

Paper Number: 115 [Version 2]

Paper Title:

Controlling Edge Organizations: Exploiting Emergence

Author and point of contact:

Patrick Beautelement

Command and Intelligence Systems,
Room 119, Alan Turing Building,
QinetiQ Ltd,
St Andrews Rd, Malvern,
Worcestershire,
WR14 3PS, UK

Telephone: (+44) 1684 896057

Fax: (+44) 1684 894492

Email: pbeautelement@QinetiQ.com

Paper for:

2005 10th International Command and Control Research & Technology Symposium
McLean, Virginia. June 2005.

Submitted for Track / Stream:

Edge Organizations

Controlling Edge Organizations: Exploiting Emergence

© Copyright 2005 / QinetiQ

Abstract

During 2003, CCRP published "Power to the Edge" which described a new kind of organization, an 'Edge Organization' (EO), which would display exceptional agility. Documents published by the Office of Force Transformation raise questions about how dispersed assets in Edge-like force structures could be controlled, especially where mechanisms such as self-synchronization and self-organization are at work.

Solving these questions is vital if EOs are to be implemented and deployed effectively. For true agility to be displayed, EOs must be capable of supporting a continuous process of dynamic execution where many effects and behaviours will be manifested as emergent phenomena.

Hence, this paper presents an holistic approach to biologically-inspired 'control' for EOs in network-centric situations. The paper looks at the challenges in terms of the dynamic interactions between all the entities and the environment. Critically, it considers the 'run-time' properties of the artefacts, actors and interactions, as well as the dynamic adaptive mechanisms, as being the key focus of attention - as opposed to the static, design-time engineering of their parts. This paper considers approaches to exploiting the emergent properties of complex systems to influence and ensure the collective, adaptive and secure behaviour of EOs and offers ideas to the research community for discussion.

Structure of the Paper

Firstly, the paper will look at the military imperatives which trigger the need for research into emergent phenomena in Edge Organisations. The paper will then briefly consider the nature of emergence - indicating some of the mechanisms available for exploitation. Next, the paper will consider the tension between bottom-up emergence and top-down control and will use this analysis as a starting point to consider implementation and acquisition issues. Finally, the paper indicates areas for further research before providing a conclusion.

Background and Relevance to Command Agility

During 2003, CCRP published "Power to the Edge" [1 Alberts] which described a new kind of organization, an 'Edge Organization' (EO), which would display exceptional command agility. Recent documents published by the Office of Force Transformation¹ raise questions about how dispersed assets in Edge-like force structures could be controlled, especially where mechanisms such as self-synchronization and self-organization are at work.

Solving these questions is vital if EOs are to be implemented and deployed effectively. Analysis of military operations in Iraq [2 Storr] have shown a worrying trend towards over-emphasis on planning as a way of trying to mitigate uncertainty. In contrast, EOs support a continuous process of dynamic adaptation during execution of operations where effects are manifested as emergent phenomena² arising from interactions among the parts. Forces which are dispersed, such as networked and semi-autonomous ones, are complex adaptive systems [3, 4] and so they will inevitably display many emergent phenomena. There are two responses to this situation: eradicate all phenomena and treat them as undesirable or exploit positive phenomena as a force multiplier for EOs. This paper takes the latter position and will suggest ways in which we can positively harness these phenomena to our benefit (for example, to wield against an opponent or to mitigate our own vulnerabilities).

Why is this significant now? As recent events have shown, multi-national coalitions constitute an increasing proportion of military operations yet, despite our increasing familiarity with them, they continue to be a challenge. In addition to the problems of integrating single-service and Joint capabilities, the nature of coalition operations implies some need to rapidly configure diverse, incompatible 'come-as-you-are' systems into a cohesive whole. When coalition partners are familiar, doctrine, systems and procedures are aligned in advance. In reality, there are always uncertainties about exactly which capabilities will be provided by whom and about how the forces will be configured. Hence, coalition operations trigger the need for rapid on-the-fly responses and cannot be predicated on using pre-existing co-ordinated systems - instead, we need flexible approaches that allow capabilities to be assembled at 'run time'³ and the emergent properties of the interactions to be exploited.

However, coalition warfare is just one of a range of types of conflict with which EOs would have to content. These various types can be mapped into a number of 'challenge spaces'⁴ which represent the 'envelopes' into which EOs have to morph to achieve their aims vis-à-vis opponents. For EOs to display this agility, they must have a number of features such as those shown in Figure 1. Key to achieving the necessary agility is understanding how to tune

¹ "Transformation Trends - WolfPAC Distributed Operations Experiment", 7th Dec 2004.

² Emergent phenomena arise from local interactions among components and their environment, where phenomena persist over time and cannot be deduced by examining the components in their inactive state.

³ Design-time activities relate to acquisition etc, whereas run-time configuration takes place after deployment and provides, the ability to flexibly adapt to the changing military imperatives of each unique operation - as they occur.

⁴ Considered further in one of the other papers on the topic of Edge Organisations being offered by myself and my colleagues Anthony Alston and Lorraine Dodd.

interactions and facilitate mobility of the 'loci of power' which exert influence on the environment and the other actors. Part of this capability comes from being able to exploit emergent phenomena as a means to self-organise, 'swarm' and adapt to cope with uncertainty. This can only be achieved, for each situation, by enabling an appropriate balance between control and emergence. For, without that balance, as Kirsch says [5]:

EOs will be "unable to control the novelties they cannot prevent, and will be unable to generate the novelties they need. They fall victim to the change they cannot inhibit and the change they cannot induce".

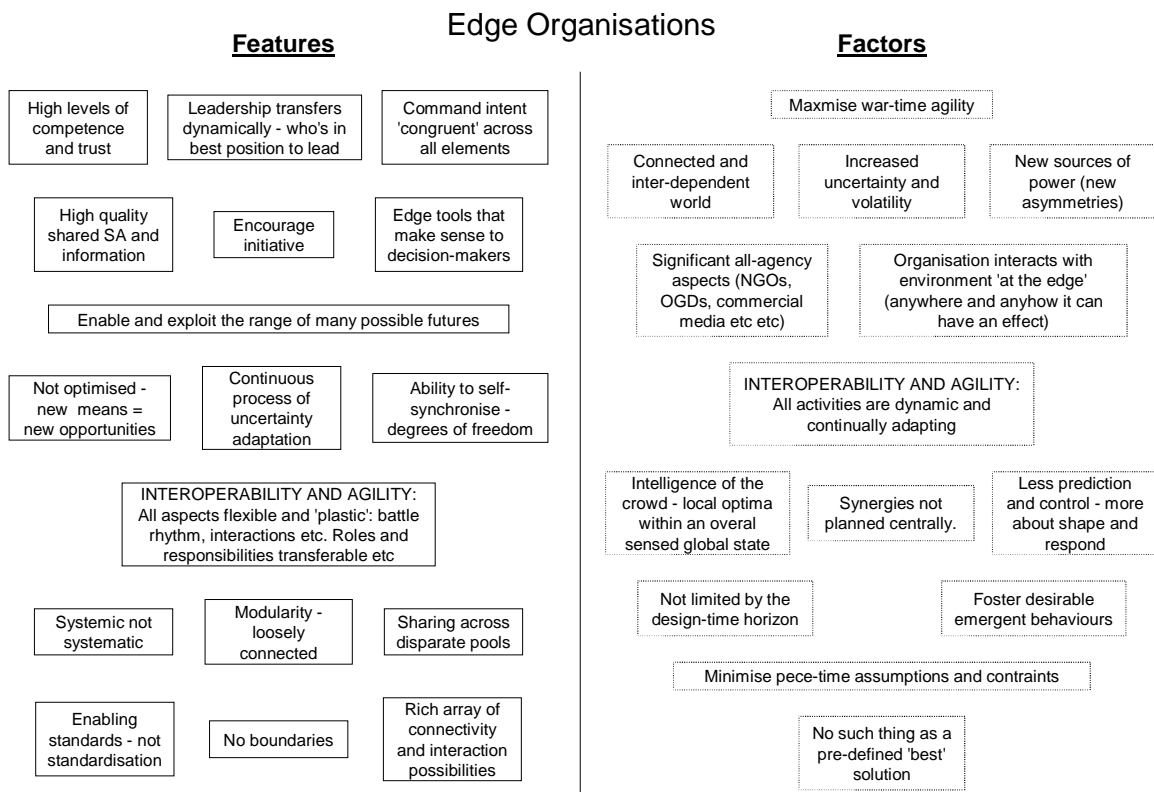


Figure 1 - Features of and Factors Relating to Edge Organisations

Hence, we cannot ignore emergence and must look for holistic approaches which employ biologically-inspired 'control' ⁵ for use in EOs and network-centric situations. A network-centric viewpoint looks at the challenges in terms of the dynamic interactions between all the entities and the environment. Critically, it needs to consider the 'run-time' properties of the artefacts, actors and interactions as well as the dynamic influences on those interactions / entities as being the key focus of attention - as opposed to the static, design-time engineering of their elements. In effect, it forces us to consider NCW as an homeostatic ecosystem - of which EOs will become a major manifestation.

There is another factor in the imperative to get to grips with emergent phenomena. Our world is now so inter-connected that we can no longer consider the battlespace to be a closed, bounded

⁵ In this sense the 'control' would not just be exerted top-down but would work in concert with the kinds of regulatory mechanisms at work in the natural world - most of which exert influence from the bottom-up.

space. Instead, we have to face the reality that we operate as part of a single, dynamic, complex adaptive system driven by emergent phenomena - the real world. It will be axiomatic to this paper that network-centric communities (and EOs in particular) are therefore:

- Complex adaptive systems (CAS);
- Open (unbounded) with distributed, yet highly interconnected, interacting elements;
- Heterogeneous, where fixed standards and procedures cannot be mandated;
- Uncertain, with ever-changing membership, events, artefacts and interactions;
- Diverse, consisting of societies of biological entities and software and hardware set in environments which span realspace, cyberspace and mindspace;
- Intelligent, being able to 'form theories of other minds' - past, present and future.

As M C Meigs [6] says: "Al Qaeda's true operational asymmetry derives from its ability to change its operational system at will in response to the methods needed to approach and attack each new target" Hence an Edge Organisation is de-facto not a fixed organisation - it's a type of adaptive, directed, hybrid swarm⁶ that can change its 'envelope of capabilities' at will - at least that's what "Power to the Edge" seems to assert.

Understanding the Phenomenon of Emergence

It is not within the scope of this paper to examine the phenomenon of emergence per se. It is discussed, as a concept, in [7 Holland and 8 Johnson], but neither book seeks to explain how emergence could be exploited as a positive tool. For that we need to look to [9 Morowitz, 10 Beaufort, 11 Davies and 12 Lewin]. Emergent phenomena may be generalised as having the fundamental characteristic of being tangible or intangible 'patterns' that persist [13 Holland] over time even though the generators of the patterns themselves may be continually changing (viz: an ant foraging party collecting food has an ever-changing membership of ants). The equivalent in the military environment would be that the organisation can continue to be 'robust' even if staff are going on and off duty. In general terms, it is well understood in that emergent phenomena arise in systems with the following characteristics, ie: with components, substrate, interactions and where synergy, antagonism and holism etc are at work. Indeed, emergent phenomena can be found in: deterministic situations, among open systems with non-linear interactions, far-from equilibrium situations, in fact, just about anywhere.

The implication here is that once the ingredients are in place emergent phenomena seem to arise 'spontaneously' (even relentlessly and unavoidably) without anyone having to do anything - but is this true? The consensus is that it is - and that emergence is a considerable force to be reckoned with and that it is something that we usually fail to exploit.

My aim here is to show that it is possible to influence and exploit emergent phenomena such that they can be mapped to trigger the kind of behaviour which would be beneficial to EOs. These features are discussed below under the headings of: conditions for emergence, features of emergence and types of emergence. To assist with the description, I have used an ants' nest as an example, as ants display so-called 'swarm intelligence' - a rich set of adaptive, emergent defence behaviours that we might wish to emulate.

Conditions for Emergence

Emergent phenomena arise spontaneously when a number of conditions are met as follows:

⁶ A diverse set of entities interacting in such a manner that they appear to be a single purposeful community.

'Substrate'. There needs to be a substrate / context / environment which supports the activities of components (viz: the ants nest, its passages, food stores etc and the surrounding environment in which the ants exist). The substrate may influence the way in which emergent phenomena arise in many ways, eg by shaping interactions (see the discussions on templates and stigmergy below). Note that the 'coupling' between the components and the environment means that inevitably co-evolve - each changing the other. In the military context this would mean that we must be able to operate in any of the substrates in which effects might need to be manifested - be that realspace, cyberspace or thoughtspace - and that 'tools' that we may use include those that enable the direct manipulation of the spaces themselves. It is inevitable that these interactions will result in profound changes back onto ourselves.

'Components'. There are some agents / elements / parts which are either assembled from other components (in a fractal / nested manner) or which function together as a part of some entity. There must be more than one component (viz: the ants in the nest) and the membership of the community is constantly changing. In the military context, therefore, we should seek to provide actors (of all types, including software and hardware elements) which are active, semi-autonomous and adaptive and which can interact and form groupings on demand.

Sensors and Effectors. Every component should have sensor(s) and / or effectors which enable interaction across their boundaries. This may mean no more than the 'ability' of a simple cell to gain and loose resources and change state. In a more elaborate example, such as the ant's nest, sensing is multimodal and involves an element of 'sensemaking' (to generate some level of internal representation to support computation and decision-making) followed by some action (again multimodal). In a military context, there will also be a range of sensing and effecting from simple (in devices) to elaborate (in people etc).

Interactions and Relationships. Interactions exchange information and are necessary if emergent phenomena are to arise. They take place between the components, their artefacts and the environment at various levels of complexity and sophistication and are mediated through many types of tangible and intangible mechanisms (in the ant's nest they involve touch, chemical / pheromone messaging between individual ants and the whole nest, individual and collective behaviours, 'crowd' movement, etc). Note that 'structure' may be achieved through communication, eg ants may be physically disconnected but use pheromones to effect communications - they are thus connected in a manner⁷. Interactions also take place between the components via the substrate (so-called stigmergy⁸ [14 Di Caro, 15 Beckers]) and between collections of components in this 'entity' and those in others (ie this ant's nest vs another ant's nest). The interactions and relationships between the participants mature as a result of the co-evolution mentioned above. In the military context there are clear equivalents, such as exchanging messages, manipulating shared artefacts (maps etc), building social relationships / coalitions etc - influencing interactions is the main tool for 'control'.

Local Rules and Templates. Following from this, in any environment where emergent phenomena are manifested, simple low-level 'local rules' are enacted which determine the nature of the interactions which take place. There are many factors which may impinge on the

⁷ It is possible to synchronise the behaviour of chaotic systems by message passing - eg: two pendulums on a wire - so called 'entrainment'.

⁸ Where communication occurs by manipulating artefacts in the environment - eg where termites build nest structures but do not 'talk' about nest building.

way that the rules are triggered and executed, though it seems that one of the most important is the substrate - including the notion of a 'template'. In the ants nest, a template may be set by the gradient of pheromone distribution around the queen. The 'rule' might then be "if I am placing earth pellets - move away from a more concentrated pheromone till the density is Y - place the pellet" - the emergent outcome is an appropriately shaped wall. In the military context, rules may be triggered by the presence or absence of events and conditions and templates may relate to, for example, gradients of bandwidth availability or force density.

Integration and Activation. Though it seems obvious to say it, emergent phenomena will not arise until we activate all the elements mentioned above and add the dimension of time. Patterns then appear, persist over time and have a manifestation which can be detected and acted upon within some context at a higher level of abstraction. However, in reality, a true information ecology can probably never be turned off (cf the Internet). In other words, all we will be doing is adding our components and tool to an existing open 'infosphere' and so we will have to hit the ground running. In the military context, we have the opportunity here to detect new phenomena and exert new effects which have never existed before.

Features of Emergence

There are a number of other factors and features that we need to take into account when considering the phenomenon of emergence. These relate to the properties of emergent phenomena whilst they are being manifested.

Observer(s) and Context. Clearly, some phenomena will emerge whether or not there are observers present (leaving aside the metaphysical argument here). However, other emergent phenomena are an artefact of the observer [16 Bass] and only have meaning in the substrate, ie the context, of the observer (viz: the perception that the ant's nest is 'angry' if poked with a stick relates to the emotions attributed down to it, from the human social world, by the observer). In a military environment a commander may perceive emergent phenomena (or abstractions of them) displayed by the opponent - even if the opponent is unaware that such phenomena are apparent - and different phenomena may be perceived in different contexts.

Lack of Reversibility and the 'Arrow of Time'. Some hold the view that emergent phenomena are not reversible - any cause and effect linkage is one-way, but this is strongly disputed [17 Bricmont]. However, even if we could reverse the 'arrow of time' we would not necessarily see emergent phenomena 'unwind', this is because a differently ordered set of interactions would now take place (in the "poking an ants' nest with a stick" example, the nest would appear to calm down for no reason just before we removed the stick). The insight here is that if the environment evolves towards an unfavourable state, influencing it back is not just a matter of unwinding - the influence required may be obscure or orthogonal to the phenomena being manifested or may rely on allowing the 'system' to self-organise back to a 'known' attractor.

Lack of Central Control. Emergent phenomena are not dictated in advance or controlled or co-ordinated centrally (top-down), instead they usually arise bottom-up and are observed at a higher-level of abstraction. To alter them, one must generally influence at the bottom - and allow the required behaviour to evolve 'upwards'. However, useful creative tension can be achieved by exploiting an observer's top-down view (at some abstraction) in concert with the bottom-up behaviours - providing a route to exert control.

Lack of Dependence on the Existence of Individual Components. Emergent phenomena will persist despite changes in components of the same 'class' - eg: the generators of the patterns themselves may be continually changing (viz: an ant foraging party has an ever-changing membership of ants, or the water molecules moving through a stationary standing wave are always changing though the standing wave remains). Indeed, components can be added and removed without the whole 'system' being decommissioned. However, what is apparent is that the diversity of the components is important, an homogenous environment is an unstable one. In the military context, diversity will provide robustness and persistence despite malicious perturbations - whereas 'normalisation' exposes vulnerabilities.

Adaptation. Emergent phenomena which arise without adaptation are like snowflakes - beautiful complex patterns, but they have no function. It seems that really useful emergent phenomena 'grow smarter' over time - as local rules and the nature of the interactions change in response to evolutionary pressures, so persistent patterns manifested in the environment mature. This infers that there is some form of hysteresis and learning - though the learning may not need to be encoded in 'data' - it may be represented by changes in trigger conditions or in the trajectory of the patterns over time. The implication here is that all our elements, interactions and the substrate itself, should be 'plastic'. A fixed, over-engineered, environment is a dead environment (see the discussion on Langton's Lambda parameter below).

Self and Non-Self. As mentioned earlier, CAS are, by default, open environments consisting of many interacting and connected elements. This means that the 'coupling' between elements across the environment makes it difficult (possibly even meaningless) to try and identify 'self' and 'non-self'. Even though one human body seems bounded and has an immune system, it is connected so closely to its environment and its community that identifying where their mutual influences stop is nigh on impossible. The implication for the military environment is profound. Enclaves (so-called compartments) and apparently bounded 'secure' communication facilities can be created - yet in fact their security can be an illusion [18 Lewin]. In a network-centric context, security will have to embrace the reality of the open environment and seize the opportunities that its complexity offers, rather than always assuming a fortress mentality.

State and Persistence. Crucially, emergent phenomena at one level may be viewed as components at a higher-level of abstraction - indeed, emergence provides freedom of action at a higher level which is denied at a lower level. However, these persistent patterns (or trajectories of patterns), must manifest themselves in such a way that mechanisms at a higher level can identify them and use them as an invariant - they may then interact and reverberate, leading to further emergent phenomena [19 Holland]. Though these mechanisms are currently poorly understood, it is in this way that sophisticated, high-level operational military behaviour could be generated from components injected at various levels within the environment.

Entropy vs 'Information' and Increasing Structure and Organisation. Emergent phenomena add to structure in the universe. Despite the Second Law of Thermodynamics stating that entropy always increases towards featureless uniformity, some see (though others disagree [17 Bricmont]) that there is opposite trend at work in the universe - that of increasing structure and organisation at ever higher levels of abstraction manifested through emergent phenomena - the so-called 'optimistic' arrow of time [20 Davies]. This may be being achieved by the fact that there are causation mechanisms at work which would not contradict the Second Law. As Donald MacKay [21] says:

"...whereas in classical physics the determination of force by force requires a flow of energy, from the standpoint of information theory the determination of form by form requires a flow of information. The two are so different that a flow of information from A to B may require a flow of energy from B to A ..."

Types of Emergence

There appears to be a view that there may be different types of emergence and that they relate to the way that they arise and to their levels of 'connectedness'. Warren Weaver [22] when writing of Shannon's work on Information Theory [23 Shannon] talked of 'simple systems' (with two or three variables), then a second group which he called 'disorganised complexity' (with millions of variables only tractable through employing statistical mechanics and probability theory) and a third type for which:

"... Much more important than the mere number of variables is the fact that these variables are all interrelated ... These problems, as contrasted with the disorganised situations with which statistics can cope, show the essential features of organisation. We will therefore refer to this group of problems as those of organised complexity"

Other viewpoints on forms of emergence relate to the degree to which the outcome can be predicted in advance. Smuts' 'Holism and Evolution' [24 Smuts] was the source for one of two 'co-emergent' notions of emergence, another view is quoted here:

"Thus, within the general framework proposed here, one must distinguish between two different kinds of emergence: A. Deducible or computational emergence. There exists a deductual or computational process or theory 'D' such that an emergent phenomena 'P' (observed in a higher-level system) can be determined by 'D' from the lower-level system. B. Observational emergence. 'P' is an emergent property, but cannot be deduced as in (A) above. [25 Baas]

Some have argued [8 Johnson, 9 Morowitz, 26 Davies] that self-organisation and the emergent properties of complex adaptive systems may be the active mechanisms in a 'theory-of-everything'. This requires us to develop acceptable representations of complexity, self-organisation, emergence etc and their relationships to the tangible and intangible worlds.

Levels of Abstraction (the Observer observed). As already mentioned above emergent phenomena may have no 'meaning'⁹ at the level at which they are generated. Though this sounds like the beginning of an endlessly infinite regress of no value, it is actually a crucial point to understand if the phenomena of emergence is to be exploited. For example, Popper [27] represents this by his description of "World 1 .. World 3 entities"¹⁰. This kind of idea might be described as follows:

That there exist higher levels of emergent abstractions at which a simplified representation of the activities of a lower level can be meaningfully manipulated

⁹ The phenomena may be real but some of the meaning attributed to a phenomena may be subjective or interpreted in a way which only makes sense to the observer.

¹⁰ Popper describes "World 1" as the physical world and "World 2" as the "world of our conscious experiences" whereas "World 3" is the world of the logical contents of books, libraries, computer memories, and the like.

by an observer *outside the system* (assuming that the observers can adapt to a world view that would let them detect the abstractions and make sense of them).

This is discussed at some length by [28 Hofstadter] and beautifully illustrated by the M C Escher lithograph "Print Gallery" [1956] which shows a young man in a gallery observing a picture which includes himself observing a picture of himself we, because we are outside the system (at a different level of abstraction), can observe and reflect on this paradox - possibly itself an emergent phenomena. The "so-what" about this is that if we do not learn how to exploit phenomena (such as that of emergence) then:

we may never know about some of the capabilities that are waiting to be used because they exist at a higher level of abstraction - of which we are currently unaware - one that conventional approaches will never reveal to us, but that an opponent could exploit against us.

Conversely, we may be able to outwit opponents by observing the emergent phenomena displayed by the opponent and then reasoning about it at an higher level of abstraction - which will be invisible to the opponent (and therefore, as an aside, infinitely secure).

Certainly there are phenomena which emerge directly from interaction at the physical level (standing waves etc), but there also appear to be conceptual emergent phenomena observed at the level of human consciousness. It seems possible therefore that there may be a continuum of types of nested emergent phenomena (depending on their tangibility or intangibility, the 'systems' from which they arose, the mechanisms of interaction / underlying theories at work, the nature of the observer / observation, the level of abstraction or other factors) and that characterising and classifying the types and their differentiating features may just be a time-wasting exercise - or extremely valuable.

There may well be some 'universal shorthand' for characterising all types of interactions but at this stage it would be foolish to assume this. Indeed, it is clear that it is currently very difficult to 'transform' a viewpoint showing interactions at one level of abstraction into an equivalent one at another level and that this difficulty is a reflection of the incompatibility of the theories and mechanisms underlying the interactions and their attendant representations at each level.

Though this seems a long way from where I started, I consider it relevant to try and embrace this broad theme as the notion of sensemaking¹¹ (for example) covers, de-facto, the physical, information, mental and social / cultural domains. Indeed, it is necessary to see consciousness as being part of the information processing capability of the biosphere, rather than something unconnected and esoteric. Hence, at some point, we will need to address theories which can span these domains if we are to employ them to provide effects 'across the board'. In studying Edge Organisations, emergence should be a topic for further research.

Mechanisms of Emergence

This section of the paper will now discuss the technologies, tools, techniques and strategies (T3S) which could be employed to enable us to exploit the emergent properties of complex systems to achieve the required behaviours in EOs. From the discussions above, I hope it is

¹¹ Used by Alberts for Network-centric Warfare (see [32]) - "... a process that builds deep understanding ...".

evident that we need several types of T3S¹² (see Figure 2), employed under different circumstances, to provide the following:

- Design-time Properties: we need to provide elements which have features which, when activated, display properties that are desirable at run-time - such as being adaptable and plastic - these form our T3S Foundations;
- Tools: we need to have tools and mechanisms available to us at run-time to employ to influence, among other things, the substrates and the behaviour of the elements and their interactions - these are the Run-Time T3S;
- Evolutionary Mechanisms: we need to be able to work with, at different levels of abstraction, the self-organising and emergent phenomena themselves to evolve beneficial behaviours (and countermeasures) within the Run-time Environment.

Each of these areas is considered in more detail below.

T3S Foundations - Design-time Properties

In terms of the kinds of T3S Foundations, over the past few years a number of technologies have begun to arrive that are individually addressing issues of information sharing and access across the distributed communities of virtual organisations. These include agile and resilient networking, peer-to-peer computing (P2P), grid computing, Web services, Semantic Web (knowledge technologies), and human-computer interaction technologies. To implement EOs, we will have to harness and intercept the latest technology developments in these areas to enable decision makers to work in a secure, flexible and agile manner. Each of these 'contender technologies' is discussed further in the Acquisition Section later in the paper.

We will need to adopt a number of design principles to ensure that we can meet the requirements demanded for EOs and net-centric approaches. The key design rule that must be applied is to minimise design-time assumptions in order to maximise run-time flexibility. This is an over-arching rule that includes the principles discussed below.

Separation and Loose Coupling. In computer technology terms, this includes the separation of interfaces from implementation (separating what an element does from how it does it), and the removal of hidden dependencies and interconnections (using the interface and nothing but the interface). Loose coupling can be achieved using dynamic discovery rather than using predefined connections and interactions - there are analogies in organizations;

Exposing Interfaces. As it is not possible to predict how elements will need to be used in the future, systems and components must be designed in an open fashion providing interfaces that allow their functionality to be invoked by new, perhaps unanticipated elements. This enables reuse and greatly improves interoperability. Again, sharing and altruism are stated to be a key enabler in Edge Organizations;

Separation of Action from Information. We should not assume that information is embedded in our active elements and nowhere else. The more widely available the information the better the interaction and the greater the flexibility. In operational terms, generic data formats and ontologies facilitate this, allowing domain knowledge to be expressed in standard flexible ways, and extracting assumptions, processes or doctrine that would otherwise be hard-coded into

¹² For a more detailed description of T3S Foundations and Run-Time Tools see "T3S for Adaptive Systems" from <http://www.tbt.org.uk/>

devices and thus very difficult to change. These approaches use open standards - strongly supported by off-the-shelf tools - such as wrapping heterogeneous data so that it can be manipulated as shared information.

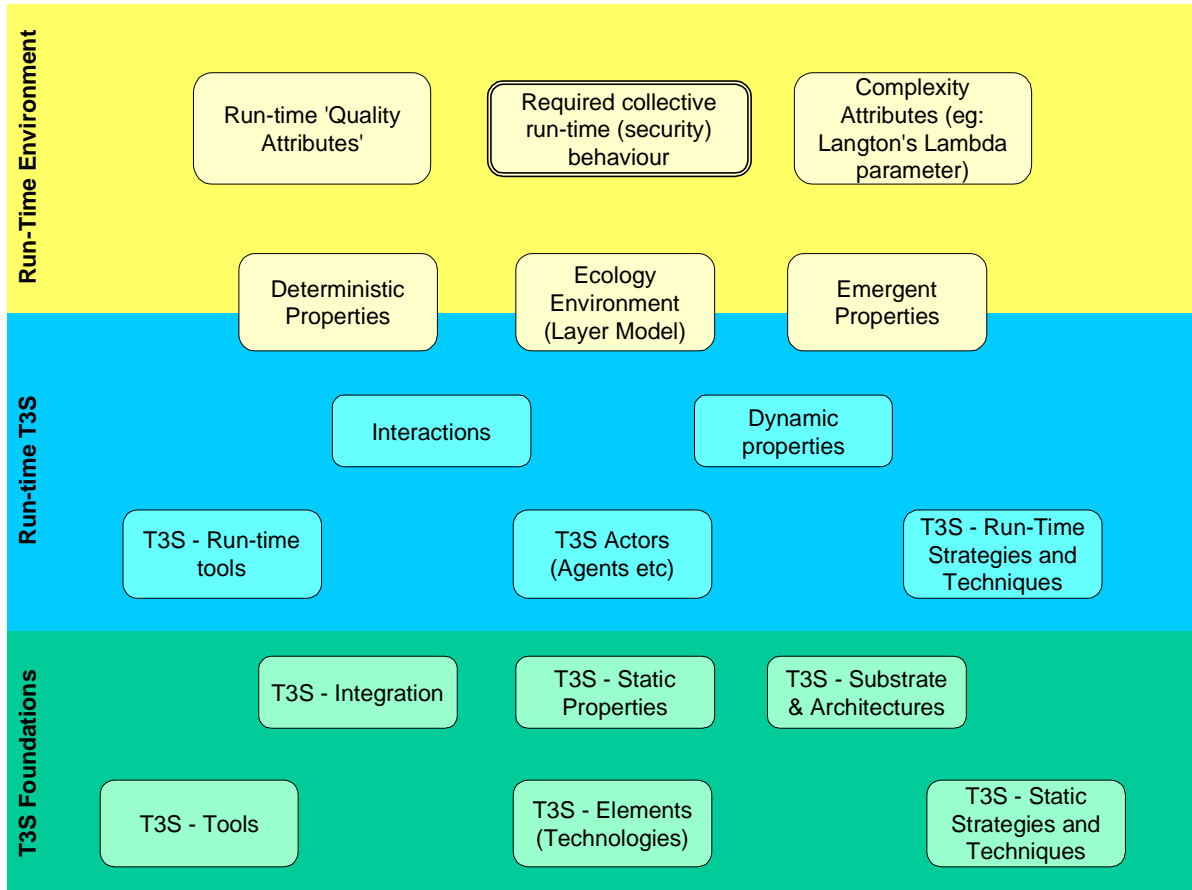


Figure 2 - The Net-centric Environment

Run-time T3S Tools

Let us now consider the run-time issues. Paradoxically, you can't begin to establish the exact performance of components until they interact within the environment and, even then, all that one can establish is a probabilistic estimate of performance. This is because it is impossible to examine every state under which the components would have to operate as many of the states are emergent and *cannot be part of the formally specified design*. It is the appearance of these emergent phenomena which leads to the failure of many of the attempts to create national computer systems and is why 'unwanted' emergence is treated as something to eradicate - leading to a retreat from embracing complexity towards design-time 'certainty'.

This paper turns this view on its head and aspires to maximise novelty at run-time. Hence, we have to accept that we cannot rigorously test the 'system' before deployment because, de-facto, the boundaries of the 'system' cannot be defined. Instead, we will provide ourselves with run-time tools through which we can enforce obligations, provide authorisation and influence interactions and social behaviour towards our requirements. The real novelty proposed in this

paper will come from the innovative integration and application of these T3S capabilities to provide the collective run-time properties that have been discussed earlier.

Evolutionary Mechanisms for Shaping Behaviour

Though it seems obvious to say it, emergent phenomena will not arise until we activate all the elements mentioned above and add the dimension of time. What now occurs is that patterns appear, persist over time and have a manifestation which can be detected and acted upon within some context at a higher level of abstraction. In the military context, we have the opportunity to detect new phenomena and apply new effects which have never existed before, but how do we do this, how do we evolve the behaviour that we want without trying to impose it top-down? We have already started on the process by altering the design-time features of our T3S Foundations and by creating Run-Time Tools - these will now take effect as they join the active environment. The rest of this section will look briefly at some of the aspects of run-time evolution over which we will have some effect and indicate strategies we could employ:

Collective (Gross) Indicators - Langton's Lambda Parameter. As emergent phenomena arise from interactions (which implies an exchange of 'information') there do seem to be principles at work which would change the nature of the onset and the 'fierceness' of the propagation of the effects. Through his work on cellular automata (CA) and artificial life Chris Langton [29] noticed a pattern to the types of CA which were derived as shown in Figure 3. The pattern also related to Stephen Wolfram's [30] system where he had classified CAs into four classes. Langton noticed the following:

- It became apparent that in situations where the 'information' (ie resources) used by the CAs was constrained to be sparse or 'frozen' few CAs existed - and those that did displayed a narrow 'fixed' behaviour.
- At the other extreme where 'information' moved too quickly to be captured it 'boiled off' and the few CAs here displayed a chaotic behaviour and had hardly any form.

He found that at a set value (about 0.273, where the complexity of the system was at a maximum, available entropy was at an optimum and where there was a phase change) the most diverse and vibrant CAs were to be found - he dubbed this "life on the edge of chaos".

It appears that the phenomena of emergence is also related to this factor and that, in general, CAS co-evolve and self-organise towards the point of optimum information flow. In principle, this means that we now have a measure which can be applied to the features of a CAS such that some assessment can be made of the likelihood that: nothing interesting will happen, 'interesting things' may happen (eg: at phase transition), or that the 'system' will generate 'chaotic' emergent phenomena. We would aim, therefore, to manipulate the opponent into one of these spaces which is mal-adaptive for their current organisational form.

As well as sensing the collective state of the information ecology, there are other evolutionary measures which we can employ (at many levels of abstraction) to change the environment and the specific behaviour manifested which include:

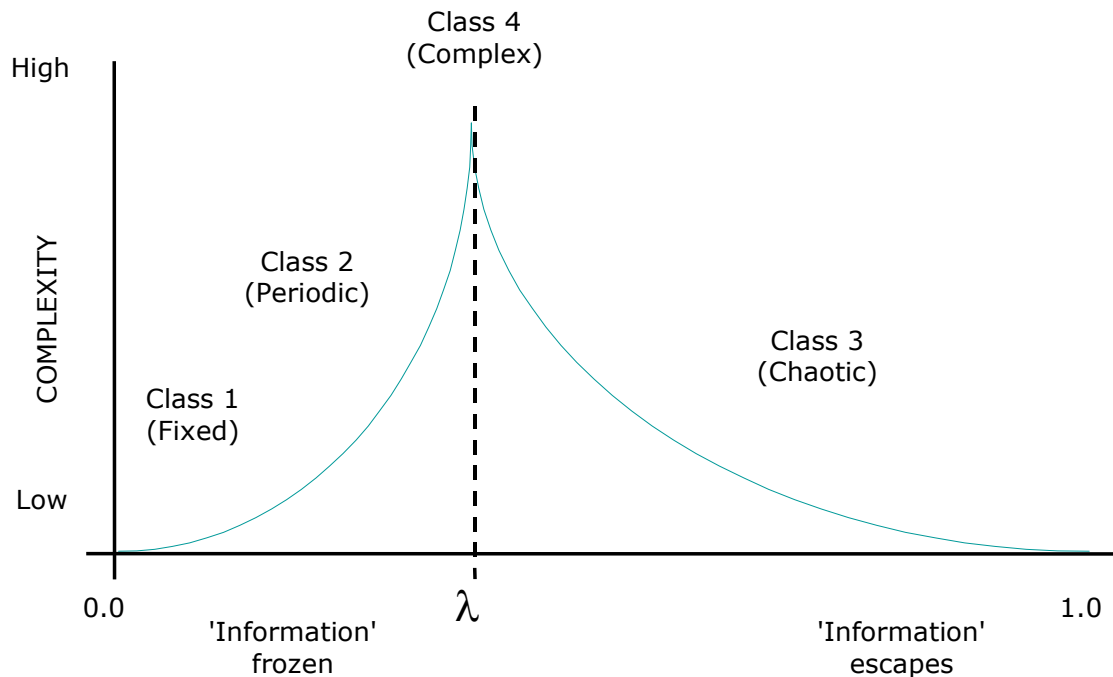


Figure 3 - Langton's Lambda Parameter

Unfitness and Unfitness Landscapes. Organisations will always have some level of 'unfitness' in relation to the environment and the other components. This unfitness is determined by a number of factors such as their degree of generalisation or specialisation, the presence of predators, the number of competitors, the 'richness' of the environment, the nature of the interactions etc. All these factors can be influenced and, when 'unfitness landscapes' are connected, behaviours can be driven towards or away from desired situations. A great deal of academic work has already been done, but it has not yet been employed to influence the behaviour of military organisations directly. However, good commanders will always seek to understand what would trigger 'discomfort' in themselves or their opponents as they know that discomfort generates an imperative for action. Understanding how different types of discomfort might drive Edge Organisations to adapt and reshape will be an important line of enquiry;

Population Variation. The balance of members of a population (trophism) is a measure of the diversity and balance within an ecology. Population membership can be altered to trigger a collective change of state to a more beneficial one or to stress the adversary (military examples would include the proportion of commanders to shooters, defenders to attackers, the number of 'embedded' media personnel, the diversity of weapon systems etc);

Downward Causation. Through the connectedness of CAS, reverberation or oscillation can occur where higher-level behaviours affect lower level elements (as with the Laser). This mechanism could be employed to induce a kind of top-down 'control';

Environmental Manipulation and Templates - Triggers and Thresholds. Many of the 'local rules' activated by the components in our ecology are triggered by various phenomena in the

environment. As has already been mentioned, components may manipulate the environment themselves as a kind of message passing (stigmergy) and we can interfere with this process to alter behaviour. In addition, in an ant's nest, pheromones are used to: leave trails, indicate presence or absence of threat, affect building behaviour (see templates discussed above) etc. This is called chemotaxis. In the military environment, messages passed between agents are like pheromones; we can alter how long they persist, how quickly stimulation causes fatigue (so messages are ignored) etc. These could be manipulated to effect the required behaviours;

Interaction and Relationship Tuning. A variation on this is interaction tuning. We can change the pattern of interaction by altering the frequency and amplitude of messages, apply damping to slow the collective information exchange and move the community from one side to the other of the Lambda Parameter (eg: force opponents into the 'freezing' zone to deny them access to resources) or by clamping certain freedoms. In addition, we can alter the 'connectedness' of elements (how much one depends on / is affected by another) and affect the extent to which changes and perturbations cascade through the environment. However, it is not clear how the different types of tuning affect the classes of phenomena that emerge;

Learning. Learning is expressed in many ways and may be represented by both static elements (artefacts in the substrate) and dynamic elements (reverberating patterns or trajectories of interactions / strange attractors). As learning is part of the way CAS adapt, we can alter the learning artefacts to accelerate adaptation to respond to damage by an opponent or to cause 'forgetfulness' to deceive an opponent. This learning can be unsupervised (eg, as in birds flocking) or supervised by a controlling mind (eg, by the commander). When unsupervised, there may not be a visible artefact of the learning and so it may be hard to understand how the learning came about or to justify its effectiveness. Learning which is too prescriptive may suffocate the generation of novelty and may inhibit the ability to adapt;

Social and Cultural. Lastly, there are also social and cultural effects to be exploited - though these may be at the level of the human decision-makers. Our options here relate to influencing: styles of command (centralised or dispersed), decision-making strategies (a-priori formal or adaptive, continuous and informal) and cultural (autocratic or mob - leading to rapid, chaotic responses). We would currently call this information or psychological operations.

As yet, these approaches are immature and need to be researched and investigated further, yet there is clear potential here, ready to be exploited to help us when deploying EOs.

Control vs Emergence in the Context of EOs

The issue of control versus emergence is discussed in several places in "Power to the Edge". It is the author's opinion that, in military enterprises, there will always be a degree of centralised thinking. At the minimum, the Commander-in-Chief has to set the context for the operation - either as part of foreign policy or to achieve specific goals. Without this single, clear vision (communicated clearly to all concerned) we are but an ant's nest - apparently wandering aimlessly without coherent purpose.

However, if we over-emphasise strong centralised control (soviet-style) and attempt to plan uncertainty out of the future we are heading for disaster for several reasons. Firstly, as Dan Quayle is supposed to have said "The trouble with the future is that we don't know what is going to happen". Despite effects-based approaches and operational network analysis, we will never

be able to predict the future. Instead, we need agile tools which will enable us to shape events to our will as they unfold - this cannot be done at design-time - whatever people say. From this simple fact flow many profound implications - not least of which is that EOs are 'on the money' when it comes to considering military operations in the future. Conversely, the current trend towards copying commerce, with its just in time (JIT), optimised approaches, is not the correct model for military operations. JIT thinking, though useful day-to-day, has shown itself to be brittle in the face of change - time and time again failing catastrophically during serious dislocation (the lack of capacitance and damping being key weaknesses). Instead, the new thinking must understand the true nature and range of military operations and adversaries that we face and, with this knowledge to the forefront, proceed accordingly.

Acquisition Issues and Contender Technologies

The overall aim of Network Centric Warfare (NCW) is to enhance military capability (enabling EOs) by exploiting information networks more effectively. The demands of coalition operations reflect the aspirations of NCW's Core Themes for flexibility, adaptability, robustness etc. Increasingly, this involves interoperability between NGOs, OGDs and commercial organisations - a very heterogeneous mix over which it is not possible to impose a single set of standards. NCW (and its UK equivalent, Network-Enabled Capability) have been widely discussed [31, 32] elsewhere, yet their implications are not fully understood.

As already mentioned, diverse, open and distributed structures of this type are, de-facto, complex adaptive systems whose boundaries cannot be rigidly delineated. Hence, it is impossible to fully define, at design-time, all possible states, configurations and the interactions that could occur. This is both an indisputable fact and a major challenge as, if we cannot define it, how can we acquire it and implement it? Currently, the tool of choice is systems engineering, yet experience has shown that, especially when dealing with information systems, its use can lead to brittle, inflexible solutions which lack adaptive capabilities. In extremis, these tools can unduly constrain the ability of commanders and their forces to act with agility and swiftness in response to unforeseen military imperatives.

So, as operational circumstances change in the real world then information systems at the boundaries of, and within, cyberspace must respond to these changes - but we have *not* procured capabilities which can do this. We have been able to achieve airspace superiority, for example, but we have not achieved the same in cyberspace. Hence, it is vital that we acquire capabilities which can act decisively, adapt at run-time and provide synchronised effects across all the domains shown in Figure 4 - this is discussed next.

The Coalition Agents Experiment (CoAX)

How do we go about achieving this coherence? Research carried out by the Coalition Agents Experiment (CoAX) [Kirton, Beutement - 33, 34] between 1999 and 2002 has addressed some of the issues discussed here and can be seen as having been a de-risking exercise. CoAX was an international collaborative research effort which involved 26 military, academic and commercial partners, funded partly by the UK MoD, but primarily by the US Defense Advanced Research Projects Agency (through the \$60M Control of Agent-Based Systems (CoABS) [35] Programme). The principal research hypothesis of CoAX was that emerging technologies such as software agents, information grids, the Semantic Web and agent control techniques could be used to construct coherent, flexible and agile 'command support systems' for coalition operations - in effect, a prototype for Edge Organisations. The technologies that CoAX employed will now be discussed in outline.

Contender Technologies

CoAX started by embracing two principles. First, it took the open, heterogeneous, diverse and dispersed nature of the Coalition environment as a given - no single standard was mandated. Secondly, requirements were not all defined in advance - neither for the 'challenge spaces' nor for the envelope of required effects. CoAX then put in place a range of technologies and tools, which are summarised briefly below:

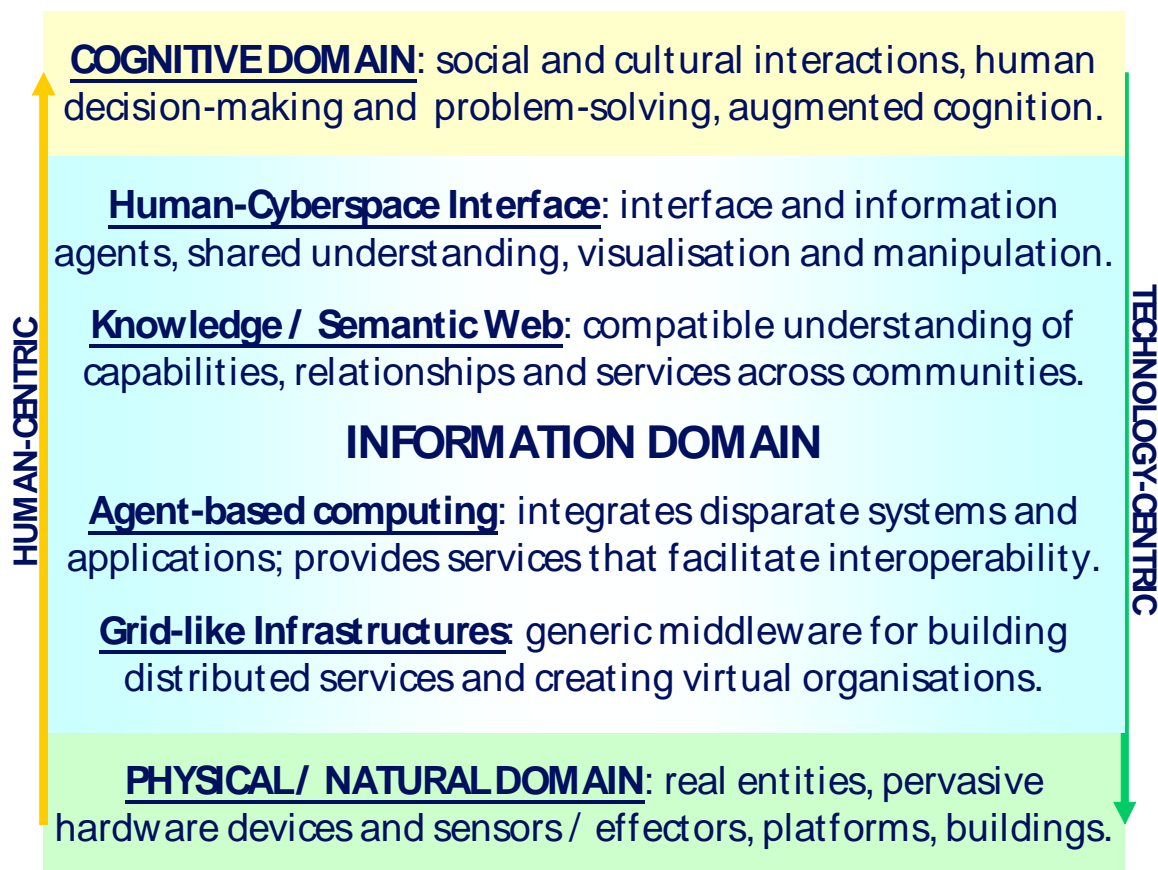


Figure 4 - This chart shows the mix of contender technologies required to achieve synchronised effects across all domains of distributed military enterprises.

Information Visualisation and Manipulation Tools. As we all know, we make sense of our world by interacting with it, then creating various kind of mental models and then by manipulating the models (so called shared situational awareness - part of 'distributed cognition' [36 Hollan, 37 Hutchins, 38 Chin]). CoAX looked at tools which enabled decision-makers to representing the world of coalition operations in the manner which made most sense to them (as typified by the Decision Desktop [Allsopp 39]). Underlying the tools were interface agents which noted what the humans were doing and, as a result, would reach out to information agents for updated intelligence about the items of interest. CoAX took a pragmatic approach (agents working invisibly in the background to support human-driven activities), whilst others are overstating the capabilities of so-called 'personal assistants' where dialogue occurs between humans and an

animated persona. However, these animated assistants are at a primitive stage and we should not expect significant development in the short term.

Trust, Persistence and Learning. In the real world we build relationships which persist over time and on which we build notions of trust. We expect our colleagues to learn our preferences (as we do theirs) and this leads to efficient social and professional interactions. It will be essential that the information age devices which support us in our work also acquire a limited version of these capabilities. This is not a trivial task - it has been said that a key indicator of intelligence is "the ability to form mental models of other minds" - software agents and robots are decades away from being able to achieve this. However, progress is being made on providing agents with persistence (a memory of previous interactions), machine learning and the ability to construct (and add to) models of users and other agents with which they come into contact. In this way we will be able to build limited relationships and develop some notion of trust - essential if we are going to delegate tasks to them [40, 41 Bradshaw].

Semantic Web. Currently, tools such as Web pages are geared towards the visual presentation of information for humans, with no support for machine understanding and reasoning, severely limiting the automated processing of information on the Web. The Semantic Web aims to provide general-purpose structures for semantics, to have data on the Web dynamically defined and linked in such a way that it can be used by machines for automation, integration, inference and re-use across various applications. The Semantic Web starts from knowing that it not possible to define a single master ontology for the whole planet. In the CoAX demonstrations, XML was one of the languages used for inter-agent messaging and DARPA Agent Mark-up Language (DAML) was used to encode and reason about domain entities, domain policies, tasks and agent message content.

Software Agents and Services. Agents can be viewed as semi-autonomous entities that help people cope with the complexities of working collaboratively in dispersed information environments. A community of agents works as a set of distributed, asynchronous processes communicating and sharing information by message passing. They work with users to make this information available whenever and wherever they need it, and can be organised to support individuals, military commands or virtual function teams. Moreover, the agent paradigm [42 Jennings] provides the modularity and abstraction required for building large, distributed and complex active information networks such as those required for EOs. Included with agent technologies are objects, components and services. Essentially, objects can be grouped into a component and provided with a defined 'interface' through which they can be activated. When a component advertises its capabilities so that other entities can find it and employ its capabilities, then it is a service. Agents are the active elements that dynamically invoke services etc in response to the tasks they are carrying out (of course agents can be composed from components and objects too!) - all this supports the achievement of agility.

Agent Domains. CoAX used the Knowledgeable Agent-Oriented System (KAoS) [43 Bradshaw] domain management services to dynamically organise agents into logical agile mission groupings corresponding to real-world organisational structures, administrative groups, and task-oriented teams - including allowing for complex hierarchical and overlapping structures. Domain managers administered agents and the specification, conflict resolution and enforcement of policies, represented in ontologies such as the DAML. Domains can be used to respond to the requirements of national sensitivities or to deal with the need for different security groupings or zones of resource availability etc.

Control Policies. The increased intelligence that software agents provide is both a boon and a danger - agents can perform tasks inside cyberspace that would be impractical or impossible using traditional software applications. However, this autonomy, if unchecked, could also severely impair military operations. In CoAX, KAoS provided services to influence the run-time behaviour of agents (even if the agents came from different developers and were running on diverse platforms). KAoS services and tools permitted policies to be written and deployed within complex military organisational structures. In addition, a system called "NOMADS" provided support for agent mobility, enabling information and computational capabilities to be deployed across the battlespace to where they were required at short notice.

Grid Infrastructure. Agents and services require infrastructures which enable the discovery of other agents and resources and which support dynamic changes in the interactions between them. CoAX used the CoABS grid middleware, but there are currently many initiatives underway for sharing resources as varied as high-end computation and digital media. The notions of information grids are seductive and seem to suggest that we can achieve perfect intelligence, so-called Predictive Battlespace Awareness, but this is not true [44 Rosenberger, 45 Lwin]. In terms of logistics, grids enable us to move away from supply chains (which can break) to robust, adaptive and self-healing [46] supply networks.

Dynamically Reconfigurable Synthetic Environments. Finally, there is the issue of turning all these capabilities into a coherent fielded force. Training and exercising, as well as support to campaign execution and mission rehearsal, cannot proceed with the simulation tools that we have now. Our current Synthetic Environments (SEs) cannot support these new battlespaces and ways of working because existing doctrine, procedures and force capabilities are embedded in proprietary data, or worse, in the code. At its most basic, the information age approach would be to move the doctrine, rules and data out of the code and expose it - so that it can be accessed and manipulated by the 'cheapest' relevant mechanism available - a significant challenge. Ideally, we need dynamically reconfigurable SEs which can support training and exercising in the kind of agile operations to which we aspire. In addition, there tends to be an assumption that in any 'what-if' analysis of effects we will wish to seek out the optimum solution, but this is false. If we can mechanistically identify the single, optimum solution, then so can our opponent, who will then neutralise our force. Instead, we need to be able to flip across an horizon of relevant options - confusing our opponent, whilst maintaining our own coherence of purpose.

Pointers for Acquisition

In many ways, COAX can be seen as a prototype of the kind of active information environment that could be acquired to support the operation of EOs. Information networks are usually seen as simple communication pipes between computers, but this is wrong. Cyberspace is an active battlespace to be dominated as part of Full-spectrum Dominance and so new tools and techniques will be required to enable coordination of effects across all battle spaces. Some initial pointers for acquisition can be drawn from CoAX - these are presented below:

Reduce the Emphasis on Specifying Design-time Requirements. Acquiring agile information networks with the necessary capabilities requires an approach which goes away from specifying everything in advance. In addition, we should look to make the design-time properties of our devices such that they can communicate, be assembled into more complex systems and then

be reconfigured dynamically at run-time. It has been shown that minimising peace-time assumptions enables greater conflict-time flexibility;

Understand how to Employ Tools in the Run-time Environment. Some uncertainties can only be dealt with at run-time. Acquisition will need to provide tools which can alter run-time behaviour on-the-fly to meet the changing demands as operations are executed. In addition, we need to understand better how to exploit the phenomena which arise from these complex environments as a force multiplier - for example, evolving and deploying emergent phenomena (such as a cascade of denial of service attacks) against an opponent [47 Beaufement];

Embrace Heterogeneity and Complexity - not Constraining Standards. CoAX accepted the reality of coalition operations - that it is not possible to mandate a single standard - and showed that this was a strength which led to greater flexibility, security and robustness. CoAX identified that some standards could be seen as constraining and others as enabling (as they can be used as building blocks) and explored the possibilities for active interoperability negotiation at run-time - a different approach from the usual acquisition of tools with pre-defined information exchange requirements - an anathema to Edge Organizations;

Exploit a Mix of Novel Architectures. CoAX showed how, within security constraints, shared information pools can easily be created by employing publish and subscribe mechanisms within a service-based architecture. In addition, legacy systems can be integrated if they are agent-enabled (exposing interfaces and data) - leading to better coalition-wide shared understanding. These information pools arise dynamically at run-time and are very robust - however, they are not database applications procured in the conventional sense;

Indeed, because we want agility, flexibility and the ability to cope with uncertainty, then by definition we cannot specify all our requirements at design-time. So, procurement must adapt so that it is able to provide capabilities which can be employed 'plug-and-play'. Consequently, a pragmatic, yet innovative, implementation of EOs would enable us to carry out dynamic operations, to change our 'operational systems' at will to achieve our aims.

Research Issues

The human race seems obsessed with overcoming challenges, shaping our world and striving for certainty by conceiving of future states and then enacting them with dramatic effect. However, we are still surprised at the many, varied and apparently unexpected outcomes which occur when we transition our schemes from their models into reality. But should we be so surprised? Foremost among our techniques is deterministic modelling based on a Newtonian view of the world. There is a view, however, that determinism is a myth [48 Prigogine]. Indeed, the deterministic experimental conditions of the science laboratory are not a microcosm of the real world - they are atypical of it - Joseph Ford [49] makes the point somewhat whimsically:

"Unfortunately, non-chaotic systems are as scarce as hen's teeth ... algorithmic complexity theory and non-linear dynamics together establish the fact that determinism actually reigns over quite a finite domain; outside this small haven of order [the 'laboratory'] lies a largely uncharted land ... where determinism has faded into an ephemeral memory."

So, even if the universe behaves like a machine in the strictest mathematical sense, it can still happen - indeed it is inescapable (as Paul Davies [50] easily proves) - that genuinely new and in-principle unexpected phenomena will occur. The conclusion must be that determinism is a myth and that we need to look beyond classical science and deterministic models and methods to understand, and then be able to harness, emergent phenomena as a positive tool in a net-centric environment. Hence, except in certain situations and despite protestations to the contrary [51 Brook], we must therefore dispel the myth that systems engineering is the tool of choice to use to produce the secure, net-centric environment we need. With its emphasis on the a-priori definition of requirements and on establishing properties and features of bounded systems of systems at 'design-time', systems engineering is clearly only suitable for engineering some of the foundations of a net-centric environment.

A third myth is that the security agencies and / or the military will not accept systems with behaviour which cannot be exactly determined in advance. The whole NCW / NEC thrust would not exist and be receiving considerable funding and support at the highest levels if this were true. Consequently, we can proceed towards a research program with confidence.

The hypothesis of this paper has been that agility in EOs is not something which is layered over the organisation, applications and infrastructure as an afterthought, but is a collective property which emerges from the interaction among the deployed elements. However, despite the positive reception given to the CoAX experiments, only the tip of the agility iceberg has been touched - self-organisation and emergence are largely unexplored topics. A multi-disciplinary team could carry out the following tasks:

- Classify the types of emergent phenomena and their differentiating features and develop a consensus on terminology of emergent properties within the scope;
- Characterise the necessary conditions under which emergent phenomena arise and relate those conditions to the characteristics and classes of phenomena;
- Understand how to trigger (and subsequently nurture) certain forms of emergent phenomena 'on demand' to support a specific 'task' and also investigate precursors, constraints and possible 'paths to emergence' for EOs;
- Use Edge Organisations (which seek to actively exploit the phenomena of emergence) as a challenging test case and use the multi-agent systems prototypes available from CoAX as a test-bed for experimentation at the conceptual (campaign) level.

Research Outcomes. The expected outcomes of the research would be as follows:

- Improved understanding of the factors and principles relating to positively exploiting the phenomena of emergence as a 'tool' to support the deployment of EOs;
- The provision of a set of formal descriptions encompassing characteristics, classifications, invariants and representations relating to the exploitation of the phenomena of emergence;
- The provision of an initial set of tools for employing and exploiting the phenomena of emergence in military and commercial contexts (the 'socio-technical' military context of net-centric conflict and in the commercial context of competition for 'resources' in distributed information systems) with some examples derived from the test-bed research into campaign analysis, military capability management, and information operations with software agent teams;
- Report on the potential applicability of the research information derived from the above to military operations and commercial activities and advice on how problems of this type should be approached when implementing EOs.

Conclusion

Increasingly, innovators look to nature for inspiration when considering the problems of our complex and interconnected world - so-called 'biomimetics'. Yet, as has been indicated above, one area seems to have been neglected - emergent phenomena. Almost without exception the natural world is formed, driven and evolves through the interaction of emergent phenomena which manifest themselves at different levels of abstraction.

However, in stark contrast, we do not usually employ the same mechanisms as the natural world in the creation of our devices, systems and human artefacts and so miss out on the potential benefits that are waiting to be accrued. The exponential growth in the use of communication, mobility and information technology is creating an ever more uncertain, highly interconnected, complex and heterogeneous world. Conventional approaches to the 'design-time' engineering of systems, which rely on the systems being closed, linear, optimized, hierarchical and 'static', do not work on complex systems.

This paper has asked: Can the principles employed in natural systems help us deal with our need for certainty in our increasingly complex human world? Clearly, understanding how to wield the phenomenon of emergence as a potential tool is part of meeting this challenge, but what part can other phenomena, such as self-organization, play and what relevance are the other theories relating to complexity and chaos? Indeed, human intelligence, sociology and culture are relevant factors too - so, how do we employ them to shape the behaviour of EOs?

The hypotheses of this paper is that control is not something which is layered over the applications and infrastructure as an afterthought, but can be seen as a collective property which emerges from the interaction among the elements which have been deployed. The paper has identified some of the technical challenges of this approach and has examined some of the mechanisms which could be employed to ensure the collective, adaptive and secure behaviour of functioning, homeostatic information ecosystems such as those in EOs.

The assertions in the "Power to the Edge" book indicate that emergence will be a key driver in Edge Organisations. It is clear that interactions and relationships (and resulting emergent phenomena displayed by organisations) are routinely influenced by effective commanders, yet there has been little recognition of this. This paper concludes that, for Edge Organisations, we should not treat this as magic but instead base the employment of positive emergent phenomena on growing understanding of complex adaptive systems. Edge Organisations cannot be implemented without a better understanding of the roles that complex phenomena such as emergence play in their operation. The aim of this paper was to stimulate debate on this topic within the defence community and the author welcomes feedback and suggestions.

References

- [1] David S Alberts, Hayes, R.E. *Power to the Edge*. CCRP Publication Series, 2003.
- [2] Lieutenant Colonel J P Storr PhD at 9th ICCRTS, *The Command of British Land Forces In Iraq, March to May 2003*, Paper 068 - British Army
- [3] See extensive research material on complexity and related topics on the Santa Fe Institute site: <http://www.santafe.edu/sfi/research/focus/networkDynamics/index.html>
- [4] Complexity Science - Network of Excellence: <http://www.complexityscience.org/>
- [5] Guy Kirsch in Chap 1, *Artificial Life and Virtual Reality*. Wiley 1994.
- [6] Comments by MC Meigs, former Commander of SFOR in Europe, see <http://carlisle-www.army.mil/usawc/Parameters/03summer/meigs.htm>
- [7] John H Holland. *Emergence: from Chaos to Order*. p7 Oxford University Press. ISBN 0-19-286211-1. 1998.
- [8] Steven Johnson. *Emergence. The Connected Lives of Ants, Brains, Cities*. Penguin.
- [9] Morowitz, Harold J. *The Emergence of Everything : How the World became Complex*. Oxford University Press, 2002
- [10] Patrick Beautement. *Exploiting the Phenomenon of Emergence as a Force Multiplier*, at <http://www.tbt.org.uk/>
- [11] Ibid. Paul Davies.
- [12] Roger Lewin. *Complexity - Life at the Edge of Chaos*, Phoenix 1993.
- [13] Ibid. John H Holland. p215.
- [14] See Di Caro, G., and M. Dorigo. *AntNet: Distributed Stigmergetic Control for Communications Networks*. J. Art. Int. Res. 9 (1998): 317-365.
- [15] Beckers, R., O. E. Holland, and J.-L. Deneubourg. *From Local Actions to Global Tasks: Stigmergy and Collective Robotics*. In *Artificial Life IV*, edited by R. Brooks and P. Maes, 181-189. Cambridge, MA: MIT Press, 1994
- [16] In Nils A. Baas and Claus Emmeche. *On Emergence and Explanation*. Intellectica 1997/2, No.25, pp.67-83 (also SFI Working Paper 97-02-008. 1997.
- [17] Bricmont, J: *Science of Chaos or Chaos in Science?*. In: *The Flight from Science and Reason*, Annals of the New York Academy of Sciences 775, ed. by P.R. Gross, N. Levitt, M.W. Lewis (The New York Academy of Sciences, New York 1996) pp. 131--175
- [18] Ibid. Lewin p210
- [19] Ibid. Holland p227
- [20] Ibid. Paul Davies. pp20 and 76.
- [21] Donald M MacKay. In *Nature* 232, p679. 1986.
- [22] Claude E. Shannon, Warren Weaver. *Mathematical Theory of Communication*. Publisher: Univ of Illinois Pr (Pro Ref);. (December 1963). ISBN: 0252725484
- [23] Claude E. Shannon. *A Mathematical Theory of Communication'*. July and October 1948 editions of the Bell System Technical Journal. 1948.
- [24] J.C. Smuts. *Holism and Evolution*. MacMillan 1926.
- [25] Ibid p5. Nils A. Baas and Claus Emmeche.
- [26] Ibid. Paul Davies. p138 onwards
- [27] Karl Popper, John Eccles. *The Self and its Brain*. Berlin. Springer International. 1977.
- [28] Douglas R Hofstadter. *Godel, Escher, Bach: an Eternal Golden Braid*. p713. Penguin. ISBN 0-14-00.5579-7. 1980.

- [29] Christopher G Langton. *Life at the Edge of Chaos*. On pp 41-49 of A-Life II. Addison Wesley. ISBN 0-201-52571-2. 1992.
- [30] Stephen Wolfram. *Cellular Automata as Models of Complexity*. Physica 10D. 1984.
- [31] Network Enabled Capability: <http://www.mod.uk/issues/nec>
- [32] Network-Centric Warfare: <http://www.dod.mil/nii/NCW>
- [33] M Kirton et al. *The Coalition Agents Experiment: Network-enabled Coalition Operations*. In Journal of Defence Science (Special edition on NEC), Vol 8, No 3, Sep 2003.
- [34] Coalition Agents Experiment Web site: <http://www.aiai.ed.ac.uk/project/coax>.
- [35] Control of Agent-Based Systems Web site: <http://coabs.globalinfotek.com>.
- [36] James Hollan, Edwin Hutchins and David Kirsh. *Distributed Cognition: Toward a New Foundation for HCI Research*. ACM Transactions on Computer-Human Interaction, 2000
- [37] Edwin Hutchins. *Cognition in the Wild*. MIT Press, Cambridge, MA, 1995
- [38] Chin M, Clothier J, Carthigaser M. *Command and Control Capability Assessment and the Criticality Issue*. DSTO. Presented at the 1997 ICCRT Conference. 1997.
- [39] Dave Allsopp and Patrick Beutement. *Decision Desktop Data Sheet*. See: <http://www.aiai.ed.ac.uk/project/coax/presentation/posters-29oct02/2002-qinetiq.ppt>.
- [40] Bradshaw, J M, Boy, G, Durfee, E, Gruninger, M, Hexmoor, H, Suri, N, Tambe, M, Uschold, M, and Vitek, J., editors. *Software Agents for the Warfighter*. ITAC Consortium Report. 2002. Cambridge, Massachusetts, AAAI Press/The MIT Press.
- [41] Bradshaw et al. *Making Agents Acceptable to People*. In Intelligent Technologies for Information Analysis. N. Zhong and J. Liu (Eds.). Springer 2004.
- [42] Jennings, N R. *An Agent-based Approach for Building Complex Software Systems*. 2001. Communications of the ACM, 44:4, 35-41.
- [43] Bradshaw, J M, Duffield, S, Benoit, P, and Woolley, J D. *KAOs: Toward an industrial-strength open agent architecture*. 1997. In Software Agents, AAAI Press / MIT Press, Cambridge, Massachusetts, editor J M Bradshaw. Pages 375-418.
- [44] Colonel John D. Rosenberger. *The Inherent Vulnerabilities of Technology: Insights from the National Training Center's Opposing Force*. U.S. Army, 2002
- [45] Michael R Lwin. *General Tzu's Army: OPFOR of the Future*. JFQ, Spring 1997.
- [46] See the so-called 'autonomic computing' research work being undertaken by IBM at <http://www.research.ibm.com/autonomic/overview/>
- [47] P Beutement. *Network-Centric Security - Approaches to Collective Run-Time Adaptation*. Adaptive and Resilient Computing Security Workshop, Santa Fe Institute. Nov 2003. Available from <http://www.tbt.org.uk/library/lib-index.htm>
- [48] Ilya Prigogine. *From Being to Becoming: Time and Complexity in the Physical Sciences*. p214. San Francisco - Freeman. 1980.
- [49] Joseph Ford. *How Random is a Coin Toss?* p4. Physics Today. April 1983.
- [50] Paul Davies. *The Cosmic Blueprint*. p31. Penguin ISBN 0-14-024362-3. 1995.
- [51] Brook, Arnold, Jackson and Stevens. *Systems Engineering - Coping with Complexity*. p94. Prentice Hall ISBN 0-13-095085-8.