995 event, *Can. J. Phys* 84, 1029.

Dobias, P. (2008a – *in press*), Self-organized Criticality in Asymmetric Warfare, *Fractals*.

Dobias, P. (2008b – *in press*), Complexity-based Assessment in Crowd Confrontation Modeling, *Journal of Battlefield Technology*, 11 (2).

Dockery, J.T. and Woodcock, A.E.R. (1993), The Military Landscape: Mathematical Models of Combat, Cambridge, England: Woodhead Publishing Limited.

Gell-Mann, M. and Lloyd, S. (1996), Information Measures, Effective Complexity, and Total Information, *Complexity*, 2 (1), 44-52.

Goldberg, D.E. (1989a), Genetic Algorithms and Walsh Functions: Part I, A Gentle Introduction, *Complex Systems* 3, 129-152.

Goldberg, D. E. (1989b), Genetic algorithms and Walsh functions: Part II, deception and its analysis, *Complex Systems*, 3, 153-171.

Goldberger A.L., Amaral, L.A.N., Glass L, Hausdorff J.M., Ivanov P.Ch., Mark R.G., Mietus J.E., Moody G.B., Peng C-K, Stanley H.E., (2000), PhysioBank, PhysioToolkit, and PhysioNet: Components of a New Research Resource for Complex Physiologic Signals. *Circulation* 101 (23), e215-e220, Circulation Electronic Pages, http://circ.ahajournals.org/cgi/content/full/101/23/e215 (Accessed February 2008).

Holland, J.H. (1975), Adaptation in Natural and artificial systems, Ann Arbor, MI: University of Michigan Press.

Hurst, H.E. (1951), Long-term storage capacity of reservoirs, *Transactions of the American Society of Civil Engineers*, 116, 770-808.

Ilachinski, A. (1997), Irreducible Semi-Autonomous Adaptive Combat (ISAAC): An Artificial-Life Approach to Land Warfare, CRM 91-61, Center for Naval Analyses, Alexandria, Virginia.

Ilachinski,, A. (2000), EINSTein: An Artificial-Life "Laboratory" for Exploring Self-Organized, Emergent Behavior in Land Combat , CRM D0002239.A1. Alexandria, Center for Naval Analyses.

Ilachinski,, A. (2003), Exploring self-organized emergence in an agent-based synthetic warfare lab, *Kybernetes*, 32 (1), 38-76.

Ilachinski, A. (2004), Artificial War: Multiagent-based simulation of combat, World Scientific.

Jones, C.L., Lonergan, C.T., Mainwaring, D.E. (1996), Wavelet Packet Computation of the Hurst Coefficient, *Journal of Physics A: Math. Gen.*, 29, 2509-2527.

Kaplan, I., Estimating the Hurst Exponent, http://www.bearcave.com/misl/misl_tech/wavelets/hurst (Accessed December, 2007).

Kolmogorov, A.N. (1965), Three approaches to the quantitative definition of information, *Problems of Information and Transmission*, 1 (1), 1-7.

Lanchester, F.W. (1914), Aircraft in Warfare: The Dawn of the Fourth Arm—No V. Principles of Concentration, *Engineering*, 98, 422-423.

Lauren, M.K. (1999), Characterising the Difference between Complex Adaptive and Conventional Combat Models, DOTSE Report 169, NR1345, ISSN 1174-3387, Defence Operational Support Technology Establishment, Auckland, New Zealand.

Lauren, M.K. (2001), Fractal Methods Applied to Describe Cellular Automaton Combat Models, *Fractals*, 9 (2), 177-184.

Lauren, M.K (2002), Fractals and Combat Modeling: Using MANA to Explore the Role of Entropy in Complexity Science, *Fractals* 10 (4), 481-489.

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

Lauren, M.K. (2003), On the temporal distribution of fatalities and determination of medical logistic requirements, DTA Report 187, NR1377, ISSN 1175-6594, Defence Technology Agency, Auckland, New Zealand.

Luscombe, R., Mitchard, H. (2005), Exploring Simple Human Behaviour Representation Using Agent Based Distillations, SimTect Conference Paper.

McIntosh, G.C. and Lauren, M.K. (2006), Genetic Algorithms Applied to Course-of-action Development Using the MANA Agent-based Model, *Journal of Battlefield Technology*, 9 (3), 27-32

Moffat, J. (2006), Mathematical Modelling of Information Age Conflict, *Journal of Applied Mathematics and Decision Sciences*, Volume 2006, Article Id 16018, 1-15.

Nanjo, K., H. Nagahama, E. Yodogawa (2000), Symmetry Properties of Spatial Distribution of Microfracturing in Rock, *Forma*, 15, 95-101.

Nanjo, K., H. Nagahama, E. Yodogawa (2001), Symmetropy and self-organized criticality, *Forma* 16, 213-224.

Nanjo, K., H. Nagahama, E. Yodogawa (2005), Symmetropy of fault patterns : Quantitative measurement of anisotropy and entropic heterogeneity, *Mathematical geology*, 37 (3), 277-293.

Narteau, C. (2007), Classification of seismic patterns in a hierarchical model of rupture: a new phase diagram for seismicity, Geophys. *J. Int.* 168, 710-722.

Nix, A.E., Vose, M.D. (1992)**,** Modeling genetic algorithms with Markov chains, *Annals of Mathematics and Artificial Intelligence* 5 (1), 79-88.

Parunak, H.V.D., Bisson, R., Brueckner, S., Matthews, R.S., Sauter, J.A. (2006), A model of emotions for situated agents, *AAMAS*, 993-995.

Peng C-K, Buldyrev S.V., Havlin S., Simons M., Stanley H.E., Goldberger A.L. (1994), Mosaic organisation of DNA nucleotides, *Phys Rev E* 49, 1685-1689.

Peng, C.K., Havlin, S., Stanley, H.E., Goldberger, A.L. (1995), Quantification of scaling exponents and crossover phenomena in nonstationary heartbeat timeseries, *Chaos 5* (1), 82-87.

Prügel-Bennet, A., Shapiro, J.L. (1994), An Analysis of Genetic Algorithms Using Statistical Mechanics, *Phys. Rev. Lett.* 72, 1305-1309.

Prügel-Bennet, A., Shapiro, J.L. (1997), The Dynamics of a Genetic Algorithm for Simple Random Ising Systems, *Physica D* 104**,** 75-114.

Richardson, L.F. (1941), Statistics of Deadly Quarrels, *Nature* 148, 598.

Roberts, D.C., Turcotte, D.L. (1998), Fractality and Self-Organised Criticality of Wars, *Fractals* 6, 351-357.

Rogers, A. and Prügel-Bennett, A. (1999), Genetic Drift in Genetic Algorithm Selection Schemes, *IEEE Transactions on Evolutionary Computation*, 3 (4), 298-303.

Rogers, A., Prügel-Bennett, A. and Jennings, N. R. (2006), Phase Transitions and Symmetry Breaking in Genetic Algorithms with Crossover, *Theoretical Computer Science* 358 (1), 121-141.

Samson, J. C., Rankin, R., and Tikhonchuk, V.T. (2003), Optical signatures of auroral arcs produced by field line resonances; comparison with satellite observations and modeling, *Annales Geophysicae* 21, 933.

Scherrer, J.H. (2003), Risks and vulnerabilities of network-centric forces: Insights from the science of complexity, Newport RI, Naval War College.

Shalizi, C.R. and Crutchfield, J.P. (2001), Computational Mechanics: Pattern and Prediction, Structure and Simplicity, *J. Stat. Phys.*, 104, 817-879.

Shalizi, C.R., Shalizi, K.L., Haslinger, R. (2004), Quantifying Self-Organization with Optimal Predictors, *Phys. Rev. Lett.*, 93, 118701.

Voronkov, I.O., Donovan, E.F., Dobias, P., Prosolin, V.I., Jankowska, M., Samson, J.C. (2004), Late growth phase and breakup in the near-Earth plasma sheet, in Proceedings of the 7th International Conference on Substorms, eds. N. Ganushkina, T. Pulkinnen, Finish Meteorological Institute, 140.

Walsh, J.L. (1910). A Closed Set of Normal Orthogonal Functions. *Mathematishe Annalen* 69, 331-371.

Yang, A. and Abbass, H. A. and Sarker, R. (2004), Landscape Dynamics in Multi-agent Simulation Combat Systems. In: Proceedings of 17th Joint Australian Conference on Artificial Intelligence, LNAI 3339, Cairns, Australia, Springer-Verlag, 4722.

Annex A Simulation Settings

The MANA personality settings for RED and BLUE squads in Sims I, II and III are provided below. Ranges are given for evolved parameters. 'X' indicates a MANA default setting was used (in all cases, the default setting was zero).

Table A.0: MANA personality parameter descriptions

Personality Trait	Description
psEnemies	Attraction/repulsion to enemies detected personally.
psFriends	Attraction/repulsion to Comrades detected personally.
psNextFlag	Atttraction/repulsion to the next waypoint.
psOrgThreat2	Attraction/repulsion to enemies of threat level 2 known through squad SA.
psEnThreat2	Attraction/repulsion to enemies of threat level 2 detected personally.
psOrgThreat3	Attraction/repulsion to enemies of threat level 3 known through squad SA.
psEnThreat3	Attraction/repulsion to enemies of threat level 3 detected personally.

A.1 Sim I

The MANA personality settings for RED and BLUE squads in Sim I (Section 4.1) are provided in Table A.1 below. Note that all RED agents were considered 'threat level 3' in this simulation.

Table A.1: Settings for Sim I.

Personality Trait	BLUE squad	RED site defender squad	RED patrol squad
psEnemies	-100 to 100	10	100
psFriends	-100 to 100	X	-50 (squad only)
psNextFlag	-100 to 100	X	20
psOrgThreat3	-100 to 100	X	X

A.2 Sim II

The MANA settings for trigger states of the BLUE squad in Sim II (Section 4.2) are provided in Table A.2 below. RED settings are constant throughout as per Sim I. Note that BLUE attraction/repulsion to RED agents was refined to allow BLUE to react differently to RED patrol agents (threat level 2) and RED site defenders (threat level 3). Furthermore, BLUE was empowered to respond differently to RED agents personally encountered versus those detected through squad (organic) SA.

Table A.2:BLUE Settings for Sim II.

Personality Trait	Trigger 1 State	Trigger 2 State
psEnemies	X	X
psFriends	100	100
psNextFlag	50	50 to 100
psOrgThreat2	-100 to 100	X
psEnThreat2	-100 to 100	X
psOrgThreat3	X	-100 to 100
psEnThreat3	X	-100 to 100

A.3 Sim III

The MANA settings for RED and BLUE squads in Sim III (Section 4.3) are shown in Table 3 below. Note all RED agents were considered threat level 3.

Table A.3: Settings for Sim III.

Personality Trait	BLUE Trigger 1	BLUE Trigger 2	RED site defender squad	RED patrol squad
psEnemies	X	X	10	20
psFriends	100 (squad only)	100 (squad only, cluster=2)	X	60 (squad only, cluster=8)
psNextFlag	50	100	X	20
psOrgThreat3	-100	100	X	X
psEnThreat3	X	-100	X	10

This page intentionally left blank.

List of symbols/abbreviations/acronyms/initialisms

ABD	Agent-based Distillation
ABM	Agent-based Model
BLUE	The Friendly Force (our side)
C2	Command and Control
CA	Cellular Automata
CAS	Complex Adaptive System
CMOE	Complex Systems Measure of Effectiveness
CR-entropy	Carvalho-Rodriguez Entropy
CORA	Centre for Operation Research and Analysis
DFA	Detrended Fluctuation Analysis
DND	Department of National Defence
DRDC	Defence Research & Development Canada
EINSTein	Enhanced ISAAC Neural Simulation Toolkit
GA	Genetic Algorithm
H	Hurst coefficient
IDF	Indirect Fire
ISAAC	Irreducible Semi-Autonomous Adaptive Combat
MANA	Map-Aware Non-Uniform Automata
MOE	Measure of Effectiveness
R&D	Research & Development
RED	The Enemy Force
SA	Situational Awareness
SOC	Self-organized Criticality
SSP	Self-similarity Parameter
WISDOM	Warfare Intelligent System for the Dynamic Optimization of Missions

This page intentionally left blank.

Distribution list

Document No.: DRDC CORA TM 2008-044

Internal

Dr. Kevin Sprague (hardcopy, CD)

Dr. Peter Dobias (hardcopy, CD)

DG DRDC CORA (e-mail)

DRDC CORA Library (hardcopy, email)

Chief Scientist DRDC CORA (e-mail)

Section Head, Land and Operational Commands (e-mail)

Section Head, Air OR (e-mail)

Section Head, Maritime OR (e-mail)

LFORT (hardcopy, CD)

CEFCOM ORT (e-mail)

LCDORT (e-mail)

CANSOFCOM ORT (e-mail)

CanadaCOM ORT (e-mail)

Deployed ORA Team, Kandahar (e-mail)

External

ADM(S&T) (for distribution) (CD)

DRDKIM 3 (CD)

Director S&T Land (e-mail)

DG DRDC Valcartier (e-mail)

Library and Archives Canada (hardcopy, CD)

CF College Library (hardcopy, CD)

Fort Frontenac Library (hardcopy, CD)

CFANS Library (hardcopy, CD)

DLCD (e-mail) LCol Rettie	rettie.jm@forces.gc.ca
DLR (e-mail) Col Lanthier	lanthier.jm@forces.gc.ca
DLSE (e-mail) LCol Bassarab	bassarab.rr@forces.gc.ca

International

Document Exchange Manager (CD)
DSTO Research Library
Defence Science & Technology Organisation
PO Box 44
Pyrmont NSW 2009, AUSTRALIA

Ms Anne Marie Grisogono (CD)
Land Operations Division
Defence Science and Technology Organisation
PO Box 1500 Edinburgh 5111
South Australia

Dr. Neville J Curtis (CD)
Research Leader Land Operations Research
75 Labs
Land Operations Division
PO Box 1500
Edinburgh SA 5111, AUSTRALIA

Michael Gillman (for dist'n and library) (CD)
Chief Technologist
Land Battlespace Systems
Dstl Integrated Systems
Room 31, Bldg A3, Fort Halstead
Sevenoaks, Kent, UK, TN14 7BP

Dr Niki Jobson (CD)
Joint Systems Department
Defence Science and Technology Laboratory
Rm C45, Grenville Building
Dstl, Fareham
Hants, PO17 6AD , United Kingdom

Director, US AMSAA (CD)
ATTN: AMSRD-AMS-S)
392 Hopkins Road
APG, MD 21005-5071

Dr. Michael Lauren (CD, email)
Research Leader, Human Systems Group
Defence Technology Agency,
HMNZ Naval Base Auckland,
Private Bag 32901, Auckland, New Zealand

Dr David Galligan (CD)
Head Operations Analysis,
Defence Technology Agency,
HMNZ Naval Base Auckland,
Private Bag 32901, Auckland, New Zealand

Alex Ryan (CD)
Land Operations Division
Defence Science and Technology
Organisation
PO Box 1500 Edinburgh 5111
South Australia

Dr. Jason Field (CD)
Land Battlespace Systems
Dstl Integrated Systems
Fort Halstead
Sevenoaks, Kent, UK, TN147BP

Dr. John Parmentola (CD)
Director of Army Research Program
Suite 9800
2511 Jefferson Davis Hwy
Arlington, Virginia 22202
United States

Mr. Patrick O'Neill (CD)
Chief, Combat Support Analysis Division
USAMSAA (ATTN: AMSRD-AMS-S)
392 Hopkins Road
APG, MD 21005-5071

Dr. James T. Treharne (CD)
OCA Division
Center for Army Analysis
6001 Goethals Road
Fort Belvoir, VA 22060-5230

Mr. John Hughes (CD)
HQ, TRADOC Analysis Center (TRAC)
Programs & Resources Directorate (PRD)
255 Sedgwick Avenue
Fort Leavenworth, Kansas 66027-2345

Mr. Paul A Shelton
Applied Physics Laboratory
Johns Hopkins University
11100 Johns Hopkins Road
Laurel, MD 20723-6099
United States

Mr. Robert Barrett (CD)
Chief, International Activities
Center for Army Analysis
6001 Goethals Road
Fort Belvoir, VA 22060-5230

Mr. Bob Barbier (CD)
TNO Defence, Security and Safety
Information and Operations
P.O. Box 96864, 2509 JG
The Hague, The Netherlands

This page intentionally left blank.

DOCUMENT CONTROL DATA

(Security classification of title, body of abstract and indexing annotation must be entered when the overall document is classified)

1. ORIGINATOR (The name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g. Centre sponsoring a contractor's report, or tasking agency, are entered in section 8.) DRDC CORA Department of National Defence Ottawa, Ontario K1A 0K2	2. SECURITY CLASSIFICATION (Overall security classification of the document including special warning terms if applicable.) UNCLASSIFIED

3. TITLE (The complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S, C or U) in parentheses after the title.)

Behaviour in Simulated Combat: Adaptation and Response to Complex Systems Factors (U)

4. AUTHORS (last name, followed by initials – ranks, titles, etc. not to be used)

Sprague, Kevin; Dobias, Peter

5. DATE OF PUBLICATION (Month and year of publication of document.)	6a. NO. OF PAGES (Total containing information, including Annexes, Appendices, etc.)	6b. NO. OF REFS (Total cited in document.)
November 2008	64	49

7. DESCRIPTIVE NOTES (The category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.)

Technical Memorandum

8. SPONSORING ACTIVITY (The name of the department project office or laboratory sponsoring the research and development – include address.)

DRDC CORA
Department of National Defence
Ottawa, Ontario K1A 0K2

9a. PROJECT OR GRANT NO. (If appropriate, the applicable research and development project or grant number under which the document was written. Please specify whether project or grant.)	9b. CONTRACT NO. (If appropriate, the applicable number under which the document was written.)
N/A	N/A

10a. ORIGINATOR'S DOCUMENT NUMBER (The official document number by which the document is identified by the originating activity. This number must be unique to this document.)	10b. OTHER DOCUMENT NO(s). (Any other numbers which may be assigned this document either by the originator or by the sponsor.)
DRDC CORA TM 2008-044	N/A

11. DOCUMENT AVAILABILITY (Any limitations on further dissemination of the document, other than those imposed by security classification.)

Unlimited

12. DOCUMENT ANNOUNCEMENT (Any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution (beyond the audience specified in (11) is possible, a wider announcement audience may be selected.))

Unlimited

13. ABSTRACT (A brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), (R), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual.)

Tactical combat has been demonstrated to exhibit properties of complex adaptive systems (CAS). In some cases, CAS dynamics can lead to large-scale (possibly catastrophic) impacts that are not in line with expectations under traditional thinking. In this paper, recognizing and exercising some degree of influence over combat CAS dynamics is investigated. In particular, we discuss approaches to selectively “drive” a conflict towards more favourable regions of the available phase space. The issue is addressed at the fundamental level of entity behaviour and interaction. Two key factors for consideration towards such a goal are 1) optimization of combatant behaviour and 2) awareness of and response to complexity within the system. A set of complex systems measures of effectiveness (CMOEs) appropriate to combat, drawn from various complex systems factors available in literature, are proposed and investigated. These measures provide a window into the dynamical progression of a combat CAS, while behaviour modifications offer the means to adapt to its changing conditions. The interplay between the two factors comprises the underlying theme of this paper. Candidate CMOEs discussed include: the fractal dimension, Shannon entropy (two forms: Carvalho-Rodrigues and spatial entropy), the Hurst coefficient, the self-similarity parameter, and symmetropy. Simulations are used to illustrate how a CAS mindset and adaptive behaviour can be leveraged to improve (simulated) mission results. CMOEs are first analyzed in repeated simulations of a simplistic combat scenario to help develop optimal behaviours. The derived behaviours are later combined with real-time awareness of a CMOE to assist in decision-making during a simulated mission of comparable nature to the original combat scenario. It was found that, overall, this new degree of freedom improved the quality of the mission outcome. Also, several CMOEs were found to be of apparent use only in repeated simulations (as opposed to real-time) due to data requirements and high variances in individual runs.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (Technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus, e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

combat, behaviour, command and control, C2, fractal, complex adaptive systems, self-organized criticality, MANA, complexity, genetic algorithm

DEFENCE DÉFENSE

DRDC CORA