2006 CCRTS
THE STATE OF THE ART AND THE STATE OF THE PRACTICE

The Implications of Complex Adaptive Systems Theory for C2

Topics: *C2 Concepts and Organizations (preferred)*
C2 Analysis
C2 Architecture

Anne-Marie Grisogono

Land Operations Division,
Defence Science and Technology Organisation,
PO Box 1500, Edinburgh 5100
South Australia
Australia
phone: +618 8259 6532
fax: +618 8259 5055
cell: +614 0907 6684
Anne-Marie.Grisogono@defence.gov.au

The Implications of Complex Adaptive Systems Theory for C2

Dr Anne-Marie Grisogono
Research Leader Land Systems, LOD, DSTO
Anne-Marie.Grisogono@defence.gov.au

Abstract

The study of Complex Adaptive Systems (CAS) has developed within a wide range of subject domains over the last couple of decades, spanning the biological sciences, economics, organisational science, public policy, environmental sciences, computer science, cognitive and social sciences, and lately, defence sciences. We have been researching how application of a CAS perspective to the most pressing and complex problems that defence faces can provide more effective tools and techniques to enable higher levels of success in dealing with these challenging problems. This approach has proved very fruitful and has generated insights that could lead to implementable and testable strategy options in a wide range of defence areas – from strategic policy, the capability development process, and defence enterprise management to the design and evolution of complex defence systems and the command and control of tactical to strategic levels of operations.

In this paper we will focus on the implications of CAS theory for C2, drawing on the understanding we have developed of what it is possible to do in the face of complexity, how adaptive mechanisms arise spontaneously in complex systems, how we may recognise them and influence their operation to better align with our purposes, and how we may develop additional adaptive mechanisms to foster more effective outcomes. The CAS we will address include not just the complex networked systems within our own forces, but also those of our allies and adversaries, and those existing in the overall environment in which we operate. All these systems influence both what we are expected to do and what we are able to do, therefore understanding how the adaptive mechanisms already operating in them shape their behaviour and how to harness those mechanisms to our purposes is potentially a very valuable and powerful strategy.

Introduction

Complex Adaptive Systems – and what they might mean for defence – has been the theme of a Long Range Research program in DSTO's Land Operations Division since early 2003.

The program was initiated for two reasons.

Firstly, complexity was a recurring source of the most challenging defence problems that we had to deal with, and the adaptivity of Complex Adaptive Systems seemed to hold the promise of dealing more successfully with them, in a way that resonated with the increasing attention being paid in defence circles to the notions of agility, robustness, resilience, responsiveness, transformation – all of which are forms of adaptivity.

Secondly, a series of attempts to exploit the rapid advances in information technology developments of the preceding decade were resulting in one techno-utopian initiative after another[1], each of which placed a premium on some aspects of increased connectivity, interaction, information and knowledge, but arguably have so far delivered mixed results at best and have not yet produced the sought-after transformative benefits in a cost-effective way. The link with Complex Adaptive Systems is evident when one considers that increasing connectivity, interaction and information also inevitably increase complexity, in an exponential fashion. We suspected that the rising complexity was making things harder at a faster rate than the technology-based solutions could alleviate. If this is indeed at least part of the reason for the disappointing progress being made, then learning more about how adaptivity deals with, even thrives on complexity, could help us find our way.

The approaches taken in this program draw on several parallel strands:

- learning from successful natural CAS, primarily biological – including evolution of species, the immune system, epidemiology, genomics, learning mechanism, and ecosystems
- mathematical modelling and analysis, formal methods
- conceptual development of CAS-based techniques and approaches transferable to real problems, and
- experimentation with real defence problems.

[1] Revolution in Military Affairs (RMA), knowledge engineering, the Knowledge Edge, Network Centric Warfare (NCW), etc

It has quickly become apparent that studying CAS was a very productive and profitable direction to take[2], and that the insights and concepts that are produced are relevant, applicable, often challenge established thinking, but do so in an integrative and constructive way. The consequences are far-reaching and touch almost every aspect of defence – how we think about the roles of technology, people, organisation and process, how we develop the force, train and maintain it, but most especially, how we use it – from the highest strategic level of examining the role of defence as an instrument of national power on the global stage, to developing campaign plans in complex situations that will produce the long term outcomes we seek, to tactical decisions and operations.

This paper will concentrate on developing the key implications of CAS for Command and Control, so we begin with a brief discussion of some essential aspects of CAS and C2 concepts which we will draw on in the rest of the paper. The following section then presents a summary of what we have learned so far about the adaptivity of CAS, in the form of a conceptual framework for adaptivity. In the final section we discuss the extent to which aspects of this framework are already practised in defence, and identify opportunities to engender more adaptivity into the command and control of defence forces and to increase their success in dealing with complex situations. A particular nexus is drawn with Effects Based Operations, and the agenda for further research and development is discussed in the concluding remarks.

Some Essential CAS and C2 Concepts

1. Complexity and Complex systems

There are many definitions of complexity ranging from very abstract and mathematical to descriptive and pragmatic. Precise definitions are often difficult to apply and justify, particularly at the boundaries (exactly what is or is not complex?), and different rigorous definitions may imply different boundaries. Moreover, formal approaches may seem obscure to the non-specialist and may not readily illuminate the salient features.

Therefore from a pragmatic point of view we adopt an operational approach – we consider a system to be complex when:

i. *Causality is complex and networked*: i.e.simple cause-effect relationships dont apply – there are many contributing causes and influences to any one outcome; and conversely, one action may lead to a multiplicity of consequences and effects

ii. *The number of plausible options is vast*: so it is not possible to optimise (in the sense of finding the one best solution in a reasonable amount of time),

iii. *System behaviour is coherent*: there are recurring patterns and trends, but

iv. *The system is not fixed*: the patterns and trends vary, for example, the 'rules' seem to keep changing – something that 'worked' yesterday may not do so tomorrow, and

v. *Predictability is reduced*: for a given action option it is not possible to accurately predict all its consequences, or for a desired set of outcomes it is not possible to determine precisely which actions will produce it.

Another way of putting it is that dealing with a complex system generally is a problem that has high *task complexity* – a concept we define as the ratio of the number of ways of getting the wrong outcome to the number of ways of getting it right.

Complex Systems with these properties generally consist of many interacting elements, and the system behaviour which results is more than just a linear[3] aggregation of element behaviours. The additional aspects of the system behaviour are often collective properties – i.e. properties which describe some aspect of a set of elements, but which is not a property of any one element. Moreover, Complex Systems often have a layered hierarchical structure, each successive layer arising or emerging from interactions between the dynamic

[2] In its two and a half year lifetime to date, the task has produced a large number of publications, launched a "CAS for Defence" workshop series, engaged with real defence problems in a number of areas, and fostered the development of a broad Australian and international research network across defence, other government research organisations, industry and academic institutions, within which we are now engaged in a number of active collaborations.

[3] The term non-linearity is often used in relation to complex system, but has two distinct senses: effects not being proportional to causes, and a multiplicity of choices to be made at points in a process. For a fuller discussion please contact author – paper in preparation.

patterns of the layer below. These dynamic and collective properties are described as *emergent*[4]. There are thus two complementary senses in which a property may be emergent:

a. it may emerge over time, producing new behaviours or structures that were not there before – dynamic emergence, and

b. it may emerge at a macro-level of description from what is happening at the micro-level – for example superconductivity is a collective property of a material resulting from the quantum mechanics of the constituent particles.

In the real world, Complex Systems are open systems, that is, they exist in a *context* with which they interact – it is not possible to draw a sharp boundary around them through which there will be no interactions at all. Attempts to do this are misguided, and can only lead to ever larger and larger 'systems' which therefore become harder and harder to comprehend – ultimately the entire universe must be considered! Choosing the right boundary is therefore an important issue in dealing with complex systems, as is understanding and dealing with the inevitable interactions with the context through the boundary. These often have significant influence on system behaviour.

Complexity should not be confused with chaos (where changeability is very high, but there is no coherence) or complicatedness (where changeability is very low, but there is high coherence). Complex systems occupy the intermediate ground between these two extremes, and so are sometimes described as being 'on the edge of chaos'.

2. What are Complex Adaptive Systems?

This term is used to describe those complex systems which have the additional important property of being adaptive – i.e. the stucture and behaviour of the system changes over time in a way which *tends* to increase its 'success'.

This requires that

i. there is a concept of 'success or failure', (technically known as '*fitness*[5]'), for the system in the context of its environment;

ii. there is a source of *variation* in some internal details of the system, and

iii. there is a *selection* process, i.e. the system preferentially retains/discards variations which enhance/decrease its fitness, which requires…

iv. some way of evaluating the impact of a variation on the system's fitness – generally achieved through some kind of external *interaction and feedback*.

v. Thus over time the system generates and internalises variations which tend to increase its fitness or success – amounting to incorporation of information[6] into the system.

In the most general sense, such a system is interacting with aspects of its environment through taking in 'inputs' or sensing, creating 'outputs' or taking actions, and some kind of internal processing in between the sensing and the acting. The details of how these three basic functions operate change over time as a consequence of the system being adaptive.

So a system which has the property of being adaptive is a system which is always changing by virtue of this adaptive process which is executing. We note that the process is a closed loop and that, because introducing variation will introduce harmful errors much more frequently than useful innovations, the selection process must serve two purposes: the elimination of fitness-decreasing variations most of the time, as well as the retention of the occasional useful fitness-enhancing variations.

Complex Adaptive Systems (CAS) have all the properties of Complex Systems, and in addition, display some characteristic hallmarks of adaptivity:

---

[4] Emergence in complex systems is a large field of research and is being studied in DSTO's CAS program in collaboration with CSIRO to understand how it may be harnessed to foster more desirable outcomes and properties in designing, managing and interacting with complex systems.

[5] The term 'fitness' comes from the study of evolutionary systems, but is generalised in the study of CAS to cover whatever it takes to define the system's success or failure. This is a critical element of adaptivity and will be discussed in more detail below.

[6] Over time the incorporation of the selected variations improves the system's abilities to succeed in its environment and therefore the nature of the variations tell us something about the nature of the environment – it is in this sense that we say that adaptation results in the system incorporating useful (fitness-relevant) information; there does not need to be any overt representational information in the system although more complex and sophisticated systems may certainly exploit more explicit forms of information encoding.

- ‘intelligent’ context-appropriate behaviour – discovery and exploitation of advantages available in the system’s environment, and recognition and appropriate response to threats to the system;
- resilience – quick recovery from shocks and damage;
- robustness to perturbations – core functionality is maintained;
- flexible responses – the system has a range of different strategies towards any given end;
- agility – rapid change of tack to more effective behaviours when needed;
- innovation – leading to creation of new strategies and new structures; and
- the system learns from experience – relevant information about past contexts is incorporated into the system in such a way that the system’s future behaviour is likely[7] to be more effective.

This last property, of learning from experience, is a defining characteristic of CAS, and distinguishes complex adaptive systems from those which are simply reactive.

There are many parameters that characterise how adaptation is operating in a particular system. The context that the system operates in also plays a significant role in the behaviour of a CAS. Not only is there a range of interactions between system and context resulting in the exchange of energy, materials, and information, but the development and nature of its adaptivity properties result to a large extent from the pressures the context places on it. These are critical issues for understanding how to exploit adaptivity and how to increase its effectiveness, and are therefore current research thrusts in DSTO’s CAS program.

Many of the properties of CAS come in complementary pairs with some tension between them, and may seem paradoxical at first sight, for example:

- robustness to damage requires that some changes in the CAS are inhibited and repaired while innovation requires that some changes in the CAS are amplified,
- coherent behaviour requires that some aspects are stable and persistent while context-appropriate behaviour requires some aspects that are sensitive to influences which may tip the balance one way or the other,
- ‘intelligent’ behaviour in a complex environment may require specialised system elements, while flexibility and resilience may put a premium on multi-purpose system elements,
- effectiveness requires both competition between elements of a CAS to refine individual strategies, and cooperation between elements to produce collective strategies,
- effectiveness in dealing with the current environment requires conformity of system properties to those that are most relevant and effective to present challenges, while effectiveness in dealing with a changing environment requires diversity of system properties.

Developing a better understanding of the dynamic processes which control and balance such paired properties will contribute to our growing ability to work effectively with and within CAS.

There are also a number of other concepts and terms (for a fuller discussion see general introductions[8] to CAS theory) that are used in discussing CAS, such as ‘tags’ which allow CAS elements to differentiate and specialise, ‘internal models’ which allow a CAS to do some prediction, and so on, but in the context of the present paper we will focus on those elements most relevant to adaptation only.

3. What We Have Learned so far from Studying CAS

One of the most fertile areas of CAS study has been the systematic analysis of how adaptivity works in naturally occurring CAS, and the identification of its principal features and mechanisms. This has allowed us to formulate a generic model of adaptation which could in principle be applied to defence systems in a number of ways:

- firstly as a template for identifying the informal adaptive loops that arise spontaneously in any complex sociotechnical system, coexisting with, and often undermining the deliberate formal adaptive mechanisms;

[7] Being adaptive does not guarantee success of course. There are many variable aspects of how adaptivity is implemented and the degree of success or failure depends on how well these are chosen in relation to the challenges presented by the context.
[8] For example *Complexity: The Emerging Science at the Edge of Order and Chaos* Waldrop, M. Mitchell 1993 (1992) London: Viking.

- secondly for analysing the factors determining the effectiveness of both formal and informal adaptive mechanisms, and suggesting leverage points for modifying them; and
- thirdly, for designing new adaptive mechanisms to deal with anticipated future pressures and changes.

Successfully applied, all three of these approaches would lead to significant insights into how defence systems should be architected, and how decisionmaking, C2 processes, information and policy should be organised and managed for the required degrees of robustness and adaptivity.

In order to facilitate these analyses, the generic model is being further developed in several directions to create a conceptual framework[9] which generates a rich set of detailed adaptive models mapping more directly onto real world applications.

While this framework is still at an immature stage, the eventual intention is to create templates, tools and processes to assist in the recognition and tuning of existing adaptive mechanisms, and in the design and implementation of effective new ones. In the following subsections we sketch out an overview of the elements of this conceptual framework – comprising a number of classifications of adaptive mechanisms (types, classes, levels and scale), and a set of factors that characterise the health of an adaptive mechanism and influence its effectiveness.

i. *Types of Adaptive Mechanisms*

The elements of adaptivity listed at the beginning of Section 2 – variation, interaction with feedback, and selective retention of fitness-enhancing variations in the sensing, processing and action functions of the system – are derived from the generic model of adaptation.

Studying real world examples of natural adaptivity shows us how those generic elements translate into working adaptive systems, and brings to light both the strengths and some limitations of the two different basic types of adaptivity that occur in nature – evolutionary adaptation (which works on populations) and learning adaptation (which works on individuals). The two types are complementary in many ways. Each can be applied in designed adaptation, and each presents different problems for the designer.

A major design problem that is common to both is how fitness is to be defined and measured.

In natural evolutionary systems this is not a problem since fitness is equivalent to surviving the selection process, but learning systems operate over a much faster timescale, generally too fast for the impact on actual fitness to be observed and provide the feedback needed to drive adaptation. Learning therefore generally makes use of proxies for fitness – observable consequences of the variations being tried which materialise fast enough to drive the adaptive cycle, and correlate well with eventual impact on fitness, but are not in themselves actual measures of fitness. The possibility exists therefore of proxies being used which correlate well only over the short term, but diverge widely from longer term fitness impacts, or proxies which may have correlated well at one point of time, but no longer do so because of external changes in the system's context.

Designed evolutionary systems (eg employing genetic algorithms) face a similar problem of defining a meaningful and measurable fitness function if they are intended to inform real world systems. We stress this point because it is one of the most critical issues to address if we wish to exploit the power of adaptation. Having a sophisticated adaptive mechanism with a poorly defined fitness measure which doesnt correlate well with what will actually be judged as success in the longer term, is like having a high tech car with all sorts of drive and control systems, but only a crude map or the wrong destination marked on it. One is hardly likely to end up in the right place! We will return below to discuss the implications of this issue for C2.

It is also possible to conceive of other types of adaptive mechanisms which combine the strengths of the natural types and augment them with features that do not occur in nature. This opens up the possibility of much more powerful adaptation than we have so far tapped.

These natural, and hybrid or augmented types of adaptation will be further discussed elsewhere but a brief summary of their essential characteristics is tabled below:

	Evolutionary	**Learning**	**Hybrid or Augmented**
applies to	populations	individuals	either

[9] Discussed in more detail in *Co-Adaptation* AM Grisogono, invited paper 6039-1 Complex Systems Conference, SPIE Symposium on Microelectronics, MEMS and Nanotechnology, Brisbane December 2005

outcome	produces new design features	improves use of fixed design	may be able to do both, or do either better
time for one loop to execute	period between generations – generally slow compared to timescale of actions	period for one action (sense-process-decide-act) loop, plus the associated learning (observe action consequences – process – make changes) loop.	could be accelerated
parallelism of processing through interaction	highly parallel – every member of the population is a simultaneous experiment 'evaluating' the fitness of one set of variations	serial – an individual system or organism experiments with one strategy at a time	could use learning mechanism to create directed evolution, and evolutionary strategies to improve learning. Could also parallelise learning through either parallel processing in single individual, or through networking a population of learning systems.
context sensitivity	in retrospect only – through some variations turning out to be fitter in the context than others	in anticipation – i.e. before choice of action or response, as well as in retrospect through feedback from consequences of action.	could extend context sensitivity to influence design choices as well as action choices.
alignment of fitness and selection mechanism	100%	highly variable	could improve alignment in learning systems by developing better proxies for fitness to drive selection.

ii. *Classes of Adaptive Mechanisms*

The overall result of any adaptivity is that the system learns from experience – information about the past is encoded into the system and is used in future to be more effective in the face of changes and stresses. The next direction in which the generic model of adaptation has been developed is in identifying a natural classification of adaptive mechanisms based on the type of changes that they are intended to deal with. We distinguish them by the timescale over which they operate, the effects-scale concerned and whether the changes are internal to the system, external or both.

This analysis produces four classes, summarised below:

1. ***Flexibility***: Effectiveness across different contexts (the ability to maintain effectiveness across a range of tasks, situations, and conditions) – i.e. the structure and capability of the force can be reconfigured in different ways to do different things, under different sets of conditions

2. ***Resilience***: Overcoming losses, damage, setbacks (the ability to recover from or adjust to misfortune/damage, and the ability to degrade gracefully under attack or as a result of partial failure) – i.e. the core functions of the force continue to achieve essential levels of capability when individual elements are disabled one or more at a time

3. ***Responsiveness***: The ability to react to a change in the environment in a timely manner – i.e. in the context of a set of strategies that are already being executed, and a change occurring that creates a new threat or opportunity, responsiveness is the capacity to recognise and deal with the new threat or opportunity as effectively as if there were ample time to plan and prepare for it.

4. ***Agility***: The ability to recognise when to shift from one strategy to another and to do so easily – producing a rapid change of tack to more effective behaviours when significant external and/or internal changes arise requiring significant and different responses.

iii. *Levels of Adaptive Mechanisms*

Adaptation can be applied at five levels, each successive level adding more power to deal with challenging complexity:

1. at the ***'action'*** level - creating a more contextually appropriate series of adaptive actions in the world, but within the constraints of existing sense, process and act capabilities
2. at the ***'learning system'*** level - using adaptation to expand sense, process and act capabiltiies in useful ways
3. at the ***'learning-to-learn'*** system level - using adaptation on the learning mechanisms themselves and thus improving the way in which adaptive actions are produced
4. at the ***'defining success'*** level – applying adaptation to the difficult problem of articulating sufficiently precise and actionable measures of success, (in order to achieve this it will be necessary to have access to some more accurate measures of success - usually only available in slowtime and in retrospect), and
5. at the ***'co-adaptation'*** level – addressing the interactions between multiple adaptive mechanisms and applying adaptation to the parameters describing the distribution of roles, resources, authorities and responsibilities between them.

   Co-adaptation needs to be further understood in two ways:

   - the fact that our systems are interacting with each other and that there are therefore additional degrees of freedom for creating adaptivity and more effective behaviour at the SoS level – this is a system management and design issue, and
   - the fact that our systems are interacting with other systems that are also adaptive, and so there is a more complex and intelligently reactive and adaptive context that our systems have to deal with, which might require more sophisticated treatment.

These five levels of adaptation are more extensively discussed elsewhere[10], but the essential difference between the levels is what part of the system is being subjected to adaptive change. As we move through the five levels, we are introducing variation and fitness-linked selection successively into

1. the parameters that characterise the operation of the existing sense, process, and act capabilities of the system;
2. the parameters that determine the scope of the sense, process, and act capabilities of the system;
3. the parameters that determine the effectiveness of the learning mechanism;
4. the internalised selection criteria; and
5. the parameters that determine the distribution of roles, resources, authorities and responsibilities between the component systems of a system of interacting CAS.

iv. *Scale of Adaptive Mechanism*

Each of these types, classes and levels of adaptation can also be applied at different scales – from the level of an individual human, or a single self-contained system, through small teams and larger groupings up to organisations, distributed systems-of-systems and full scale enterprises.

Adaptation at the scale of an individual is generally very effective because of the inherent adaptivity of humans – provided that the adaptivity being called for is within the scope of what the individual is equipped to deal with, in terms of authority, motivation, capability and information.

A similar argument can be made for the adaptivity of a tight hardened team, but when we move up the scale to larger and looser couplings of capability elements and people we often find greater rigidity and inertia.

If we examine the adaptive mechanisms that are intended to be operating at those scales and compare them with our generic model, the reasons for the lack of adaptivity are evident – the elements are generally now distributed over different parts of an organisation or even across different agencies, there may be inconsistencies between how these operate, the loop may be broken or the flow restricted, and so

[10] *ibid*

on - all the problems discussed in ii. above become more likely to occur as the scale of the adaptivity required increases.

v. *Health of an Adaptive Mechanism*

Each element of an adaptive mechanism needs to be just right in order for adaptation to be successful at improving fitness, and maintaining it in the face of environmental change, both with respect to its own functioning and in its relationship to the other elements. If any one element is defective or not properly linked to the others, the adaptation will fail in various ways, for example:

- if there is not enough variation (both in scope – what gets varied, and extent – how much it gets varied), the system may not be able to find new strategies or structures which would maintain or improve its fitness relative to competitors; however if there is too much variation it may not get to recognise the value offered by a new variation before it gets replaced by yet another version;
- if the variations that do arise are not experimented with through fitness-relevant interactions that create feedback then the adaptive mechanism is not able to distinguish between those that enhance fitness and those that dont; timeliness of feedback is in tension with comprehensiveness here – fast specialisation may miss some downsides, but the compromises forced by robustness to many stresses may reduce some benefits;
- if the feedback from the interactions is not correctly linked in to the selection process then the system may fail to retain the winners and discard the losers; and
- if the selection process is not well-aligned with the system's fitness then it may incorrectly label the winners and losers.
- Moreover, adaptation is a dynamic process and the various timescales need to be well-matched. The total cycle time is the sum of the times needed for each step, and in order for the adaptation to cope with changes in the environment, this cycle time needs to be much faster than the time constant of change in the environment – *how* much faster depends on other parameters that determine the difficulty of finding areas of greater fitness, the topology of the fitness landscape, and the rate at which it can be traversed.

These are some of the factors that influence the effectiveness of an adaptive mechanism at producing context-appropriate system behaviour and outcomes. Analysing all of them in a systematic way and developing diagnostic tools to assess the health of an adaptive mechanism is a current active area of research. One of the insights that is emerging is that there is a large and complex set of interrelated parameters that characterise an instantiation of an adaptive mechanism and that the task of determining their values is itself a complex task which is best handled through an adaptive process.

4. Command and Control

There are probably as many definitions of C2 as there are of complexity. Given that the intended audience of this paper is the C2 research community that attends CCRTS, it is somewhat presumptous to explain or define C2 to that audience.

However, since we aim to explore a broad range of CAS implications for C2, we need to base that exploration on the most generic interpretation of what C2 is, consistent with an understanding of the CAS that our own and others' forces consist of, and to avoid predicating it on specific current C2 concepts which might foreclose some interesting options.

In an earlier paper[11] we asserted that in the most general sense:

> *C2 is about determining objectives, and under constraints, choosing and causing actions to happen which increase the probability of achieving those objectives over time.*

Elaborating what we mean by objectives entails an exploration of strategic intent in what we might call the 'space of possible futures' and the assignation of value (from 'very desirable' to 'avoid at all costs') to different parts of that space, and then the identification of the roles that defence must or might play and the outcomes it should produce or avoid in its contribution to the generation or prevention of the preferred or dangerous futures.

Moreover, recognising that ultimately armed conflict is about a clash of wills, and that the objectives of each party are to exert a shaping influence on the will of the other parties[12], i.e. to create an outcome in the cognitive domain, we further assert that:

---

[11] *A Generic Framework for Generating and Exploring C2 Concepts* AM Grisogono, 9th ICCRTS, Copenhagen, 2004

[12] Each party also seeks to align the will of the various levels in their own force.

The challenge of C2 is to choose and cause to happen those actions which, through shifting perceptions and assessments in the cognitive domain, bring about outcomes which contribute effectively to shaping the future that unfolds in the directions set by higher strategic intent.

This suggests the following primary roles for the command and control of defence:

C2 Role 1 having understood the higher strategic intent in the space of possible futures, negotiating the defence role in implementing that intent,

C2 Role 2 determining the outcomes defence needs to produce or avoid to implement that role, alone or in cooperation with other agencies

C2 Role 3 choosing defence actions to undertake to produce or avoid those outcomes,

C2 Role 4 causing those actions to happen, and

C2 Role 5 monitoring and continuously re-assessing all of the above in the light of unfolding events.

We will now in the next section bring together these threads of articulating a generic concept of C2, a generic model of adaptation, and a theoretical framework for thinking about how adaptation may be implemented and exploited, and discuss some of the implications and consequences that result for C2.

Putting it all together… Implications of CAS for C2

5. Complexity and CAS in Defence

From the above descriptions of complexity and CAS it is not difficult to recognise that the main business of defence is dealing with complexity in various forms, and that there are adaptive mechanisms in defence systems and operations wherever we look. Many of these are spontaneously arising out of informal interaction loops rather than deliberately designed as CAS, since any grouping of interacting elements which has feedback loops that tend to stabilise particular patterns of interaction and to trigger changes over time in response to external stresses and stimuli, can be thought of as a CAS.

Not only are the systems of our own and others' forces and agencies themselves CAS, but various subsystems and groupings of them are also CAS. In fact every individual human is a CAS in his or her own right – and going in the other direction of scale, interactions between defence forces can also be seen as occurring within various larger CAS containing other CAS such as political, dilpomatic and economic systems.

Once we start to see the world through these perspectives, we begin to see more and more adaptive mechanisms operating in any complex organisation. Individual humans and capability elements besides being CAS in their own right, are also likely to be part of many different higher level adaptive mechanisms, simultaneously contributing to different classes, levels and scales of adaptivity.

Each one of these adaptive mechanisms operates in its own way, and has its own implicit 'fitness' or internalised success and failure measures, which can be deduced by observing how selection *actually* operates in the system – i.e. what 'criteria' the system seems to be using in selecting whether to adopt or reject internal variations of the system that are experimented with – which may often be at variance with the formal stated objectives or success measures.

How do these insights help us?

We have many options for dealing with complex systems – familiar ones we have used traditionally or instinctively for a very long time, but also some new ones that we are just beginning to grasp the potential of.

CAS theory helps us to construct a conceptual framework which both encompasses the existing approaches and suggests the new possibilities, but importantly also helps us to formulate and address questions about what works best under what conditions, how to increase the effectiveness of any particular approach, how to choose the best approach for particular situations or problems, and how to actually implement it at any particular scale.

In this paper we have only lightly skimmed over the theory, and the applications to these questions, but we can highlight a few major implications:

- On a coarse level our options for dealing with a particular CAS (whether our own, or not) include

- observation without intervention – simply seeking to understand it better and to have some insight into why it does what it does, or what it might do in the future;
- shaping its behaviour through indirect influences of the context it operates in;
- more direct interventions in the system, by altering some its component elements or their interactions, or by control of resources or information available to it, and
- completely destroying it – but to do that successfully requires a good understanding of its sources and mechanisms, else we risk simply stimulating them to recreate it and possibly replace it with more resistant forms.

CAS theory can assist us in each of these options – recognising and mapping adaptive mechanisms, identifying influence pathways and also in putting in place an adaptive approach to gradually refine the quality of our understanding of the CAS in question, or the effectiveness of our shaping or direct intervention actions. It can also help us understand the process of emergence that spawns new CAS within a larger CAS, so that we can identify the limits to our ability to eliminate or contain CAS that behave in particular ways.

- Many of the spontaneously arising adaptive mechanisms in our systems may not even have been recognised as such, yet they exert influence over our systems' behaviours. CAS theory can help us identify them, map them out and deduce their 'fitness' functions.
- Many of the spontaneously arising adaptive mechanisms in our systems may have implicit 'fitness' which is not well-aligned with stakeholder intent. CAS theory may assist us in identifying interventions to modify the workings of those mechanisms to improve alignment.
- Many of the deliberately introduced adaptive mechanisms in our systems are dysfunctional. Analysing them in terms of the conceptual framework can help us understand the reasons and how to redress them.
- A particular example of this is the recognition that if we chose to adopt an intentionally adaptive stance in our own organisations, then wherever projections, predictions or estimates were made (generally with the intention of supporting adaptive action) we would instigate a learning mechanism that would gather information as the future unfolded about how accurate (at an appropriate level of detail) those projections turned out to be, and with hindsight, would analyse how a better prediction might have been made, and the impact of doing so, and would institute the required changes where the cost and benefit would warrant it.
- Many of the spontaneous and deliberate adaptive mechanisms in our systems are in conflict with each other – CAS theory may help us understand how to achieve stakeholder intent better through 'tuning' the existing formal and informal adaptive mechanisms and replacing conflicts with synergy.
- Understanding the adaptive mechanisms operating in a particular system can reduce the baffling unpredictabilities of complex systems – and moreover such understandings can themselves be developed in an iterative adaptive way – indeed one can argue that this is the *only* way it can be done.
- When what constitutes longer-term success and failure are well understood, and are difficult to achieve because of the complexity of the situation then we may be able to do much better by exploiting adaptive approaches. The benefits are that one can start with a modest appreciation of the situation and improve by exploiting level 2 and 3 adaptation.
- Even when the longer-term success measures are not well-understood we can make good progress by invoking level 4 adaptation.

We will discuss some of these points in more detail in the following sections.

6. What are the implications for control, influence and predictability?

A common reaction to hearing about complexity and CAS is concern that it is all too hard, that complexity is equated with chaos and loss of control, and is to be avoided at all costs. Yet it is clear that we cannot avoid it, and it is also clear that in spite of complexity in all but the simplest systems around us, things do work reasonably well a lot of the time, and sometimes exquisitely so. Complexity is therefore obviously not the same as chaos. Furthermore, we are all competent (albeit to varying degrees) in our day to day dealings with other complex adaptive systems, i.e. people, and although we are not able to fully and deterministically control them (for which we should be relieved rather than frustrated!), we have many strategies for influencing their behaviour. The point being made here is that in addressing CAS in defence, we are not talking about something totally new and threatening, but rather about extending our already extensive capabilities in dealing with CAS into domains where we were previously, but no longer are, able to cope with simpler strategies.

We have argued that in dealing with CAS we cannot expect the high degree of predictability and control that an engineered system usually provides, but we can still have influence on outcomes and *some* confidence in our estimates of what may happen. Predictability is not an all (strictly deterministic) or nothing (totally chaotic) quality – it may exist in varying degrees along a continuum, and how much we have depends on what we are asking for. There are a number of ways of sliding along that scale of predictability, for example by varying:

- the level of detail of the prediction being sought, or precision of effect we require;
- the choice of which particular effects we are seeking to predict;
- the time horizon over which prediction is made;
- the required confidence of prediction; and
- generalising the notion of predictability to include inferences in different directions in time, whether
  - forwards (= predict in advance i.e. what happens next?)
  - backwards (= explanation in retrospect i.e. what must have happened?), or
  - sideways (= retrospective 'prediction' or inference of existence of aspects that had not yet been observed, i.e. what else must be happening now?)

So in other words, a given complex system or complex adaptive system may seem quite unpredictable, for example if are seeking high confidence, highly detailed predictions a long way into the future about system behaviour that is influenced by multiple contingent factors, or the same system may seem quite predictable, if we are instead seeking predictions for example about aspects that are stable, or short term, or only require a probabilistic assessment. The predictability of a complex adaptive system also depends to a great degree on the quality of understanding of the system that we bring to bear on it. A good example is the confidence that a marketing analyst may have about the average change in buying patterns in a well-known target market that might result from a particular advertising campaign, while being completely unable to predict what an individual consumer will buy in six months. In essence, the better we understand the adaptive mechanisms operating in a CAS the less baffling its behaviour will appear.

Predictability is therefore as much a property of the one doing the predicting, as it is of the system whose behaviour is being predicted. The inherent complexity of the system will limit the range of predictability, while the choices made by, and the quality of understanding of the system achieved by, the predictor will influence the degree of predictability that is actually obtained.

In simple systems, predictability results from understanding cause-effect relationships and being able to say that if this cause is removed, then that effect will not occur, or if this cause is changed by a known amount then the resulting effect will be changed by a calculable amount. But in reality even simple cause –effect relationships do not exist in isolation. One can ask what caused the "cause" and what are the effects of the "effect" and one quickly recovers a complex causal network, and with it, the need to distinguish between what have been described as proximate causes, remote causes, precipitating causes, necessary causes, sufficient causes, contributing causes, reciprocal causes, and so forth[13].

To restore some clarity to the discussion, one should focus on what is really sought, the ability to anticipate the consequences of actions that might be taken, and conversely, the ability to choose actions to produce desired effects, *and* not produce undesired effects. We note that even in simple systems our ability to do this is achieved through restricting the focus of attention to particular classes of effects, to a limited timescale over which we address consequent effects, etc.

Our ability to apply similar approaches in systems of higher complexity[14] is reduced in a number of ways:

[13] A number of authors, starting with Aristotle, have written about these and other types of causes, and the whole topic of causation and control will be discussed in more detail in a forthcoming publication. Please contact author for more information if required.

[14] Recalling the operational definitions of complexity given in the first section of this paper.

- even when all the relationships (both within the system and between the system and its context) are deterministic, it does not take very many interacting influences on system behaviour for the computation of their net effect on behaviour to become intractable;
- some relationships may be stochastic, or deterministically chaotic – i.e. extremely sensitive to the exact values of parameters which can never be known with sufficient precision to provide any more than short-range estimates;
- some parameters that influence outcomes may be inaccessible for measurement;
- since complex systems are characterised by complex networked causation (as opposed to simple cause-effect relationships) with multiple interacting causes and influences leading to multiple interacting effects, even just identifying all the contributing factors to a given effect becomes very difficult, let alone understanding the consequences of any variations on them; similarly, for the reverse task of deducing what actions may lead to particular outcomes. To a large extent these difficulties stem from the high branching ratios at every point in the evolution of a complex system – it is simply impossible to trace through every possible future path or to determine which branch will be taken because complex systems are essentially open systems that are subject to contingent influences from their context, which are essentially uncontrollable;
- a further consequence of complex networked causation is that the effect of adding or removing one element will be highly dependent on what else is there, so that the order in which elements are added or removed is important, and the adding or removing of elements may perturb the system so that adaptive mechanisms in the system may kick in to compensate for the change. This leads to the possibility of causation hysteresis, where simply undoing the changes that have been introduced does not restore the starting state of the system.

On the plus side however, a major advantage of working with complex *adaptive* systems is that their adaptivity actually *increases* their predictability and influencability when the adaptive mechanisms operating within them are properly understood.

It becomes a critically important task therefore to identify and study the adaptive mechanisms operating in CAS that we wish to infuence, whether they are our own systems, benign or neutral systems that we need to interact with, or the hostile systems of our adversaries.

Recalling that the essential elements of any adaptive mechanism are:

- internalised measures of success and failure (= ‘fitness’),
- a source of variations in the internal details of the ‘sense’, ‘process’ and ‘act’ aspects of the system,
- a means of real or modelled interaction with the context to provide fitness-relevant feedback,
- and a fitness-linked selection process which retains and propagates, or discards and inhibits those variations,

these four elements constitute four primary *targets for influence* in CAS that we do not own, or *levers of change* in those that we do, suggesting a broad range of potential tactics for influencing outcomes.

In brief, what springs from this analysis is the idea that we can increase the probability of obtaining the outcomes we wish by:

- finding and identifying existing adaptive mechanisms in the systems that contribute to the outcomes we care about,
- selecting those adaptive mechanisms that are the best targets of our interventions, and
- shaping perceptions of what constitutes success and failure in those mechanisms, or by modifying aspects of their three other elements so that the system will as a result, but of its own ‘volition’, take a different course.

There are two reasons why this is a more effective strategy than trying to force the system to undertake particular actions, or to prevent it from undertaking others:-

- if the system and its context truly are complex, then trying to work out the right actions in advance becomes a futile exercise, and

- the principle of economy of effort – a subtle but effective intervention to align the system's own adaptive mechanisms with our intents produces the desired outcomes with less apparent interference and exertion and more precision in effects.

Note that the reduced predictability of complex systems should not be interpreted as implying that there is less value in doing analysis at all – by taking an adaptive (and therefore iterative) approach to how analysis is done an initially poor predictive ability will both improve over time as feedback is used to refine the models of the causal and influence networks operating in the system, and also will progressively refine the accuracy of particular outcomes of interest.

Before concluding with a brief discussion of new options, we will first examine in the following section to what extent we are already practicing a CAS-based approach in the way we do C2 now.

7. C2 as an Adaptive Mechanism

Unsurprisingly, the C2 roles listed at the end of section 4 map quite well onto an adaptive mechanism at the action-in-the-world level, with the first two roles together determining the measures of success or failure for adaptive action in the remaining roles.

The concept of mission command also fits well into the construct of an adaptive mechanism since the focus is on articulating the intent, i.e. how success and failure will be judged, rather than specifying in great detail what is to be done, and empowering the mission team to make their own adaptive action choices in the implementation of that intent. In fact we would assert that the reason that mission command is a successful strategy in complex, dynamic and unpredictable situations is precisely because it adopts the stance of focusing on articulating measures of success and deliberately empowers individual and team adaptivity. Conversely, in tasks that are very predictable, although possibly complicated, it might be more efficient to employ a directive approach – as for example in the cooperative building of a bridge.

Adaptivity at the action-in-the-world level is ingrained into many aspects of how defence operates – for example every element of the Intelligence Preparation of the Battlefield (IPB)[15] and the Military Appreciation Process (MAP)[16], can be mapped into an action-in-the-world adaptive loop.

The Battlespace Analysis, Mission Analysis and Enemy Analysis phases which begin the overall planning process are a combination of 'sensing' (in that information inputs from various sources are actively sought) and 'processing' to develop specific mission-relevant internal models[17] and leading into the generation of 'acting' options in COA Development. These action options are then tested against the internal models in the COA Analysis phase, which is still part of the 'processing' phase, leading to refinements of the action options until a decision is made and the system moves into the Decision and Execution phase.

Of course the real process is never as clean and linear as this very brief description suggests, and the fact that real command teams make many jumps back and forth between different parts of these processes is evidence that they are actually using this adaptive mechanism in an appropriate way given the complexity of the task they are undertaking, in other words, they are executing many rapid iterative trajectories through the process, gradually refining their understanding of the relevant aspects of the situation.

In doing so, individual command team members are also operating at the second (learning-system) level of adaptation because they are learning about the situation, and their accumulating experience and insights enriches their adaptive action capabilities in subsequent stages of the mission, and in later missions. The quality of the learning of course depends on (among other factors) the extent to which their developing mental constructs are actually tested in success-relevant ways through real or simulated experience.

But does the team also operate *as a team* at the second level of adaptation? This would require that one or more of the 'sense', 'process' and 'act' capabilities of the team are changed as a result of team learning, over and above the individual learning that its members achieve. For example the team might develop its own version of the MAP and IPB process, might augment it in various ways, might introduce innovations into their planned actions or in how they are executed etc. These team-learned advances of course are

---

[15] http://www.fas.org/irp/agency/army/mipb/2001_10.pdf

[16] ATIB#74. Military Appreciation Process. Army Training Information Bulletin, No 74.

[17] The term 'internal models' in the context of CAS refers to the implicit or explicit representation of the system itself and its context that more sophisticated CAS develop and that are used to test action options during the 'processing' stage in order to screen out potentially dangerous actions and to increase the probability of selecting successful action options for implementation.

also learned by the individuals that are involved, but the innovations belong at the team level because they change how the team operates, and new team members would have to pick up these local variations of the doctrinal processes on arrival in order to become integrated into the team.

This example highlights in quite a stark way one of the many trade-offs that have to be handled in exploiting adaptivity – in this case between *standardisation* of processes, which enables more rapid team integration and interchangeability of components, and *diversification* of processes, which enable more locally effective and context-appropriate behaviours to develop, and which may also generate learning that might be more widely relevant and could therefore be fed into learning at higher scales of organisation and thus propagated to other command teams.

8. Opportunities for Increasing Adaptivity in Defence

The above discussion of the first two levels of adaptation as currently operating is intended to acknowledge the extent to which these concepts are consonant with existing C2 practice, but that does not imply that the potential power of adaptivity at these first two levels is being fully tapped. A more detailed and careful analysis will suggest many avenues for enhanced learning and adaptive action at the scales of individuals and of teams – this would be another paper in its own right.

A particular thread in this regard that would be enormously valuable to explore arises from the work of Dietrich Dörner[18]. Through a combination of historical analyses of massive failures (such as Chernobyl) resulting from the inherent limitations of human capacity to deal with complex decision-making, and years of well-designed human-in-the-loop experimentation, Dörner has identified a number of specific factors[19] that distinguish those few individuals that perform better at handling complex situations over a long enough time-frame for the 'chickens to come home to roost', from the majority that are only able to achieve initial short-term goals but in doing so, sow the seeds of later catastrophes. There are significant lessons for defence here, and interestingly, though not surprisingly, these lessons are very well-aligned with the conclusions we draw from a CAS theory perspective.

Another thread worth exploring comes from the organisational science and policy science literature, where applications of complexity and adaptivity have been studied for a couple of decades, and from where such commonplaces as the 'learning organisation' arise. A particularly insightful analysis of the opportunities inherent in complex organisations (which surely includes defence) is provided in a treatise[20] by Robert Axelrod and Michael Cohen, two of the giants in the field. They identify three design issues for complex systems which map precisely onto the three elements of an adaptive mechanism: variation, interaction and selection. While not explicitly addressing the fourth element which we have discussed here – fitness (perhaps because it is not as difficult to articulate what constitutes success and failure in business organisations as it is in enterprises whose purpose is not about turning a profit) – they do discuss in a lot of useful detail a wealth of intervention options available to managers of an organisation for modifying the way in which those three elements of adaptivity are implemented. Most of their options are transferable to defence, and illustrate the point made above, that there are a large number of degrees of freedom, or choices to be made about how to be adaptive – which also implies that there is a lot of scope for error and that the task of engendering adaptivity is a complex task in its own right (recalling our definition of task complexity in Section 1). Hence the need to employ adaptivity once again to tackle this complex task, which is what the third and fourth levels which we have described attempt to do.

When we turn our attention to higher scales (combined arms team, task force, joint force, organisation, enterprise etc) the scope for exploiting these first two levels of adaptation is arguably even greater than at the team and individual scales, and more difficult to implement effectively because the sense, process and act elements at these scales are likely to be much more widely distributed and possibly influenced by conflicting agendas.

In our framework there are three more levels of adaptation beyond learning. To what extent are these higher levels exploited in current defence practice at any scale?

The third level asks us to critically examine, and apply adaptation, to how we do learning in its most general sense of whatever leads to improved sense, process and act capabilities for adaptive action. Possible practices that would qualify as Learning-to-Learn adaptation include

---

[18] see *The Logic of Failure*, Dietrich Dörner, New York 1996

[19] These include the abilities to tolerate uncertainty, manage time, evaluate exponentially developing processes, assess side-effects and long-term repercussions, test their hypotheses, maintain focus and reflect on their own behaviour.

[20] *Harnessing Complexity: Organizational Implications of a Scientific Frontier* R Axelrod and MD Cohen, 2001, Simon and Schuster.

- after action review of how effective individual and command team learning were during operational planning and mission execution, introduction of experimental changes to increase learning effectiveness, and adoption of those changes that in fact do improve learning;
- review of the effectiveness of training programs and introduction of new training technologies, methods and processes;
- review of how effectively lessons learned in the field are captured, analysed and implemented and introduction of improvements in these functions;
- review of how effective technology insertion and science and technology support to operations are at remedying identified issues, and developing better ways of handling them; and
- review of how effective the process of capability development is at expanding the sense, process and act capabilities in useful ways for adaptive action, and hence introducing improvements into how it functions.

While current examples of these Learning-to-Learn adaptations could certainly be found, it is probably fair to say that they would fall a long way short of what is possible at this level. The aspects of these various functions that could or should be subject to adaptation include

- how variations are created and their rate, extent, scope and concurrency;
- how the variations produced are experimented with and evaluated and at what rate, cost, accuracy and concurrency; and
- how the selection process operates – how tolerant it is, how fast, how completely it eliminates discarded variations and how widely it propagates retained variations, how wide or narrow the unit of selection is and so on.

A full discussion of these options would require a much longer paper, but clearly there is considerable scope for exploiting Learning-to-Learn adaptation in many aspects of defence operations. We note also that these aspects not only include operational C2 functions, but also raise, train and sustain functions and are therefore broader in their scope.

Level 4: Defining-Success, falls more closely into the remit of C2 since it addresses how objectives for action are determined, which is what we have described as C2 Roles 1 and 2 in Section 4.

The issue here is to improve the alignment of local intents (as expressed by the measures of success and failure that are used for driving adaptation in multiple levels of scale vertically, and in many different systems horizontally across each level) with the overall and ultimate measures of success and failure for defence as they will be eventually judged. Since these ultimate measures are not accessible to the multiple adaptive action loops that are operating throughout the force, these loops are instead driven by proxy measures for success and failure, which are observable, and which are intended to correlate well with the ultimate measures. Developing good proxies for success is one of the important roles of C2.

What Level 4: Defining-Success offers is the possibility of deliberately employing an adaptive process to refine the proxies that are used in levels 1 and 2 so as to create more successful long-term outcomes. This requires that hindsight measures are available to provide the yardstick against which the outcomes from using the current or hypothesised alternative proxies can be judged.

In order to do this we have to first define some ways of characterising the ultimate success of the outcomes:

- Did we get the ‘right’ outcomes? Would the resources have been better allocated to another course of action instead? i.e. in retrospect, could we have achieved a better outcome with the same cost / risk / time if we had had better processes and made better use of information?
- Did we make the best possible use of the defence and non-defence elements of capabiity that were available? i.e. in retrospect, could we have achieved a <u>*better*</u> outcome within the <u>*same*</u> overall cost / risk / time constraints if we had had better processes to address the interoperability and integration issues for utilising the capabilities that were available?
- Were the outcomes efficiently produced? i.e. in retrospect, could we have achieved the <u>*same*</u> outcome with <u>*less*</u> cost / risk / time if we had had better processes and made better use of information?

- Did we adequately protect ourselves from risk? i.e. in retrospect, could we have achieved the same or better outcomes with a lesser exposure of our capability elements to risk?

We do this ‘in retrospect’ because generally it is how the outcomes are perceived and judged over a longer time period that is the acid test, and we need to learn from hindsight. Because of the inherently long latency in arriving at those judgments these measures cannot be used to drive the adaptive processes directly at the first two levels. However they could be used to drive the slower higher-level adaptive processes (Learning-to-Learn and Defining-Success) so that overall our ability to effectively and efficiently produce successful outcomes improves with time.

Implementing Level 4 adaptation implies that not only are hindsight measures developed when sufficient time has elapsed to allow a confident assessment, but that the proxies that were used at the time to make local decisions are compared against hypothesised variations on them to judge whether the variants would have led to better decisions and better hindsight outcomes. As evidence accumulates that a varied proxy produces better alignment than an existing one it can be introduced to replace it for current decisions. The proxies in use need to be continuously tested and evolved in this manner.

Finally Level 5: Co-Adaptation is core business for C2. At this level we are addressing how to best use the system of systems that we have, in order to achieve success, and to do this, we are experimenting with different roles, relationships and resourcing of our component systems. Every time we do this we necessarily are changing the boundaries and proxies for the systems in question and they then need to cycle through their own adaptive processes in response. This is routinely practiced to some extent – constrained by the decision rights of the agent who is doing it – a theatre or mission commander for example. One of the possibilities being opened up by wider networking is of performing this co-adaptation process across wider groupings. The question it begs is how the hypothesised changes to the system-of-systems parameters are to be evaluated – other than the clearly impractical route of trying every set of changes in the real world. The current extent to which co-adaptation is practiced is limited, for good reason, to the domain over which there is confidence in judging the utility of different options.

9. How can we use our understanding of CAS to better deal with opponents and allies?

When we ask what interventions are possible in dealing with CAS we do not own, we have to decide where to operate on the spectrum of

destroy it ← → *seek to shape and influence it* ← → *accept it as it is*.

The middle ground is the interesting territory here and CAS theory is offering the possibility of becoming better at it.

Moving into this territory requires us to understand more about the system, eg what are its adaptive mechanisms? what are its success and failure criteria? what levels exist in it and how do they interact? what information do they seek and have? what response options do they have? how do they process information and how do they decide? how much variety or diversity do they allow? what is their interaction structure? how are new ideas generated and evaluated? how are policies enforced? What conflicts are there between its adaptive mechanisms and how could they be exploited? and so on.

The task now is to refine the crude models of the system that we may have started with so that we get better and better insights over time. These are difficult questions to answer – granted – but the power of five levels of adaptivity can of course be used to deal with those difficulties. At the Action level we would be probing the system and observing how its response supports or differs from our expectations. At the Learning level we would be reflecting on the gaps in our understanding and devising new probes to test its behaviour in those aspects. At the Learning-to-Learn level we reflect on how effective we are at the Learning level, and experiment with variations in how we devise new probes and learn from them. At the fourth level, Defining-Success, we would be rethinking how we judge the quality and adequacy of our growing understanding, and evolving that into more useful directions. And at the fifth level, Co-Adaptation, we would be considering what resources and systems and elements we are bringing to bear on the problem of developing our understanding of the CAS in question, and whether we could be more effective by using our assets differently.

As the quality of understanding of the system matures, we would be using our understanding to test and refine our interaction options with the system. The ‘probes’ we would have been employing in the earlier understanding-building phases are on a continuum with more overt interactions that have their own purposes, but of course the latter should also be seen as probes that provide further tests and refinements of our models.

The picture painted here is of a process that starts lightly and tentatively and gradually builds confidence and effectiveness. Interestingly, and not surprisingly, this picture is entirely consistent with Dörner's profile[21] of the rare decisionmaker who is successful at dealing with highly complex systems over the long term.

10. What are the implications for how we may better foster desirable adaptivity properties in our own force?

In other words, what interventions are possible in our own CAS? – what are we trying to achieve? what are our degrees of freedom? and what are the mechanisms available to us?

By now it should be clear that a CAS-based approach suggests many opportunities for interventions to influence the behaviour and properties of our own systems. We have briefly touched upon classes, levels, types of adaptation, and so on.

We have indicated that techniques are emerging from our research in this area that could help us address adaptivity at multiple simultaneous levels, to look at the interactions between levels, identify conflicting adaptive mechanisms, debug them, tune them, and introduce additional adaptive mechanisms as needed.

The single most critical issue we have to tackle to allow us to apply these techniques is developing a clearer articulation of our measures of success and failure, which we turn to in the next section.

11. The Importance of Understanding Success and Failure

The paucity of good examples of 'fitness' articulated in a sufficiently precise and useful way to drive adaptation is evidence that it is a very difficult task – but why is it so hard to specify what we want?

We can attempt to define success and failure at various levels – but measures of success at very high levels tend to be vague, ambiguous, easy to state but not easy to measure, and operate over very long timescales, while very low level measures of success (really performance specifications) tend to be too detailed and prescriptive and to have already pre-empted many of the choices that really should have been made adaptively, so that there are insufficient remaining degrees of freedom fro adaptation to work on.

The latter are easier to action but may not meet the higher level goals, while the former are impossible to action without making many further value judgements. Mid-level measures are both more difficult to state and more difficult to relate upwards and downwards, yet this is the level most relevant to the big decisions that have to be made so it is important to come to grips with the challenge. Ultimately what one would like to have is a hierarchy of measures which are coherent so that achievement of lower ones build towards achieving higher ones.

However lower level measures are probably best left to local actors to determine adaptively, and the real challenge of adaptive C2 at the strategic level in complex situations is to pitch the articulation of success and failure at the right level to both guide local actors, and allow them the maximum possible latitude to exercise their inherent adaptivity – provided of course they are appropriately motivated and supported.

Then there is the issue that there is a disconnect between rhetoric and action, between what we say we value and how we actually make choices. To see what fitness function or what measures of success we *actually* operate with we need to look at how the selection processes of our adaptive mechanisms work. One can ponder to what extent this disconnect is an unintended consequence of the difficulty of articulating useful mid-level measures, or a cynically deliberate stance which maximises the opportunity to let factional agendas drive decisionmaking?

Some further reflections on the difficulties and importance of understanding what constitutes success and failure:

- It is more difficult to articulate success measures for defensive postures than for offensive ones. What are the consequences of this? Does this bias our decisionmaking towards more offensive strategies?
- Measures are also important in concept development, because we need to know what is required to make a concept work, and how to tell when it is working. We also need a framework that the measures fit into so we know how much is needed of any particular measure, how to influence an objective measure, what unintended effects may be generated, and so on.

[21] See 18 and 19 above.

- It is important for the set of measures to explicitly include failure measures since these are not simply the absence of success.
- It is also important for the measures of success and failure to be comprehensive if they are to drive adaptive behaviour, else we risk favouring the explicit measures over the implicit ones, with possibly disastrous consequences.
- We also need to distinguish and know how to use various levels and types of measures that relate to success and failure at a number of levels and timescales. This is important because as dicussed above, the ‘real’ success or failure that is created is generally only apparent after quite long periods of time, too long to be useful for driving the much faster adaptive mechanisms we require. So we have introduced the notions of:
    - ***Symptoms*** of the problem – the negative consequences of not being adaptive enough
    - ***Measures of success and failure*** – in reducing the negative consequences of the problem and creating the expected benefits
    - ***Proxy measures*** of success and failure – these are shorter-term observable outcomes which are not themselves the primary outcomes we care about, but that correlate well with those longer-term ‘real’ success and failure measures and that can be detected over a timescale which is fast enough to allow adaptive action in response; and
    - ***Indicators*** – these are observable outcomes which do not qualify as proxy measures of success and failure, but which give an indication of what path the unfolding chains of interrelated events, factors and influences are taking and in particular, whether we are on a path towards success (so stay with current CoA and keep monitoring) or whether the path is veering towards failure (so some corrective action is called for). The relationship between the indicators and the success/failure outcomes may be heuristic or based on a theory or model of how the causal and influence network operates.

  Obviously these proxies and indicators are related to what we called symptoms above – and one would expect to be able to generate one or more measures corresponding to each symptom.

Concluding Remarks

If we agree that the potential value of understanding and exploiting adaptivity in defence is immense, then we must also recognise that the danger in failing to rise to this challenge is even greater. Complexity is not about to go away – rather it is continuously and inevitably rising, and we must deal with it.

Clearly there is still much work to be done to operationalise our growing understanding of CAS and to inrease our ability to effectively exploit it. There are also some difficult implementation issues arising from traditional ways of thinking and doing things, and from the inherent nature of human cognition and the history and structures of our organisations. However the overall take-away lesson from this discussion should be an optimistic one – if we make proper use of adaptivity we do not need to know all the answers, we do not need to make better plans, or get everything right the first time. We just need, as nature has ever done, to start with a straw man and know how to grow it to a better one.