CHAPTER 4

ON THE ROLES OF DESIGN IN DEFENCE

Anne-Marie Grisogono

How should we think about design in defence? This publication is an opportunity to take stock of what has been said, tried and learned so far, and also to explore some paths not taken. Ideally the outcome will be to recalibrate the way forward and reenergise implementation of design approaches in the service of improving defence effectiveness.

The origin, history and practice of design thinking in defence will no doubt be amply covered by several of the other chapters in this publication. What this chapter proposes is a fresh perspective on the value and roles of design in defence, based on first principles, systems thinking and complexity science.

But a situation appreciation is called for first. One cannot expect that the introduction of new ideas into large organised enterprises with long histories and proud hard-earned traditions will be easy or smooth. One of the founders of quantum physics, Max Planck, claimed that science progresses one funeral at a time.[1] That is much too slow for defence.

About new ideas in defence

Defence is risky business—the costs and stakes are very high, while the benefits are often elusive, difficult to measure and even more difficult to attribute. There is no end point, just an ever-evolving competitive struggle for an edge that can only be temporary. This dynamic creates an imperative for innovation, for an unflinching, realistic appraisal of what works and what doesn’t, what should be retained, reinforced and proliferated, and what should be pruned, eliminated and replaced—in theory.

In practice, there are a few things that get in the way. Human nature is one, for a start—with all the well-known cognitive biases and heuristics, which channel our thinking in familiar ways and raise anxiety about change, contradiction, loss, ambiguity and the unknown, unconsciously distorting judgement as a result.[2] And, then there are inherent difficulties in separating the wheat from the chaff in the marketplace of ideas.

A reason for this difficulty stems from language itself. Words are powerful, evocative, and make it possible to discuss and reason about abstract and hypothetical ideas—but they are also often vague, ambiguous, and lacking precision where nuanced distinctions are necessary, hindering rather than enabling effective dialogue. The attachments people feel to their own interpretation of a word’s meaning, deriving from associations accrued through their own experiences, contribute to the amount of effort wasted in semantic arguments. Korzybski told us we should not conflate the map and the territory.[3] The same goes for words and the concepts they represent. A map needs to be precise enough to enable accurate exploration of the territory. And, the words we use should do the same for the conceptual territory.

Then there's complexity, a commonplace motif these days in fora that discuss defence concepts—but rarely go beyond superficial nods to non-linearity, fog and friction, and black swans— acknowledging the challenges posed, but often short on practical support. We will have more to say about this in the following section.

Throw in time and workload pressures, and natural suspicions about the hidden agendas of those who stand to benefit from the latest 'big new thing', and the unfortunate consequence of all these factors is that new concepts in defence are often subject to a predictable life-cycle from which little of real value is ultimately extracted.

A new idea is launched, with some fanfare and hype, and a web of jargon grows around a catchy phrase or acronym. Then, some people drink the Kool-Aid and become enthusiastic proponents while a backlash slowly mounts from those who see it as no more than established wisdom in new clothes; or worse, see it as an assault on established wisdom and an attempt to replace it with a slick but unsubstantiated new fad. So, battle lines are drawn and the cycle plays out. If the new idea has high-ranking support then the tactics are to pay lip service as required and wait it out till the next posting cycle, perhaps white-anting its implementation at ground level. If the new idea lacks high-level sponsorship, then it's easier to ignore, and ridicule.

Meanwhile, both sides may attract new adherents but often through the aforementioned lenses of heuristics and biases, which draw us emotively to one pole or the other rather than through thoughtful deliberation and judgement. As a result, the debate is often polarised and, therefore, generates more heat than light. Over time, as people jump on the bandwagon, the concept and its associated jargon are heard more frequently. Then, its popularity peaks, eventually slides into obscurity, and it becomes passé or even off-limits. Think RMA (Revolution in Military Affairs),[4] EBO (Effects Based Operations),[5] NCW (Network Centric Warfare),[6] and COIN (Counter-Insurgency).[7]

But was there a baby in the bathwater? Of course, some new ideas do deserve to be discarded, or to be correctly identified as revamped old ones. But equally, new language is sometimes necessary to make a new distinction, or a new concept, accessible, and it pays to invest in the effort required to discover and develop its value.

The potential value of a new idea will not generally be immediately obvious or available. It needs to be discovered, developed and protected from premature rejection, until enough is understood to make a proper assessment of its benefits against its costs and risks. Just because they have different views, both proponents and opponents have important contributions to make. The former because they envisage its possibilities; the latter because they recognise its hazards. These two frames need to be appreciated as complementary not contradictory and woven into a more robust, comprehensive and nuanced view. Opponents whose concerns are heard and addressed can become supporters. And, fans whose enthusiasms are tempered by reality checks can work more effectively towards realising an idea's potential.

Admittedly, design thinking is not exactly new, even in defence, but one thesis of this chapter is that despite its successes so far there is much potential value yet to be developed. In the meantime, it is arguably still at risk of being prematurely discarded in favour of some yet-to-be-named new idea.

The other thesis, to be developed below, is that accessing the full potential value of design requires understanding that design is the logical and necessary complement of an adaptive approach to navigating complex situations successfully; that it is applicable across many domains and scales, not just at the operational level; and that much greater value can be derived when they are judiciously applied in concert, because each enables and depends on the other.

Complex problems

Complexity stems from interdependence between the elements of a situation. Because these interdependencies create many possible pathways through which consequences of changes and events can propagate, it becomes impossible to correctly anticipate all the relevant consequences of a proposed action or a hypothetical event. Conversely, it is similarly impossible to be sure of exactly what conditions to set and actions to take to achieve specified desired outcomes while also averting unwanted outcomes. This observation applies not only to the operational context but also to capability development and to the broader contexts within which defence is embedded.

Since complex problems by definition do not lend themselves to obvious solutions, it is necessary to pay attention to *how* approaches and solutions are going to be developed or, in other words, to a strategy for dealing with complexity. The outcome or product of such a strategy is also a strategy, but a specific one—the concept of how the actual situation in question will be handled.

To avoid confusion, we use the word *stratagem* to denote the strategy for the particular situation, and *meta-strategy* to denote the strategy used to develop the stratagem. While a stratagem is necessarily situation specific, a meta-strategy is couched in more generic terms.

For example, one obvious aspect of a meta-strategy for dealing with complexity is the need to take an adaptive approach,[8] which translates to an iterative evidence-based and goal-oriented engagement with the problem situation, enabling one to continuously learn about its essential dynamics and to develop, trial and adapt approaches to deal with it. In other words, the situation specific operational stratagem needs to be developed in an adaptive fashion. This is discussed in more detail below, but note here that this is only one part of an overall adaptive approach: there are several other aspects that need to be developed adaptively including goals, capability and execution of the stratagem.

There are many advantages to an adaptive approach – it does not depend on having complete information about the situation (rarely possible); when comprehensively applied it is inherently a risk-management approach, and it is the only way known to cope with situations which are rapidly evolving. Furthermore, if a complex organisation that is dealing with a large scale problem situation adopts what we have called an Adaptive Stance,[9] then it is able to effectively leverage the detailed local knowledge, insights and intelligence of all its distributed agents in the situation—a huge effectiveness multiplier. Western defence forces have increasingly recognised these benefits and the consequent need to embrace adaptivity as a response to the complexity challenges they face.[10]

However, while it is necessary to take an adaptive approach in dealing with a complex situation, it is not sufficient; and indeed, there is a hidden danger in it, if taken to an extreme and in isolation. To illustrate, consider a humanitarian relief operation. Those contributing to international assistance efforts after a natural disaster might individually all be taking an adaptive approach to identifying local actions they can take to help victims. Nevertheless a whole situation assessment might find areas of desperate need without any assistance, and over-servicing in other areas of lesser need. Worse, the actions taken by one group may clash with those of another (e.g. one laying a road through another's building site), and natural competition between aid agencies and between individuals may lead to outright conflict, and bidding wars for local staff who can facilitate their projects. Despite best intentions locally, the overall outcome may be quite negative. Similarly, any organisation that enthusiastically embraces the need to be adaptive by instructing all its personnel to go forth and adapt without any constraint runs the risk of rapid disintegration if they do. Evidently an enterprise that seeks to deal with a complex situation has to confront the inherent tension between the alignment necessary for coherence, and the diversity necessary for adaptivity.

The answer is not to abandon the adaptive approach but to place it within a design construct that makes intelligent use of adaptive approaches and, at the same time, frames them so as to provide coherence. A major mechanism whereby this can be done is through design of the enterprise's operational stratagem, and especially two of its aspects: the adaptation architecture of the enterprise and its hierarchy of intents and associated metrics framework.

In order to flesh out this claim, we turn next to looking at the nature of design from first principles and then, in the subsequent sections, examine its importance and applications to defence operations in complex situations.

Fundamentals of design

Design takes many forms in many domains, but what is the core essence they have in common?

> *[Design is a] specification of an object, manifested by an agent, intended to accomplish goals, in a particular environment, using a set of primitive components, satisfying a set of requirements, subject to constraints.*[11]

This proposed definition recognisably refers to the outcome of design as being something functional, which aligns closely to the everyday intuitive meaning of design. However for our purposes in defence, it is perhaps too restrictive in implying that the designed object is material, that it is the product of a singular agent, that it is intended for a particular environment, and that its component parts are primitive. Even the implication of a well-defined set of requirements may be too optimistic in a genuinely complex situation. So we will need to generalise most of these provisos but retain the essence of what design is—the creation of something functional, something that works—and allow variability in how goals, options and constraints are determined. The outcome of design could be any combination of processes, tools, networks, strategies, tasks and so on. But, for simplicity we use the term 'system' to imply them all. Herbert Simon probably said it best: 'Everyone designs who devises courses of action aimed at changing existing situations into preferred ones'.[12] To investigate the process of design we now need the following concepts:

Design parameters are the variables that specify the essential characteristics of the system that enable its useful functions. Changing the values of design parameters results in changing what functions can be performed, how it performs them, maybe even whether it performs them, and can also change the performance outcomes. Changing the design of the system by removing or adding features also changes the set of design parameters that are needed to describe the system.

Design space is the set of all possible combinations of values of all the design parameters. The design parameters are therefore the dimensions of this virtual or conceptual space. Since the number of design parameters is very large for real systems, design space is a hyper-dimensional space. Innovation that leads to new design parameters changes the dimensionality of the design space. A particular instantiation of the system is represented by a single point in the space, i.e. a set of values specifying each of the design parameters, and an evolving system therefore carves out a trajectory through design space.

In practice, not all design parameters are equally available or suitable for modification, for example because any change in them is too costly or too harmful, or because the designer does not have the authority to change them. So it can be more useful to consider a subspace defined by those design parameters that actually are design degrees of freedom. Even with this useful reduction in the number of possibilities to be considered, it is still an immensely vast hyperspace, permitting a vast number of possible designs.

Therefore choosing values for the design degrees of freedom is not something that can generally be done in a sequential or deductive way such that there is clearly only one right answer. Rather, *design is a creative process drawing on many sources and can therefore appropriately be described as an art*.

Action decisions have their primary intended effect on/about the actual situation of interest, whether to change something in it or to discover something about it. Examples include: decisions to build, modify or destroy something in the situation; to measure or observe events in it; and to give information to, or ask for information from a participant in the situation.

Design decisions, in contrast, have their primary intended effect on what future action decisions will be possible by enabling or constraining future sensing options, action options, decision rights and responsibilities. In other words, design decisions are primarily intended to affect what outcomes are going to be possible to affect in the future, by affecting what action decisions are going to be available.

Design decisions therefore determine objectives, stratagems, capabilities, distribution policies for information and material resources, distribution of responsibilities, authority and autonomy, task procedures and doctrine, training programs, and so on. Specific examples include: decisions to develop improved sensors for finding survivors; to acquire or deploy new search equipment; to hire or fire individuals with specialised skills; or to change delegations of authority.

Such design decisions clearly impact on the subsequent action decisions that can be made about employing the capabilities that are actually available as a result of those design decisions, to do particular things in the situation of interest; for example, where to deploy rescuers and how to distribute the available emergency supplies. The projected timeframe of design decisions may be the near future in a particular situation or it may be very far-reaching and intended to impact in a wide range of future situations that are yet to emerge. Of course real-world decisions may not neatly fall into one category or the other. But, even though there will often be impacts on both the complex situation and on what later decisions will be possible, one can in principle think about these two types of consequences separately.

Thus design decisions are about making choices that remove some degrees of freedom in the system (by assigning specific values to them so that they are no longer free) in such a way as to modify the degrees of freedom available for action decisions later. Good design then, translates into improved chances of success in those later action decisions.

There is an important design issue here—those making the action decisions later do not necessarily need as many degrees of freedom as possible, since (a) that would place too great a decision burden on them as they deal with an actual situation, (b) they may not have sufficient decision support (e.g. relevant information) and (c) the timescales required to implement some of the necessary design decisions may be too long for operational timescales. Therefore many design decisions need to be made well in advance of any intended application timeframe. This observation highlights a special class of design decisions, which we will call *meta-decisions* because they *are decisions about decisions*; in other words the decisions about *what* design decisions need to be made, *when* they need to be made, and *how* and *by whom*.

Meta-decisions are further discussed below but here we simply observe that from the perspective of those making action decisions in a particular operation these meta-decisions, and the design decisions that are consequently made prior to the operation, contribute to determining the range of options that will be available to them in the operation. If that range is too narrow and prescriptive their effectiveness may be hampered but, equally, if it is too wide and open-ended it may hamper efficiency and responsiveness. Good meta-decisions mean good design, which will translate into getting that balance right.

Why design is important—Part 1: Managing the risk in the art

The concept of design decisions presented here is much broader than a traditional interpretation of design, implying that many decisions routinely made in an organisational or operational context are in fact design decisions. There are two reasons for claiming it is important to explicitly identify such organisational decisions as design decisions:

a. design decisions may have a large impact on future outcomes

b. design decisions are based on conjectures and therefore entail risk.

The first point follows from the definition of design decisions: since the impacts of each design decision will be generated by all the subsequent action decisions that it makes possible (for as long as the design decision remains in force) plus, indirectly, the consequences of the inability to make the action decisions that it precluded.

The second point follows from the earlier observation that choosing the values of design parameters out of a vast hyperspace of possibilities is necessarily a creative not a deductive process.

We can also frame this in the converse: for a decision to qualify as a design decision (in this conceptualisation) there must be an element of creative choice being made. If there is not, for example if the value of that design parameter is constrained such that there is only one acceptable value for it, then setting it to that value is not a design decision but rather the logical consequence of earlier design decisions (which created the constraints). Similarly, action decisions can be seen as either calling for choices to be made or as logically deducible from existing facts.

To the extent that there is an element of necessarily free choice in either action or design decisions, the decision-maker is exercising what we can call 'art'. In the domain of defence operations, this is called 'operational art' for both operational design and operational execution/action. As intimated earlier, the use of the word 'art' is very appropriate because it is a creative process and cannot be rigorously defined, described or prescribed.

However, to the extent that such an element of free and necessary choice exists, the quality of the decision made will depend on the quality of the decision-maker's understanding of:

a. the complex situation of interest

b. the relevant high-level objectives or measures of success and failure

c. the available options (for action decisions this means the capabilities that can be called upon, and therefore the available actions that can be taken; for design decisions this means the possibilities implied by the design degrees of freedom)

d. their own experience, intuition (generalised experiences), education and creativity.

Since (a), (b) and (c) are all complex, multi-dimensional and multi-scale their understanding will necessarily be imperfect (incomplete, and incorrect in parts) and so the choice that is made as a result (through exercising operational art) is a *conjecture about what will work*, rather than a logical consequence of known facts.

Making conjectural decisions is unavoidable when one must deal with sufficiently complex problems and situations. But it is equally necessary to recognise them as conjectures and not mistake them for certain knowledge. Failing to do so amounts to taking possibly very significant but unknown risks. On the other hand, acknowledging and managing those risks invites one to ask (and act on) two very important questions:

- How would I know? i.e., what plausibly observable evidence would indicate the conjecture was wrong?

- How much would it matter? i.e., what would have to be changed if evidence implied the conjecture was wrong?

Expressing design decisions as conjectures means that instead of just stating what the decisions are (with or without supporting evidence, such as analytical studies) there is an additional statement expressing the underlying conjectures about how the design choices are anticipated to result in observable benefits without unwanted outcomes, thereby also specifying what would constitute contradictory evidence, a critical enabler of a fully adaptive approach.

When the design choices are implemented, this then permits not only implementation monitoring (*were the design decisions correctly executed?*), but also conjecture testing as proposed above (*does the implementation result in the expected positive impacts and not in the unwanted negative ones?*).

Rather than waiting for good or bad news, agility calls for rapid discrimination of which of the possible paths to both wanted and unwanted outcomes are being activated, so that appropriate adaptive action can be taken promptly. The design process should therefore also identify the earliest indicators and proxies that can be monitored for such evidence, and the relevant timescales over which they can be expected to develop.

Obviously monitoring just the implementation is not sufficient to avert the risks inherent in the conjectures but just monitoring for impacts is not sufficient either—one needs both to be able to conclude that, *yes, the design was correctly implemented*, and, *no, it did not pan out as expected in the situation*, and therefore there is something important to be learned. Seeking such feedback from the complex situation and being prepared to act on it by adapting the design and action decisions that depended on the newly-refuted conjectures amounts to taking an adaptive approach to both design and action/execution.

This means that the design process never ends—it has to continue throughout an operation, continually setting up and evolving the conditions for successful action. Doing so provides some insurance against unknown risks inherent in acting on imperfect knowledge about complex situations. But, like real insurance, it depends on full disclosure—being able to identify all the conjectures (and hence risks) that are made. An unconscious conjecture, implicit in an unrecognised (design) decision, is unlikely to be critically examined and therefore the decision-makers, and the enterprise they are part of, are blind to the consequent risk it creates—until it materialises into what might by then be an emerging disaster. By that time, earlier windows of opportunity to avert the disaster, which might have been taken if an adaptive approach had been in place, are often closed.

The combination of design and adaptive approaches therefore provides an effective meta-strategy for dealing with complex situations; fostering rapid learning about what is relevant in the situation, at both the individual and enterprise level; and increasing the probability of success by enabling the operational design to continually adapt to the relevant realities of the situation.

In summary, our arguments so far for explicitly taking a broad view of design are: that it provides an integrating framework within which adaptive approaches can be exploited; and that the design process can be harnessed to elicit the otherwise implicit underlying conjectures and subject them to explicit risk management through an adaptive design process.

But there is a third argument that further strengthens the nexus between design and adaptivity. The extent of a system's ability to exploit the power of adaptation across the full range of conditions requiring it depends to a large degree on design decisions that are generally not recognised as such, and therefore often not oriented towards enabling adaptivity. To set the context for developing this third argument, we first address the following question.

What needs to be designed?

A system is a set of interdependent elements forming a complex whole that act together in a common purpose. An operation is a coordinated set of activities undertaken (by a system) to achieve a common purpose. Operations are systems in action and systems are the means by which operations can be implemented, each providing both context and constraints for the other. This deep duality suggests that systems and operations should ideally be co-evolved in an iterative fashion.

They also both have multi-scale structure. Systems consist of functional subsystems connected in various ways, and those subsystems may also be resolved into networks of component systems with more limited functionality, etc., thereby creating a multi-scale system-of-systems view. At each scale, it is both the properties of the components and the topology of their interconnection that are important in determining the functionality that is thereby enabled. Similarly, operations can be broken down into functional lines of operations connected in different ways, which can then be broken down into networks of interlinked tasks, then procedures etc., so also creating a multi-scale and multi-dimensional view of a complex operation.

But design proceeds in the opposite direction, building up complex multi-scale networks of simpler elements. And, it is evident that the number of possible ways to select subsets of components and to assemble them into multiscale networks of various topologies is vast, for both systems and operations. Most of those possibilities will not make any sense, nevertheless there will be many possibly effective solutions, so design is essentially the problem of finding ‘good enough’ needles in the haystack-hyperspace of all possible designs.

Applying an adaptive approach to the multi-scale nature of both systems and operations helps structure complex design problems and make them more tractable. The system and the operation can be co-evolved through a combination of top-down sketching out of the stratagem (the highest level of the operational design) and of the system-of-systems concept (the highest level of the system design), and middle-up build and test of the possibly useful contributing plans and capabilities, with many up and down iterations between them to resolve discrepancies, alleviate identified problems, and exploit identified opportunities.

Since the lower scales of elementary procedures and functions are less situation specific and are likely to have been evolved over a long period and many previous experiences, it will generally not be useful to revisit them unless there is some innovation or identified problem at those scales calling for a reappraisal.

There are intimate linkages between successive scales: the functionality of the lower scales enables the range of possibility in the higher; the higher scales set goals for, and constraints on, the functioning of the lower. Thus changes at any scale have potential impacts at all other scales. The design process therefore needs to be iterative, not only between the system and the operation but also up and down the scales of each.

Moreover, and as discussed above, because of the necessarily conjectural nature of design decisions and the complex dynamics of the situation itself the design process needs to be placed within an adaptive construct that seeks relevant evidence to test and evolve the underlying conjectures; and hence, evolve the designs to better fit the real situation. This amounts to a continuous learning process about the complex situation but one that is focused on the aspects that are most relevant to the success of the operation. Since success also requires avoiding many possible failure modes, which are not simply the absence of success, the relevant aspects of the situation that need to be continuously learned about are much wider in scope than those that can be directly deduced from a proposed stratagem. A devil's advocate appraisal of designs, for example of a new policy initiative, is an important tool therefore for identifying such failure modes and evolving the robustness of the design. In the defence domain, where hostile elements of the situation can be expected to look for and foster failure modes, red-teaming is employed to identify and mitigate those risks and hence develop more robust operational and system designs.[13]

It is not just the design process that needs to continue throughout an operation, co-evolving the two 'pillars' of design: the system design and the operational design. There also needs to be continuous learning about the situation, which in general will be very rapidly evolving, partly in response to the operation but also because of its own inherent dynamics. The conceptual model of the situation is the third pillar that needs to co-evolve with the operational and system designs. Finally, one more co-evolving pillar is implied by the approach described here: the collection plan to provide the information needed from the situation to feed the required adaptive processes.

Since the adaptive processes are testing and evolving the conjectures about the situation (the conceptual model) and also the conjectures about what will lead to success and avoid failure (the operational and systems designs) and since these are all multi-scale and multi-dimensional, it follows that the information that needs to be continuously collected to provide the relevant feedback for those adaptive processes can also be organised into a multi-scale and multi-dimensional framework of measures and information; or, for brevity: a metrics framework.

Putting all these threads together, we arrive at a meta-strategy that co-evolves four multi-scale pillars: the operational design, the system design, the conceptual model of the situation, and the supporting metrics framework, as illustrated in Figure 4.1. The diagram also shows some detail about the co-evolutionary cross-impacts.

Thus both design and execution are very dynamic processes, intimately engaged with each other and with the complex situation. This approach encompasses everything we would want to include in design. It relates design to operational consequences (for operational design) and to capability consequences (for system design), and helps clear thinking about design. It also helps focus the effort to learn about the situation and, hence, make best use of limited information resources.

While the diagram in Figure 4.1 illustrates the co-evolutionary relationships, it does not show how these are implemented in an ongoing engagement with the situation nor how the co-evolution proceeds temporally. The simplified schematic diagram in Figure 4.2 attempts to paint that picture.

Figure 4.1: The continuous co-evolution of situational understanding, system design, operational design and design of the supporting metrics framework[14]

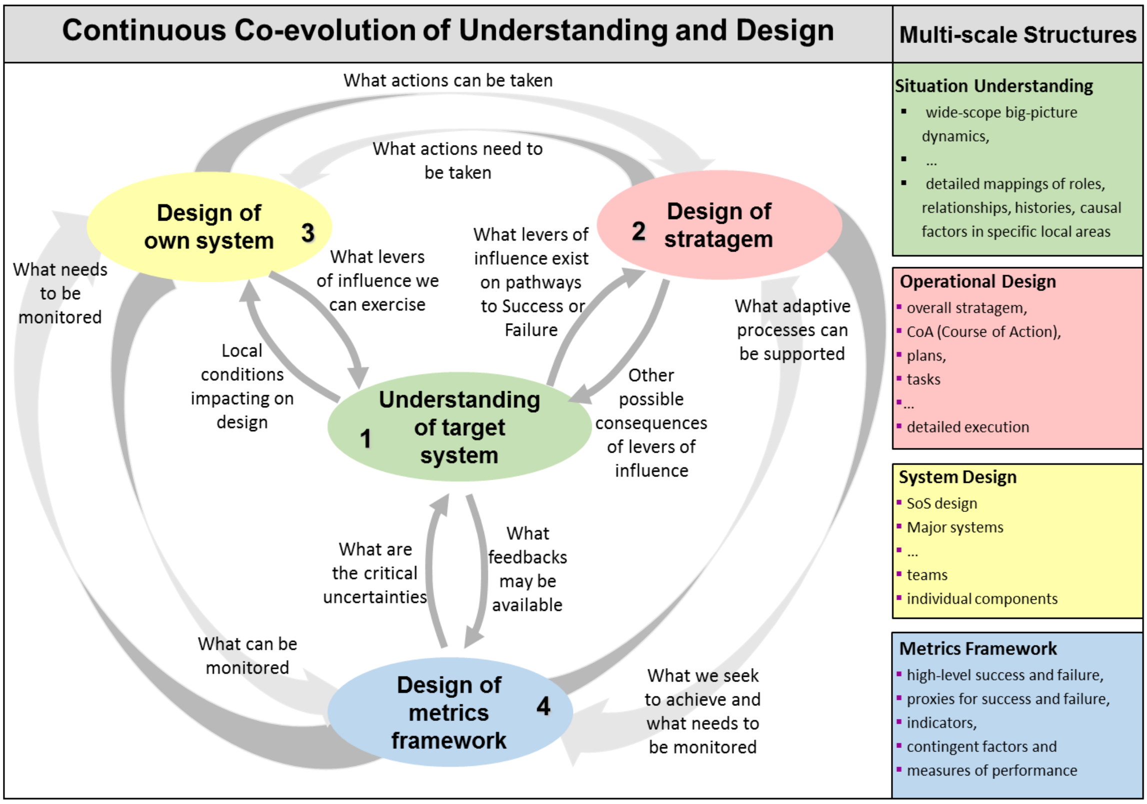

Figure 4.2: Three nested loops of co-evolution of (1) the conceptual model, (2) the system and operational designs and (3) the adaptive implementation and execution of the designs[15]

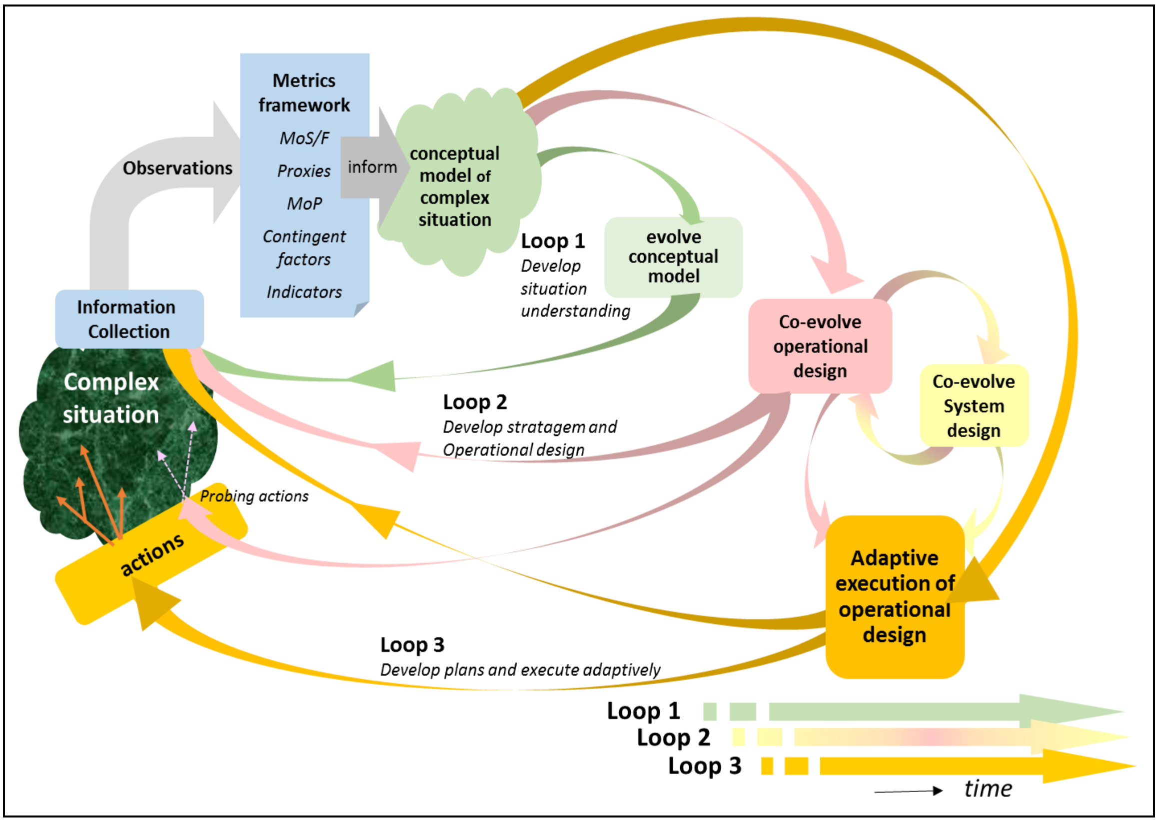

In Figure 4.2, Loop 1 begins with observing the complex situation and developing a conceptual model of it by identifying significant aspects, the causal and influence pathways leading to them, and hence relevant contingent factors and indicators (observables which convey useful information about the dynamic state of the situation). The cycle implied by the loop is one of using the current understanding of the situation to interpret the information collected, and hence forming expectations about how the situation will develop, adapting the information collection plan as a result to ensure that the subsequent relevant developments are actually observed, and then comparing them with expectations to inform the adaptive evolution of the conceptual model. One of the goals of this process is also to reduce the initially vast number of situational measures that might be observed into a succinct manageable set of indicators that is adequate to support all the necessary adaptive processes.

Loop 2 uses the conceptual model to explore how the desired high-level goals will translate into observable success and failure measures, and what intervention opportunities exist on the causal and influence pathways that lead to them. This is the cycle in which the co-evolution of operational design and systems design occurs. As design conjectures for feasible elements of the stratagem are identified, their possible positive and negative consequences through various causal and influence pathways need to be explored. Any critical uncertainties resolved through probing actions and further targeted information collection as well as cross impacts between the options being considered also need to be explored. At the same time the system design takes shape and thereby informs the feasibility assessment of options generated. The goal here is to arrive at a stratagem whose elements work together in a synergistic way to identify and mitigate the consequent risks and to identify the additional metrics (measures of performance and proxies for success and failure) needed to support the testing and adaptive refinement of the operational and system design conjectures.

When sufficient confidence is reached to begin execution of the designs, Loop 3 is initiated. To support adaptive execution, monitoring of performance and proxy measures begins in earnest. Monitoring of any additional indicators required to flag conditions arising that call for an adaptive response also begins, looking in particular for reductions in capability calling for a resilience response, and significant changes to the conceptual model (either because of new understanding, or because the situation has changed in a significant way) that would call for an agility response, i.e. a change in approach or stratagem.

As the time arrows in the diagram suggest, the three loops begin in fairly quick succession and all continue in parallel, since much of the evidence needed to test and evolve the conjectures has to come from actual interaction with the situation.

We are now in a position to give a more detailed and systematic reply to the question of what exactly it is that needs to be designed. There are two parts.

- The first part conforms to a traditional interpretation of design, and relates to what is intended to be done, and can be done. Thus at the highest scale, on the system side it calls for specifying the overall structure and functions of the system, including the relationships architecture (with partners, external stakeholders, service providers, etc. that all form parts of the extended 'system'), while on the operational side it calls for a description of the higher level success and failure measures for the situation, and the stratagem—what aspects of the situation it is intended to influence, and how doing so is expected to transform the situation towards success and away from failure. At successively lower scales, more detail is filled in about the subsystem components and the tasks and procedures needed to use the system to implement the stratagem in the particular situation. *Together, these design aspects define the kinds of outcomes that* ***can*** *be generated in a given situation.*

- The second part relates to the adaptive properties of the systems and operations: their ability to understand themselves (their purposes, their capabilities and their current state) and their current context well enough to determine what outcomes they should be generating, when they should change what they are doing, how they need to change themselves in order to do so, and their ability to make those changes smoothly and quickly enough in the context of the situation. In other words, it is how the system and operational designs adapt in a given situation. *Together these design aspects define what outcomes* ***will*** *be generated over time, and hence how well the overall high level goals are achieved.*

Although these two design aspects are equally important, we will pay more attention now to the second part because although the need to adapt and learn is widely acknowledged, adaptibility is the part that is generally much less well understood and addressed, particularly from a design perspective.

Why design is important—Part 2: Design *for* adaptivity

Being adaptive is part of the evolutionary heritage of human beings, just as it is for all life forms to varying degrees. The extent of human adaptivity is amply demonstrated by our successes, including our remarkable recoveries from various failures. But adaptivity is clearly not limitless, and it has proved difficult, for many reasons, to engender the necessary levels of adaptivity in larger-scale human organisations.

The basic algorithm underlying adaptivity is very simple: iteration of a cycle of introducing change, testing it in real-world interaction (or realistic enough simulation) generating success-relevant feedback, which then informs a selection decision—to retain or discard?—and is followed by implementation of the selection decision.

The capability to execute the interaction testing and feedback steps of such a cycle is vested in the system design. But, the processes and authorities to introduce changes into the designs, to direct that success-relevant tests are performed, and to make judgments and implementation decisions as a result, are vested in aspects of the operational design. Most particularly, they are vested in what we term the *meta-decision architecture of the designs*.

As introduced in the previous section, meta-decisions are those design decisions that determine what kinds of decisions should be made when, and by whom, and how they will be supported. Examples of meta-decisions include the allocation and delegation of authorities to individuals and the mandating of particular decision processes; for example: how defence procurement decisions are to be made, or how decisions about recruitment and selection of staff in an enterprise are to be made. So, meta-decisions are decisions about future decisions. Some of the future decisions that will be needed are those that the adaptive processes require i.e. decisions about what changes to introduce, how to evaluate them, whether to retain them and how to implement that decision. For that reason, meta-decisions have the power to both limit and enable the adaptive potential of the enterprise.

The reason for calling it a meta-decision *architecture* is that decisions are rarely independent. As we have seen, both action decisions and design decisions create and destroy possibilities, necessities and opportunities for later decisions, thus creating networks of interconnected decisions. The connection pattern of decisions can be important, especially in the context of adaptive processes, which call for both action and design decisions in particular patterns. The meta-decisions will determine what decision patterns are possible.

This is particularly relevant in the context of design as an integrating concept that can resolve the tension described earlier between, on the one hand, the diversity that adaptation needs and the divergence that can cause and, on the other hand, the coherence that is needed for a complex enterprise to achieve complex inter-dependent outcomes in a complex situation. The way in which design can be used to resolve that tension is through two key avenues: the design of the enterprise's multi-scale metrics framework, and the design of its adaptation architecture.

The multi-scale metrics framework was introduced above as the fourth pillar of a co-evolutionary meta-strategy for dealing with a complex situation. At its highest scale it articulates the values-based measures of success and failure in the situation, i.e. what is really important, in and of itself, to achieve or avert in the situation. Successive scales flesh out the conjecture-based proxies that become the progressively more detailed positive and negative objectives for action, until we arrive at the scale of measures of performance that indicate progress and quality of task implementation, and indicators that convey something significant about the state of the situation and may trigger the need for an adaptive response.

As such, the multi-scale metrics framework is a distillation of the essential structure of the stratagem and its subsidiary courses of action, plans and tasks, identifying:

a. the aspects of the situation that need to be monitored to provide feedback for the various adaptive processes

b. the intermediate goals or proxies at each scale that are therefore built in to the relevant adaptive processes as their selection criteria.

In other words, the elements of the metrics framework provide the direction-setting selection criteria for the adaptive processes as well as guiding the collection of feedback data for the adaptive processes—feedback data in the form of both implementation measures for adaptive refinement of action and effects measures for adaptive refinement of the design conjectures. It is through these two important roles that the design of the metrics framework is such a powerful means of enabling and focusing the power of adaptation in a coherent way onto the enterprise's overall high-level goals.

The second key avenue for doing that is through the design of the adaptation architecture, which describes how all the many adaptive processes, operating in parallel, are put together.

It is evident that adaptive processes in a complex situation can interact in many different ways. For example, they may interact on particular elements of the situation either synergistically (producing cumulative aligned effects) or antagonistically (their separate effects are opposed), one may indirectly modify the impacts or the operation of another, they may interact spatio-temporally to produce oscillations or more complex patterns and so on. In the absence of an organising architecture, the potential for chaos and disintegration is significant. The role of the adaptation architecture is therefore to establish the necessary relationships between the adaptive processes, by such means as nesting them temporally and hierarchically, automating necessary linkages, e.g. to create a trigger for deliberate resolution if one adaptive process results in changes to the operating conditions of another, and creating additional adaptive processes where needed to resolve inconsistencies arising from the separate operations of parallel adaptive processes. Of course, many such relationships, linkages and processes already exist in any complex organisation or enterprise, but the point of calling them out as elements of an adaptation architecture is to pay them the attention they deserve and thus more explicitly shape and enable the adaptive potential of the enterprise in a coherent way.

In most enterprises the actual meta-decision architecture is likely to be largely an accidental by-product of decisions made for reasons of local efficiencies, tradition, administrative convenience, short-term cost-cutting and so on, all of which may have been justifiable within the frames of those decisions, but which do not necessarily combine to produce the most agile and effective use of resources to address overall goals and higher objectives in complex dynamically changing situations. Unfortunately, organisations also have the tendency to invest heavily in the structures and processes they develop to implement their meta-decisions, which then become costly and difficult to modify, resulting in organisational inertia. The complexity of existing enterprises will generally be too great for pure top-down restructuring to be effective, therefore changing the meta-decision architecture will usually be best addressed adaptively as structural and business process design issues at all scales of the enterprise.

Concluding remarks

In summary, what needs to be designed is not just the operation to be performed in the complex situation and the system capabilities that are needed to do it, but also, very importantly, the adaptation architecture i.e. what adaptive processes operate at every scale and throughout the enterprise, how they are linked up, how they themselves can evolve, and the closely related meta-decision architecture.

Together the meta-decision architecture and the adaptation architecture form a significant aspect of the stratagem and determine how the enterprise (both its systems and its operations) are able to evolve over time. So, as environments and social and organisational contexts develop and change, the enterprise can keep learning and adapting itself to achieve and maintain high levels of effectiveness in the eyes of its stakeholders. The overall process that integrates all these aspects is the co-evolution of the four pillars as depicted in Figures 4.1 and 4.2.

A key task at the enterprise level is therefore to empower and support its own systems and operations to develop their own meta-decision and adaptation architectures so that they can keep learning about their own complex changing environments, and keep adapting their own operational and system designs and the metrics frameworks needed to support them.

The benefits of design richly applied in this fashion are many. Not only does it provide an antidote to the dangers of a mechanistic approach, through stimulating systems thinking and eliciting more comprehensive analyses of risks and their mitigations but, even more importantly perhaps, it legitimises the essential characteristics of adaptation: self-correction, tolerance of ambiguity and the restraint necessary to simultaneously pursue a particular approach in a complex situation while also resisting the temptation for total commitment to it, the better to be poised for strategic agility in response to rapid situational learning.

On the other hand, we should not underestimate the inherent difficulties in fully implementing a design approach in an organisational culture that has not yet learned to value these essential characteristics, which might easily (if superficially) be seen as contrary to traditional doctrine. Small steps in the right direction may have more chance of success than attempting culture re-appraisal on a large scale.

To close, and in keeping with these observations, we reflect on the prospects for this broader conceptualisation of design to inform defence thinking and for its full potential to eventually be realised. The inroads that have been made so far under the banner of operational design are important but, from the perspective offered here, they are only part of the story. And in order to escape the bandwagon treadmill and transition into enduring capability and practice they still have to navigate the hazards discussed in the opening section. The other part of the story rests in what the arguments presented here can contribute to buttressing the case for both design writ large and for the nascent practice of operational design, in particular.

There are two obvious opportunities. Firstly, the current practice of operational design relies extensively on discourse to explore multiple perspectives, surface assumptions and generate a richer systemic picture within each of its frames. But, if we acknowledge that the products of these dialectical processes are conjectures then it is evident that there is scope to increase their effectiveness and accelerate learning by more systematic application of an Adaptive Stance to robustly test the conjectures and evolve them in the light of real-world evidence.

Secondly, and conversely, creating the degree of adaptivity required of the defence enterprise at every scale could be greatly enhanced by more deliberate application of operational design thinking, to explore the opportunities and the risks through multiple system frames, and thereby arrive at a better prioritisation of the requirements for adaptivity, and more robust initial designs for them that can be more rapidly evolved to harvest the opportunities, and to be hardened to the risks.

Design thinking and the Adaptive Stance can each stand alone but taken together they are much more powerful. It is worth the effort to save them both from the fad cycle.

Notes

[1] 'A new scientific truth does not triumph by convincing its opponents and making them see the light, but rather because its opponents eventually die, and a new generation grows up that is familiar with it'. Max Planck, *Scientific Autobiography and Other Papers*, referenced in: P. Azoulay, C. Fons-Rosen and J.S. Graff Zivin, 'Does Science Advance One Funeral at a Time?', *NBER Working Paper No. 21788*, December 2015, revised October 2016, p 1.

[2] See, for example: D. Kahneman, *Thinking—Fast and Slow,* New York: Farrar, Straus and Giroux, 2011.

[3] 'A map *is not* the territory it represents, but, if correct, it has a *similar structure* to the territory, which accounts for its usefulness'. Alfred Korzybski, *Science and Sanity: An Introduction to Non-Aristotelian Systems and General Semantics* [1st ed. 1933], Institute of General Semantics, 1994, hardcover, 5th edition, ISBN 0-937298-01-8., p. 58. (Emphases in the original).

[4] E. A. Cohen, *A Revolution in Warfare*, Foreign Affairs, March–April 1996, pp. 37–54.

[5] E.A. Smith, *Effects-Based Operations* Command & Control Research Publications (CCRP), Washington, 2003; P.K. Davis, *Effects-Based Operations (EBO): A Grand Challenge for the Analytical Community.* RAND, Santa Monica, 2001.

[6] D. Alberts, J, Gartska and F. Stein, *Network Centric Warfare*, DOD Command and Control Research Program, Oct. 2003

C. Wilson, "Network centric operations: background and oversight issues for congress", DTIC Document, 2007.

[7] K. Friis *Revising COIN – The Stakeholder Centric Approach.* in S. Gates and K. Roy (eds.): War and State-Building in Afghanistan: Historical and Modern Perspectives. London: Bloomsbury, 2014.

[8] A. M. Grisogono & A. J. Ryan, *Operationalising Adaptive Campaigning*, 12th ICCRTS (International Command and Control Research and Technology Symposium), Rhode Island, 2007; A. M. Grisogono, M. Spaans, M. Spoelstra, E. Douze, & R. Pieneman, *Learning to be Adaptive*, 14th ICCRTS, Washington, 2009; The Technical Co-operation Program, Joint Systems and Analysis Group, Action Group 14: Complex Adaptive Systems Science for Defence. *Synthesis Report: A Conceptual Framework for Adaptation,* DOC–AG14 - CFA #3-2010, 2010.

[9] In brief the Adaptive Stance consists of four elements: (1) ambiguity tolerance; (2) willingness to have one's ideas proved wrong and changing them as a result; (3) turning conjectures and assumptions into explicit expectations, and deliberately seeking disconfirming evidence; and (4) supporting peers and subordinates in doing the same. For more details, see: A. M. Grisogono & V. Radenovic, *The Adaptive Stance– steps towards teaching more effective complex decision-making*, International Conference on Complex Systems, Boston, 2011.

[10] Examples include: US Office of the Chairman of the Joint Staff, *Joint Publication 5-0 Joint Planning*, 2017; Australian Army, *Adaptive Campaigning: Future Land Operating Concept*, Canberra: Defence Publishing Service, 2009; and various implementations of Systemic Operational Design.

[11] P. Ralph & Y. Wand, 'A proposal for a formal definition of the design concept' in: K. Lyytinen, P. Loucopoulos, J. Mylopoulos & W. Robinson (Eds.), *Design Requirements Workshop,* Cleveland, OH, USA, June 3-6, 2007: Lecture Notes in Business Information Processing,.Springer-Verlag 2009, pp. 103–136.

[12] H. A. Simon, *The sciences of the artificial*, Cambridge, MA: MIT Press, 1996.

[13] For example, see: Ministry of Defence, *Red Teaming Guide* (2nd ed.), Shrivingham, UK: Development, Concepts and Doctrine Centre, January 2013. Available online:

https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/142533/20130301_red_teaming_ed2.pdf, accessed 28 May 2018.

[14] *Notes about Figure 4.1:* Labelled arrows indicate the primary impacts of changes in one area on another. The multi-scale structure of each is illustrated in the tables on the right.

[15] *Notes about Figure 4.2:* The supporting metrics framework evolves throughout in response to the needs for feedback and information that arise in the three loops. Each loop is shown to provide direction to the information collection plan, and to be informed by the information collected and structured by the evolving conceptual model. The sequence of initiation is shown but once commenced, all three adaptive loops continue concurrently in an ongoing engagement with the complex situation through both actions taken in it, and information collection from it.