Identifying Cyber Defence Challenges in Acquisition: A Soft-Systems Approach

Prepared by Andrew Beard
of the Defence Science and Technology Laboratory (Dstl)
for the 30th International Symposium on Military Operational Research (ISMOR)

30th July 2013

Abstract

The UK Ministry of Defence (MOD) is embracing cyber as both a capability and operational environment. Military systems are, by necessity, increasingly digitally-enabled and interconnected and subsequently more vulnerable to cyber attack. MOD requires the confidence that cyber defence risks are appropriately managed in the acquisition of new systems, and applies the principles of Information Assurance (IA) to ensure the integrity, availability and confidentially of its information systems.

This paper describes the approach to, and findings of, an independent review of MOD's IA process from a cyber defence perspective. The customer requirement for a high-level, stakeholder-informed review of policy and procedure was met through the application of soft methods. Elements of the Appreciative Inquiry Method (AIM) and Soft Systems Methodology (SSM) were employed to engage stakeholders across organisational boundaries, and at various levels of influence, to explore the problem space.

The review identified that, whilst the principles of the IA process are generally sound, its implementation is ineffective in managing cyber defence risks. Follow-on work will begin to address the most critical issues through a coherent, cross-Government approach.

ISMOR
The International Symposium
on Military Operational Research

[dstl]

Ministry
of Defence

Introduction

In writing this paper I thought carefully about how best to approach it and, as you the reader may have already remarked upon, I elected to write it in the first person. I imagine this is somewhat unusual for conference papers (I have certainly never come across any written in that way) but there are, I believe, good reasons for doing so in this instance.

The work I describe herein has been – and still is, at the time of publication – a personal journey into the realms of soft systems, and I was keen to capture and share my new experiences with these methods in a manner that befitted that journey.

My background is one of ‘traditional’ science: I am a physicist by discipline and most of my career has been spent undertaking so-called ‘hard’ operational analysis in a variety of different contexts. It is through the breadth of application, rather than academic study, that I gained an appreciation of systems problems, which then led me into softer systems work.

Thus, I am a relative newcomer to the practice of soft methods, although I have come to realise that my work has, for some time, warranted such an approach. Had I known what I now know about these approaches when I began this work, I probably would have applied them more rigorously, and certainly earlier on.

Fortunately for me – and the study to which I applied these methods – I have learned that soft methods are not an exact science, and my somewhat inconsistent application of them has not been detrimental to the work. On the contrary, even their intermittent use has brought clarity and structure to aspects of the problem space.

I also discovered elements of these methods that did *not* lend themselves well to the problem space, and whilst my unfamiliarity with the methods may have influenced my success with them, I nonetheless explore these observations in the paper.

It is my intention for this paper to serve as an insight into the practical application of several soft methodologies to a large, complex, and unwieldy systems problem – that of defence acquisition. I have written it in such a way so as to appeal (hopefully) to newcomers to the field of soft systems and experienced practitioners alike. The work itself is by no means complete and there remains considerable scope for further application of systems methods to embed cyber defence in acquisition.

Background

Much has been written on the subject of defence acquisition and, more recently, cyber and cyber defence. However, the integration of cyber defence into the acquisition process has remained curiously absent since the adoption of cyber, despite widespread concern that military capabilities were being acquired without due consideration of their potential cyber vulnerabilities. Of particular concern were the platforms and so-called ‘non-networked’ systems that tend not to be subjected to the same level of information security scrutiny that, say, networked-systems might during their acquisition.

The UK Ministry of Defence (MOD) has launched a variety of studies to investigate the extent of the cyber vulnerabilities of in-service platforms and systems. However, to date there had been no examination of the *process* through which our future capability is acquired and the cyber vulnerabilities that might be introduced unknowingly or otherwise mismanaged. It was feared that, without intervention, our existing acquisition practices could introduce capability into service with undisclosed cyber vulnerabilities that might be identified and exploited by our adversaries.

Orientating the Study

To address this I was asked by the MOD customer to conduct a short, high-level review of the policy and process of *Information Assurance* (IA). This is the process through which cyber defence, as we now refer to it, is purported to be addressed in acquisition, although there is no reference to 'cyber defence' in any of the IA documentation. (The concepts of IA were established long before anyone was talking about cyber or cyber defence).

Information assurance, as MOD interprets it, is predicated on the principles of *integrity*, *confidentiality* and *availability* of information. That is, an assured system is one in which the information it handles is not susceptible to unauthorised manipulation (integrity), is protected from unauthorised disclosure (confidentiality) whilst remaining accessible to authorised users (availability).

The customer believed that any change to the IA process (and by extension the acquisition process) would require a joined-up approach between several organisational entities, and saw my work culminating in a senior-level cross-Government workshop to achieve this.

I had, at the time, only been working on cyber for a short while and had no knowledge of IA whatsoever. This was perceived to be an advantage as I was not in a position to impress any bias upon the review, which therefore should, in theory, make it easier to remain impartial.

It also coincided with my recent development into the role of *technical consultant* – a term Dstl uses to encapsulate the skills, competencies and personal attributes necessary to provide high-impact and timely science and technology advice to decision-makers. (Establishment of the technical consultancy roles is symptomatic of Dstl's deliberate rebalance away from performing original research towards managing and integrating the delivery of research placed elsewhere, a transformation that will become more pronounced over the coming years).

In this mindset I characterised the study in a consultative fashion which, crucially, would enable me to make progress despite my unfamiliarity with the problem domain. Whilst I realised that soft approaches would be necessary, I didn't feel I had sufficient grasp of the discipline to confidently apply problem structuring methods to the *totality* of the problem space or govern my approach.

Exploring Soft Systems Methods

The customer was keen for me to engage with a large community of stakeholders, whom played different roles in acquisition and would no doubt offer unique insights into the challenges for cyber defence. Having recently undertaken some training in both Soft Systems Methodology (SSM) and the Appreciative Inquiry Method (AIM) I saw this as an opportunity to provide some semblance of structure to my early engagements.

It is not the purpose of this paper to regurgitate SSM for the reader, and I forgo such a digression for the sake of brevity. A substantial body of work exists in the field, notably by Checkland [1] and Wilson [2], which will introduce the reader to the methodology far better than I could accomplish in this short paper. I will, however, devote some discussion to the principles behind SSM, what AIM is and how it differs, and the reasons why I elected to employ these methods as I did.

SSM is predicated on the design of a conceptual system – one which may be explicitly and defensibly derived from our understanding of the problem itself – that is then used to test against the real world. It describes a series of steps that are best done *alongside* one's stakeholders, rather than in isolation or through individual consultations. Figure 1, from Checkland and Scholes (1990) [3], depicts the seven stages of SSM (although, notably, in later texts Checkland prefers not to sequence the method quite as rigidly).

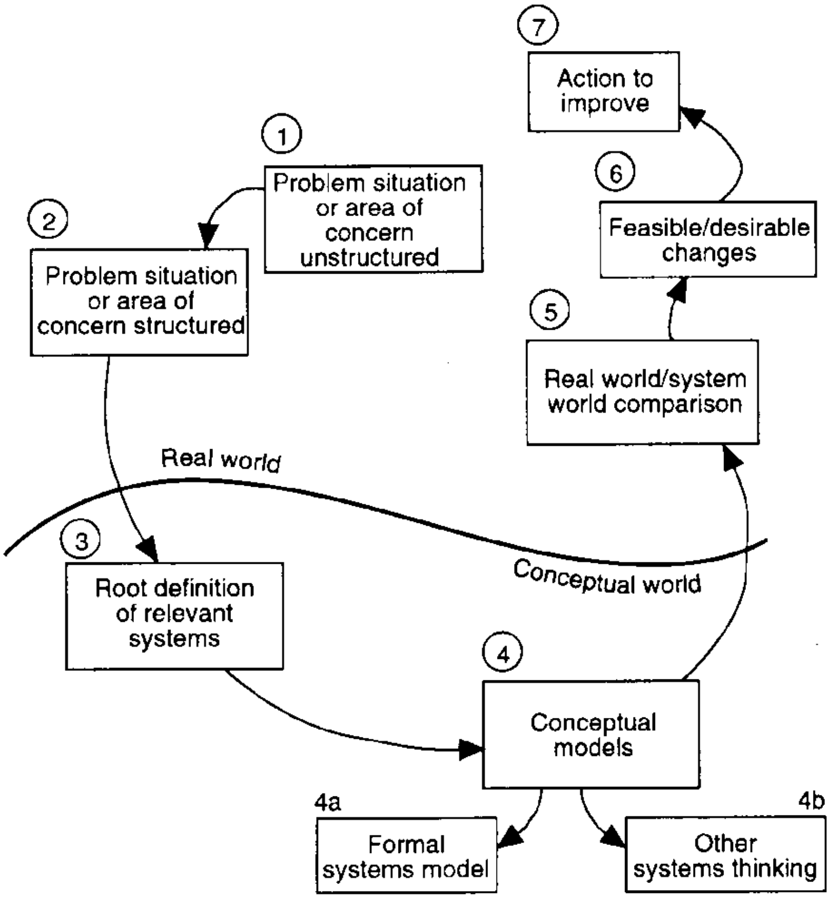
7
Action to improve
1
Problem situation or area of concern unstructured
6
Feasible/desirable changes
2
Problem situation or area of concern structured
5
Real world/system world comparison
Real world
3
Conceptual world
Root definition of relevant systems
4
Conceptual models
4a
4b
Formal systems model
Other systems thinking

Figure 1: The seven stages of SSM, from Checkland and Scholes (1990) [3]

AIM, on the other hand, is well-suited to individual consultations (at least in the early stages) and is predicated on exploring the *as-is*, rather than the *as-should-be*, from the outset. AIM, developed by Stowell and West [4], is a specific technique in the field of Appreciative Inquiry (AI), first coined by Vickers in 1968 [5]. (The American-English *inquiry* is preferred over the British-English *enquiry*). Instead of asking what the *problems* are with a situation – which tends to focus one's attention in a particular direction – AI asks what *is* working. By focusing on what an organisation does well stakeholders are encouraged to understand its strengths and the positive contribution they bring, and to build on these.

AIM takes the principles of AI and provides a structure with which to tackle systems problems. The process starts by asking individual stakeholders to draw a conceptual model of the system as they understand it and, as such, was an ideal approach to take with my initial stakeholder engagements to learn about the problem space. Figure 2, from Stowell (2012) [6], depicts the three stages of AIM.

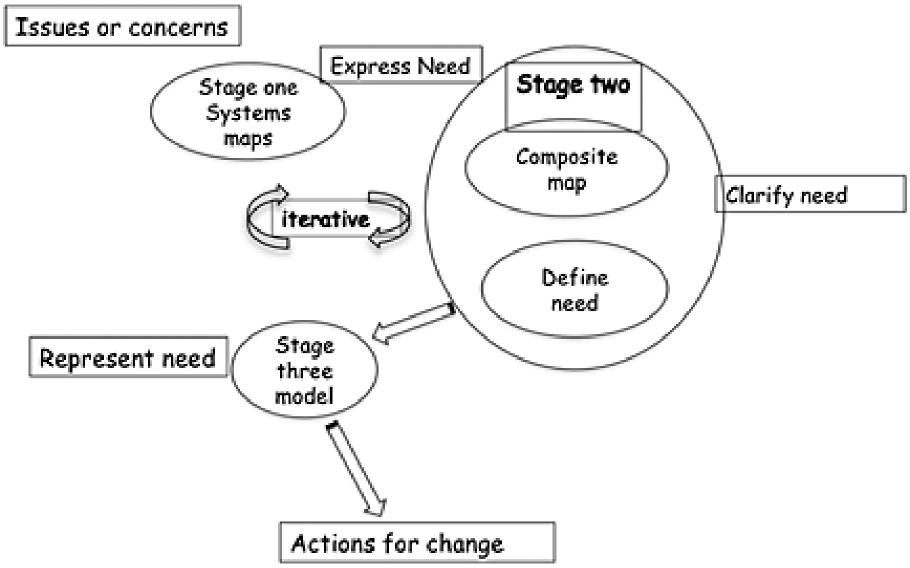

Figure 2: Schematic of the three stages of AIM, from Stowell (2012) [6]

Entering the Problem Space

I anticipated that few of the participants would be comfortable with the abstraction of their business to a 'system', so instead I simply asked them to explain how information assurance worked in Defence, from their perspective, by drawing it. In some instances the participant led the drawing; in others, where they were less confident articulating the process diagrammatically, I reflected back their observations in my own drawing. In some cases the pen exchanged hands frequently over the course of the session.

As stipulated by the method, I approached each consultation with a blank sheet (or whiteboard) on which to express the participant's unbiased perception, resorting to pre-prepared material only when absolutely necessary to engage the participant. This maximised the participant's ownership of their product.

Interestingly, despite the 'appreciative' nature of my consultations, participants readily opened up about the problems with the IA process, often without prompting. Participants were, almost without exception, intent on distinguishing the *theory* of the process – that which they had invariably portrayed in their system diagram – from the *reality*, which was, they argued, somewhat inconsistent.

Thus, over the course of the discussion, our systems diagrams tended to evolve into 'rich pictures', whereby the well-defined illustration of the process was annotated with words, symbols and links depicting the friction or conflict between entities. Figure 3 depicts two examples of the output produced when stakeholders were asked to draw the model of how the IA process worked. Through the course of my consultations I began to feel more confident in the problem space, with expressions of both the theory and practice being increasingly repeated across sessions, albeit independently.

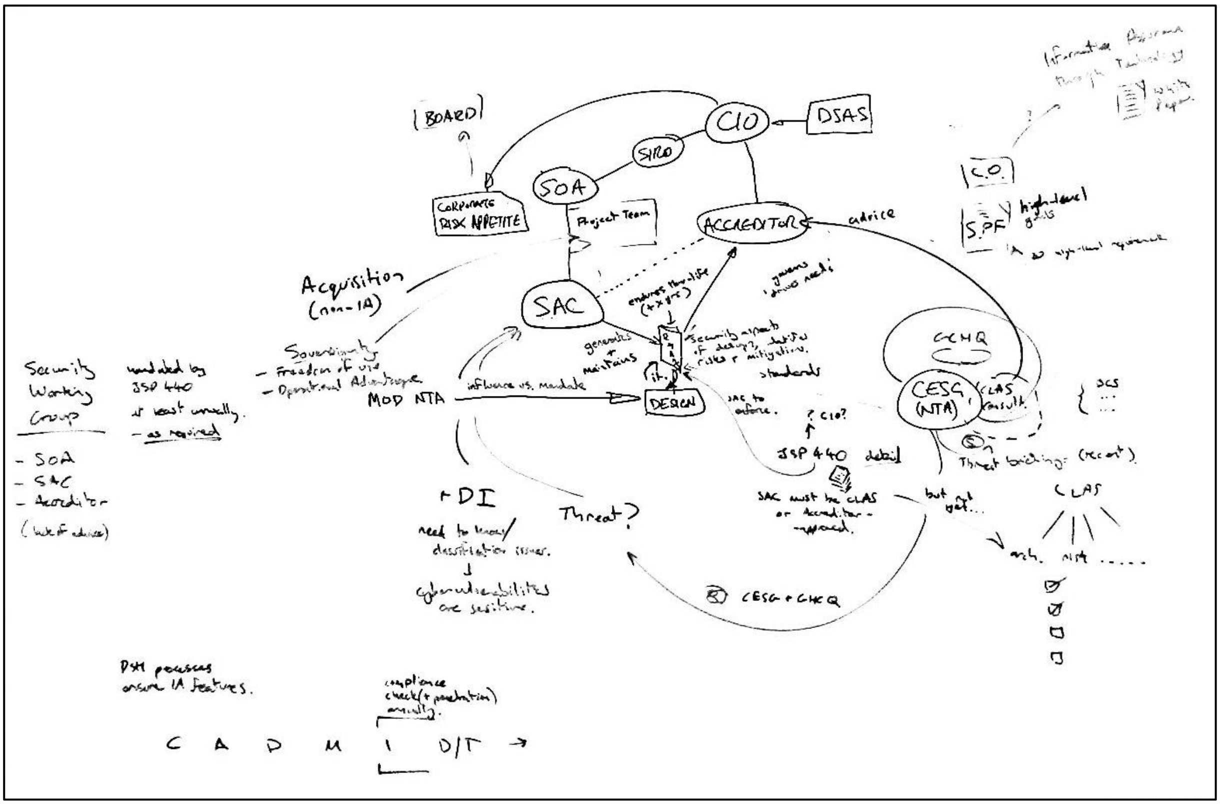

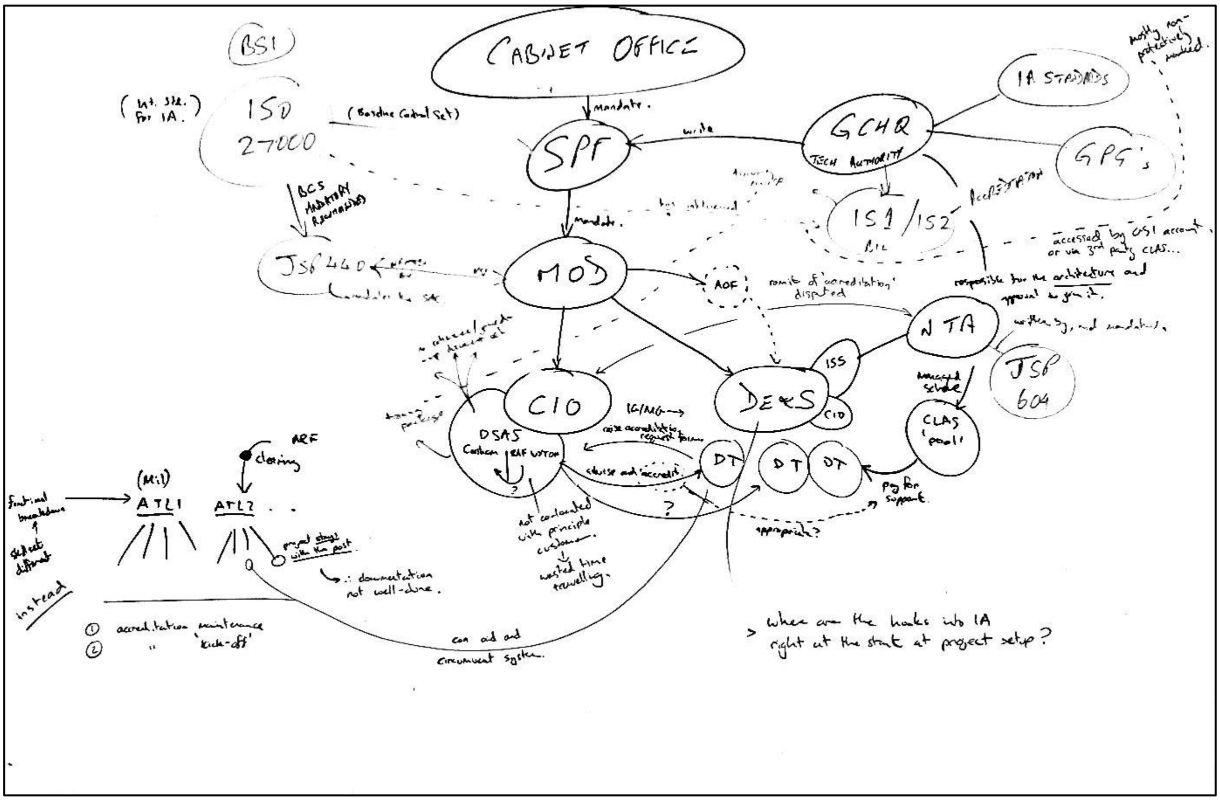
BSI
ISO 27000
(Int. Std. for IA.)
(Baseline Control Set)
CABINET OFFICE
mandate.
SPF
write
GCHQ
TECH AUTHORITY
IA STANDARDS
GPG's
IS1/IS2
ACCREDITATION
JSP 440
MOD
mandate.
AOF
remit of 'accreditation' disputed
NTA
responsible for the architecture and approval as given it.
written by, and mandated,
JSP 604
ISS
CIO
DE&S
CIO
DSAS
DT
DT
DT
CLAS 'pool'
pay for support.
appropriate?
not contacted with principle customer.
wasted time travelling.
(Mil) ATL1
ATL2
ARF
clearing
documentation not well-done.
can aid and circumvent system.
instead
① accreditation maintenance
② " 'kick-off'
> Where are the hooks into IA right at the start at project setup?
accessed by CESG account, or via 3rd party CLAS...

Figure 3: Examples of rich pictures generated through one-to-one consultations with stakeholders.

Emerging was a picture of a process that was theoretically sound and fundamentally unarguable, but that suffered substantially in its implementation. In almost every facet of the process revealed to me, an example had been offered as evidence of how it had failed in a particular instance. Many problems concerned poor behaviours: sometimes defensible from a particular perspective, sometimes not – although always driven by local incentives and organisational culture. Others concerned confusing process documentation, lack of governance, insufficient resources, improper incentives and myriad other problems that one might expect to find in a poorly implemented process.

I felt obliged to reflect on whether this was indeed indicative of a serious problem for acquisition or merely 'business as usual' for a large process-driven organisation such as MOD. Perhaps stakeholders were over-emphasising the scale of the problem? Ultimately, I decided, it all came down to *impact*. A few years earlier, in 2009, Charles Haddon-Cave QC had published his damning independent review of the Nimrod incident in 2006 [7], a report that made very sobering reading for many people in Defence. Here was an example of a catastrophic failure that had been the product of a sequence of seemingly innocuous mistakes, erroneously assured by a 'lamentable' process. With this in mind it was easy to envisage how a failure to conduct IA effectively might lead to undiscovered cyber vulnerabilities that would jeopardise a platform system, potentially leading to loss of life.

Moreover, the IA process – even when conducted properly – was not capable of addressing all aspects of cyber defence in acquisition. Issues surrounding the risk management of cyber vulnerabilities, particularly for complex systems-of-systems and systems on networks, was non-trivial. Identifying the boundary around the 'system' to undergo accreditation was difficult. These were the challenges facing MOD in addition to the improper conduct of IA.

Structuring the Problem Space

The second stage of AIM requires the facilitator (me, in this instance) to synthesize the individual stakeholder's maps into a single composite map that best-reflects the consensus view. I found this rather difficult, and made several failed attempts before concluding that a different approach was needed. Upon reflection, I believe my difficulty arose because I attempted to incorporate too many concepts into the systems model. Because the stakeholders' maps were essentially rich pictures it was not trivial to separate the pure systematic components from the non-functional 'issues'. This underscored an intellectual struggle I had been wrestling with since undertaking my SSM learning; that of the difference between issue-based and functional/process-based modelling.

In practising SSM, particularly at the conceptual modelling stage, it is important to make the distinction between a description of the system (or process, in this instance) of interest and that of activity or issues associated with achieving the desired system state. Both are valid approaches (depending on the nature of one's problem space), yet combining the two (whether consciously or not) can quickly confuse matters. This exercise reinforced the importance of this distinction.

Instead of persevering with the composite map I made a conscious decision to step out of systems thinking for a period and turned to the written word to describe the system as I understood it. What emerged from this was a short report comprising a 'pen-picture' of the information assurance process, followed by a discussion of the issues raised by my stakeholders.

I found this to be an easier medium in which to document my findings, although – as a strictly linear construct – I struggled with where to begin the 'story'. (Re-applying SSM later revealed that this was probably because of the considerable interdependency between the issues; a feature of the problem space that was not possible to explore in prose).

The purpose of writing the report was two-fold. Firstly, to demonstrate (not least to myself) that I understood the problem space and had achieved sufficient coverage and depth to make sound judgements. Secondly, to provide some substance upon which to engage stakeholders on the *totality* of the problem.

From my consultations it was clear that the stakeholders were well-informed and knowledgeable about their local issues and in some cases more widely. However, no single stakeholder had visibility of the entire process or the issues affecting it. I believed it important to establish a common, baseline understanding of the whole system and its issues amongst my stakeholders and tailored my report accordingly. It was concise, (fairly) high-level and – above all – short (circa ten pages).

The customer and I had hoped that stakeholders would take ownership of the report itself, and I encouraged particular stakeholders to contribute passages they were best-placed to write. I envisaged – perhaps naïvely – a situation whereby stakeholders identified their own solutions to the problems identified and documented these in the report for re-circulation, thereby progressing the collective thinking from the problem space into the solution space.

Upon reflection I should, perhaps, be unsurprised that this second objective was not met through the report alone. In order to make an impact with stakeholders who shared concerns about cyber defence in acquisition but were otherwise pre-occupied with delivering defence core business I needed more than just a description of the problem.

Several stakeholders correctly identified the absence of any evidence in support of the assertions in the report. I had deliberately refrained from citing specific evidence for two reasons, explained in the report. Firstly, the mere suggestion of a particular cyber vulnerability associated with an operational platform, device or other system would be highly classified, and I wanted to ensure my report remained accessible to all stakeholders, not just those with appropriate access.

Secondly, the majority of the issues identified by stakeholders were that of process implementation; that is, the way in which IA was being conducted. Establishing a causal relationship from poor acquisition performance to a specific cyber vulnerability, or even mis-management of cyber risks, would be tenuous at best. For these reasons, I elected not to engage in sourcing and documenting supporting evidence – notwithstanding its importance – and returned to systems thinking to reconsider my approach.

Rediscovering Soft Systems Methods

It was in the re-application of SSM – in particular a 'CATWOE' analysis of the problem space – that I discovered a potentially critical obstacle. CATWOE is a mnemonic for **C**ustomer (or **C**lient), **A**ctor, **T**ransformation, **W**orldview (loosely derived from the German 'Weltanschauung', meaning 'wide world perception'), **O**wner and **E**nvironmental Constraints. It is typically used within SSM to validate so-called 'root definitions' – statements encapsulating the what, how and why of a given system's function – to ensure they are rigorous and comprehensive.

In applying CATWOE, one seeks to answer each of the following questions (which have been re-arranged to match the order of the 'six questions' in systems thinking practice [8]):

1. Transformation — What transformation does this system bring about? (What are the inputs and what transformation do they go through to become the outputs?)

2. Worldview — What particular worldview justifies the existence of the system? (What point of view makes this system meaningful?)

3. Actors — Who are responsible for implementing this system? (Who would carry out the activities which make this system work?)

4. Customers — Who are the beneficiaries or victims of this particular system? (Who would benefit or suffer from its operations?)

5. Owners — Who has the authority to abolish this system or change its measures of performance?

6. Environmental Constraints — Which external constraints does this system take as a given?

CATWOE, being a simple yet powerful means of insight, has been adopted by a variety of other soft systems methods (including AIM), and may even be employed independently to characterise human activity systems – as I was doing in this instance.

In attempting to identify the customers, actors and owners of the 'system' – which, by now, I had begun to refer to as 'cyber-hardened acquisition' – I realised that it was not obvious whom would be taking my recommendations forward and integrating cyber defence into acquisition.

Acquisition itself is 'owned' by the Chief of Defence Materiel (CDM), although governance of the process seemed to sit elsewhere, under the Director General Finance (DG Fin). Furthermore, cyber as a capability – which includes all aspects of cyber defence – sits within the newly-established Joint Forces Command (JFC).

Not only that, but, turning CATWOE onto the study itself, I realised I had no senior-level champion for the work whom could drive it forward and rally engagement at senior levels across the various organisations involved. Thus far I had been engaging with stakeholders at my level to understand the problem space, but without a senior champion any solution I or my stakeholders proposed would fail to have an effect on the acquisition process.

I now had a reason – an imperative, really – to convene *senior*-level stakeholders from across the relevant organisations: the identification of problem ownership. I knew enough about the problem space to articulate the challenges to cyber defence in acquisition in front of these senior stakeholders (I had already approached several of them directly beforehand) to stimulate a debate about whose problems they were. This was the next, necessary step to gain traction on the issues I had uncovered.

I reconvened a small group of informed stakeholders – people whom I had already consulted with on the study – alongside some systems-minded colleagues whom had no prior exposure to the work. This union was deliberate to provide a fresh perspective on the work, with its newfound direction, and adopt a systems approach to the workshop. It also took advantage of my department's evolving ways of working on high-impact systems consultancy tasks.

I facilitated an internal workshop with this group, using rich pictures, so-called 'PQR' statements (essentially asking what, how and why), CATWOE or the 'six questions', and a root definition of the idealised process I had developed alongside the customer.

Our root definition, reproduced here for completeness, emerged not from the rigorous application of SSM, but the assembly of the concepts we wished to convey into a meaningful statement of vision. It does not, therefore, stand up to the scrutiny of CATWOE as it should be applied within the method, and I leave it as an exercise for the reader to identify, through CATWOE, how our imprecise root definition of 'cyber-hardened acquisition' could be improved.

> *"An established process for cyber-hardened acquisition and through life support that protects our digital assets against the evolving cyber threat. A process underpinned by a culture of security, with the right governance, structures and ways of working, that incentivises the right behaviours across government and industry."*

With these hand-drawn materials I introduced newcomers to the problem space and validated the thinking that led me to conclude a senior workshop was necessary to address the apparent lack of ownership. Figure 4 shows the rich picture used as a basis on which to bring my workshop delegates up to speed.

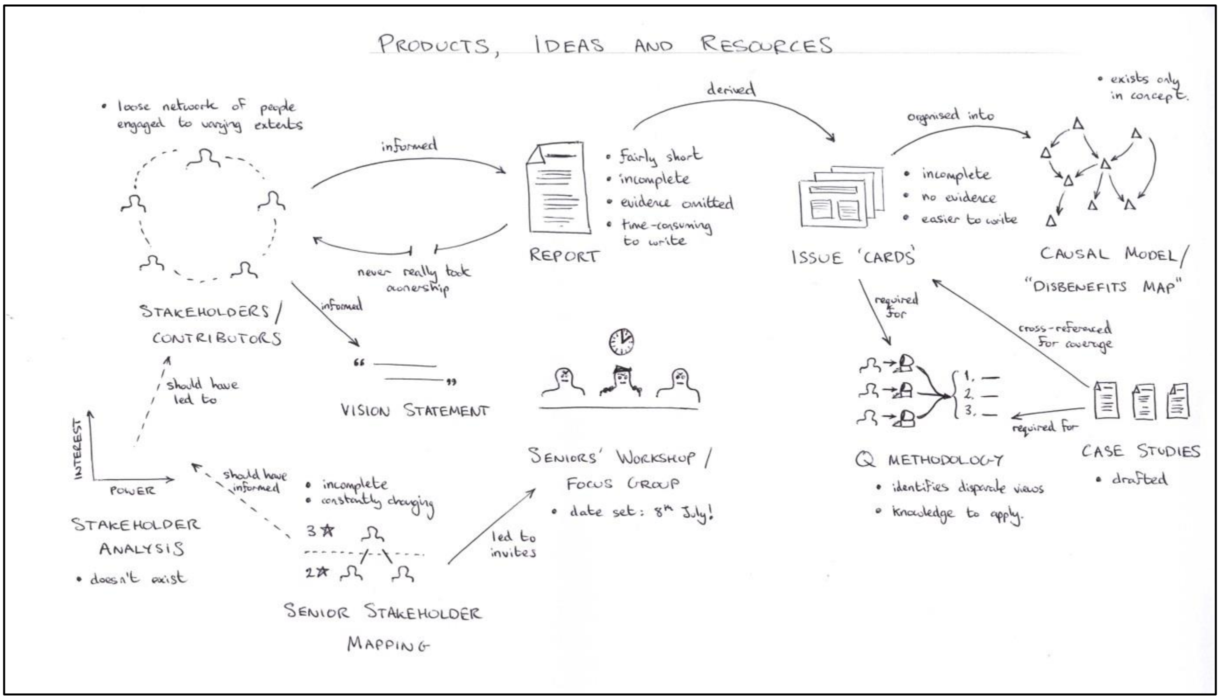

Figure 4: My issue-based rich picture of the situation prior to workshop planning.

This time around I remained acutely aware of the distinction between the functional/process aspects of the problem space and the issues associated with achieving the desired end state. In this instance it was the issues that warranted attention, in particular that of establishing senior-level ownership.

Key to this session was a problem structuring chart I had developed for the occasion, combining PQR analyses with my answers to Facione's 'IDEALS: Six Questions for Effective Thinking' [9]. These six questions, each beginning with a letter from the mnemonic IDEALS, may be used to help arrive at and justify a decision. In this sense they are not dissimilar to the 'stages of soft systems thinking' [8].

I	Identify the problem	*What's the real question we're facing here?*
D	Define the context	*What are the facts and circumstances that frame this problem?*
E	Enumerate choices	*What are our most plausible three or four options?*
A	Analyse options	*What is our best course of action, all things considered?*
L	List reasons explicitly	*Exactly why are we making this choice rather than another?*
S	Self-correct	*Okay, let's look again. What did we miss?*

Figure 5 shows this problem structuring chart, unaltered from the original I used in the workshop. It was readily apparent from this framework that a gathering of senior stakeholders was necessary to identify respective owners of the process.

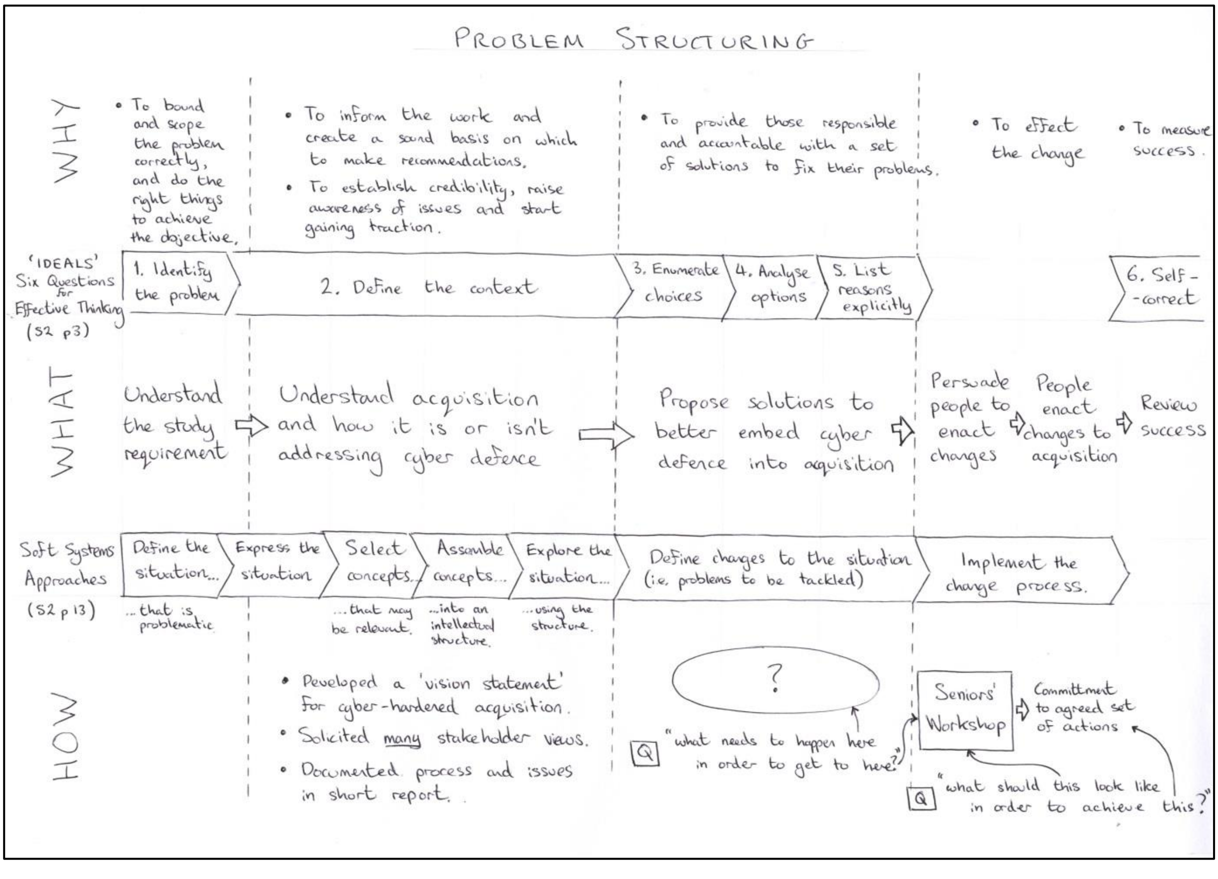
PROBLEM STRUCTURING
WHY
• To bound and scope the problem correctly, and do the right things to achieve the objective.
• To inform the work and create a sound basis on which to make recommendations.
• To establish credibility, raise awareness of issues and start gaining traction.
• To provide those responsible and accountable with a set of solutions to fix their problems.
• To effect the change
• To measure success.
'IDEALS' Six Questions for Effective Thinking (S2 p3)
1. Identify the problem
2. Define the context
3. Enumerate choices
4. Analyse options
5. List reasons explicitly
6. Self-correct
WHAT
Understand the study requirement
Understand acquisition and how it is or isn't addressing cyber defence
Propose solutions to better embed cyber defence into acquisition
Persuade people to enact changes
People enact changes to acquisition
Review success
Soft Systems Approaches (S2 p 13)
Define the situation... ...that is problematic
Express the situation
Select concepts... ...that may be relevant.
Assemble concepts... ...into an intellectual structure.
Explore the situation... ...using the structure.
Define changes to the situation (i.e. problems to be tackled)
Implement the change process.
HOW
• Developed a 'vision statement' for cyber-hardened acquisition.
• Solicited many stakeholder views.
• Documented process and issues in short report.
?
Q "what needs to happen here in order to get to here?"
Seniors' Workshop
Committment to agreed set of actions
Q "what should this look like in order to achieve this?"

Figure 5: My problem structuring diagram, created to facilitate an internal systems workshop.

Having achieved consensus agreement of *what* needed to be done I devoted the rest of the session to determining *how*. I was fortunate to be in the company of experienced consultants whom had tackled similar challenges before – that of conveying complex systems issues to senior non-specialists – and we soon identified the need for a simple visualisation to expose the problems with cyber defence in acquisition. This visualisation was required to communicate the *impact* that the problems, if left unaddressed, would have on military capability, using language that senior decision-makers could identify with.

In my rich picture I had identified the close-coupling of several of these problems, and expressed how the complete set of problems might be interrelated through a complex cause-and-effect network, possibly even with causal loops. (This is visible in the top-right of the rich picture in Figure 4). I had begun exploring this approach prior to my internal workshop, having being introduced to the concepts of system dynamics a few months earlier.

Although the problems I had identified were not strictly functional components of the system (rather, they were known deficiencies) they were readily expressed in an influence diagram, and I suspected they would lend themselves well to causal loop modelling. Together we concluded that this was an appropriate foundation on which to build the visualisation, as it would present a traceable route from the problems in acquisition to their ultimate impact on capability.

Expressing the Problem Space

My first attempt at such a model was successful only insofar that it expressed the problem space; a communicable version would thwart me for a little while longer. After deconstructing the known problems into indivisible statements and recombining them in an all-encompassing causal model I discovered the true nature of the problem space. A space containing more than one hundred discrete problems and exhibiting complex interdependencies far beyond what I had anticipated.

Having undertaken the exercise I felt confident that the problem statements featured at the nodes (or ‘vertices’) of my diagram and their connections (or ‘edges’) were apposite and representative of the problem space. To decompose the problem statements any further was to introduce redundant nodes into the model, whilst to assimilate any of the problem statements together severed key interdependencies.

Because of the sensitivity of many of the problems described by the model it is not possible to reproduce it here in its entirety. Instead I hope to satisfy the reader with an extract thereof containing one of the causal loops I identified. In Figure 6 I depict the cyclic relationship between three issues pertaining to the retention of information assurance specialists.

In this instance, a lack of suitably-qualified and experienced IA specialists causes the few available to become overcommitted to projects. High workload and fatigue were cited as reasons for IA specialists not remaining in post for long, which in turn exacerbates the shortfall, creating a feedback loop. Not shown in this extract is the multitude of inbound and outbound dependencies to these nodes alone. Extrapolating across the totality of the problem space then presents a substantially complex picture which, whilst nonetheless useful, was not conducive to conveying the required messages to non-specialists.

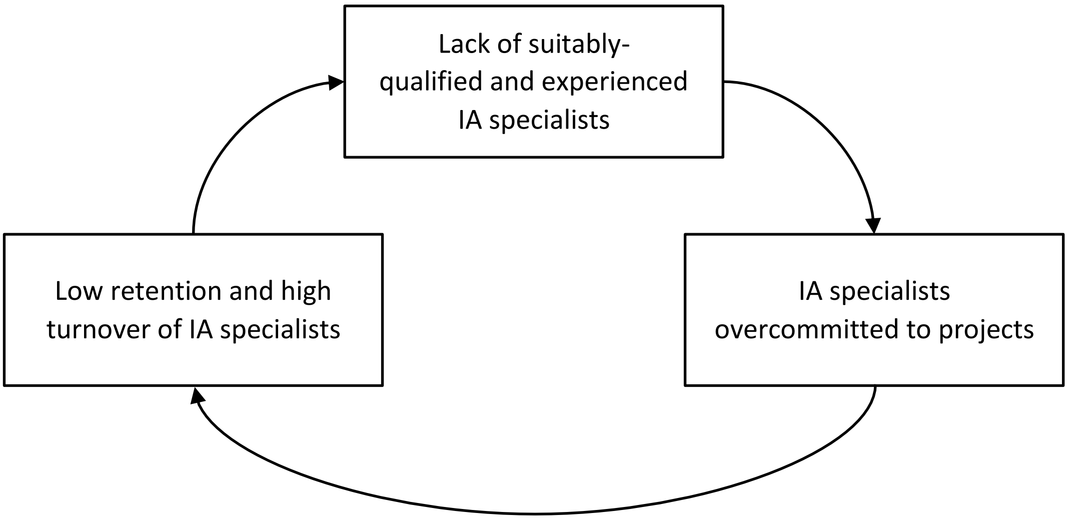
Lack of suitably-qualified and experienced IA specialists
IA specialists overcommitted to projects
Low retention and high turnover of IA specialists

Figure 6: Extract from the causal model of issues affecting the information assurance process with extraneous connections removed for clarity.

Working in software I spent many hours repositioning the nodes of the graph in an attempt to reduce the apparent complexity of their interrelationships, hoping that a natural clustering would emerge by which to simplify the picture. When this failed I resorted to grouping the statements by 'theme' as it was apparent that certain problems pertained to matters of talent management (such as in the example above) and others to governance, culture, behaviours, risk management, and so on. Whilst this proved to aid understanding of the nodes themselves, it exacerbated the intersection of edge lines, and was only comprehensible with these edges hidden.

It was at this point that I circulated the product for internal review, fully-aware that the complete picture was not fit for consumption at the stakeholder workshop. By now, though, I was keen to expose the complexity of the problem space to my systems colleagues and identify a solution to visualise it more effectively, as I had exhausted all the techniques I was aware of.

I was, then, unsurprised by the feedback I received regarding the networked picture – it was perceived as far too complex. What did surprise me was how well the 'themed and grouped' version was received, despite having no cause-and-effect relationships portrayed. Exploring this I discovered that these relationships were themselves not necessary to convey the impact of the problems, provided that the impact remained prominent in the visualisation.

I had previously rejected this stylisation, believing the problem statements and assertions would lack *provenance* in the absence of an explicit cause-and-effect construct. This not being the case, I proceeded to aggregate the hundred-odd issues into a more manageable number; a task made substantially easier by the absence of any interdependencies. A few iterations later I arrived at a simple yet informative illustration of the problem space, comprising approximately twenty issues, distributed fairly evenly into six high-level groups.

Even at this level of abstraction the problem statements and their impact remain too sensitive to disclose herein, however the essence of the visualisation and the top-level themes that emerged are replicated in Figure 7. This pictorial representation of the challenges to cyber defence in acquisition contained sufficient detail to stimulate rich discussion at the senior-level workshop, whilst remaining simple and accessible enough to engage all stakeholders.

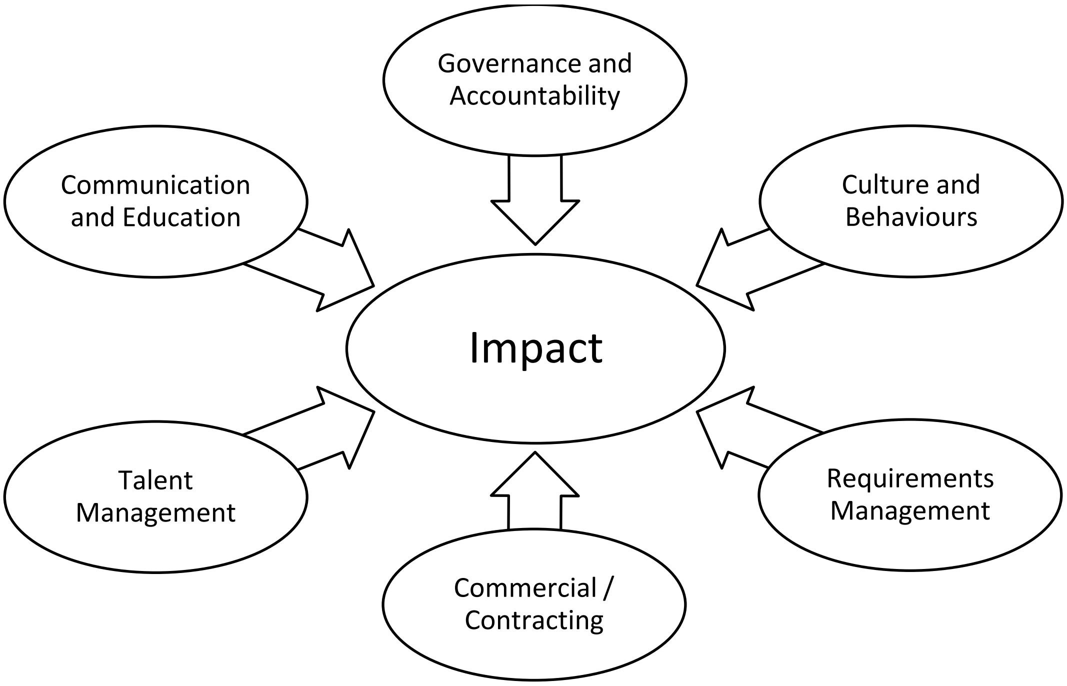

Figure 7: Sanitised abstraction of the challenges to cyber defence in acquisition and their impact.

I described earlier how the problems exposed by stakeholders with whom I consulted were not technical in nature. Rather, they were predominantly issues of process implementation, in particular that of the softer 'people' aspects. Irrespective, it was necessary to tease out how these issues, if left unaddressed, might affect the risk of cyber vulnerabilities creeping into military systems, and the impact that would have on military capability.

To articulate this simply and plainly I aggregated all the impact nodes of the original causal model into two distinct high-level issues, and related these to the impact on military capability. Figure 8 depicts this concept as featured in the issues diagram in Figure 7.

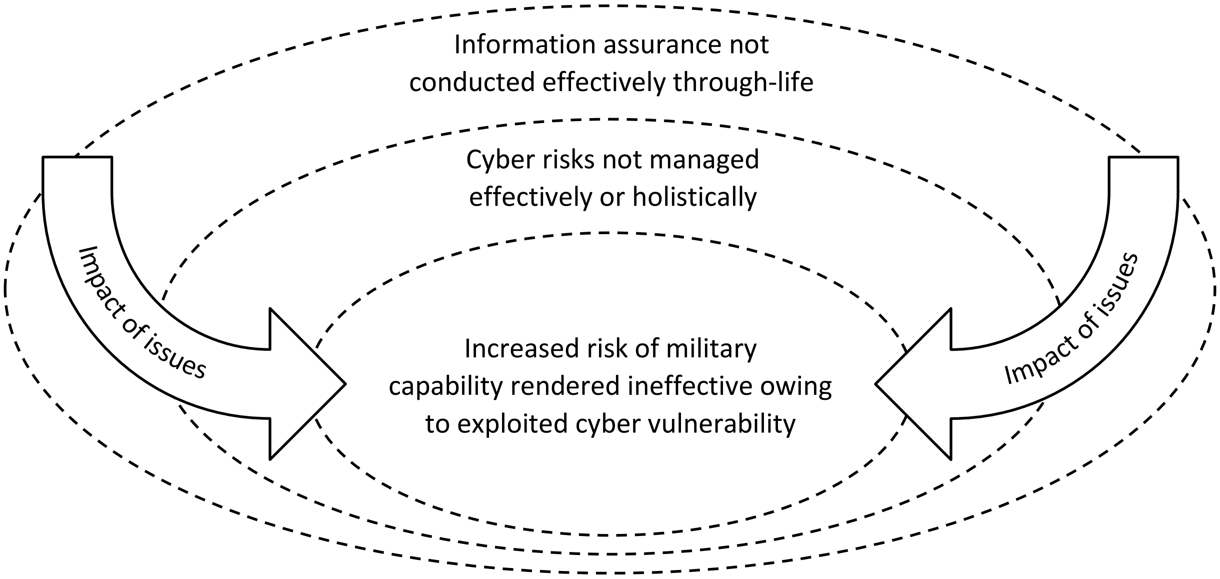
Information assurance not conducted effectively through-life
Cyber risks not managed effectively or holistically
Increased risk of military capability rendered ineffective owing to exploited cyber vulnerability
Impact of issues
Impact of issues

Figure 8: Conceptualisation of the two main issues for cyber defence in acquisition and their impact.

Making an Impact

With these visualisations I facilitated the senior stakeholder workshop and succeeded in exploring the problem space with the delegates. The pictorial representation of high-level issues, whilst (in my view) a gross simplification of the underlying complexity of the problem space, proved to be effective in engaging senior stakeholders in the aggregated issues. A reference to the underlying causal map was all that was necessary to convey the provenance of the issues as I presented them.

In syndicates the delegates were asked to prioritise the twenty-odd issues according to the impact they would have on future military capability if left unaddressed. Then, borrowing from the Delphi structured communication technique [10] (albeit outside the realm of forecasting), the delegates were invited to discuss the syndicates' prioritisations and given an opportunity to reprioritise in order to arrive at a (modest) consensus. Common across the syndicates were high-impact issues of governance and accountability, culture and behaviours, and communication and education. Issues associated with the other three themes were perceived with varying impact, although all (with only one or two exceptions) were considered to be relevant across the syndicates.

A responsibility assignment matrix, commonly referred to as a *RACI matrix* [11] was then employed to identify ownership of the highest-impact issues. RACI is a mnemonic for **R**esponsible, **A**ccountable, **C**onsulted and **I**nformed, and is often used to clarify the roles and responsibilities in organisational structures. Smith [12] defines these as:

R Responsible *Those who do the work to achieve the task. There is at least one role with a participation type of responsible, although others can be delegated to assist in the work required.*

A Accountable *The one ultimately answerable for the correct and thorough completion of the deliverable or task, and the one who delegates the work to those responsible. There must be only one accountable specified for each task or deliverable.*

C Consulted *Those whose opinions are sought, typically subject matter experts, and with whom there is two-way communication.*

I Informed *Those who are kept up-to-date on progress, often only on completion of the task or deliverable, and with whom there is just one-way communication.*

For the purposes of the senior stakeholder workshop the delegates were asked to identify only those roles *accountable* and *responsible* (in that order) for addressing the issues, as I viewed these as fundamental to 'ownership'. In several cases accountability was identified at the highest levels of the MOD, up to and including top-level budget holders (typically 'four-star' positions in MOD language). Where responsibility could be identified this was typically held two or more levels down in the MOD organisational hierarchy.

As I had anticipated, there were several issues for which responsibility – and even accountability – could not be readily identified. It was in the discussion of these issues that the need for inter-organisational cooperation to address them became self-evident. By the end of the workshop, the delegates had identified a number of mitigating actions to the highest-impact issues, spanning the three themes described above. Thus, the workshop had achieved its purpose, and I was most satisfied that the delegates had agreed – of their own volition – to re-convene the community at a later date to review progress.

Reflection

It is at this point that I now bring this paper to a close, not because the work I describe herein is in any way complete; rather that further progress is, at the time of writing, a future endeavour. There remains much to do to build on the momentum gained through the workshop, to identify solutions to further embed cyber defence into the acquisition process, pursue these solutions through to implementation and ultimately realise the benefits for Defence.

Upon reflection I remark that this study was – and still is – somewhat unique, in that Dstl was commissioned to review and influence a process over which we have no formal authority. My observation that there was no one 'customer' within MOD to champion the work and take any recommendations forward was pivotal, and fundamentally defined the scope of the workshop.

To conclude this paper (at least for now) I present a summary of the insights I have made herein that pertain to my first foray into the application of soft systems methods:

a. How one 'orientates' oneself around or within a project from the outset can prejudice how one treats the project throughout. (I was asked to conduct a review of IA process as a starting point for embedding cyber defence in acquisition, and the final workshop kept a strong focus on IA even though many of the issues fell outside the traditional remit of IA). Such convergent thinking should be balanced with convergent thinking through the judicious application of systems methods.

b. Soft methods are 'frameworks for learning' so it does not matter if the most appropriate method was not applied from the outset, or indeed 'correctly'; an appropriate method can still be employed to bring clarity to the problem space, even late in the project.

c. Whilst it is unlikely that any single method will serve to govern the project from start to finish, there is no harm in selecting a (seemingly) appropriate method to guide one's thinking, particularly through uncertain times. This should of course be moderated through peer review, and a method may be rightfully abandoned (in favour of another) if it is not helping anyone understand the situation.

d. As a technical consultant one does not require any prior knowledge of the problem space or domain to add value (simply asking the right questions can help stakeholders identify solutions to their problem). However, complex systems problems are best explored after one has been immersed in the detail, and it's of benefit to learn fast.

e. More is learnt about a method each time it is applied, so perceived unfamiliarity with a method should not preclude one from applying it, and learning from the experience.

f. Adoption of an appreciative inquiry approach (AIM or otherwise) is useful when consulting with stakeholders for the first time, and lends itself well to problems that may be readily expressed as a system (and less so for those that do not).

g. The evidence offered by stakeholders in the form of examples or even anecdotes should be captured in addition to their assertions, as it may be necessary to add credence to one's insights and/or serve as 'existence theorems' when communicating with other stakeholders.

h. There exists a fundamental distinction between the functional description of a system and that of its constituent parts to the activity or issues associated with achieving the desired system state.

Both are valid approaches, depending on the lens through which the problem space is being tackled, yet combining the two unintentionally can confuse matters.

i. What might be perceived as the most insurmountably complex problem at first can probably be made some sense of if one spends enough time exploring it. It may be difficult to appreciate just how much of the problem space one has learned from stakeholders whom, independently, have limited visibility of the complete picture. Communicating one's understanding of the problem space to stakeholders can aid everyone's understanding.

j. When expressing the problem space, or an aspect thereof, initial progress is likely to be greater with pen and paper (as opposed to computer-aided methods). It is a more direct form of expression, and frees the mind to concentrate on the subject matter. When, and only when, the hand-drawn expression becomes unwieldy on paper should one consider transferring it to a digital medium.

k. It is likely that any expression of the problem space, whether a systems model or otherwise, will require several iterations, so it is not necessary to seek perfection from the outset. No articulation will be 'right', however subsequent iterations are likely to be more useful, especially if others have been involved in producing them.

l. Where problem structuring techniques fail to lend any structure to the problem, simply writing down what one knows about the problem, in prose, will at least capture one's knowledge about the problem space, and in so doing a structure may emerge.

m. By involving others in the learning process one can explore divergent options as well as establish direction at times of uncertainty. If stakeholders are not readily accessible one's peers may offer unique insights.

n. Answering the 'six questions' of systems thinking practice [8] can help reveal potential obstacles to success, in particular the absence of customers, actors or owners. It's imperative that one identifies one or more champions for the work within the organisation under review, otherwise it is unlikely that any recommendations for change will be taken forward.

o. Decision makers are rarely concerned with the detail or complexities underlying a problem, although they will likely seek assurance that it's been explored and analysed. Communicating the essence of a complex problem in a manner conducive to senior decision makers is itself a non-trivial exercise and may require considerable resource investment.

p. Reflecting back on the journey one took to understand the problem can offer insights into the learning process and perhaps improve future success. Insights into what worked and what didn't may be of value to other systems practitioners, perhaps in the form of a symposium paper.

q. It can be easy to underestimate the time and effort required to write a symposium paper.

References

[1] **Checkland**, Peter B. (1981) *Systems Thinking, Systems Practice*, John Wiley & Sons Ltd, ISBN 0-471-98606-2

[2] **Wilson**, Brian (1984) *Systems: Concepts, Methodologies and Applications*, John Wiley & Sons Ltd., ISBN 0-471-92716-3

[3] **Checkland**, Peter B. and **Scholes**, J. (1990) *Soft Systems Methodology in Action*, John Wiley & Sons Ltd., ISBN 0-471-92768-6

[4] **Stowell**. F. A. and **West**. D. (1991) *The Appreciative Inquiry Method, A Systems Based Method of Knowledge Elicitation*, Systems Thinking in Europe, Plenum, New York. p. 493-497

[5] **Vickers**, G. (1968) *Science and the Appreciative System*, in Human Relations 21: 99-11

[6] **Stowell**, F. A. (2012) *The Appreciative Inquiry Method – A Suitable Framework for Action Research?*, Systems Research and Behavioural Science, Volume 30, Issue 1, p. 15-30

[7] **Haddon-Cave** QC, Charles (2009) *An independent review into the broader issues surrounding the loss of the RAF Nimrod MR2 aircraft XV230 in Afghanistan in 2006*, London: The Stationary Office, ISBN 9780102962659

[8] **Dodd**, L. and **Hilton** J. (2013) *Systems Thinking Through Practice* (Training Course Material), College of Management and Technology, Cranfield University / Defence Academy of the United Kingdom

[9] **Facione**, P. (2011) Think Critically, Pearson Education, Englewood Cliffs, New Jersey

[10] **Helmer-Hirschberg**, Olaf. (1967) *Analysis of the Future: The Delphi Method*, RAND Corporation, Santa Monica, California, P-3558

[11] **Jacka**, Mike and **Keller**, Paulette (2009) *Business Process Mapping: Improving Customer Satisfaction*, John Wiley and Sons, p. 257, ISBN 0-470-44458-4

[12] **Smith**, Michael (2005) Role and Responsibility Charting (RACI), Project Management Forum, p. 5