

THE INTERNATIONAL
C2 JOURNAL

VOLUME 2, NUMBER 2, 2008

SPECIAL ISSUE

*Representing Human Decision Making
in Constructive Simulations for Analysis*

GUEST EDITOR

James Moffat
Dstl

Modeling the Complexity of Combat
in the Context of C2

K.B. Sprague
P. Dobias

THE INTERNATIONAL C2 JOURNAL

David S. Alberts, Chairman of the Editorial Board, *OASD-NII, CCRP*

The Editorial Board

Berndt Brehmer (SWE), *Swedish National Defence College*

Reiner Huber (GER), *Universitaet der Bundeswehr Muenchen*

Viggo Lemche (DEN), *Danish Defence Acquisition and Logistics Organization*

James Moffat (UK), *Defence Science and Technology Laboratory (DSTL)*

Mark Nissen (USA), *Naval Postgraduate School*

Ross Pigeau (CAN), *Defence Research and Development Canada (DRDC)*

Mink Spaans (NED), *TNO Defence, Security and Safety*

About the Journal

The International C2 Journal was created in 2006 at the urging of an international group of command and control professionals including individuals from academia, industry, government, and the military. The Command and Control Research Program (CCRP, of the U.S. Office of the Assistant Secretary of Defense for Networks and Information Integration, or OASD-NII) responded to this need by bringing together interested professionals to shape the purpose and guide the execution of such a journal. Today, the Journal is overseen by an Editorial Board comprising representatives from many nations.

Opinions, conclusions, and recommendations expressed or implied within are solely those of the authors. They do not necessarily represent the views of the Department of Defense, or any other U.S. Government agency.

Rights and Permissions: All articles published in the International C2 Journal remain the intellectual property of the authors and may not be distributed or sold without the express written consent of the authors.

For more information

Visit us online at: www.dodccrp.org

Contact our staff at: publications@dodccrp.org



Modeling the Complexity of Combat in the Context of C2

K.B. Sprague and P. Dobias (Defence Research & Development Canada, Canada)

Abstract

Tactical combat has been demonstrated to exhibit properties of complex adaptive systems (CAS). In this paper, recognizing and exercising some degree of influence over CAS dynamics is investigated in the context of command and control (C2). In particular, approaches to selectively “drive” a conflict towards more favourable regions of the available phase space are discussed. Two features of key importance to such a goal in a CAS environment are combatant behaviour and measures of effectiveness that incorporate complex systems factors. The measures provide a window into the dynamical progression of the system, while behaviour modifications offer the means to adapt to it. The interplay between the two factors comprises the underlying theme of this study. Candidate measures of effectiveness in a complex systems environment are discussed, including: the fractal dimension, Shannon entropy (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 achieve better C2 and improve (simulated) mission outcomes.

Introduction

Motivation

Traditional command and control (C2) in warfare draws upon results based on the analysis of classical combat dynamics such as, for example, the Lanchester equations describing attrition rates (Lanchester, 1914). Most models are implicitly based on the assumption of normal (Gaussian) underlying statistical distributions for salient characteristics such as kill probabilities, and consequently attrition. However, it can be shown that in many cases the dynamics of combat obey fractal, rather than normal, statistics (Lauren 2003, Dobias 2008b)—not only with regard to the distribution of casualties, but also regarding the spatial distribution of forces, the distribution of radio traffic and the frequency of conflict intensity. Thus expectations that are in-line with traditional thinking may not reflect certain realities of a conflict. In addition, combat dynamics have been shown in some cases to exhibit signatures of self-organized criticality (SOC) (Ilachinski 2004, Lauren 2001, Dobias 2008b) when viewed as a complex adaptive system (CAS)¹. This means that once the system reaches its critical point² (typically an attractor), a rapid transition can occur, possibly leading to catastrophic large-scale events having a tremendous impact on the outcome of the conflict under study. Ignoring these factors can lead to a perilous misconception of the risks involved in a combat operation (Ilachinski 2004).

Thus recognizing and exercising some degree of influence over CAS dynamics at or near critical points is worthy of investigation in the context of C2—as is avoidance or migration towards such a

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 is consistent with underlying fractal statistics.

state (depending on the circumstances), if possible. Furthermore, simply being able to characterize the dynamics of a given CAS, even in more mundane settings (i.e., no evident SOC), may provide insights into useful dynamical patterns or symmetries within the system not readily apparent using more traditional analyses. In particular, it is of great value to develop methods that selectively “drive” a conflict towards more favourable regions of the available phase space (Ilachinski 2004). Since the CAS dynamics depend on the underlying rules of behaviour adopted by the interacting entities (e.g., training, tactics), it follows that reaching a given objective within a conflict requires the parallel consideration of behaviour and the system response to behaviour measured with respect to achieving that objective. Thus two factors of key importance for favourably influencing a conflict are behaviour and complex systems measures of effectiveness (CMOEs). The measures provide a window into the dynamical progression of the system, while behaviour modifications offer the means to adapt to it. The interplay between these factors comprises the underlying theme of this paper.

Relevance to C2

If a conflict scenario, or specific aspects thereof, were analyzed beforehand in consideration of the relationship between behaviour (friendly, foe, or neutral) and the resulting CAS dynamics, then it is conceivable that command could leverage this information to prescribe optimal behaviour ‘rules’ that probabilistically improve mission outcomes. In particular, a robust set of behaviour profiles can be adopted that apply to general circumstances one might encounter, or alternatively a highly optimized set of situation-specific behaviours can be adopted to deal with an exclusive scenario of interest. In any case, the commander is endowed with a heightened awareness linking behaviour to consequences and extended capabilities in directing those he/she commands in a CAS combat environment.

Previous Work

Others have demonstrated the value of improving mission success rates by optimizing the behaviour of combatants under fitness criteria tied to the achievement of the mission objectives in a simulated combat environment (Ilachinski 2004, Lauren 2002, Tolk 1995, Hofmann 1995). Measures of effectiveness (MOEs) associated with overall mission success drive the optimization process. Common examples of MOEs for a friendly (BLUE) force include ‘number of BLUE casualties’ and the Boolean variable ‘BLUE reached the desired location’. It seems reasonable to move towards extending the above notions of optimization to include the use of CMOEs, especially with regard to *triggering* a desired behaviour (see the Section *Genetic Algorithms in MANA*, below).

Aim

The aim of this paper is to illustrate how knowledge of complex systems factors in combat can be characterized and how it may lead to a tactical advantage within a few conceptually simple combat situations. Moreover, advantages are to be shown both in the analysis of aggregate CMOE results (i.e., repeated simulations of a given scenario) and real-time CMOE tracking (i.e., response to one or more CMOEs within a single run).

Scope

In summary, experiments are conducted through simulation and address the following:

1. Exploration of behaviour optimization in a difficult combat scenario;
2. Use of CMOEs in repeated simulations to improve mission success probabilities for the given scenario; and

3. Successful use of a real-time CMOE coupled with optimized, situation-specific behaviour profiles in a second scenario having similar elements to the first.

The principle metrics employed consist of the following:

1. Degree to which CMOEs contribute to the success and/or efficiency of the optimization process (in repeated simulations); and
2. The interpretive value of CMOEs in a real-time combat simulation environment.

The scenario employed was 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 dynamics.

Simulation Environment

The arena employed was an agent-based distillation (ABD) called ‘Map-Aware Nonlinear Automata’ (MANA) (Lauren and Stephen 2002), described below. ABDs form a subset of the more general class of agent-based models (ABM). ABMs in this regime are generally based on the philosophy of cellular automata (CA). 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 situational awareness (SA). ABMs have been successfully utilized to model a variety of scenarios where emergent behaviour rather than specific technical properties were to be 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.

Traditionally, the opposing force (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.

Before delving into simulations, potential CMOEs are briefly described (see Appendix A for more detailed descriptions) and the framework for behaviour representation and development is reviewed. Furthermore, genetic algorithms are portrayed as a search tool designed to ‘find’ the optimal state of behaviour for BLUE agents in various circumstances given limited SA.

Complex Systems Measures of Effectiveness in Combat

Heuristics

To understand the dynamics of a complex system requires that appropriate measures be established. We identify CMOEs with variables that capture features of particular interest in complex sys-

tems, on presumption that they might be applicable as measures of effectiveness for combat. A more general term might be ‘complex system factors’. The list includes various entropies, variables capturing long-term correlations, measures of complexity, and others. In this paper the CAS is described via a few carefully chosen indicators normally associated with disorder; namely, the fractal dimension, Shannon entropy³ (two forms: Carvalho-Rodrigues and spatial entropy), the Hurst coefficient, the self-similarity parameter and symmetry. Most of these measures have been applied to combat dynamics previously and each measure is described, in turn, below (also see Appendix A). 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 potential upcoming events. Note that whether disorder is increasing or decreasing in a CAS can be a matter of perspective. For example, 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. 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., The Second Law of Thermodynamics)⁴. 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

3. Note that when measuring the degree of self-organization in a complex system, entropy is somewhat contentious as a ‘measure of complexity’. See (Shalizi 2004), for instance. Nevertheless, in general various forms of entropy are relevant to understanding the dynamics of complex systems.

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

it progresses to a criticality holds potential value for future analyses (see, for instance, Shalizi 2004 for a promising candidate measure).

In traditional combat models and wargames the primary measure of effectiveness is often attrition—whether measured directly (number of killed, loss-exchange ratio⁵, etc.) or indirectly (attrition-based definition of mission success). However, in some cases the focus on attrition actually ignores the complexity of combat (Dobias 2008a). 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 measures of complexity are described—one attrition-based and 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 scenario-specific—none are universally relevant. A key determinant seems to be how disorder unfolds in the system. Therefore, the main focus of most measures is on the progression and degree of disorder within the system from various perspectives. 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.

Individual CMOE Descriptions

Several measures have been proposed to capture the complexity of combat. Possibly the oldest of them is entropy, of one form or another. Entropy is a measure of disorder in a system from a particular perspective that varies depending on the application.

5. Ratio between RED and BLUE killed.

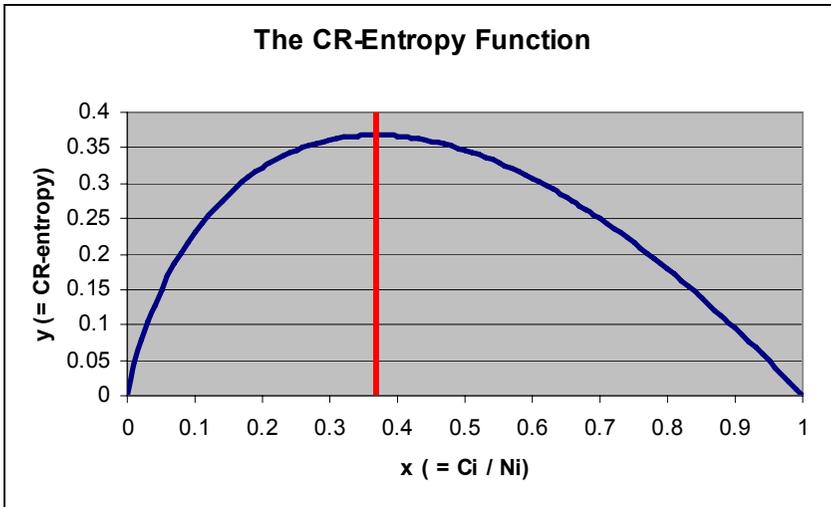


Figure 1. The CR-entropy function. A plot of $y = x \ln(1/x)$.

Carvalho-Rodrigues (1989) (CR) proposed an attrition-based definition of combat entropy applicable to each force individually (e.g., RED or BLUE) or as a whole through reporting the difference in force entropies. It is a form of Shannon entropy (also referred to as ‘information’ entropy) (Shannon 1949). One of the main features of CR-entropy is that, prior to attaining a maximum value, a higher CR-entropy for a force translates to a more disordered combative state. Thus prior to the engaged forces reaching maximum CR-entropy, the force having the lower CR-entropy is considered to have the advantage. Once the CR-entropy of a force breaches the maximum value, it enters into a *disintegration phase*, wherein combat capabilities are assumed to have declined substantially (Ilachinski 2004). For formulae related to Figure 1, see Appendix A.

Ilachinski (2004) suggested a specific form of Shannon entropy based on the spatial distribution of soldiers relative to a regular grid covering the battlefield area. The resulting ‘spatial entropy’ is closely related to the fractal dimension when the latter is computed via the ‘box counting’ method (see Appendix A for a comparison).

Rather compact, non-dispersed patterns display low spatial entropy whereas disorganized, spread-out patterns display high spatial entropy.

The fractal dimension can also be used as a measure of the spatial distribution of combat units (e.g., crowd, BLUE force) (Ilachinski 2004). It is a statistical quantity that quantifies the self-similarity of the distribution of units on the battlefield from the large distance scales of the system to finer and finer scales. In particular, it describes the clustering properties of force units (which has been related to firepower concentration (Lauren 2000)) and can act as a rough discriminator between laminar and turbulent classes of behaviour (Ilachinski 2003). The fractal dimension also relates to scaling relationships typical of SOC.

Temporal and spatial correlations in agent velocity (speed and direction) are other characteristics that could possibly provide additional insights into complex system dynamics. Such correlations can be calculated independently for each velocity component of moving entities. Correlations can be described in terms of the Hurst coefficient H and/or also the self-similarity parameter (SSP). For the scenario investigated, H and the SSP displayed the same basic pattern. In both cases, velocity correlations are characterized by a scaling between the number of steps and the root mean square distance traveled. When H or SSP values are at 0.5, the motion is random. If between 0.5 and 1 the motion is correlated; if between 0 and 0.5 the motion is anti-correlated; and if 0, the motion is centered about a point. There are notable differences between the two, however, most evident in computational aspects and the interpretation of SSP for values larger than 1 (see Appendix A to compare). For the scenario investigated, the two measures yielded similar apparent behaviour.

A new quantity was proposed on the basis of Shannon entropy that measures the symmetry and entropy of a given spatial pattern or shape. This measure is called symmetry and has been applied to investigations of SOC (Nanjo 2001, 2005). The spatial distribution

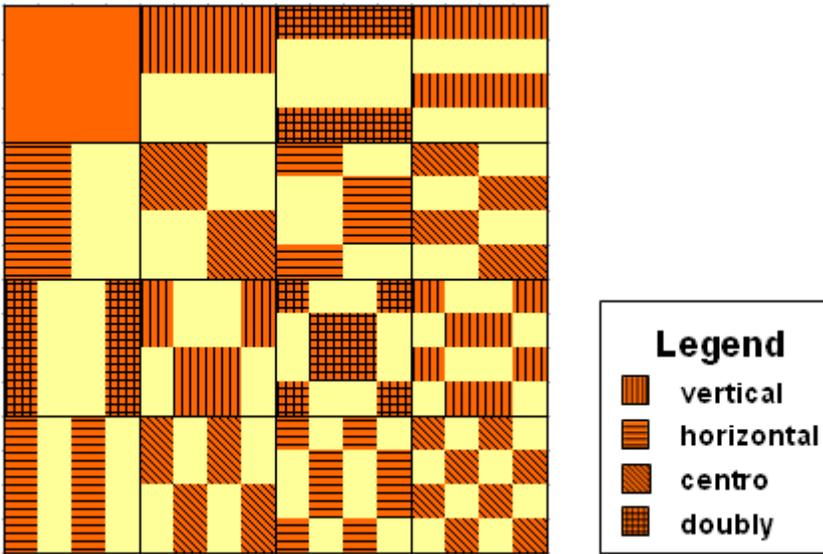


Figure 2. The first 16 members of the 2-D Walsh function kernel.

of units is projected onto a pattern basis to determine the relative contribution of reference symmetries to the observed pattern. The rise and fall of the various spatial symmetries can be tracked during a simulation. When the pattern is random, the value is high (1), and when dominating symmetries are present in the system the value is low. Figure 2 shows the first sixteen patterns of the two dimensional Walsh function kernel used to compute symmetry. The patterns are separated into four main classes of symmetries: vertical, horizontal, centro, and double.

Application, or Potential Application, to Conflicts

The fractal dimension and corresponding power-laws have been used to describe the statistical distribution of the intensities of wars (Roberts and Turcotte 1998), warfare statistics (Richardson 1941) and attack casualties (Lauren 2001, Dobias 2008b), 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 used to explain spatial properties of the battlefield 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 and the degree of disorder in a manner similar to that of the fractal dimension.

CR-entropy was first used to address logistical concerns during military exercises (Carvalho-Rodrigues 1989). Dockery *et al* (1993) employed 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 crowd control and, in particular, signal a phase transition between a group confrontational mindset and the inclination to disperse (Dobias 2008a). The SSP potentially could have been used in an analogous manner.

To our knowledge, the concept of symmetry is new to the domain of combat dynamics⁶, but holds promise in general for pattern recognition under a degree of disorder, including possibly the identification or classification of forces based on limited SA. It also holds promise for identifying the state of a complex system. Examples from the geological sciences involving earthquakes and/or acoustic transitions leverage symmetry values and corresponding symmetry projections to describe various dynamical aspects of the system in question. For a fault model with SOC, fault patterns of critical states and sub-critical states⁷ are distinguishable via symmetry—

6. The article by Dobias (2008b), in press at the time of writing, also applies the concept of symmetry to combat. Therein, symmetry is found to be quite effective in capturing system-wide changes in conflicts exhibiting SOC.

sub-critical fault patterns show nearly constant symmetry values whereas various values are taken on during critical states (Nanjo 2005). Work with microfracturing in rock indicates that the process evolves under a constraint of increasing richness in double symmetry (a trend towards low symmetry 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 symmetry might exhibit meaningful variations in certain combat systems.

Precursors and SOC

The idea that precursors might exist in certain critical and/or sub-critical states of a complex system holds obvious potential benefit if measurable and applicable to a combat CAS. The concept of SOC 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 prototype of SOC (Bak 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, and 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 com-

7. 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.

plexity, the SOC is 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 SOC.

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 the large ones.

In general, large-scale critical events in dynamical systems are often (but not always) preceded by smaller, more frequent events (precursors). The existence of precursors depends on the state of the complex system. A complex system may have multiple critical configurations—some with precursors and others without (for instance, see Narteau 2007). Examples from natural complex systems where precursors are known to occur include tremors that precede large earthquakes (foreshocks) and localized intensifications or resonances preceding the onset of magnetospheric disturbances (substorms) (Samson 2003, Voronkov 2004). Such precursors can facilitate early response to the **possibility** of a large-scale event in the near-time horizon. Precursors should be interpreted as probabilistic indicators, and can be extended to incorporate factors such as the chance of observing multiple large-scale events in a short time frame (based on the fractal distribution of events in the system and knowledge of the current state of the system). Note that the large-scale events themselves may indicate the onset of more of the same.

As an example related to combat, Figure 3 shows the results for the entropy for a hostile crowd-control scenario modeled using MANA. Entropy shows a slight increase and then a dip preceding the main increase due to crowd dispersal (time steps $\sim 150-250$). This is consistent with the change in the system's state corresponding to a phase transition (Dobias 2008a). Similar precursors are common in other dynamical systems such as those mentioned above.

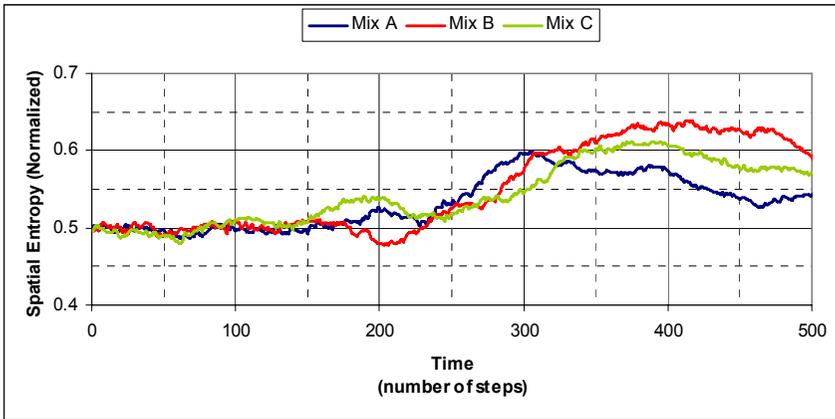


Figure 3. Average Shannon (spatial) entropy for three different Blue configurations.

The existence of the precursors (pseudo-breakups) in a magnetospheric system was explained by the lack of sufficient free energy in the system to trigger a full-scale event (Dobias 2006). This is consistent with a sub-critical system near a critical point (discharge event systems). Thus it is reasonable to expect that in military combat systems, a major system change (phase transition) likewise will be signalled by precursors of a similar nature in certain cases. 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.

Evolving the *Agility* of Combatants within an Agent-based Distillation

Motivation

With enough simulations, a description of the relative frequencies of possible outcomes for a virtual battle scenario emerges. 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 the operation’s likelihood of success?”

Given fixed forces, weaponry and equipment, the success probabilities of simulated combatants in an agent-based distillation can vary widely dependent upon 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 is referred to here as *agility*. Agility, in the context of the simulations employed, amounts to manoeuvre tactics and involves different ways of traversing the environment in response to obstacles, friends, opponents, neutrals, and preferred paths (e.g., towards a waypoint). Unless otherwise stated, hereafter the term ‘behaviour’ refers to the agile variety.

ABDs such as MANA (Lauren and Stephen 2002), ‘**I**rreducible **S**emi-**A**utonomous **A**daptive **C**ombat’ [*ISAAC*] (Ilachinsky 1997), ‘**E**nhanced **I**SAAC **N**eural **S**imulation **T**oolkit’ [*ELNSTein*] (Ilachinsky 2000, 2003) and ‘**W**arfare **I**ntelligent **S**ystem for the **D**ynamic **O**ptimization of **M**issions’ [*WISDOM*] (Yang et al. 2004, 2005) 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 realism⁸;
- They provide controls to vary the behaviour of combatants through in-built agent parameters including personality, weaponry and sensor capability;

8. 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.

- 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 genetic algorithm, allowing one to explore a large parameter space of behavioural possibilities for an optimal solution.

Genetic Algorithms in MANA

The MANA genetic algorithm has been employed by several authors to investigate the effects of behaviour on combat (Luscombe and Mitchard 2003, McIntosh and Lauren 2006, Parunak et al. 2006). Simply stated, a genetic algorithm (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 and allows one to rank them from best to worst⁹. GAs maintain a population of candidate solutions which evolve over successive generations in response to the fitness function. 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*:

1. 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., fit-

9. In what follows, high fitness solutions are assumed to be better, to within the limitations of measurement, than low fitness solutions.

- ness-proportional pairing), often to the exclusion of unfit pairs. These pairs become the ‘parents’.
2. Crossover – ‘child’ solutions are generated by randomly combining genes of ‘parents’. The children represent new, possibly unexplored solutions.
 3. 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 (e.g., imagine the potential impact of crossover on chromosomes in a problem having two equally good, paired solutions that are polar opposites of one another). 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 know, minimally, the essence of a good solution in addition to how it will be measured.

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’. In other words, 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, chro-

mosome 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 **final** 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 once opponents of a given type have been detected.

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 repre-

sentation and structural scoping (i.e., defining a minimal list of parameters, anticipating architectural ‘building-blocks’, or recognizing ‘genetic drift’ (Rogers et al. 1999), which amounts to gene evolution in the absence of fitness criteria).

Although numerous involved methods are available for evaluating and monitoring the performance of GAs (Bornholdt 1998, Goldberg 1989a,b, Holland 1975, Prügel-Bennet and Shapiro 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 and Vose 1992) requires knowledge of chromosome transition probabilities). Nevertheless, even a small subset of these methods is enough to infer important characteristics about the dynamical progression of GA runs.

In the section that follows, GA progression was evaluated and monitored by following the genes of the fittest member of the population as the generations proceed. Additionally, mean population fitness was tracked and the effects of varying attribute settings of the genetic operators themselves examined, including: fitness criteria (MOEs), gene set, population size, number of repetitions, mutation rate/size and the use of trigger states. For the sake of brevity, however, non-essential details are omitted.

Scenario

As alluded to in the *Introduction*, the purpose of this section was two-fold 1) to demonstrate application of the GA in MANA, and 2) 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 complexity measures appropriate to the given circumstances were chosen and used to plan and execute a challenging (virtual) mission. Three simulations were conducted: *Sim I*, *II*, and *III*. In these simulations focus was on the

‘control’ aspect of C2. 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 of BLUE was to reach a waypoint ‘B’ from their starting position ‘A’ on a billiard table battlefield. It was a tough scenario for the BLUE force—conditions were tailored to make their situation extremely difficult, with the hope that behaviour would emerge that significantly improved chances of success.

In *Sim I*, the optimal BLUE force strategy was found without complexity SA and without using triggers. The scenario was treated as a single obstacle for BLUE to overcome—meaning that only one set of optimal behavioural parameters were sought. 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 was found utilizing MANA triggers (state changes altering behaviour) aided via analysis of CMOEs monitoring the system dynamics. 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, and the fitness for each potential solution (chromosome) was determined by averaging over 10 runs (using the multi-run feature). 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 CMOE 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 behavioural profiles. 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) on or off.¹⁰

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 kills per shot) and slightly longer range sensors (25 versus 20 distance units). The site defenders remained proximal to position B, moving randomly within a confined area. The setup is displayed in Figure 4. The measure of success (fitness) was defined as the number of BLUE combatants within the first cluster of agents to reach waypoint B under a time constraint of 500 steps. Baseline MANA settings and ranges for evolved BLUE traits for this simulation can be found in Appendix B.

A single set of personality traits was evolved for the BLUE team that optimizes the situation described above. 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 an overly large set:

- Attraction/repulsion to friends (range -100 to 100)
- Attraction/repulsion to enemies detected personally (range -100 to 100)
- Attraction/repulsion to enemies detected by others [SA] (range -100 to 100)
- Attraction to waypoint B (range 50 to 100)

10. 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.

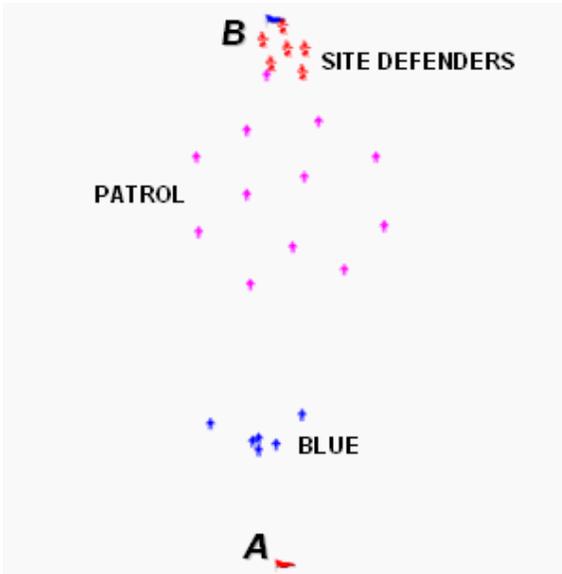


Figure 4. MANA Scenario. BLUE force consisted of 1 squad with 6 members heading towards the blue flag (top). RED force consisted of 2 squads: a 6 member squad defending the blue flag and a 12 member dispersed patrol heading towards the RED flag (bottom). The RED patrol is shown in magenta. Battlefield dimensions are 200x200.

Simulations indicated that the evolved behaviours had at most an 18% success rate (1000 repetitions). 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 of defeating the RED site defender squad using optimized tactics, which represents an approximate theoretical upper limit in this scenario.

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 listed in the *Discussion* section (below).

Sim II) Evolving Combatant Behaviour in Stages

It is easy to see that some benefit could be gained by partitioning *Sim I* into different stages. It is of interest to explore how the CMOEs might help to understand how best to partition the scenario. To achieve as complete a picture as possible, complexity is viewed from various perspectives and scales. Note that in the figures that follow, the two squads of the RED force are combined unless otherwise stated. To begin with, the fractal dimension is plotted for BLUE and RED forces at two different scales: 1) entire battlefield and 2) minimal containment (see Figure 5). The latter restricts the evaluation space to a minimal, axis-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.

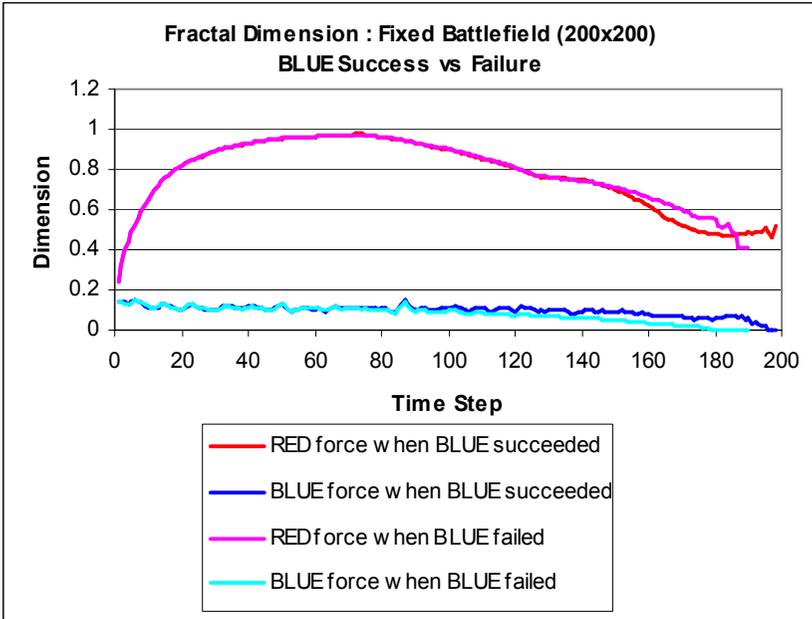
The fractal dimension plots in Figure 5 are time step averaged over many simulations—168 for the case of BLUE success and 832 for BLUE failure (1000 total simulations conducted)¹¹. At battlefield scale (Figure 5a), the RED force starts off tightly clustered (dot-like) as evidenced by the low fractal dimension. The dimension then

11. Note that the Legend for Figure 5 applies to subsequent figures, where appropriate.

increases as they spread out (near line-like) 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 5b, 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 *inside* 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.

Spatial entropy for the minimal containment case was also computed. It displays nearly identical behaviour to the corresponding fractal dimension plot (not shown here – refer to Appendix A).

a)



b)

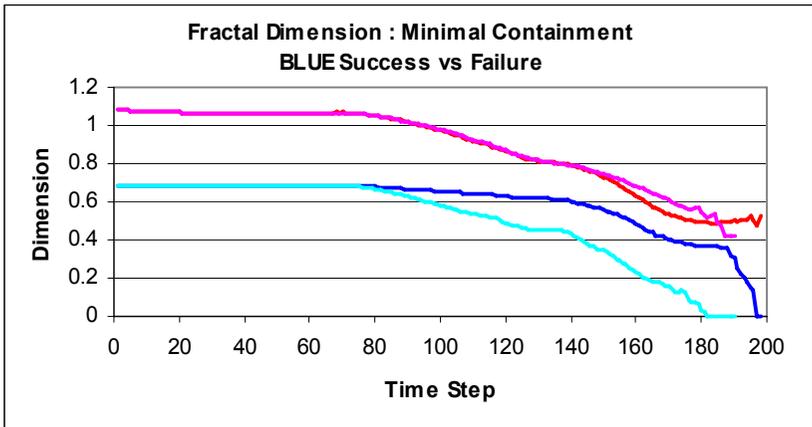


Figure 5. Fractal Dimension from a) battlefield and b) local perspectives.

Figure 5b suggests that perhaps something important (spatially) is happening to BLUE around time step 70, and that something important happens to RED around time step 150. The ‘important’ events merely correspond to the timing of significant changes in the combat strength of the various forces; the computed fractal dimension and spatial entropy are more-or-less counting the surviving number of agents for each force following an encounter. So, although the two quantities 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. To shed some light on the complex nature of the dynamics, further CMOEs must be explored.

Recall that symmetry is a spatial measure taking into account entropy with respect to a kernel of two-dimensional symmetries. Symmetry plots for the BLUE and RED forces are shown in Figure 6. Symmetry was computed using minimal containment.

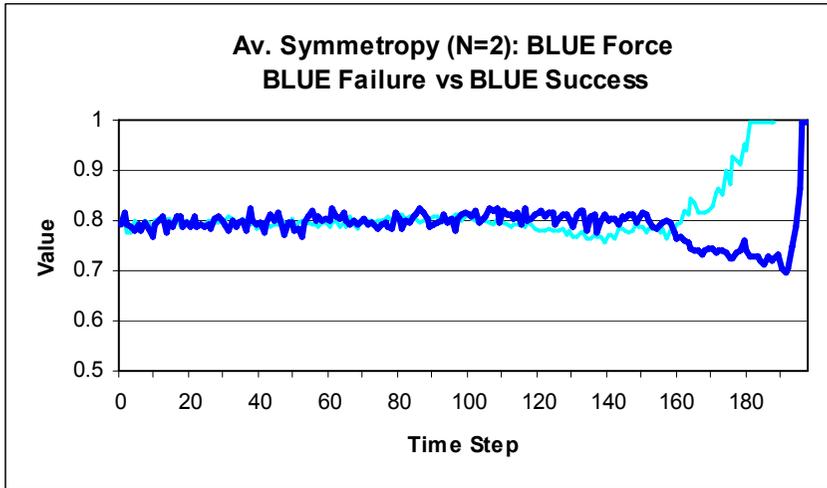
In Figure 6a, it can be seen that BLUE symmetry paths diverge around time step 100 (in the midst of the battle with the RED patrol). At this point the successful branch maintains its course while the unsuccessful branch falls below it. Then near time step 153 the successful curve makes a gradual downturn, crossing the unsuccessful curve at ~ 160 . Thus BLUE success favours steadiness with regard to the distribution of symmetries during this interval, owing to the fact that BLUE attrition coupled to formation tendencies causes the squad to spatially reorganize itself. Afterwards, the successful branch proceeds downwards while the unsuccessful branch shoots upwards. Note that the steep climb in BLUE symmetry after time step ~ 160 (failure case) seems to be weighted by the timing of the elimination of the BLUE force. 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 7a). This decrease in double symmetry (P_d in the figure) is accompanied by a convergence of the other symmetries (P_v , P_h , and P_c) 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 7b).

For RED symmetry (Figure 6b), 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 symmetry curve first rises above the unsuccessful curve, and then falls below it.

Figure 8 shows the symmetry 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 the BLUE symmetry 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., perhaps by this time it can be determined whether BLUE had a good or bad encounter with the RED patrol). The remaining CMOEs are now examined to provide additional insight.

a)



b)

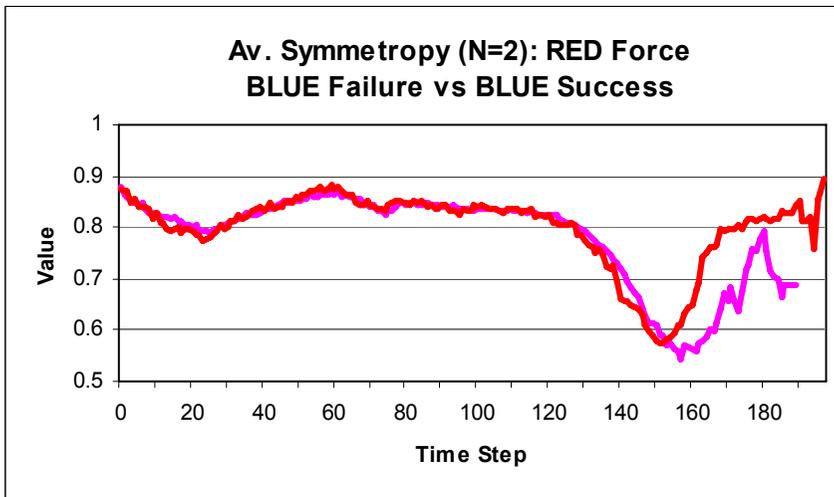
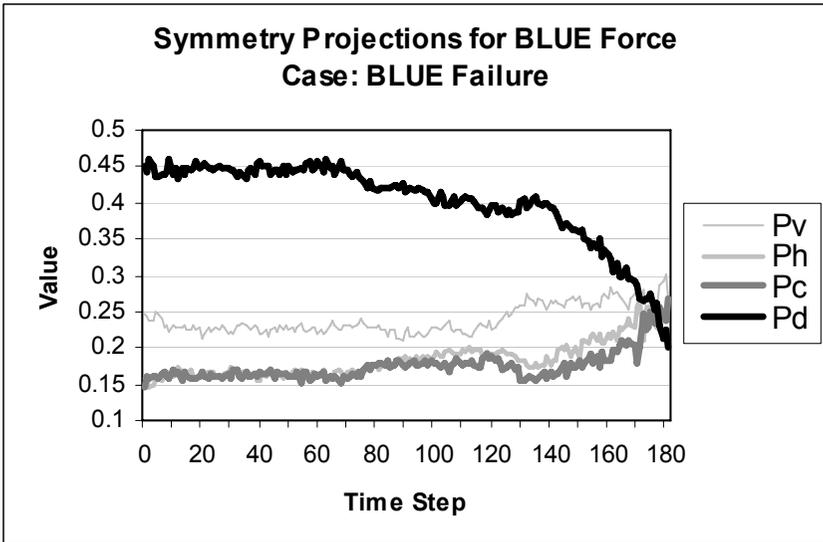


Figure 6. Symmetry for a) BLUE and b) RED forces (Legend of Figure 5 applies).

a)



b)

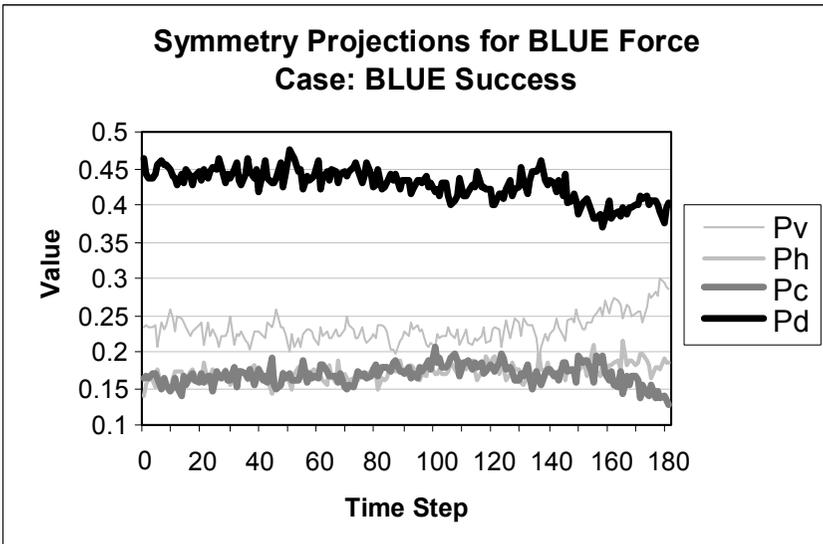


Figure 7. BLUE force symmetry projections.

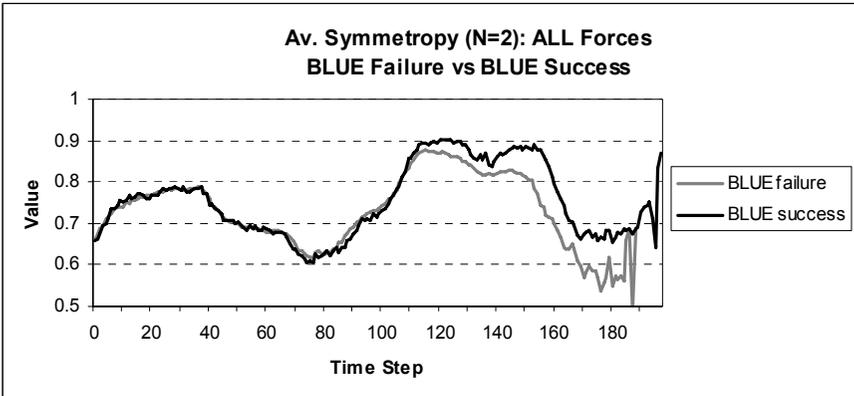


Figure 8. Combined forces symmetry.

Figure 9 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 point is that the plots for SSP in the X and Y directions are quite different. This could be an indication of self-affinity (see Appendix A). In the Y direction (Figure 9b), the crucial time step occurs around step 75—the instant that 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 fails when motion at this juncture tends toward randomness or even anti-correlation.

In the X direction (Figure 9a), time step 135 is where the branching occurs. In opposition to Y, the better path for BLUE is one of anti-correlated motion. Together, the two seem to suggest that a higher

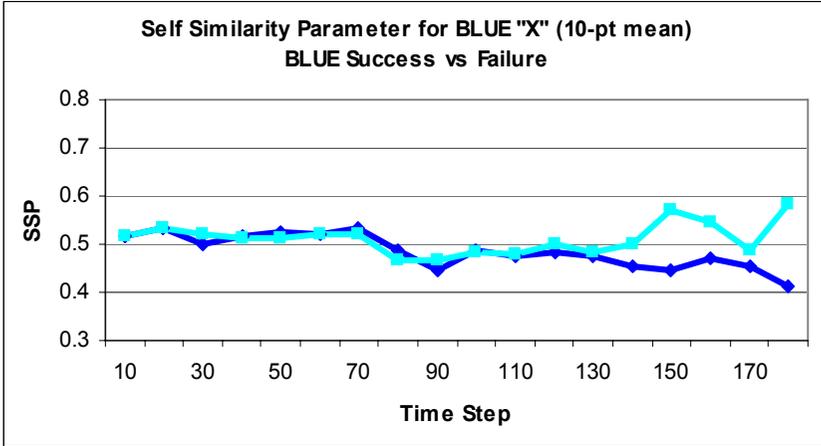
degree of self-affinity is beneficial (i.e., the fractal scales differently in X and Y).

The Hurst coefficients for BLUE X and Y were also computed, showing similar patterns to those of SSP. Appendix A contains the Hurst coefficient plot for the Y direction.

For comparison, Figure 10 shows the SSP for the RED site defender squad. This squad randomly moves about a small area proximal to waypoint B until the BLUE force arrives, at which time BLUE is pursued. It is interesting to note that the X and Y correlations are nearly identical to one another, suggesting non-self-affinity.

The branching of the Figure 10 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 (combat strength). Although the timing is roughly coincident with symmetry 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 likened to a precursor.

a)



b)

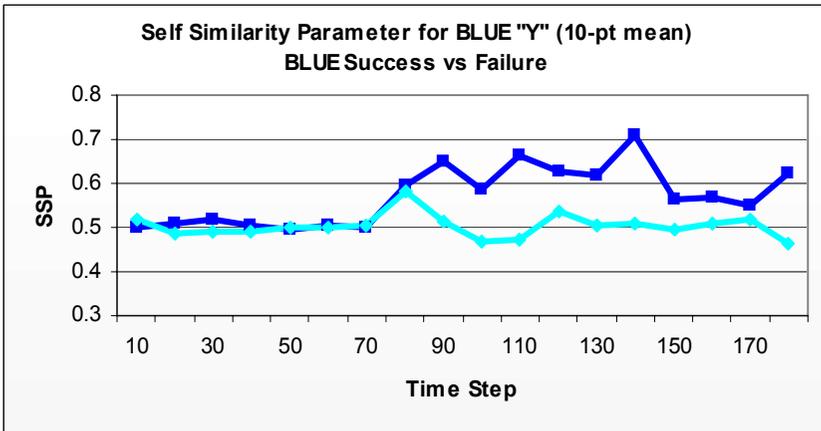
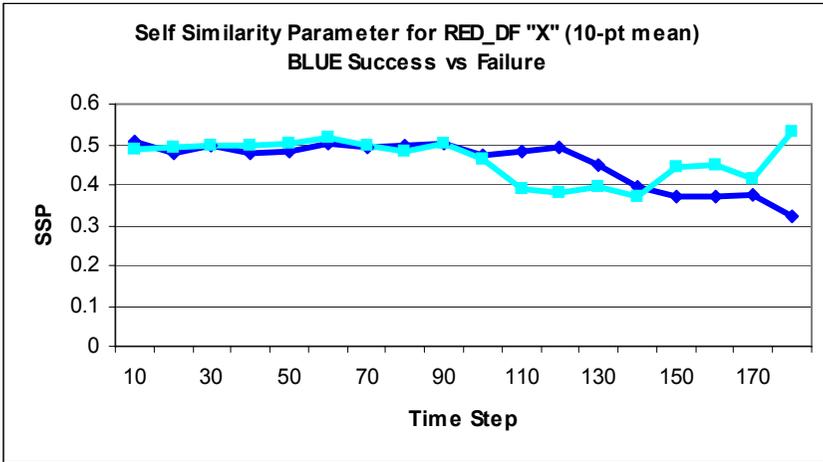


Figure 9. Self-similarity parameter for BLUE force (Legend of Figure 5 applies).

a)



b)

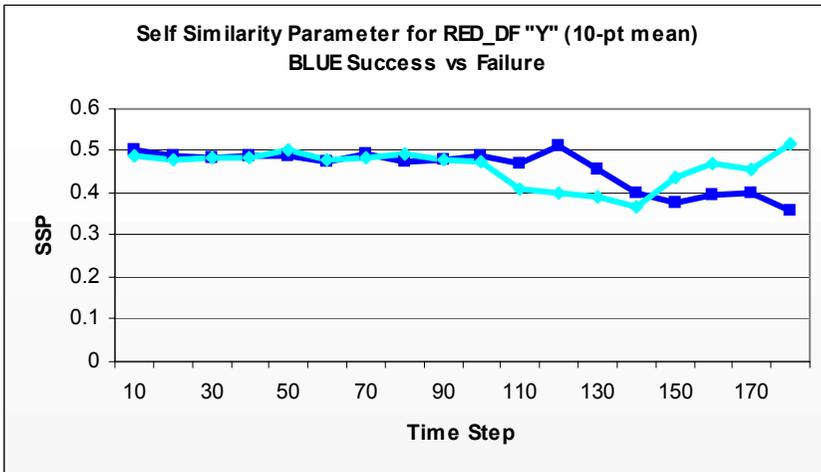


Figure 10. Self-similarity parameter for RED 'site defender' force (Legend of Figure 5 applies).

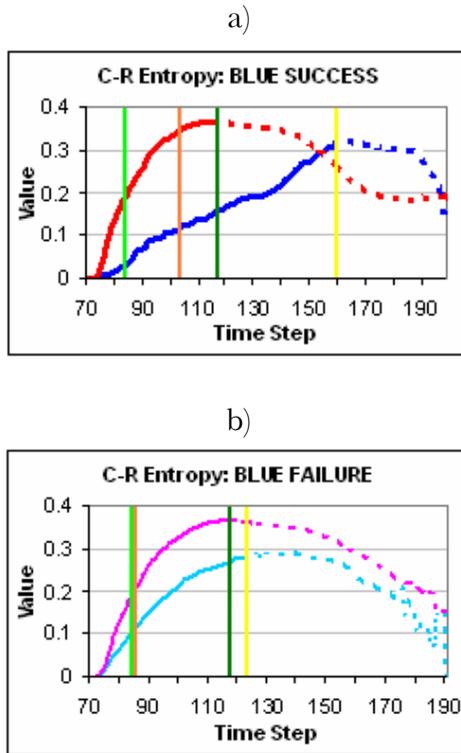


Figure 11. CR-entropy when a) BLUE was successful and b) BLUE failed (Legend of Figure 5 applies).

The final CMOE, CR-entropy, is evaluated beginning at the first sign of attrition (around time step 70 – Figure 11). The disintegration phase (as noted previously and in Appendix A) for each case is represented by a dashed curve. In many ways, the indications of CR-entropy at particular time steps seem to underlie the patterns observed in other measures. In Figure 11, the curves when BLUE succeeded (Figure 11a) and failed (Figure 11b) are divided into intervals corresponding to CR-entropy states (see Appendix A). The states provide a rough indication of how the battle is progressing. Prior to the light green marker (\sim time step 83), 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 on BLUEs' successful curve (Figure 11a), the advantage is maintained

for around 20 time steps, whereas on their failure curve (Figure 11b) 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 of BLUEs' curves (\sim time step 113), transitioning to one of BLUE advantage. This advantage is due to the disintegration of the RED force, which occurs simultaneously for both cases (the 0.37 threshold has been passed). On BLUEs' failure curve (Figure 11b), the advantage is once again short-lived and transition to the full disintegration phase follows (\sim time step 123), whereas on their success curve the transition is delayed until approximately time step 160 (yellow marker).

Note also the standard deviation (or spread) of BLUE's CR-entropy data was highly variable throughout, whereas RED was tightly controlled, reducing to almost nil near the onset of REDs' disintegration phase (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 perhaps in many simulations BLUE circumstances fell somewhere in the midst of the two paths.

It could 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 pivotal (perhaps critical) and distinct dynamical events in the system. Since time step 110 falls near the end of the first encounter, 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. Since the degrees of freedom required to characterize beneficial behaviour rules at this juncture fall outside the bounds of the original behavioural parameter set, this aspect was not investigated, despite its obvious potential.

We proceed by partitioning the scenario according to the intuitively obvious transition events (i.e., by encounter type) which are further 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 guesses of what the optimal GA behaviour settings might be for the BLUE force. The point is to help assess the added value of employing a GA in this scenario, as opposed to simple (human) reasoning. Both guesses 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*.

Guess 1: 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)

Guess 2: 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).

Personality settings for BLUE while running the GA are provided in Section B 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 significantly over the solutions without triggers and slightly over guesses made 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 1, 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 in motion, between the reaction to personal versus squad (SA) detections of enemies. In retrospect, the tactic for confronting the RED site

defenders could have been anticipated as it merely quantifies a tendency to attack a superior (single) foe as a group rather than individually.

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 1 for reference: 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’. Standard errors are included.

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 were 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 would be expected.

Table 1. 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

Sim III) Real-time Response to CMOEs for Tactical Advantage

In the previous section, it was shown that the CMOEs were able to distinguish between successful and unsuccessful behaviour. Therefore, it follows that actual knowledge of real-time complexity might be beneficial to improve the likelihood of success. The next 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. Various C2 options were exercised by monitoring and responding to the temporal evolution of a chosen CMOE.

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) via entropic heterogeneity or otherwise 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 the Section *Precursors and SOC*, above). 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—the pre-cognitive signatures sought are definitely not evident in the averaged results for this scenario.

On the other hand, computing the fractal dimension and symmetry 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. Use of this quantity is further cautioned by its association with simple casualty counting in the simulations examined.

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 waypoint B). Nonetheless, as stated in Appendix A, computing the SSP or Hurst coefficient is data intensive and the real-time scenario is not likely to be capable of producing the required data support (several hundreds to several thousands of data points).

Therefore, since symmetry alone is not overly constrained in the case of sparse data, it is the only measure investigated herein as a prospective CMOE for real-time determination of the encounter-type in this situation. Like the Hurst coefficient and SSP, it also requires at least a slight SA advantage to be particularly useful. A symmetry 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 symmetry patterns of detection preceding an encounter should provide a reasonably accurate cue about what to expect, and

will demonstrate the feasibility of real-time response to one facet of complexity in the system.

To begin with, a new, but similar challenge for BLUE is designed (Figure 12). 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. To get to the waypoint, they expect multiple encounters with a RED force similar to the RED patrol above (however, in this case, it is more of an occupying force than a patrol). 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 not 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*).

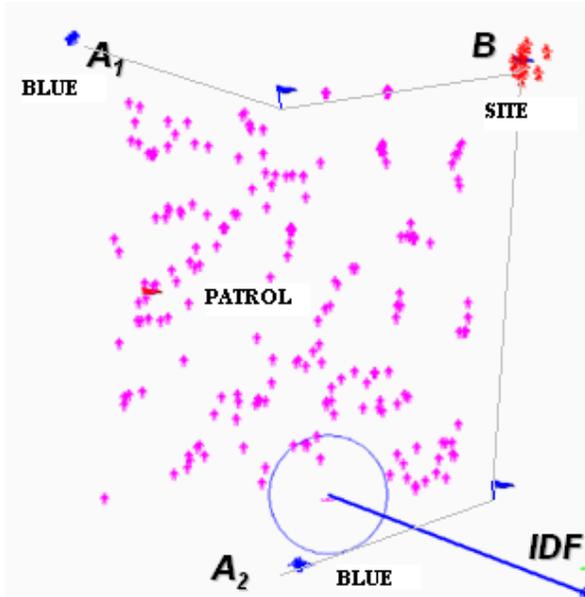


Figure 12. MANA scenario for Sim III.

The first potential symmetropic signature investigated was the mean (or average) force symmetry. 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 symmetry exponent q (grid matrix is $2^N \times 2^N$), and the spread is high enough to blur any distinctiveness¹². The mean symmetry data are presented in Table 2 (format is ‘value, spread’). 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

12. 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 2 because 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.

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 2 that real-time use of the computed symmetry means is of no value in this case.

There is, however, another option worth exploring. The detection data can be separated into distinct symmetry ‘modes’. These modes are a reflection of commonly encountered patterns in the symmetry matrix that characterize entropic heterogeneity¹³. In Figure 13, the frequency of symmetry modes is shown for the two RED encounter 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 symmetry 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 2. Mean symmetries of encounter types for *Sim III* mission reference.

RED Force Constituent	Force Strength	N	Symmetry Of Entire Force	Symmetry 3 Detections	Symmetry 5 Detections
Site Defenders “SITE”	10	2	0.81 , 0.12	0.86 , 0.05	0.80 , 0.11
		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 “PATROL”	200	2	0.82 , 0.04	0.86 , 0.04	0.82 , 0.12
		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

13. Symmetry ‘modes’, in this context, refer to recurring measured symmetry values. A mode roughly indicates that some particular combination of patterns seems to occur in the system repeatedly over time.

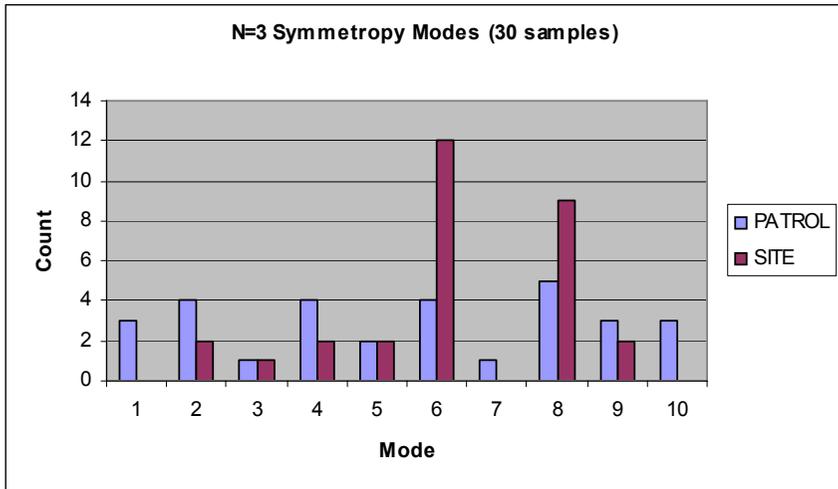


Figure 13. Symmetry ‘modes’ of the RED force.

Now it is time to capitalize on the distribution of symmetry 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 good match to the sought distribution, for example a preponderance of Mode 6 and, to a lesser extent, Mode 8—above what would normally be expected when encountering a PATROL. Then limit 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 switch into the trigger state found in *Sim II* to be most successful against this group. A level of risk tolerance could be fixed before the simulation, essentially defining the cutoff 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 cutoff, the more certain BLUE is that the target site defenders have been correctly identified.

However, due to the sparseness of detection data in this scenario, a more conservative approach was taken. By default the IDF is set on, then simply turned off whenever Mode 6 appears. All other detections are ignored (including Mode 8)¹⁴. Note that this slightly magnifies BLUEs' exposure to the risk that the IDF might be turned off too soon. To minimize this exposure three steps were taken: 1) a minimum symmetry bounding box length, equal to the BLUE detector range of 50, was introduced—this helps to ensure that the symmetry signature is spatially no smaller than the observed scale, 2) the initial start-up of the scenario was not processed (first 500 time steps) 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 MOE for a run was defined as the number of RED site defenders alive at the time of the earliest correct Mode 6 discovery¹⁵. The overall mission was deemed 'completed' if BLUE reached the waypoint B. The results of the simulations and default runs are displayed in Table 3. '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. 'Random' runs (10) employ the same set of trigger states, except that in this case the 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. The maximum number of time steps for any run was set to 2000.

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 difference between the Default and CMOE-triggered simulations (num-

14. Note that this 'switch' had to be implemented manually, since MANA does not have the capability to trigger a state change based on CMOE values.

15. When there were no SITE detections by the end of the run (e.g., simulation #6 in Table 3 in addition to all 'Default' and 'Random' cases), ten subtract the number of RED site defenders killed by the IDF was substituted for the numerator of the MOE.

bered 1-10) lies in the value of the MOE. The MOE was improved upon by over 12-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.

Table 3. 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
<i>SimIII Avgs</i>	<i>1030 (52%)</i>	<i>35%</i>	<i>3.3%</i>	<i>5.1 /10 (51%)</i>	<i>100%</i>
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 comrades 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.

Discussion

The simulations conducted in the previous section revealed several interesting features concerning the use of a GA to help define behaviour patterns in combat operations. It also hints at possible advantages of endowing agents with an awareness of complexity in a combat system when the system is viewed as 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 a 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, 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*, CMOEs based on BLUE SA were shown to have distinctive features for different encounter-types (RED patrol or site defenders), inducing a partitioning of the behavioural pattern search space into two distinct groups—one for each encounter type. The result was a significant improvement in mission success rates. Regarding the GA used in *Sim II*, a surprising, unanticipated result was generated. That is, the opposing movement pattern that BLUE evolved for dealing with personal versus squad detection of enemies. 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.

The methodology applied in *Sim I* mapped out a framework for running, analyzing and interpreting genetic algorithms used in MANA to optimize agent behaviour and tactics in a difficult scenario. The main lessons learned from this simulation were as follows:

1. The Multi-runs option should be used when running a GA to buffer against the effects of randomness in the outcomes of a conflict (10 were used in this study).
2. 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 a practitioner.
3. Testing for genetic drift and evolving the system using a high mutation rate may help to eliminate extraneous variables (genes) thus improving performance and simplifying the interpretation.

In *Sim II*, the concept of a phase transition (as per SOC) was used as a framework to improve upon the findings of *Sim I*. These concepts were not, however, strictly demonstrated to be applicable to the specific scenarios investigated. In particular, the presence or absence of SOC was not demonstrated in any of the simulations. Nevertheless the *idea* of approaching possible 'criticality', adapting to a 'phase transition' and looking for 'precursors' fit well as an approach to framing the improvement of mission success. Intuitively, it parallels a sensible and careful approach to optimal mission planning. Thus, despite the lack of rigor, it was demonstrated that the complexity indicators introduced in the section *Complex Systems Measures of Effectiveness* can display recognizable and distinctive patterns during encounters, and that these patterns can be leveraged to help partition the application of the GA into streams that deal with different

types of encounters separately. It was also shown that dividing the solution up in this way can yield significantly better mission success rates compared to undivided solutions. The partitioning was achieved using MANA triggers that caused state changes in agents' behavioural profiles dependent upon the current situation and agent's intent.

Note, however, that the results of *Sim II* **do not** indicate that partitioning chosen without the aid of CMOEs coupled with behaviour chosen without using the GA would have fared much worse. The CMOE-driven selection of a trigger pointed to the 'most obvious one', and (educated) guesses at optimal behaviour profiles performed nearly as well as GA-optimized ones. Perhaps the scenario was too simple, or the force sizes too small, to demonstrate any practical advantages gained in this scenario. In the end, the CMOEs really only helped by providing insight into the underlying dynamics of the *Sim II* runs, and the GA only helped by presenting an alternative to an already good option. It is worth noting though that a potentially useful possibility for catching the RED site defenders off-guard was exposed via the SSP, however the available degrees of freedom did not allow BLUE to take advantage of the opportunity.

In *Sim III* it was found that precursory-like signatures derived from CMOEs could constitute an early warning in real-time via entropic heterogeneity, hinting at the nature of an imminent near-future encounter. These precursors, recognized via limited situational awareness, were successfully used by the BLUE force to distinguish between RED encounter types and to call off IDF support when appropriate. Furthermore, the use of precursors was combined with state changes and partitioned, evolved behaviour as demonstrated in *Sim II*. Although mission completion rates were not radically improved upon (all missions were completed), those that were completed 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 3). 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 CMOEs 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 CMOEs can provide unique capabilities relevant to C2 in the general case.

Conclusions

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: the fractal dimension, Shannon entropy (via CR-entropy and spatial entropy), the Hurst coefficient, the self-similarity parameter and symmetry. All but one of these measures are directly based on the spatial dynamical properties of the system rather than on attrition. Therefore they are better suited to capturing certain aspects of the complexity of combat than attrition-based measures. 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.

The combat scenario faced by the BLUE force in this paper presented a difficult challenge to overcome (under default settings, the success rate was negligible). Mission success and agent response capabilities were generally enhanced by adapting agent behaviour based on the knowledge of complexity in the system. Factors that contributed to the improvements were 1) how to partition the system on the basis of various entropies and long-term correlations, and 2) the early determination of enemy type based on an entropy/symmetry measure (symmetry).

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 support.

Future work

Improving understanding of the progression of self-organization within a combat CAS is a topic of future research interest.

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.
- Carvalho Rodrigues, F. 1989. A proposed entropy measure for assessing combat degradation. *J. Opl. Res. Soc* 40, 8: 789-793,
- Dobias, P. 2008a. Complexity-based Assessment in Crowd Confrontation Modelling. *Journal of Battlefield Technology* 11, 2: 1-8.
- Dobias, P. 2008b. Self-organized criticality in asymmetric warfare. *Fractals* (in press).
- Dobias, P., Wanliss, J.A., Samson, J.C. 2006. Non-linear stability of the near-Earth plasma sheet during substorms: 9 February 1995 event. *Can. J. Phys* 84: 1029.
- Dockery, J.T. and Woodcock, A.E.R. 1993. *The Military Landscape: Mathematical Models of Combat*. Cambridge, England: Woodhead Publishing Limited.
- Feder, J. 1988. *Fractals*. Plenum Press.
- 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, <http://circ.ahajournals.org/cgi/content/full/101/23/e215>.
- Hofmann, H.W., Schnurer, R., Tolk, A. 1995. On the Impact of Stochastic Modeling in a Rule Oriented Combat Simulation Model on Division/Corps Level. *NATO AC/243 (Panel 7) Symposium on Uncertainty, NATO Shape Technical Center, The Hague*.
- Holland, J.H. 1975. *Adaptation in Natural and artificial systems*. Ann Arbor, MI: University of Michigan Press.
- Ilachinski, A. 1997. Irreducible Semi-Autonomous Adaptive Combat (ISAAC): An Artificial-Life Approach to Land Warfare. *Center for Naval Analyses Research Memorandum CRM: 97-61*.
- Ilachinski, A. 2000. EINSTEIN: An Artificial-Life “Laboratory” for Exploring Self-Organized, Emergent Behavior in Land Combat. *Center for Naval Analyses Research Memorandum CRM D0002239.A1*. Alexandria.
- 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*. Singapore: World Scientific.
- Jones, C.L., Lonergan, G.T., Mainwaring, D.E. 1996. Wavelet Packet Computation of the Hurst Exponent. *Journal of Physics A: Mathematical and General* 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. The Principle of Concentration *Engineering* 98: 422-423.
- Lauren, M.K. 1999. Characterising the Difference between Complex Adaptive and Conventional Combat Models. *DOTSE Report* 169, NR 1345.
- Lauren, M.K. 2001. Fractal Methods Applied to Describe Cellular Automaton Combat Models. *Fractals* 9, 2: 177-184.
- Lauren, M.K. 2000. Firepower Concentration in Cellular Automata Models - An Alternative to the Lanchester Approach. *DOTSE Report* 172, NR 1350.
- Lauren, M.K., Stephen, R.T. 2002. Map-aware Non-uniform Automata (MANA): a New Zealand Approach to Scenario Modelling. *Journal of Battlefield Technology* 5, 1: 27ff.
- Lauren, M.K. 2003. On the temporal distribution of fatalities and determination of medical logistic requirements. *DTA Report* 187, NR1377, ISSN 1175-6594.
- Lloyd, S. and Pagels, H. 1988. Complexity as Thermodynamic Depth. *Annals of Physics* 188: 186-213
- 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.
- Nanjo, K., Nagahama, H., Yodogawa, E. 2000. Symmetry Properties of Spatial Distribution of Microfracturing in Rock *Forma* 15: 95-101.
- Nanjo, K., Nagahama, H., Yodogawa, E. 2001. Symmetry and self-organized criticality. *Forma* 16: 213-224.

- Nanjo, K., Nagahama, H., Yodogawa, E. 2005. Symmetry 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., S. Havlin, H.E. Stanley, A.L. Goldberger. 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., 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.
- Rogers, A. and Prügel-Bennett, A. 1998. Genetic Drift in Genetic Algorithm Selection Schemes. *IEEE Transactions on Evolutionary Computation* 3, 4:298-303.
- Samson, J. C., R. Rankin, R., V. T. Tikhonchuk, V.T. 2003. Optical signatures of auroral arcs produced by field line resonances; comparison with satellite observations and modeling. *Annales Geophysicae* 21: 933.

- Shalizi, C.R., Shalizi, K.L., Haslinger, R. 2004. Quantifying Self-Organization with Optimal Predictors *Phys. Rev. Lett.* 93:118701.
- Shapiro, J., Prügel-Bennet, A., Rattray, M. 1994. A Statistical Mechanical Formulation of the Dynamics of Genetic Algorithms. *Lecture Notes in Computer Science*, 865, special edition on Evolutionary Computing edited by T. C. Fogarty.
- Shannon, C.E., Weaver, W. 1949. *The mathematical theory of communication*. Urbana, Illinois: Univ. Illinois Press.
- Tolk, A. 1995. Zur Reduktion struktureller Varianzen - Einsatz von künstlicher Intelligenz in geschlossenen Gefechtssimulationsmodellen. (publ. in German) *Pro Universitate Verlag*, Sinzheim.
- Voronkov, I. O., E. F. Donovan, P. Dobias, V. I. Prosolin, M. Jankowska, J. C. 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. *Mathematische Annalen* 69: 331-371.
- Wanliss, J.A. 2004 Nonlinear variability of SYM-H over two solar cycles *Earth, Planets, Space*, 56: 13-16.
- 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.
- Yang, A., Abbass, H.A., Sarker, R. 2005. Evolving agents for network centric warfare. In *Second Workshop on Military and Security Applications of Evolutionary Computation*, Washington D.C.: ACM Press.

Appendix A

CMOEs for the Practitioner

The following potential CMOEs are described in detail:

- Carvalho-Rodrigues Entropy;
- Spatial Entropy;
- Fractal Dimension;
- Hurst Coefficient;
- Self-similarity Parameter;
- Symmetry.

Some other complex systems factors of interest that were not pursued in this analysis include the Kolmogorov Complexity (Kolmogorov 1965), Statistical Complexity (Shalizi et al. 2004), Depth (Lloyd and Pagels 1988), Effective Complexity and the related term Total Information (Gell-Mann and Lloyd 1996).

Carvalho-Rodrigues Entropy

Carvalho-Rodrigues (CR) proposed an attrition-based definition of combat entropy for the i -th force (i being RED or BLUE) of the form (Carvalho-Rodrigues 1989, Ilachinski 2004)

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

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 ($N_i = N_{i0} - C_i$). 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). Up until this point, high CR-entropy translates to a more precarious position for the force in

question. The maximum value 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 below are based on dividing the range of C_i/N_i for each force into three regions: 1) $C_i/N_i < 0.185$ – less than half way to the disintegration point, 2) $0.185 = C_i/N_i < 0.37$ – more than half way to the disintegration point, and 3) $C_i/N_i \geq 0.37$ (disintegration – past the half-way point).

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:

a. Moderate:

- i. $C_{BLUE}/N_{BLUE} < 0.185; 0.185 \leq C_{RED}/N_{RED} < 0.37$
- ii. $0.185 \leq C_{BLUE}/N_{BLUE} < 0.37; C_{RED}/N_{RED} \geq 0.37$

b. High: $C_{BLUE}/N_{BLUE} < 0.185, C_{RED}/N_{RED} \geq 0.37$

2. Advantage RED: Analogous to above, switching RED and BLUE subscripts.
3. Balanced: Neither RED nor BLUE has a notable advantage. RED and BLUE CR-entropies are comparable.

Similarly, the difference $(C_{RED}/N_{RED} - C_{BLUE}/N_{BLUE})$ is also a relevant parameter to monitor. 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.

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.

Spatial Entropy

Ilachinski (2004) suggested a specific form of Shannon entropy based on the spatial distribution of soldiers. The computation and meaning of spatial entropy are somewhat akin to the fractal dimension computed via the box-counting technique (below). Figure A1, when compared to Figure 5b in the main text, shows that the two are nearly indistinguishable for the scenario examined.

To compute the spatial entropy, 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 are in the i^{th} sub-block, the probability of finding a soldier in that sub-block is $p_i(b) = N_i(b) / N$. Then Shannon entropy takes the form

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

The expression $1/(2 \ln(B/b))$ 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.

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

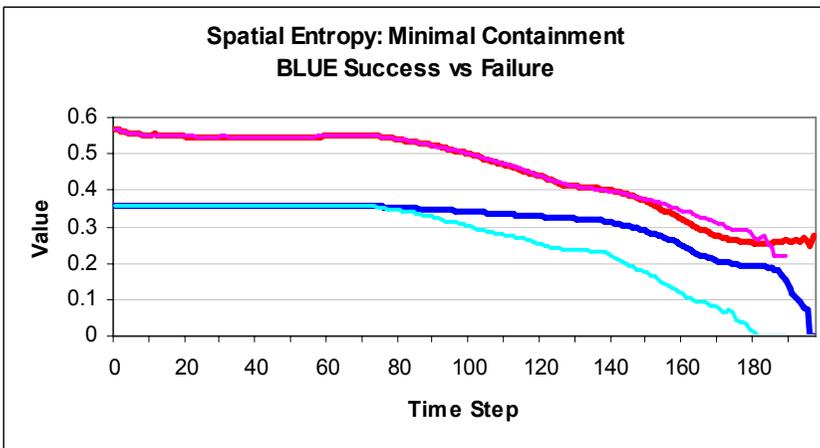


Figure A1. Spatial entropy – compare with Figure 5b (Figure 5 Legend applies).

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) (Ilachinski 2004). The most natural of many possible fractal dimensions to describe spatial dynamics of a combat seems to be the box-counting (or capacity) dimension D_F . It expresses the relationship between the size of a box ε , and the min-

imum number $N(\epsilon)$ of boxes needed to cover all the agents. Generally, the dependence is a power law:

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

In the expression above, L is the size of the battlefield. For agents uniformly distributed over a two dimensional (2D) battlefield, $D_F = 2$. Taking the logarithm of both sides of the equation for sufficiently small ϵ , a formula for D_F is obtained:

$$D_F = \lim_{\epsilon \rightarrow 0} \frac{\ln N(\epsilon)}{\ln L / \epsilon}$$

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

Hurst Coefficient

The interpretation of the Hurst coefficient bears a strong resemblance to that of the self-similarity parameter (below). The calculations for each also share similar features. In fact, for the scenario examined it is evident that they show nearly equivalent behaviour (compare Figure 9b (main text) with Figure A2 (below)).

Temporal and spatial correlations in agent velocity (speed and direction) are found via the Hurst coefficient. Such 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. 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. As H approaches 1, the RMSD becomes directly proportional to the number of steps, $L = \lambda N$. 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 2008a). 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.

A caveat needs to be included at this point. Due to their stochastic nature, the velocity correlations are relevant only for large numbers of data points (thousands and more). 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).

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

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 \mathcal{N} points (within a box) is given by:

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

where k ranges from 1 to \mathcal{N} .

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

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

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. Also, generally it is best to use values for

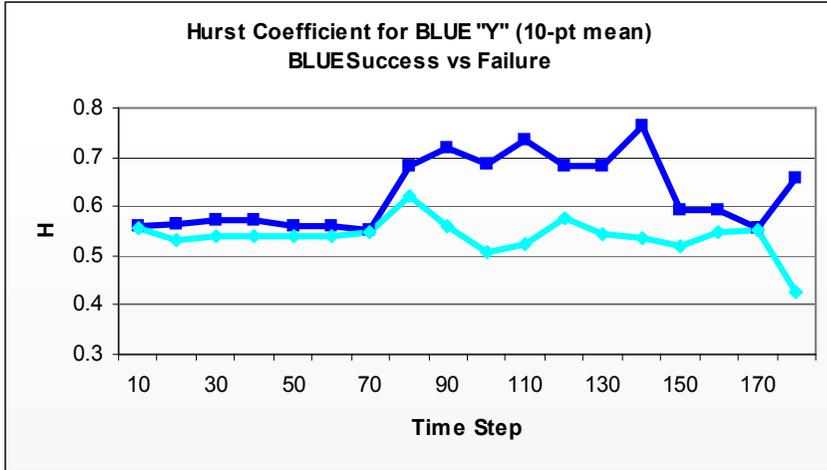


Figure A2. Hurst Coefficient for BLUE ‘Y’ – compare with Figure 9b (Legend of Figure 5 applies).

box sizes uniformly distributed in logarithmic space. This allows for a better fit of $\log R/S(n)$ as a function of $\log n$.

Self-Similarity Parameter

The self-similarity parameter (SSP) can be viewed as a measure of the ‘roughness’ of a time series (Peng *et al* 1995). Furthermore, it shares many interpretive properties of the Hurst coefficient (above). One advantage of using the self-similarity parameter over H is that it can be applied to a non-stationary time series. Now a time series is self-similar if the process $y(t)$ shares the same statistical properties as a properly rescaled process given by $a^\alpha y(t/a)$. α has the following interpretation (Goldberger *et al* 2000):

- $0 < \alpha < 0.5$: 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).

A method called detrended fluctuation analysis (DFA) is commonly used to calculate the SSP (Peng *et al* 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 can be briefly summarized as follows.

At first the entire data series is integrated and then divided into boxes of length n . The integration equation for a data series D_i of \mathcal{N} points is given by:

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

where k ranges from 1 to \mathcal{N} .

Afterwards, 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 $F(n)$ versus n is used to calculate a slope, which in turn provides the SSP. 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 $\log n$.

Symmetry

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

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

where $m, n = 0, 1, 2, \dots, M-1$, $x_{i,j}$ is the value of grey (e.g. “black” and “white” – i.e. 1 or 0) in the i^{th} row and the j^{th} column. $W_{m,n}$ is the two dimensional Walsh function defined as

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

In the above expression the function $b_k(i)$ denotes k^{th} bit in the binary representation of i . For instance, for a number $5 = (101)_2$ the values of b are $b_0(5) = 1$, $b_1(5) = 0$, and $b_2(5) = 1$. $b'_k(m)$ is a

transformed function for the binary representation of the number m . The transformation is defined as

$$b'_0(m) = b_0(m),$$

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

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 2001) (see Figure 2 in the Section *Complex Systems Measures of Effectiveness*). $W_{0,0}$ is the exception.

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

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

In the expression above, $(m,n)^{Sk}$ 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 for entropy $S = -(1/2) \sum_k P_k \log_2 P_k$ can be applied. The 1/2 factor serves to normalize the symmetry so that the maximum value is 1. The higher the pattern, the higher the symmetry. For a random pattern (randomly distributed black and white cells), the symmetry is 1.0^{16} (Nanjo 2001).

16. Note that in Nanjo (2001), the author does not use the normalization factor 0.5. Consequently, therein the maximum value for symmetry (equal to the symmetry of a random distribution) is 2.

Appendix B

MANA 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).

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

Sim II

The MANA settings for trigger states of the BLUE squad in *Sim II* 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

Sim III

The MANA settings for RED and BLUE squads in *Sim III* are shown in Table A.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