19th ICCRTS: C2 Agility: Lessons Learned from Research and Operations

Paper 081: Using causal models to manage the cyber threat to C2 agility: working with the benefit of hindsight

Topics

Primary: Topic 6: Cyberspace, Communications, and Information Network
Alternates: Topic 3: Data, Information, and Knowledge
Topic 4: Experimentation, Metrics, and Analysis

Authors

Mr A Barnett
Principal Technical Authority
Dstl
Porton Down Science Park
Salisbury SP4 0JQ, UK

Dr S R Smith
Chief Technology Officer
MooD International Ltd
University Science Park
York YO10 5ZF, UK

Dr R P Whittington
Chief Strategy Officer
MooD International Ltd
University Science Park
York YO10 5ZF, UK

Point of Contact

Dr Richard Whittington, MooD International Limited, University Science Park York YO10 5ZF, UK
dick.whittington@moodinternational.com
+44 7909 711 987

Abstract

The effectiveness and agility of an operation can be compromised through impact of adversary cyber activity on performance of key assets. The agility of adversary activity demands focused and effective investment in cyber security: there is a need for strong awareness of the military impact of threat vectors within the mission landscape, such that cyber security can be managed through a balanced, measurable, mission-driven process. With insight into the landscape and potential impact on mission outcomes, decision makers in the capability acquisition or mission control space can make informed decisions and interventions; and hence can plan and act proactively and with agility, rather than reacting to events.

MooD International has been working with Dstl to develop a *causal modelling* approach to support the mission commander in taking appropriate action in the face of overwhelming amounts of data that may be associated with cyber events. The approach projects likely consequences of cyber events by exploiting a cyber kill-chain concept, embedded in a model that aligns capabilities and components of the mission landscape.

The approach has been applied as the engine for an interactive cyber-security situational awareness and decision environment for the mission commander, supporting impact analysis and exploration of alternative interventions.

Introduction

UK Ministry of Defence (MOD) operates by being proactive and agile; in cyberspace this requires understanding of the impact of cyber events on the MOD Enterprise in terms of its ability to deliver core mission objectives. To achieve this, defence enterprises need strong awareness of the military impact of the threat landscape so that cyber security and associated risks can be managed through a balanced, mission-driven process. With insight into the landscape and its potential impact on mission success, decision makers in capability acquisition or mission control can make more informed decisions and investments, and so can plan and act decisively rather than reacting to events.

Methods and systems to support Cyber Situational Awareness within Defence enterprises have to date focused largely on data and events most directly connected with physical and electronic computer assets (networks, physical IT infrastructure, end devices, operating systems, applications), and on providing functionality related directly to managing these assets (see, for example, [1, 2, 3]).

Whilst such approaches play an essential part in tackling cyber security, there are three significant challenges that are not adequately addressed, in particular with regards to the needs of the military decision maker:

1) The military decision maker is not primarily concerned with cyber assets, but with mission impact, understood in terms of restrictions in the use of capabilities that might depend on these assets – and so information about risk to cyber assets is in itself little help without the military context;
2) The military decision maker needs time to be able to assess alternative courses of action, anticipating a problem before it becomes critical; and
3) As they are acting on the basis of information that is partly to do with assumptions about the future, and about 'layers' of cyberspace with which they may have little direct experience or intuition, the military decision maker needs to understand how much confidence to place in any information used as the basis for decisions.

The consequences of not addressing these problems include:

- A misreading of the level of risk being carried at any point in time, and so decisions being made that are either too cautious, leading to loss of mission objective, or that are too bold, leading to excessive exposure of missions to cyber-instigated disruption and loss;
- A lack of insight into the future impact of current and potential future events in cyberspace, and so a highly curtailed amount of time for alternative options to be considered concerning the overall mission; and
- Poorly focused investment in defensive assets, leading to either spend on cyber security in non-critical areas, or excessive exposure in areas that are critical to mission success

The situational awareness concept addressed by this paper responds to the challenge of enabling proactive decision making around cyber activity (whether adversary or mission-initiated), with particular regard for future impact on mission assets and objectives. Specifically, this relates to a number of key questions for the mission commander, which have set the scope and aims of the work reported here:

1. What is the status, now, of the capabilities I need in my mission?
2. What is the likely status, into the future, of these capabilities, given the likelihood of future activities that I know about, and the likelihood of events that I might have less certainty or control over?
3. What are the implications on mission status, but also financial cost, resourcing, etc. if these happened in a different way – what if the adversary took a less likely course of action? What if a less likely event did occur? What are the implications if I need to undertake additional activity in order to mitigate?
4. And how much confidence should I place in the guidance this system is providing?

Although these questions are not unique to cyber security, cyber activity does generate specific considerations for a mission that makes support for these questions more critical. In particular, cyber-related events and assets:

- cut across the physical arrangements of command, personnel and systems,
- may do so in non-obvious ways in situations where previous patterns of activity are not necessarily the complete guide to future (or even past) activity,
- and in such a way that the tempo can be rapid, especially when considered in terms of human in the loop decision making.

It is important that physical and virtual activity is seen not in the context of two different domains, but in the context of a single environment in which physical and cyber activities each impact upon elements of the other. This perspective enables the mission commander to address questions purely from a perspective of impact upon the mission.

This paper reports on work carried out jointly by UK Defence Science and Technology Laboratory (Dstl) and MooD International, around developing and exploiting an approach to apply causal relationships across the cyber landscape. The approach seeks to avoid the misreading of risk levels against acquisition or mission objectives, to enable proactive testing of potential interventions by decision makers ahead of critical events, and to focus cyber security investment in areas that are critical to mission success.

After describing background influences and methodology, this paper walks through an application of the approach to a surveillance scenario, contrasts the approach with related methods and technologies, and summarises the learning from the research to date, together with further implications and ambitions.

Background

Dstl's purpose is to maximise the impact of science and technology on UK defence and security, serving all areas of the Department from operational commands to the Defence Reform Unit, Defence Equipment & Support to the Front Line Commands or Land, Sea and Air. Dstl is also responsible for championing and developing science and technology skills across MOD and UK industry, promoting innovative and novel capabilities that may benefit UK and international partners. In the context of cyber situational awareness, Dstl's research approach takes a Business Enterprise view of cyberspace to understand the impact of cyber events on real world operations and so enable the application of timely and appropriate response.

MooD International is a UK-based, research-led software enterprise which specialises in the development of software to aid decision support and governance within large, complex organisations, within government and industry. MooD's achievements in successfully delivering generations of innovative, R&D focused technology projects that extend the boundaries of technical feasibility and knowledge have been recognised by successive Queen's Awards for Innovation.

The MooD framework for Decision Insight

The approach has adopted MooD's framework for decision insight, which layers three increasingly sophisticated perspectives (see Figure 1):

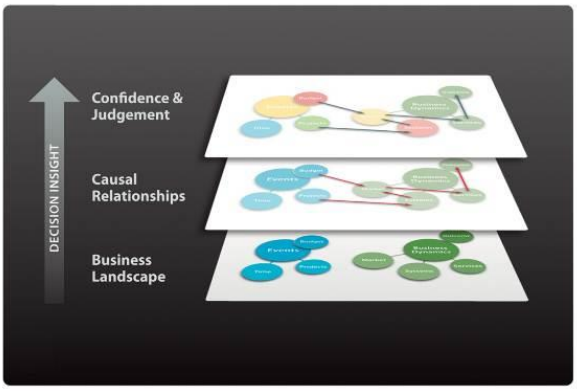

Figure 1. Levels of Insight to Support Decisions

- At the root is the Business Landscape – a connected architectural model of relevant components, moving parts and connections, populated in part through automated feeds from operational data.
- The landscape is overlaid with Causal Relationships that can be executed against live landscape content, and used to project future events, whether under the control of the enterprise or not, and their potential impact on the enterprise.

- Through a further overlay, Causal Relationships are associated with levels of confidence by means of assigned probabilities of status or events, to support judgement, including through Bayesian analysis of likelihood.

Under the approach adopted, populating the levels of the framework is not solely reliant on access to trustworthy data. The models are configured and augmented by a combination of live operational data plus judgment from domain experts.

Business Landscape

An architecture model is used to represent the relevant mission capabilities, and to define relevant dependencies across components at different levels of the cyber domain (see below). Content for the business landscape is derived from a range of external systems, structured by business analyst using a configurable, domain-specific meta-model:

- Connections across the levels of the landscape draw in part on the constructs proposed by Scott Borg in [4, 5].
- Modeling of the concepts that span cyber and kinetic activity has included an adaptation of the MITRE Corporation view of cyber activity [6, 7], elaborated to accommodate the blending of cyber with non-cyber activity, together with a mechanism for calculating and revising probabilities based on observation of events.

Causal Relationships

Causal models and associated causal reasoning have been proposed (see, for example [8, 9, 10]) as powerful methods for understanding and working critical connections between factors in a complex landscape. In this framework, causal relationships build on landscape dependencies to provide additional connection properties across components, supporting the determination of overall impact to mission capabilities across levels. Rules can be based on Subject Matter Expert (SME) judgement, validated or augmented where possible from pattern analysis of historical data. As inputs, such rules take data about current performance, status and events, which might originate from a range of sources, including non-obvious, non-structured sources.

The work has identified three principal kinds of causal relationship that can be applied together – as described in the following section - including an adaptation of the concept of 'cyber kill chain' recently proposed by Lockheed Martin [11].

Confidence & Judgement

Bayesian approaches to probability have been applied in machine learning based technologies, most prominently in search, by for example HP Autonomy [12], or in evidence-based learning with IBM's Watson [13], and applied in the cyber domain to support better detection through learning-based anomaly detection by, for example, SRI International [14].

Levels of confidence across the causal model are calculated as a function of a number of parameters, including the degree of cyber resilience of causal connections, and confidence levels in the occurrence of projected events given e.g. alternative 'kill chain' scenarios that might connect cyber events through time.

Applying Decision Insight to the Dstl view of Cyber Situation Awareness

Dstl recognise a multi-level perspective on the cyber domain, as shown in Figure 2.

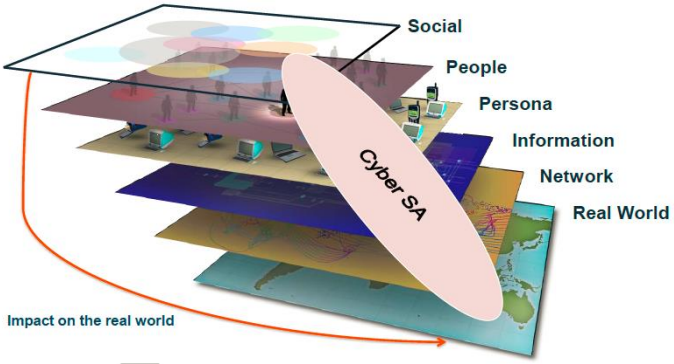

Figure 2. Levels of Cyber Domain

Within the landscape model adopted here, the connectivities across and between elements enable the impact of events to be analysed across the levels of this domain (i.e. from Real-World / Physical up to the Social level).
Going beyond the structural potential of this multi-level connectivity, the approach applies causal rules to current information about performance and events, to project the likely future impact across the levels, and hence show the business impact on assets and capabilities of seemingly unrelated or low level events, as indicated in Figure 3.

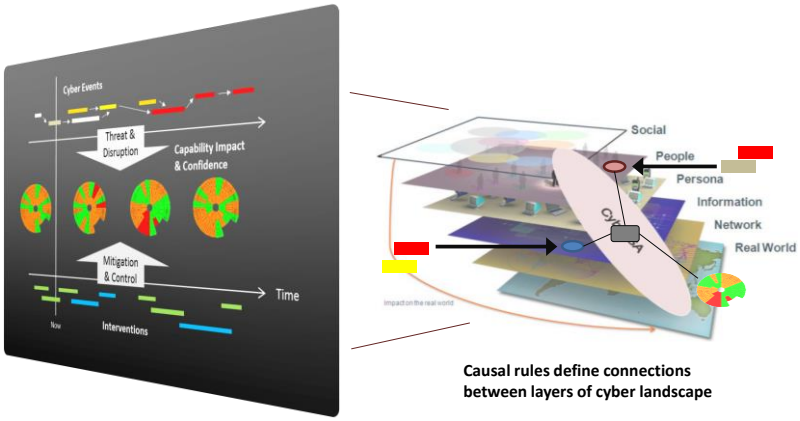

Figure 3. Projection of Impact over Time

Learning & Innovation

Although each of the constituent techniques applied in this framework may not in itself be novel, we have defined an approach that synthesises a number of proven techniques in a sufficiently tractable way so as to be workable within a software environment to support mission-level situational awareness. This topic is addressed in more detail in a later section.

The project reported here has implemented a demonstration system that takes data concerning cyber events and applies these to a causally connected model of operations to calculate levels of confidence given alternative kill chain scenarios.

The solution has been packaged as a highly visual web application with a user interface suitable for direct use by a military decision maker, providing a projection of future freedom of action that is based on levels of confidence in critical capabilities, and providing the ability to explore the effect on this freedom of action of alternative scenarios of enemy and own forces interventions.

The system:

- implements a lightweight causal model representing the dependencies between the 'layers' of cyberspace, and in particular the mission outcomes and required capabilities
- exploits data feeds concerning cyber interventions alongside own forces' interventions that influence causal rules in order to give calculated levels of confidence in future mission impact
- allows options for alternative courses of military action to be explored alongside alternative potential courses of cyber activity action, again exploiting the causal rules mechanism

It has been developed using MooD software and demonstrated using a scenario that locates cyber activity within the context of real world operations, in this case within the scenario of clearing routes from Improvised Explosive Devices (IEDs). The particular focus has been on how such an approach can provide proactive and mission-focused support for cyber situational awareness, from the perspective of a mission commander.

Concept & Methodology

The core concept reported here is the use of cause / effect relationships between capabilities, assets and events, in conjunction with algorithms for calculating levels of confidence, to provide enhanced, mission-focused cyber situational awareness.

In particular:

- The incorporation of a **cause / effect model** to capture dependencies between capabilities, assets and events that is expressive enough to provide future projection concerning the status of seemingly unrelated concepts, including abstract and non-cyber concepts of interest to a mission commander (hence addressing not only the need for anticipation, but also the cyber '**cut across**' issue).
- A **probabilistic approach** to these cause / effect dependencies that allows guidance to be given on the most likely state of affairs in the future. This guidance is revised given the observation of actual events and asset status, but is based primarily on domain-expert sourced cause / effect dependencies, from current situational awareness, rather than being predicated on past behaviours or patterns in data (hence addressing the uncertainty / **non–obvious** causal dependency issue).
- The implementation of this kind of model in a system that **takes operational data**, applies the model to this data, provides hooks for further automated analysis support, and delivers the whole as a usable, visual, mission-relevant environment that can be used to assess the impact of alternative courses of action (hence addressing the issue of the **rapid tempo** of decision making).

Principal Components of a Business Landscape for Cyber Situation Awareness

At the centre of the methodology is a concept of three fundamental types of component within the Business Landscape: Capabilities, Assets and Events. It is the connectivities within and between these that underpins the design of the decision environment.

Capabilities

Capabilities are abstract, derived, mission-relevant concepts that represent one of the ways in which the mission commander reasons about mission outcomes. The mission commander requires an up-to-date view on the current and historical status of mission capabilities. These capabilities are not necessarily directly measurable in a physical sense.

Capabilities typically represent the required effects for a particular kind of mission e.g. surveillance or logistics, and can be broken down into lower level capabilities e.g. imaging or storage, and quantified in terms of measures of effect or effectiveness, e.g. image resolution or temperature. The actual, observed measure of effect for a capability can be compared against a target or required measure of effect to produce a status for a capability. Status can be visualised using, as illustrated in Figure 4, a 'target graph' visualisation, indicating the status at different time periods (e.g. 'Now'; 'Now-1' or a time in the past'; 'Now-2' or a time further in the past).

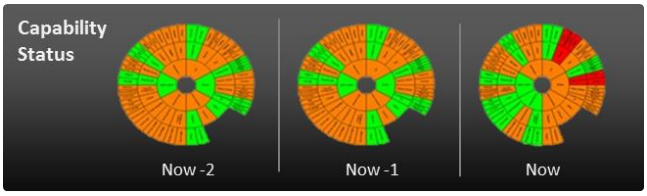

Figure 4. Target Visualisation of capability Status

In the scenario adopted within this paper, the capabilities relate to a route clearing operation, deployed in the context of a particular mission, and comprise enabling capabilities (e.g. Mission Planning, Wide-area Surveillance) as well as the specific required capability (e.g. Counter IED).

In principle, the actual performance of every capability can be independently assessed. In operational analysis, this is typically by direct observation, e.g. a post mission analysis that compares what was required with what was achieved, so determining the status of the capability at a point in time.

Assets

Where the actual performance of a capability is not directly observed (i.e. its measures of effect cannot be directly observed, or where it requires an aggregate of observations), it can be assessed in terms of the aggregation of performance data from a collection of underlying enabling assets that are expected to provide that capability, as illustrated in Figure 5.

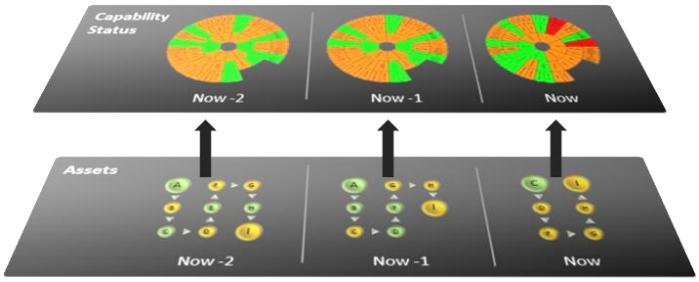

Figure 5. Capability Measures of Effect calculated from underlying Asset Measures of Performance

The enabling asset Measures of Performance are themselves either directly assessed (e.g. by human observation or automated measurement) or aggregated from actual performance of related assets. In the scenario adopted, the assets implement some or all of the required capability – e.g. a 'Gorgon Stare' asset that provides 'Wide-Area Surveillance'.

Within the landscape definition, relevant assets can be located across the levels of cyberspace. The assets and associated vulnerabilities (assets or properties that are of particular interest with regard to kill chain activity) in the scenario adopted include traditional physical assets, such as equipment, together with human assets, network components and qualitative information, as indicated in Figure 6.

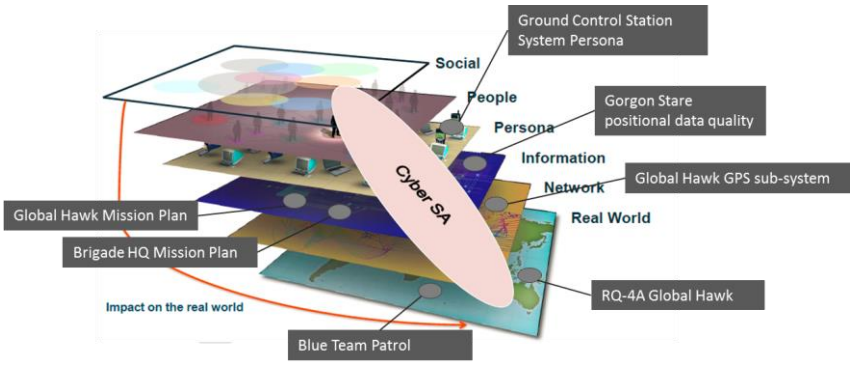

Figure 6. Assets Aligned with Cyberspace Layers

The relationship of assets to cyber vulnerability applied here is based partly on the view of cyber concepts provided by MITRE Corporation in [7]. This involves the association of vulnerabilities with assets, with these vulnerabilities then being susceptible to attack events.

Events

In addition to having a view on the actual performance of assets and capabilities in historical periods up to 'Now', the Mission Commander also requires the associated context of the actual activities and events that have occurred over the same time periods, as indicated in Figure 7.

Events have been observed and measured with a degree of confidence in actual performance, and may be red team events, such as a Denial of Service attack, or blue team events such as Mission Re-planning.

There is not necessarily a relationship between the performance of assets or capabilities and the events that have occurred, although with regard to 'calibration', it is possible to identify correlations between these.

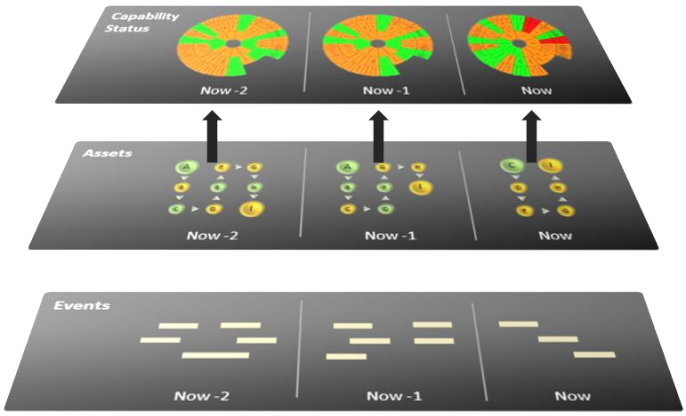

Figure 7. Events through Time

Exploiting these Components to Project the Future and Inform Interventions

The Mission Commander requires a forward view of the status of each mission-relevant capability. By definition, this is not directly observable. Techniques for predicting the future state of data are typically based on statistical techniques applied to historical data, using correlations or patterns in that data to predict future effect (i.e. extrapolation from data mining functions, or functions from 'predictive analytics').

However, we do not want to predicate future status solely on past behaviours, for three reasons:

1. The performance of an asset may depend on the context of operation, which may be different from a previously encountered context of operation. The status of a capability, calculated from the performance of its assets, may be acceptable, but the effectiveness of that capability may also be affected by the performance of other capabilities being deployed in the same mission. For example, the ability of a surveillance system such as Gorgon Stare to support the mission may depend on its capability to deliver imagery of the target area. The quality of that imagery will be dependent on the atmospheric conditions and the capability of the platform to know where it is – as imagery of a location is only useful if we know the location.
2. Seemingly unrelated concepts may be affected by composite causal chains, where these connections are not apparent in historical data. For example, a physical intrusion onto a secure site could be a precursor activity to a cyber-attack on a system inside that building that supports a UAV being controlled on another continent.
3. The probability of events occurring in the future may not be determined by their occurrence in the past. For example, a pattern may be detected from an adversary's past behaviour, but this may not hold in the future. Indeed, observing those events from the adversary's perspective might reveal that they have changed tactics or processes.

The methodology reported here answers the question of 'what drives future effect' by including three kinds of causal relationships, and so three kinds of domain knowledge:

- How one event might lead to another (adapting the concept of a cyber 'kill chain')
- How an event might affect the performance of an asset
- How a change in performance of one asset or capability might drive the change in performance of another

<u>Cyber Event Chains – from Event to Event</u>

Events are not unrelated, particularly when related to the motivation or mission of a particular actor. To be able to link events together in a 'chain' is helpful for two reasons:

1) it lets us have more confidence in the occurrence of future events, and
2) it gives a way of recognising or anticipating in advance things that may have an undesirable effect.

This method builds on Lockheed Martin's adoption of the concept of a cyber kill chain [11], characterised as a sequence of phases through Reconnaissance, Weaponization, Delivery, Exploitation, Installation, and then C2 and Actions. In the scenario adopted, a number of short, related event chains combine cyber and kinetic events, to reflect the cross-cutting nature of cyber warfare.

For example, Figure 8 shows an event chain comprising event types that begin with access to the Ground Control Station network ('Reconnaissance'), carry out persona management activities leading to the compromising of Mission Planning and operational disruption ('Actions on Objectives').

Figure 8. Event Chain - Hostile Influence on Mission Planning

Each event in the chain has a probability that is conditional on the probabilities of other events in the chain. If an observed event is identified to be an instance of one of these event types, then this may affect the probability of instances of related event types occurring. Figure 9 shows a chain where 'GCS Hostile Persona Management' has been detected, leading to a heightened assessment of the probability of subsequent events in the chain, including Operational Disruption, and also a heighted probability of access to systems and networks.

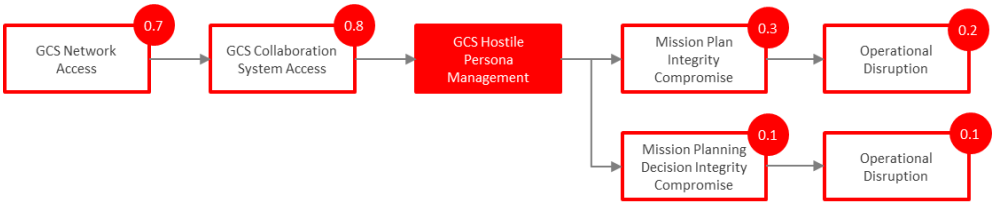

Figure 9. Event Chain – Including Probability Overlay following Detection of Events

This methodology adds three main features to the cyber kill chain concept:

- Levels of confidence in events – this is forward looking, nothing is certain, and our understanding of where we are in an event chain, or which event chain we might be in, is constantly revised on the basis of observed events
- Combinations of red and blue team events
- Flexibility in what these chains are, including the ability for events to be modified dynamically based on domain knowledge or revised courses of action

The method adopted can also be compared structurally with that of Jakobson [15], which also takes a strongly mission-centred view, although it goes further in terms both of the above factors, and also the application of causal reasoning beyond dependency analysis.

As well as providing levels of confidence to a mission commander, these features also increase the expressiveness of the kill chain: complex process chains, including branching of event chains and concurrent events, can be represented through a dynamic treatment of probabilities. Observation of an event makes future events more or less likely. So, as events are observed, the likelihood of other events may approach 1 (if X has happened, then so has Y) or 0 (if X has happened, then Y cannot have happened). For example, it is unlikely that a physical network asset will be compromised until a port has been opened on a related network. Once that port is open, we might conclude that the probability of a successful attack on a vulnerability is now more likely. Equally, if we detect an intrusion attempt and successfully stop it, we might conclude that a successful attack is now less likely.

The overall effect for the mission commander is that as events occur, so the landscape of potential future courses of action change, in potentially non-obvious or not previously observed patterns.

Cause / Effect Rules from Event to Asset

Events are not interesting unless we can infer or reason about the effect that they have. This principle is reflected in causal rules that describe the effect that a kind of event can have on a vulnerability, which in turn affects the level of performance of an asset. This approach is comparable to the MITRE view of assets and vulnerabilities, in which an event can 'activate' a vulnerability of an asset, so reducing its measure of performance.

To date, the development of the approach reported here has not been concerned with the complementary problems of event detection and classification; it assumes a feed of event data that can be matched with kill chain nodes, providing an observed event that is used to project a corresponding effect on an asset. For example, we assume that we have a data feed that provides us with the status of an asset and with events that might impact it. So for a UAV, we assume we can measure and be provided with the appropriate measures of performance and that we have access to events, such as attacks on the ground control segment or attempts to shoot down the UAV that might impact its ability to perform the mission. In the mission planning example, each of these events has an effect on the measures of performance of several assets.

Cause / Effect Rules from Asset to Asset / Capability

Once we know the probability of a future event, and the effect that the event could have on assets, the final piece of domain knowledge required is the impact of that change. This requires a different understanding of cause and effect. By way of analogy, consider a situation where the fuel usage of a number of vehicles is regularly measured. The aggregate fuel usage of the fleet of vehicles at a point in time can be calculated from the actual fuel needs of individual vehicles, and is unlikely to be directly observable independently of this calculation. This kind of aggregation function is part of the regular Management Information ('MI') needed during operations. A cause/effect relationship between assets is similarly a function, but with a different interpretation and usage. The value of one property is used to generate the value of another, but as a 'hypothesis' that can be verified independently through observation. Continuing the analogy, a fleet planning system is distinct from a regular MI system in being able to, for example, estimate at some point in the future the fuel need of a fleet of vehicles from the total payload, the total distance to be transported, and the type of terrain to be traversed.

This kind of causal rule is important for two reasons:

a) It can provide the ability to generate values of observable measurements of performance into the future, that can then be independently verified through observation
b) It can be used to capture the mission context of asset performance

The concept can be illustrated through cause / effect dependencies between capabilities. For example, assuming Wide-Area Surveillance is important to support the activities of the Patrol, we can generate a value for the effectiveness of the Patrol in the future, based in part on the future effectiveness of Wide-Area Surveillance. In particular, the effectiveness of the Patrol will degrade as the effectiveness of Wide-Area Surveillance degrades. In principle, these dependencies can also be described at the level of assets.

Inter-Working of Concepts

The three kinds of causal rule, implemented over a Landscape and ascribed with probabilities, provides a way of generating a forward-looking view of mission capability, as impacted by cyber events, that is not predicated on past behaviours. Figure 10 summarises the three kinds of causal rule and their inter-play, from event to event relationships at the bottom, through the connection of events into the effect that they have on assets in the middle, and through causal dependencies between capabilities at the top.

The interplay between these three pieces can be emphasised by the implications of not having them available:

- If we didn't link assets and capabilities there would be no basis for understanding the overall mission context of an event – e.g. the event may not matter, depending on the environment of operations
- If we could not link event occurrence to asset performance there would be no basis for reasoning about the impact of events.
- If we could not link events to events then there would be a much weaker basis for expecting or predicting future events.

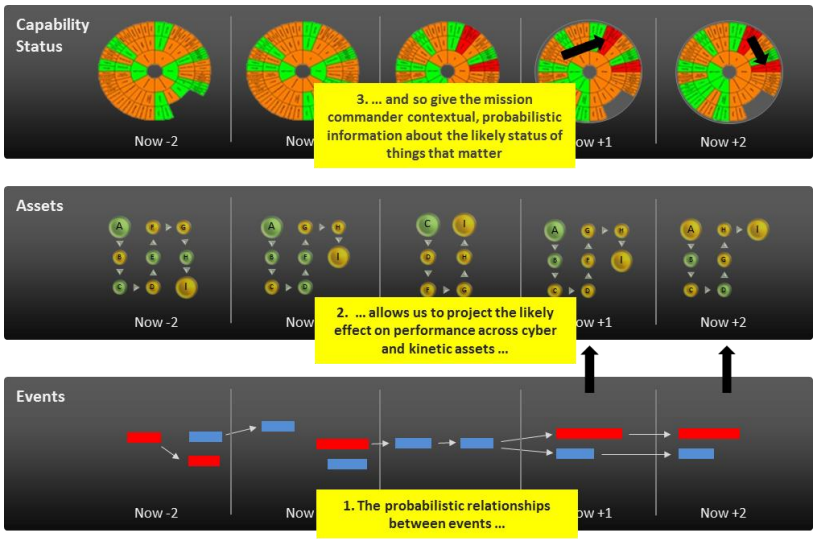

Figure 10. The Interplay between the Three Kinds of Causal Rule

Application to a Cyber Situational Awareness Scenario

Application Overview

On the basis of the methodology described previously, the broad structure of a configured MooD application to implement the concept as an interactive software environment is as shown in Figure 11.

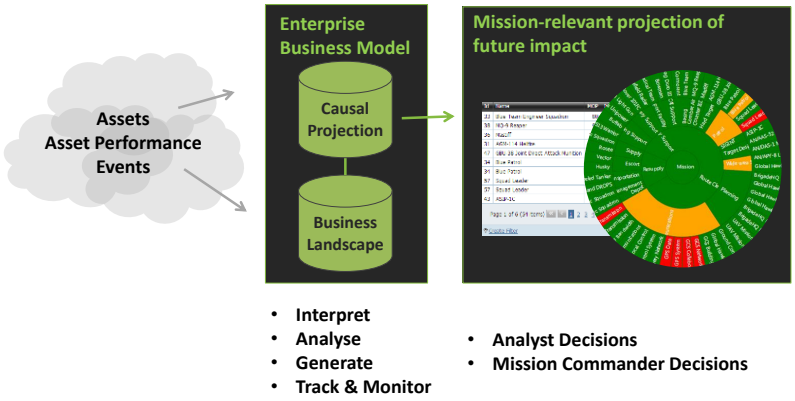

Figure 11. Summary Schematic of the Application

In essence, events and performance data are sourced externally and these are interpreted and analysed using causal rules expressed against the components of the landscape.

A user interface engages the Mission Commander with target representations of projected future impact against mission-relevant capabilities. On the basis of such projections, the Commander explores and tests feasible interventions to enhance the likelihood of mission success.

Multiple Perspectives

In the scenario assumed, the primary role is the Mission Commander, exercising concern about risk to their mission and stakeholders – loss of mission objective, unintended damage or loss, loss of future capability that might impact other missions.

The Mission Commander has a particular perspective on mission, capabilities, and intended outcomes. However, the adversary also has a mission, capabilities and intended outcomes. These two different perspectives join on assets and events. A feature of the application (allowing events to be associated with actors) enables the mission commander to switch their view to see this 'game' from their adversary's perspective – their adversary's mission and capabilities, and how the events might appear to them. Within the limits of what information is available concerning the adversary's mission, the Mission Commander can also in principle explore scenarios as they might appear to the adversary, adjusting and incorporating events that impact the Mission Commander's mission, but progress the adversary's mission. In principle, the Mission Commander can also adjust causal rules in this model, to explore how that world would look if the adversary was running to a different understanding of how the world works i.e. different sets of causal rules.

Initial Concept of Operation & Sample Screenshots

The scenario assumes an autonomous Intelligence, Surveillance and Reconnaissance (ISR) capability deployed within a larger mission. In this concept of operation, (blue) mission events owned by the Mission Commander are distinguished from (red) adversary events, and relevant event chains are shown with their 'precedence' relationships, as indicated in Figure 12. Note that the detail of the events and their impact on assets and capabilities (shown on the target panel) are not important to an understanding the

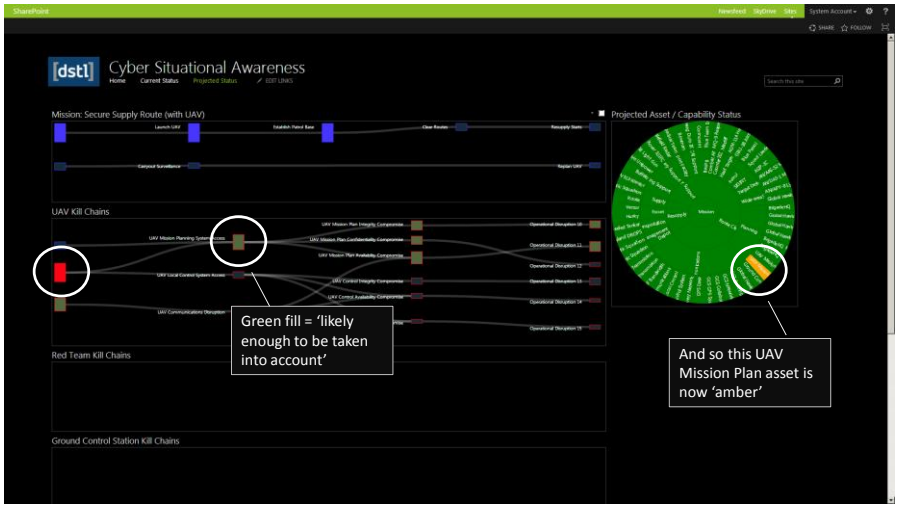

Figure 12. Sample Screenshot (1) from a Mission-Focused Cyber Situational Awareness System

approach. This figure shows a simple campaign plan (the sequence of blue nodes towards the top left); the node sizes indicate the assessed likelihood at the current time of an event occurring.

The lower event chain shows a cyber-event chain (or "kill" chain) that has been identified by detection of an event that may indicate adversary activity. The consequences of detecting that event have been calculated from relevant causal rules – in this case these have determined that subsequent events in the chain become more likely, with a corresponding status change to a mission asset.

Suppose now (see Figure 13) that a further adversary event is detected, indicating the likelihood of a second active kill-chain, with additional implications for mission-relevant assets and capability.

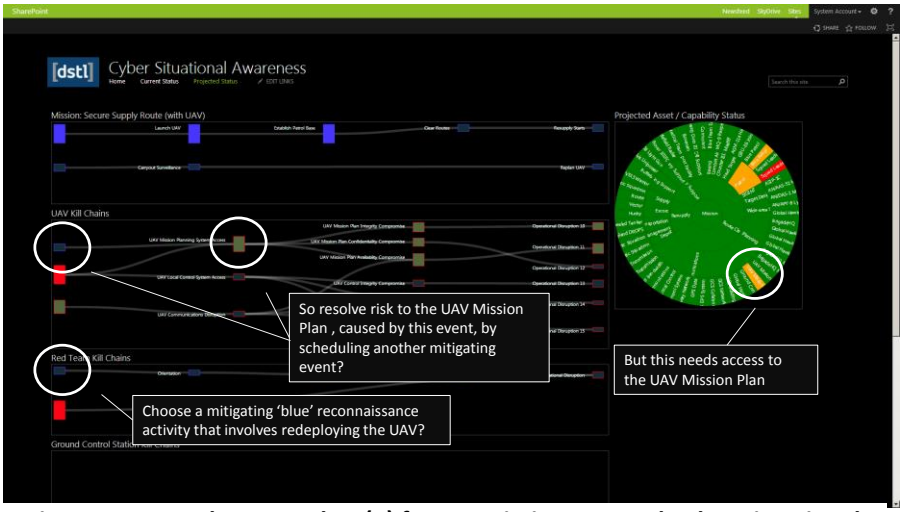

Figure 13. Sample Screenshot (2) from a Mission-Focused Cyber Situational Awareness System

In response to the projected impact on mission success, the Mission Commander identifies an intervention by means of a sequence of mitigating actions (i.e. a further blue event chain aligned with the predicted adversary kill chain) to reduce the likelihood of adversary disruption.

Figure 14 then indicates by an overlay the consequences of performing the identified intervention. Effectively, the system shows that the mitigating event chain would resolve the Mission Plan issue, but would temporarily reduce the effectiveness of UAV bandwidth, and consequently image quality, and so reduce surveillance capability.

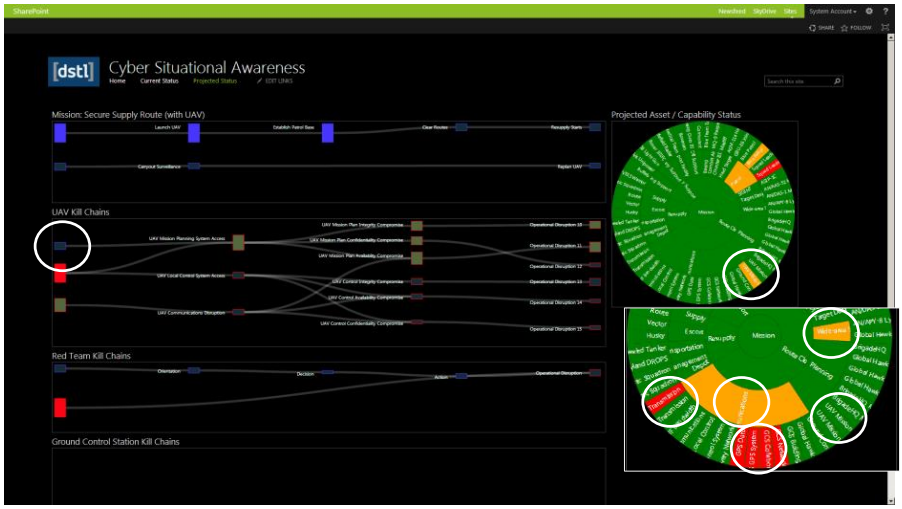

Figure 14. Sample Screenshot (3) from a Mission-Focused Cyber Situational Awareness System

This situation highlights the effectiveness of the approach in projecting future implications of detected events, of analysing these in terms of probable mission consequences, and hence supporting the Mission Commander in critical decision making.

The example continues (as shown in Figure 15) with detection of further adversary activities, indicating additional kill-chains and projected consequences on mission-relevant assets and capabilities.

In this case, having resolving the previously identified physical red team threat, evidence occurs of hostile persona management that calls into question the integrity of the mission itself. What now: search for a mitigation, assess potential impact, risk and outcome, make a decision, and observe the actual effects as they happen?

Building on the native 'optioneering' capabilities of the MooD software, the ability to drop into an exploration of how things might be different, becomes very natural. A causal rule, being generative, allows revised results to be generated given changes in any related data, and in principle based on changes to any of the rules themselves. Any changes are made in the context of a 'scenario', and remain available for comparison against the 'baseline', or against other scenarios.

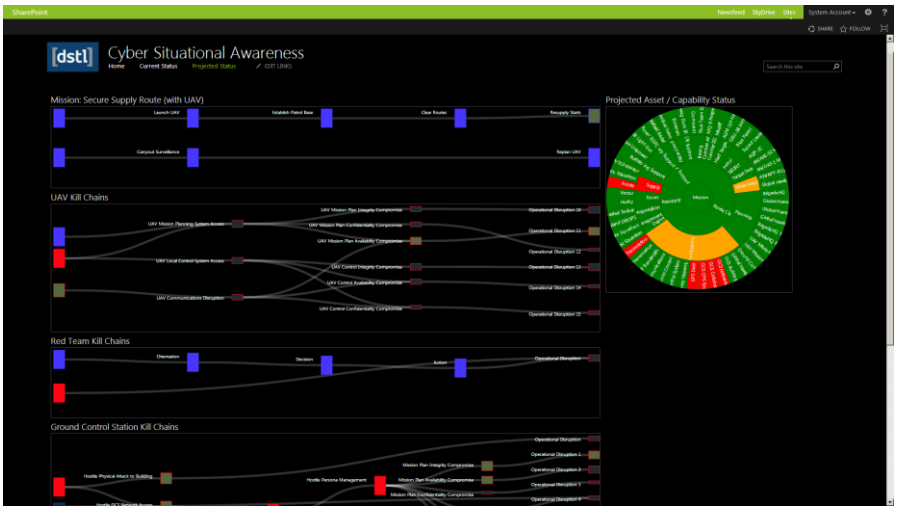

Figure 15. Sample Screenshot (4) from a Mission-Focused Cyber Situational Awareness System

The level of sophistication of this scenario is sufficient to illustrate how the approach extends the traditional envelope of situational awareness in the cyber domain:

- It adds the required intelligence about how the mission or environment is operating now, and could / should operate in the future.
- It exploits causal rules in conjunction with real data to create at least one projection of how the future might look, given the likelihood of certain events.
- The decision environment is live and always running as part of the operational environment – not just 'planning'.
- It provides the necessary support for exploration / reaching shared position on future courses of action, given the risk that might need to be carried in return for a particular outcome.

Extending the Concept of Operation

The approach also permits the modification of causal rules themselves as part of a scenario, allowing the exploration of the impact of cause and effect working in a different way to that expected. This means the ability to make modifications to kill chains, to the effect of events on assets, or capability to capability causal rules (i.e. "suppose that my world didn't work quite like that…").

The specific User Interface approach is a significant area for continuing development. For example, in one concept of operation, the Mission Commander could be presented with a collection of events of varying probability. The event chain with the highest probability is the one that is being used to drive the causal rules that ultimately determine the status of capabilities. If the Mission Commander selects an alternative event through a suitable User Interface, the causal rules can be re-executed for that scenario. Effectively, a central 'target map' showing the status of Capabilities would then be calculated based on the effect of the likeliest nodes in the kill chain.

Because of the indirect dependencies across cyber and kinetic assets, the non-obvious nature of the information being worked with, and the rapid tempo with which events occur, and with which understanding may change, confidence in the model and what it is telling us is paramount.

Being able to rely on the system for mission critical decisions under these conditions means that having confidence in, and transparency of the basis for, that decision making, becomes a primary, rather than a secondary activity for a Mission Commander. The reasons for the difference, and so failure in understanding, may be due to a number of causes, and understanding how little is known or can be relied on is a valuable contribution in its own right to the cyber situational awareness of the Mission Commander.

Complementary Methods and Technologies

The approach described here touches a number of technologies. In many cases these can be seen as complementary, either through addressing comparable needs in different situations, or through offering potential for enhancement or integration with this approach.

Increased automation of data analysis and visualisation is a necessary enabler for scalability of information exploitation across "big data" environments. And this requirement applies also to operational Cyber Situational Awareness, where the volumes and velocities of relevant cyber assets and events pose a real challenge.

However, the solutions that are being adopted in the wider data visualisation marketplace (see e.g. the research associated with the IEEE Visualization and Graphics Technical Community [16]) have been designed for a very different situation, where queries can be expressed by end-users against pre-tabulated data models, where outputs are computed by relatively simple pattern-matching, and results expressed using traditional graphic controls. They describe a viewpoint into a transactional "machine" and serve to inform about current status, potentially augmented with trending based on historical values.

By contrast, the approach described here cannot assume the existence of a predictable machine. Consequently, our concept and method involve creating an overlay across the data perspectives to provide a "business level" of decision support that is designed to enable decision makers to manage future interventions on the basis, not only of past trends, but also what is likely on the basis of inherent causalities in the domain. From this perspective, our model-based design provides a structure within which the decision-maker's 'hypothesis' can be expressed and evolved based on the effect of proactive or reactive interventions . In other words, the decision maker is able continually to explore the likely future effect of a proposed plan of activities and interventions, and hence to challenge the current understanding of assumed behaviours. To perform under operational conditions, this model needs to be integrated with lower level automation technologies.

In effect, the approach benefits from automation technologies at the 'data scientist' or analyst level, where the responsibility is to deliver the parameters of the causal model from source data. This aspect becomes more pressing as we anticipate greater volumes of data through an increasing number of cyber connected assets, especially as we progress towards the exploitation of the 'internet of everything'. This is an important area for continuing development where we hope to leverage an increasingly commercially available set of techniques (see e.g. recent work on event-based systems [17]).

Another critical requirement for the effectiveness of this approach is the identification and exploitation of causal rules that effectively describe the behaviour of the cyber domain. Several potential sources of rules have been identified, including:

- War-gaming / scenarios for event chains.
- Vulnerability analysis on equipment, operations.
- Statistical analysis of data (patterns).

What is apparent is that no one of these sources is necessarily reliable or correct, but they all have the potential to contribute to a situation picture presented for human analysis. And this is an important aspect of the approach described here: the significance of gathering and embodying domain knowledge from subject-matter experts, and applying the approach in a situation assuming human action and judgement. This contrasts with approaches (see, for example [18]) that seek much higher levels of automation across the complete decision cycle.

The use of statistical analysis techniques to compare patterns in actual data with the causal rules in the causal model can act as an additional calibration mechanism, in particular where these patterns can be extrapolated to suggest modifications that could be made to causal rules ([19] gives a useful overview of the contentious topic of causal inference), or vice versa, as indicated in Figure 16.

Rapid modification and feedback is an important part of the deployment of this approach, to enable testing of projections against actual observations, with the aim of constantly building and testing the situation

awareness picture, which feeds into the area of battle damage assessment. In this regard, aspects of the methodology suggested in [20] become relevant, given its focus on cross-domain integration.

Simulation technologies complement these techniques by providing additional information for the Mission Commander, constraining the decision space, and also identifying and validating probabilities around event chains being used in the causal model.

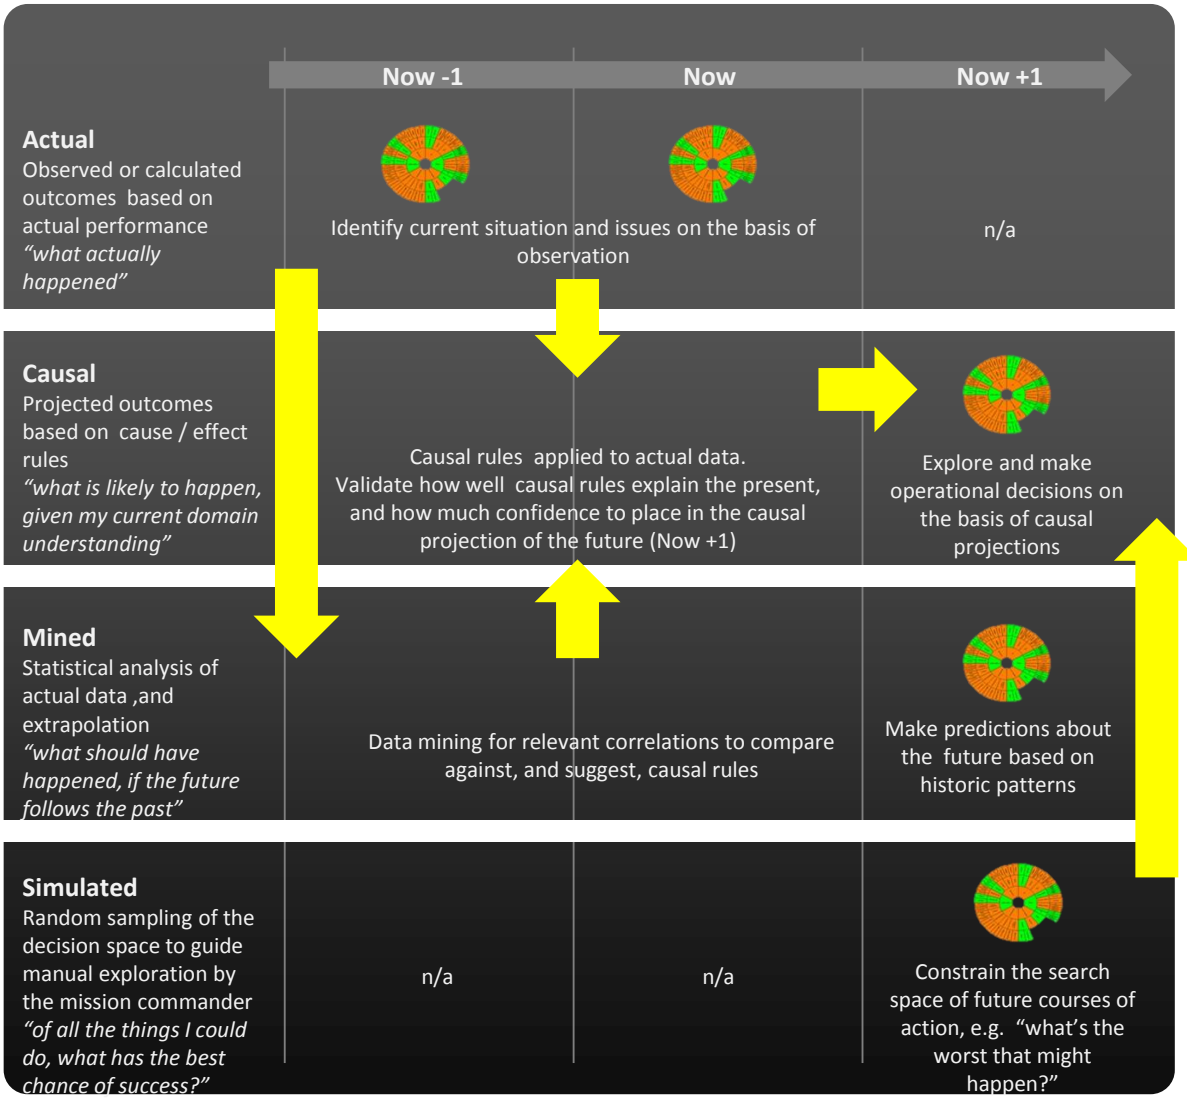

Figure 16. The Use of Data Mining and Simulation to Complement the Use of Causal Rules

Figure 16 also illustrates the complementary positioning of simulation technologies alongside the proposed approach, in providing contrasting alternative projections of future outcomes from potential interventions.

Through building on and extending these various technologies, this approach delivers a unique effect through a combination of:

- A causal model that contains dependencies across and between the key components events, assets and capabilities
- Synchronization of this model with live baseline operational data.
- A future projection method that also incorporates scenarios relating to potential interventions aimed either at mitigating risk or achieving advantage over adversarial action.

Results & Learning

This paper has described an application of causal modelling to a cyber-warfare scenario, locating cyber activity within the context of real world operations, in this case within the scenario of clearing routes from Improvised Explosive Devices (IEDs). The particular focus has been on how such an approach can provide proactive and mission-focused support for cyber situational awareness, from the perspective of a mission commander.

The key results highlighted by the work to date include:

- An autonomous ISR (route clearing) scenario – with supporting data – that exposes key types of trade-off and decisions for the mission commander. This scenario has further utility in cyber situational awareness research
- A taxonomy of causal relationships to enable modelling of the factors underlying cyber situational awareness, including an adaptation of the concept of kill chain
- A business model that reflects the essential concepts that span cyber and kinetic activity across the cyber seven layer model, and a mechanism for calculating and revising probabilities based on observation of events
- A concept of operation of a cyber situational awareness system that supports the scenario, and which gives an indication of how these concepts implemented together can provide an enhanced level of cyber situational awareness for the mission commander, including the provision of views from the perspective of other stakeholders, in particular the cyber adversary
- A positioning of causal modelling in the context of actual data reporting, data analytics and simulation, showing how a combination of actual data and causal data provides direct operational support for mission command, supported indirectly by analytics and simulation

The following observations / learning points have been demonstrated:

- A combination of three kinds of causal rule is required to account for mission-focused situational awareness in cyber security, covering the interplay between events and outcomes.

- Applying probabilities to these rules is essential to make this a usable risk focused system that enables management of uncertainty.

- The flexibility to represent events independently of actors allows for views from different perspectives – e.g. how do I think the world looks from the perspective of my adversary?

- The criticality of being able to contrast the model underlying the system with actual observations being gathered from other sources, in order to provide guidance on levels of confidence that should be placed in future projections and decisions made on that basis, and also to highlight gaps in actual observations and sources that might be required to increase that level of confidence.

Conclusions

In delivering the results and observations highlighted above, this paper has shown how a number of key questions relating to Cyber Situational Awareness can be addressed so as to achieve the particular aim of improving anticipation from a mission perspective. The approach successfully takes into account:

- The need to account for the combination of cyber and kinetic assets, and the direct and indirect (via cause and effect) dependencies that exist or need to be explored across these
- The non-obvious nature of cyber events and their effect on assets, and the need for decisions not to be based solely on analysis of past behaviours, from one perspective, or according to one way of understanding what drives behaviours and outcomes
- The need to match the rapid tempo of cyber warfare through the ability to respond to events and assess the impact of potential, alternative courses of action

We have demonstrated an approach to situational awareness embodied in a User Interface concept that provides a contextual view of events and capabilities, projecting directly onto the mission perspective – a view that brings to the fore the key information for the mission commander. This approach has been implemented using a non-trivial ISR scenario to explore a feasible concept for deployment.

There are a number of ways forward that are immediately apparent:

- Explore the required supporting tool capabilities and identify routes to automatically integrate the outputs, such as event detection and business process mapping, into the User Interface concept.
- Develop the approach further with real data to gauge the right level of granularity at which events and properties of assets need to be described (i.e. not too complex to be infeasible, but rich enough to provide realistic support) to inform the mission commander's decision needs in the required timeframe; aligned with this is the challenge to introduce increasing data-level automation.
- Scope the balance of support needed for decision space exploration by the mission commander, including the use of simulation methods to constrain the decision space at the right time so as to enable timely focus on key decisions
- Calibrate the approach using analytic methods, to validate the effectiveness of causal projections, to explore the potential for automated inference of causal rules, and to provide an alternative predictive methodology to challenge or reinforce causal projections.

It may also be sensible to now look at integration and complementarity with related challenges and approaches, as we are conscious that this work may be able to exploit advances in adjacent areas concerning e.g. detection, as well as potentially raising new challenges to be addressed in these areas. Ultimately, these new techniques are vital to enable MOD and other large scale, cyber-dependent enterprises to operate in a proactive and agile way, with effective decision making supported by mission relevant analytical evidence.

References

[1] A. Nicholson, T. Watson, P. Norris, A. Duffy and R. Isbell, "A Taxonomy of Technical Attribution Techniques for Cyber Attacks" Proceedings of the 11th European Conference on Information Warfare and Security, ESIEA, Laval, France, July 2012), pp188-197

[2] A. Barreto, P. Costa, E. Yano, "A Semantic Approach to Evaluate the Impact of Cyber Actions on the Physical Domain", Semantic Technology for Intelligence, Defense, and Security (Fairfax, VA, October 2012)
http://stids.c4i.gmu.edu/papers/STIDSPapers/STIDS2012_T08_BarretoEtAl_EvaluateImpactOfCyberActions.pdf

[3] A. Kim, B. Wampler, J. Goppert, I. Hwang, "Cyber Attack Vulnerabilities Analysis for Unmanned Aerial Vehicles", Infotech@Aerospace 2012, Garden Grove California, June 2012
http://arc.aiaa.org/doi/pdf/10.2514/6.2012-2438

[4] S. Borg, "How Cyber Attacks will be used in International Conflicts", 19th Usenix Security Symp (Washington DC, Aug 2010)

[5] S. Borg, "Economically Complex Cyberattacks", IEEE Security & Privacy (Nov/Dec 2005)
http://digitalstrategies.tuck.dartmouth.edu/cds-uploads/publications/pdf/Pub_EconCyberattacks.pdf

[6] S. Musman, Aaron Temin, Mike Tanner, Dick Fox, Brian Pridemore, "Evaluating the Impact of Cyber Attacks on Missions", MITRE Corporation 2010
http://www.mitre.org/sites/default/files/pdf/09_4577.pdf

[7] S. Barnum, "Standardizing Cyber Threat Intelligence Information with the Structured Threat Information eXpression (STIX™)", MITRE Corporation 2012;
http://msm.mitre.org/docs/STIX-Whitepaper.pdf

[8] N. Fenton and M Neil, "The use of Bayes and causal modelling in decision making, uncertainty and risk", Risk and Information Management Research Group (QMC, University of London, 2011)

[9] D. Zhang and N. Fooare, "EPDL: A Logic for Causal Reasoning", School of Computer Science and Engineering, Univ New South Wales, Australia (2000)

[10] N. Cartwright. "How to do things with Causes", American Philosophical Association, Vancouver British Columbia (April 2009)

[11] E.M. Hutchins, M.J. Cloppert and R.M Amin PH.D., "Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains", Proceedings 6th International Conference on Information Warfare and Security (ICIW 11 - Washington DC USA March 2011), Academic Publishing International Limited, pp. 113–125, 2011
http://www.lockheedmartin.com/content/dam/lockheed/data/corporate/documents/LM-White-Paper-Intel-Driven-Defense.pdf

[12] See e.g. Autonomy Technology Overview White Paper, HP Autonomy, 2009
http://publications.autonomy.com/

[13] See e.g. URL: http://www-03.ibm.com/innovation/us/watson/index.shtml

[14] L. Briesemeister, S. Cheung, U. Lindqvist and A. Valdes, "Detection, Correlation, and Visualization of Attacks against Critical Infrastructure Systems", Proceedings Eighth Annual Conference on Privacy, Security and Trust, Ottawa, Ontario, Canada, August 17-19, 2010
http://www.csl.sri.com/papers/PST2010/pst2010.pdf

[15] G. Jakobson, "Mission Cyber Security Situation Assessment Using Impact Dependency Graphs", 14th International Conference on Information Fusion (Chicago, Illinois, USA, July 5-8, 2011)

[16] See e.g. Proceedings IEEE VAST 2013 (Atlanta, Georgia, USA, 13-18 October 2013), IEEE Transactions on Visualization and Computer Graphics, Dec 2013

[17] See e.g. Proceedings of the 7th ACM international conference on Distributed even-based systems (Arlington, Texas, USA, 29 June - 03 July 2013), ACM 2013

[18] J. Mugan, "A Developmental Approach to Learning Causal Models for Cyber Security", SPIE Defense, Security, and Sensing, Machine Intelligence and Bio-inspired Computation: Theory and Applications VII (2013)

[19] J. Pearl, "Causal Inference in Statistics: An overview", Statistics Surveys Volume 3, The American Statistical Association, the Bernoulli Society, the Institute of Mathematical Statistics, and the Statistical Society of Canada, pp. 96-146, 2009

[20] R. Martino, "Leveraging Traditional Battle Damage Assessment Procedures to Measure Effects from a Computer Network Attack", USAF Institute of Technology (June, 2011).

