Ubiquitous command and control

Dale Lambert and Jason Scholz

Abstract. We assert that the current conceptualisation of military operations has itself become a legacy system, which fails to take advantage of emerging technology and evolving commercial business practice.

In response to this, we present a conceptual framework for the development of potentially extremely robust societies of human decision-makers and automated machine decision systems appropriate for future military conflict, that is commensurate with emerging technology.

The framework is based on a context of adaptation of location, through telecommunications and transportation resulting in richer human presence, and its ensuing effect on the function and structure of enterprises. Eight principles or tenets are presented which exploit this context. We present principles of decision devolution, ubiquity, automation, human and system integration, distribution and decentralization to each provide forms of robustness through diversity. Principles of social coordination and management provide unity and bounds on this diversity.

1. Introduction

Command and Control is defined in the doctrine of many Western allied nations as:

> The exercise of authority and direction by a properly designated commander over assigned and attached forces in the accomplishment of the mission. Command and control functions are performed through an arrangement of personnel, equipment, communications, facilities, and procedures employed by a commander in planning, directing, co-ordinating, and controlling forces and operations in the accomplishment of the mission. United States Department of Defense [29].

Command authority has generally been typified by access to privileged information, which historically has been disseminated hierarchically. The impact of telecommunications technologies on military and security operations has served to undermine this situation. Extending this phenomenon, we propose the development of similar and significant decision making capabilities (human and machine) throughout the organisation to provide robustness and diversity.

The conduct of military operations has a rich, long and established history, which by *argumentum ad verecundiam*, may cause some to immediately dismiss this framework as erroneous or impractical. For those entrenched in the current ways of doing business, it will be difficult to appreciate that the current conceptualisation of military operations has, *ipso facto*, become a legacy system. So as a purely academic exercise to free the reader from legacy thinking, in considering the following paper, do not consider the framework in the context of conventional military operations. Pretend instead, that you are designing a system for offensive terrorist units and defensive counter-terrorist units, a contest for which we have no established capability and so no preconceived approach. Then ask yourself which system you would adopt – the conventional military approach or the one presented here?

The concept was first described in 1999 [13], using the US "Cooperative Engagement Capability" as a foil for illustration, and later in [17] using "Network Centric Warfare" as a foil. The following draws on both of these works and further extends the latter to include more detail on regulation in social coordination.

2. Context of adaptations

We assume a "process philosophy" [3], in which identity (temporal unity) is defined on the basis of change (temporal diversity), rather than defining change on the basis of identity. In an object philosophy, identity is assumed and then you choose how to deal with change, whereas in a process philosophy, change

is assumed and you choose how to assemble identities. Consequently, adaptability is a way of understanding the world, rather than something to be understood about the world. Adaptability therefore provides a way of understanding the process that engenders the context of what is to follow.

2.1. Adaptation of location

Adaptations in transportation and telecommunications have altered the extent to which presence is influenced by distance.

- Villages were the centre of social organisation when horses were the primary personal means of transportation, because the speed of the horse constrained the rate at which one could change the location of their physical presence. Virtual presence required the physical presence of a messenger.
- The development of the telegraph in 1837 allowed a degree of influence over considerable distance with minor delays, though this species of virtual presence still offered noticeably less influence than physical presence.
- The invention of the telephone in 1876 provided a more effective mechanism for virtual presence, though telephones did not reside in almost every home in Australia until the 1960s.
- Private motor vehicles had superseded horses and public transport as the principal means of personal transportation for every home by the 1950s. This resulted in the proliferation of highways during the 1950s and 1960s. The resulting interaction between larger centres of population reduced the importance of villages and smaller towns as a social cluster of influence.
- Overseas air travel usurped ocean travel during the 1960s and 1970s to provide more accessible international presence.
- More recent advances in satellite communications, mobile telephones, text messaging, video conferencing and the Internet have all contributed to a substantially richer capacity for virtual presence.

With each of the advances in transport and telecommunications, the effect has been to increase the scope for both presence and virtual presence. The outcome of this effect has been to shift the sphere of influence from localisation to globalisation.

> The world of the industrial revolution is giving way to an era of global economy, powerful information technologies and relentless change [12, pp. 216].

2.2. Adaptation of function

The adaptation of location is engendering an adaptation of the function performed by individuals and organisations. Increased presence and virtual presence has increased both the nature and number of players that can influence a function. The effect is to increase the scope for both competition and collaboration.

- Competition increases because an external presence can be more easily imported to perform a function.
- Collaboration increases because one can more easily export functional expertise as a component within broader functions.

The outcome of these competitive and collaborative forces has produced an increased focus on competitive strength. Individuals and organisations have been compelled to understand the functions that they can competitively perform and then apply them collaboratively in strategic alliances with individuals and organisations possessing expertise in complementary functions. The management literature has been preaching collaborative advantage through strategic alliance since the mid 1980s, where a strategic alliance is a cooperative agreement on,

> . . . joint research efforts, technology sharing, joint use of production facilities, marketing one another's products, or joining forces to manufacture components or assemble finished products [28, p. 165].

The motor vehicle industry typifies the approach [4]. Companies are increasingly understanding themselves as "Lego blocks" of functional capability within a global economy. A recent growth area has been in metalevel businesses that integrate other "Lego block" businesses.

2.3. Adaptation of structure

The adaptation of function is inducing an adaptation of structure. Walker [31] observes that,

> Companies typically divide themselves into chunks according to some logical differentiation. At the same time, they seek to find ways to glue these chunks together as a coherent enterprise. Walker [31, p. 136].

Differentiation (diversity) allows the organisation to form smaller units that can be more focused and better managed. Differentiation is usually undertaken on the basis of product, function or market. *Integration* (unity)

provides competitive advantage through efficiencies, synergies and combined value adding (*ibid*).

Under the early industrial model, industries were centrally located, usually opportunistically based on geographic features, and were operated under the ideals of specialisation in production and a division of labour. As the size and complexity of firms increased, the roles became more specialised, encouraging a hierarchical integration to manage the centralisation of differentiated function. The hierarchical management structure has, in principle, remained until current times, and influences the conception of strategic planning. Through an analysis of texts, Mintzberg [21] identified the following three propositions as basic premises of the classical strategic management framework:

1. strategy formation should be a controlled conscious thought process, which is economically based and rationally determined;
2. responsibility for the conscious thought process rests with a single individual at the top; and
3. strategy proceeds in a linear top down fashion through development and then to implementation.

The effect of an increase in competition and strategic alliance is to erode the classical hierarchical structure to include networked structures. This induces a number of outcomes.

1. Command and subjective "rational" decision-making will be tempered by negotiation. Command and control will be supplemented by collaboration.
2. Single minded emphases, such as the economist's focus on profit or the militarist's focus on force, will increasingly need to be understood and applied against a broader diversity of motivating goals.
3. The presumption of control will increasingly be understood as a question of managing change in a complex environment.

2.4. Adaptation of adaptation

The adapted networked structures will increase the diversity available to an organisation, and this in turn, *will intensify the pace of change.* As the tempo of change increases, organisations must learn to adapt "Lego block" capabilities to satisfy intent, with an awareness "... equivalent to the craftsman's feel for the clay" [20, p. 109].
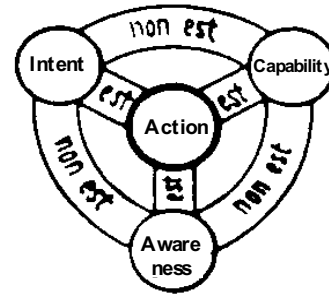

Fig. 1. The action trinity.

At work, the potter sits before a lump of clay on the wheel. Her mind is on the clay, but she is also aware of sitting between her past experiences and her future prospects. She knows exactly what has and has not worked for her in the past. She has an intimate knowledge of her work, her capabilities, and her markets. As a craftsman, she senses rather than analyzes these things; her knowledge is "tacit". All these things are working in her mind as her hands are working the clay. The product that emerges on the wheel is likely to be in the tradition of her past work, but she may break away and embark on a new direction. Even so, the past is no less present, projecting itself into the future (*ibid.*).

This also affects how we understand the identity (unity) of an organisation. Organisations have traditionally understood themselves as persistent, and as a consequence, changes in organisational location, function and structure have often been violent. A process view of organisations, as intimated by Senge's [25] "learning organisations", counters this violence.

One outcome of a process conception of organisational identity is that organisational change becomes less a centralised decision and more of an environmental effect of adaptations in location, function and structure. A second outcome is that competitive strength no longer lies solely in knowledge and strategic alliance, but also in their adaptation. Innovation and the ability to form dynamic relationships become the basis for competitive strength. In an information economy in which information is rapidly traded, innovation becomes the new means of production.

Adaptability is a fundamental tenet of the synthesis in that it is a feature of all of the following tenets, each of which successively builds upon its predecessors. The benefit of adaptation is the ability to operate in changing environments.

2.5. Action, decision and command and control

Some further conceptualization of action and decision is needed before presenting the framework. We may consider decision to be the reasoning side of action. One may view action as involving intent, capability or awareness. The trinity of capability, intent and awareness depicted in Fig. 1 is founded on mutual tension. For example, if capability is held fixed and awareness is reduced, achievable intent is also reduced, and so forth.

Intent is in some sense the foremost of the trinity, driven by an underlying will. Intent establishes future goals and forms the basis for so called 'effects-based (military) operations'.

We assert that *command* involves the expression of *intent* to another[1]. Complementing this we assert, *control* involves the expression of a *capability* (a plan is an example of a capability) to another and the monitoring and correction of the execution of that capability. Lambert [14] in effect suggests that we can understand *action* as the utilisation of *capability* to achieve *intent*, given *awareness*.

Endsley [10] provides a useful definition of situation awareness as,

> Situation awareness is the perception of the elements in the environment within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future [10].

To illustrate, at some time t an individual may harbour the intended effect $i_{t,k1}$ for future time k_1 and intended effect $i_{t,k2}$ for future time k_2. Figure 2 illustrates these intended goals on a time line. At time t, these future goals are framed against some awareness of the current situation, represented in Fig. 2 by a_t. To achieve the intended effects $i_{t,k1}$ and $i_{t,k2}$, various capability options c_j can be considered. A capability option is viewed as anything that has a capacity to change one's awareness of the world, typically by changing the world. This is a much broader sense of the term 'capability' than is commonly applied in military contexts, which often restrict the focus to military equipment. The cascading of capability options offers a number of potential states of awareness of the world, with the best course of action being the set of capability options $\{c_1, c_3, c_7\}$ that delivers the states of awareness $\{a_t, a_{t,<1>}, a_{t,<1,3,>}, a_{t,<1,3,7>}\}$ that devi-

ate least from the intended effects $\{i_{t,k1}, i_{t,k2}\}$ at times k_1 and k_2 respectively. Decision or decision-making is the choice of capability to achieve intent given awareness, and can be formulated as a dynamic programming problem if we care to do so. Of course things change, so the best course of action at time t may not remain the best course of action at some time t' greater than t.

3. A conceptual framework

The theme of the framework is 'unity with diversity'. It balances the robustness but instability of diversity, with the stability but fragility of unity.

3.1. Decision devolution

The decision devolution principle is founded upon the idea that additional individuals or entities are not always required to govern collectives. When appropriately equipped, collectives can *sometimes* govern themselves. In the context of defence systems this signals dynamic liaisons adaptively forming from operational assets without the oversight of a command headquarters.

The conduct of military operations without the oversight of a command headquarters is of course an anathema to current military practice, and might well foster allegations of heresy against the authors. But large-scale collectives can successfully operate without a ruling class as exemplified by such Internet sites as Wikipedia, eBay, YouTube, Geocaching, MySpace, Chat rooms, Instant messaging, Digg, and Second Life. Command involves the creative expression of intent to another. Control involves the expression of a capability to another and the monitoring and correction of the execution of that capability. Processes akin to these operate within eBay on a significant scale, without the oversight of a ruling class. Command resembles the vendor expressing the intent of sale, with any member of the collective potentially being a vendor. Control resembles the process by which the purchaser acquires the sale item, with any member of the collective potentially being a purchaser. Control works in eBay because the collective is largely self-monitoring and self-correcting. Customer satisfaction with each transaction is recorded and made visible to all in the collective. Ideally, this monitoring mechanism then facilitates correction, by steering prospective purchasers away from exposed historically fraudulent vendors.

Decision devolution means that decisions run across the collective,

---

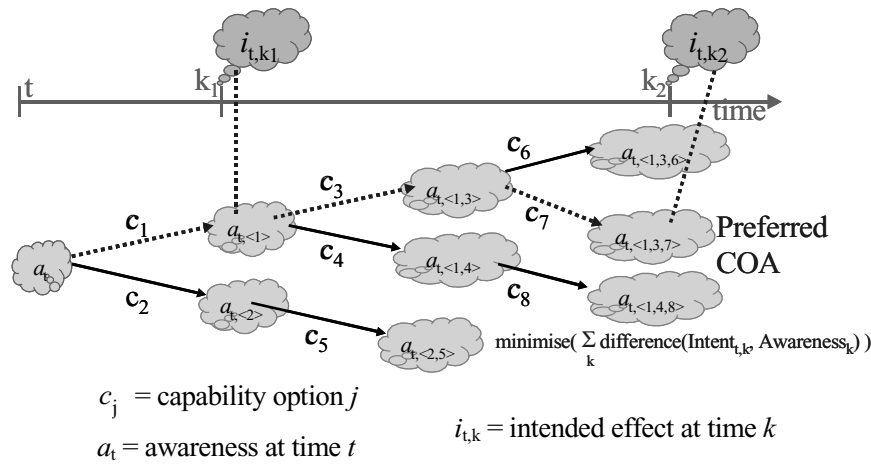[1] Of course one might include command of one's-self.

Fig. 2. A conceptualisation to aid in understanding action and decision: fiAct to use capability to achieve intent, give awareness'.

... once the environment cranks up the rate of change required for competitive performance, complex tasks must be correspondingly repartitioned, and human capital correspondingly reallocated. And absent some infinitely capable overseer, this repartitioning problem must be solved by the same individuals who have to perform the task of production. The result, in a successful firm, is a continual swirl of problem-solving activity and ever shifting interactions between the problem solvers, each of whom has information relevant to the solution of a particular problem but none of whom knows enough to act in isolation. Nor does any one person know precisely who knows what; hence, problem solving is a matter not just of forming the necessary combination of resources ... but of searching for and discovering those resources in the first place [32, p. 269].

The potential benefits of decision devolution are flexibility and redundancy. Flexibility can arise through the ability to share the load throughout the collective. Redundancy ensues because the conduct of military operations can still proceed even if its command centre becomes inoperative.

3.2. Ubiquity

The ubiquity principle argues: (i) for a C^2 component on *every* platform; and (ii) that these components should be *similar*, not identical.

3.2.1. Graceful degradation

A C^2 component on *every* platform allows command and control to degrade gracefully under strike by re-configuring C^2 among the remaining assets.

In the Information Age, C^2 centres have become the enemy's centre of gravity, and are therefore the prime targets for precision strike. In defending against precision strike, one approach is to build a duplicate C^2 centre. The neutralisation of the C^2 centre is then less catastrophic, as the duplicate centre can assume its function. But redundancy offers only one level of reprieve. By enabling C^2 functionality to re-configure as necessary, ubiquity offers greater sustainability, by enabling the quality of defence to degrade gracefully, rather than instantaneously, under the threat of surgical strike [13, p. 13].

Figure 3 illustrates the point. In frame 1 a forces of darkness fighter approaches a ubiquitous command and control (UC^2) system with malevolent intent. In frame 2 the fighter destroys the commanding ship and C^2 reconfigures in response. In frame 3 the fighter destroys a C^2 significant bomber, and again C^2 reconfigures. In each successive frame the quality of C^2 diminishes but is not extinguished. In principle defeating a ubiquitous system amounts to defeating all of its assets.

3.2.2. Agreement

Having similar, rather than identical components, offers a balance between unity and diversity. But why do we need this balance? This relates to the notion of "common".
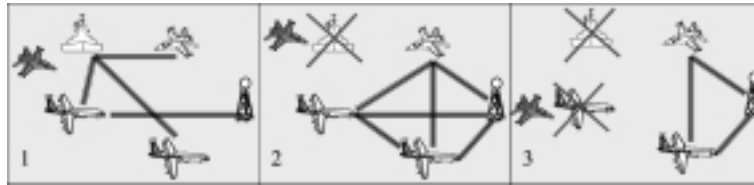
Fig. 3. Reconfiguring C2 under strike.

A first interpretation of "common" is "common as identity". This may for example involve disseminating an identical "common operating picture" to each person in the collective. This mistakes identity for unity. The Great Irish Potato Famine of the 1840s led to a significant number of deaths and refugees. It resulted from a uniform dependency on an identical food source (potatoes) that became infected. The distribution of an identical picture beckons analogous concerns, as an infected picture might ensure everyone has the wrong understanding. Biology teaches us of the fragility of uniform identity and the robustness of diversity. Another drawback with "common as identity" is that not everyone wants to see an identical picture. Different individuals are interested in different aspects of the environment and at different levels of granularity.

The second interpretation is "common as consistency". Instead of disseminating an identical picture, consistent databases and/or information feeds are disseminated. This allows different individuals to generate their own picture of interest from the same underlying consistent information. But consistency is not as desirable as it first seems. If a fusion system receives assertion α from source X and assertion not(α) from equally trusted source Y, then which assertion should be entered into the consistent database? If the wrong one is entered, then the wrong information is propagated to every individual in the environment. "Common as consistency" lacks robustness because it eliminates diversity.

The third interpretation "common as agreement" allows individuals to harbour both public and private views. Public views being a product of agreement with other individuals, while the private views retain alternatives should they be required. Under the weight of public opinion, individual Y might be persuaded to accept some statement α, but is free to privately retain his or her reasons for endorsing not(α). This might subsequently prove to be invaluable if it turns out that not(α) is in fact correct. Inconsistencies should be managed, not discarded. Agreement facilitates social unity while retaining the robustness of diversity.

3.3. Automation

The automation principle argues that some expertise should be automated through software, and indeed, that this is the mechanism by which ubiquity might be achievable. Some decision making can be fully automated. Other aspects will be better performed by humans, with the choice between the two being mediated empirically.

3.3.1. Automated decision-making

The intention is that automated software expertise should facilitate automated decision makers (we might term these "software agents") as well as automated decision aids. The proliferation of automated decision-makers within society generally was noted in Lambert [13].

> Smart decision software is steadily becoming a ubiquitous commodity throughout our society. It pervades throughout our homes, our work, our transportation, our health, and our leisure. The military environment too, has experienced the proliferation of smart decision software within its assets. The upshot is that something like a C^2 capability is steadily emerging within our assets, and the communications to efficiently link them is improving [13, p. 36].

Since then that trend has only increased. Our automobiles, for example, now integrate the driver with a society of automated decision-making components developed by a disparate group of strategically aligned global partners.

The prospect of automated decision-making in a military context is controversial. Some might contend on moral grounds that military operations should be immune from the automation progression otherwise evident in society. There are two responses to this. First, automation will proceed in military operations whether or not it should. In 2000, the US Congress stated:

> It shall be a goal of the Armed Forces to achieve the fielding of unmanned, remotely controlled technology such that –

(1) by 2010, one-third of the operational deep strike aircraft of the Armed Forces are unmanned; and
(2) by 2015, one-third of the operational ground combat vehicles of the Armed Forces are unmanned [30].

Second, there is a case for including automation within military weaponry. Automobiles rival wars as a contributor to human death, and yet the automobile industry is one of the leaders in integrating automated decision-makers. Much of the manufacturer's motivation is to make automobiles safer. A similar motivation could apply in a military context. If a missile that has been instructed to destroy a train bridge observes or is informed of a passenger train traversing that bridge as it approaches, then we would want the missile to exercise moral judgment and defer its strike on the bridge until after the passenger train has departed the scene. This might be achieved by building in Rules of Engagement (ROE) into the missile that ensure conformance with national moral intent.

3.3.2. Ubiquity through automation

The advantage of automated software expertise is that it is easily replicated, adapted and distributed. The benefit is that automated software expertise is more readily transferable, which enables the ubiquity of C^2 capability.

> The fact that we can readily duplicate software then becomes the crucial attribute, because duplicated software encoded human expertise is the mechanism that facilitates the ubiquitous capability [13, p. 36].

An illustration of the theme occurs in the movie "The Matrix" when the character Trinity, played by Carrie-Anne Moss, suddenly needs to be able to fly a nearby helicopter. The following dialogue ensues,

> Neo: "Can you fly that thing?"
> Trinity: "Not yet." *Trinity dials on her cell phone.*
> Tank: "Operator."
> Trinity: "Tank, I need a pilot program for a V2-12 helicopter. Hurry."
> Tank: *Tank downloads the program.*
> Trinity: *Trinity turns to Neo.* "Let's go."

This expresses the idea, but in the near term we would download expertise relevant to that person and the helicopter, to the helicopter or the person's personal automated agent, not the person.

The encapsulation of expertise in software will gain in currency as two mind set changes become more pro-

nounced. The first is an acceptance of *semantic machines*. Computers are so named because they were conceived during a wartime calculation boom as rapid number crunching devices. Nowadays computers are instead viewed as something akin to post office boxes that serve as repositories in which people store information, so that they or other people can access that information subsequently. The machines themselves have no understanding of the information they hold. One might be able to retrieve the statement "Saddam Hussein bought munitions from Mussoria" through a syntactic search on "Saddam", but one cannot retrieve it through a semantic search on "Iraq". Nowak and Lambert [23], for example, reports recent software applications using the five-tiered conceptual structure of Fig. 4, taken from Lambert [15]. It is symptomatic of a new shift toward a "semantic web" and semantic machines that associate meanings with the information they hold about the world by constraining possible interpretations through formal logics.

The second mind set change is an acceptance of *cognitive machines*. Computers are currently viewed as machines that hold information that people reason about. In time computers will come to be understood as machines that have agents that people reason with. Over a decade ago the first author developed a software agent system called ATTITUDE, initially for dynamically managing a phased array radar on an Airborne Early Warning aircraft. ATTITUDE is so named because it is programmed at the cognitive level in terms of propositional attitudes like beliefs, desires and expectations [14]. Figure 5 illustrates the cognitive model associated with an ATTITUDE agent's "mind".

Lambert [16] outlines a strategy for capturing cognitive routines as ATTITUDE software.

ATTITUDE adopts an extended "Beliefs Desires and Intents" (BDI) architecture. A well known example of the BDI architecture was proposed by Bratman et al. [2].

ATTITUDE has a 'learn' attitude which may be used to automatically construct cognitive routines using cognitive routines. In this way, behaviour ascribable to intentions may be formed by the machine automatically on the basis of its interactions with the environment, rather than on the basis of its programming by a programmer. Machine intentions need not be explicit. The computer program "Temporal Difference (TD) Gammon", by Tesauro [27] has no explicit intentional representations (as in a BDI architecture), and learned to play Backgammon by interaction with its environment (both human and machine players).

Social: group, ally, enemy, neutral, own, possess, invite, offer, accept, authorise, allow,
Cognitive: individual, routine, learnt, achieve, perform, succeed, fail, intend, desire, belief, expect, anticipate, sense, inform, effect, approve, disapprove, prefer,
Functional: sense, move, strike, attach, inform, operational, disrupt, neutralise, destroy,
Environmental: land, sea, air, outer space, incline, decline, number, temperature, weight, energy,
Metaphysical: exist, fragment, identity, time, before, space, connect, distance, area, volume, angle,

Fig. 4. Tiers of Semantic Primitives for Military Operations.
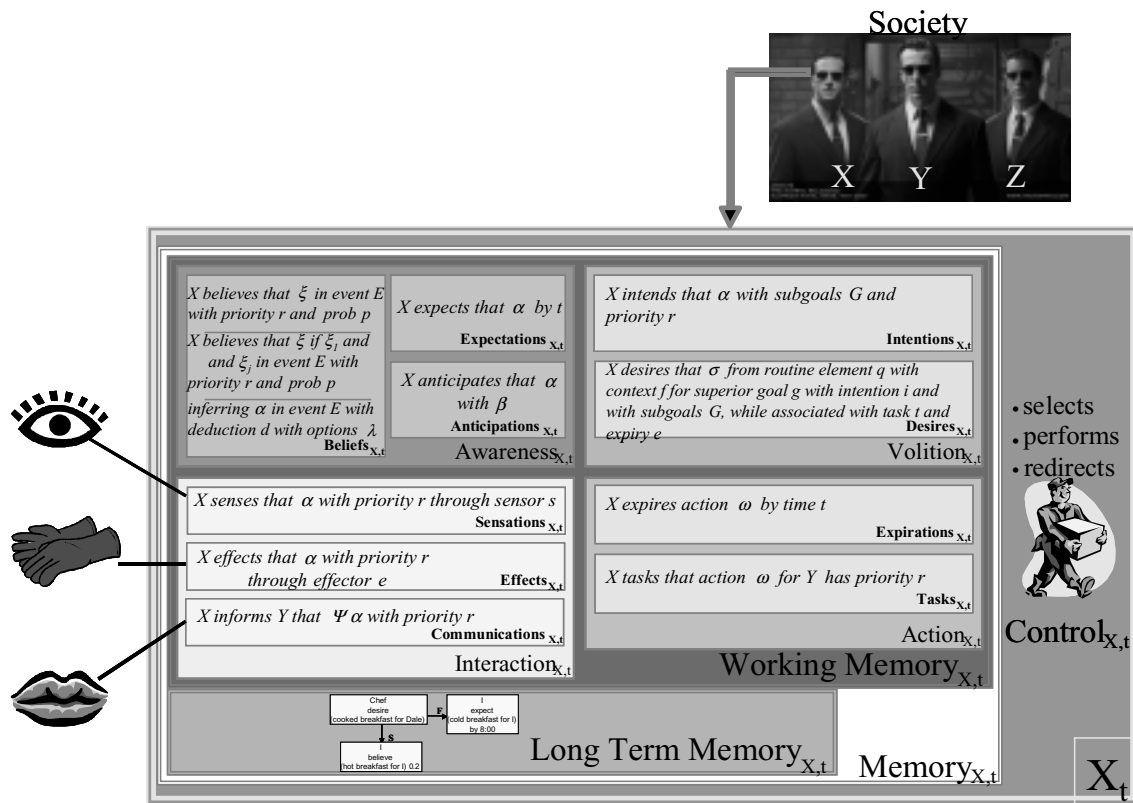

Fig. 5. ATTITUDE Cognitive Model.

(TD Gammon) plays at a strong master level that is extremely close (within a few hundredths of a point) to equaling the world's best human players [27].

Some of the plays achieved by TD-Gammon are considered by experts to be extremely novel. Kit Woolsey was ranked number 5 in the world in 2007, and is attributed to stating,

TD-Gammon has definitely come into its own. There is no question in my mind that its positional judgment is far better than mine. Only on small technical areas can I claim a definite advantage over it . . .

In the more complex positions, TD has a definite edge. In particular, its judgment on bold vs. safe play decisions, which is what backgammon really is all about, is nothing short of phenomenal. (Quote by Kit Woolsey in Teasauro [27]).

3.4. Integration

The integration tenet addresses the integration of people and machines. It makes two points, one in relation to mixed initiatives and the other in relation to improvements in interaction.

3.4.1. Mixed initiative on authority, responsibility and competency

The principle is for automated and human decision-making to be fully integrated. Integration exists to complement the weaknesses in some parts of a system with strengths in other parts of a system. This includes the division of labour between people and machines. James Reason, who has undertaken extensive research on human expertise and error, captures the intent beautifully from the human perspective, through contrasting "the human as hazard" with "the human as hero". People can exhibit great flexibility, adaptation, recovery and improvisation to perform heroic acts. Apollo 13 and Chess grand masters, who can play simultaneously against over forty Chess players while blindfolded, are examples of the incredible capability of humans. But humans also make errors, commonplace errors of little consequence and uncommon errors with serious consequence. The shooting down of the Iranian passenger aircraft by the US Navy in 1987, and the Challenger space shuttle disaster, are examples of the human as hazard. And most importantly, the heroes and hazards are not two different groups of people. The heroes are sometimes hazards.

The division of labor between people and machines should be developed to leave human decision making unfettered by machine interference when it is likely to prove heroic, and enhance human decision making with automated decision aids, or possibly override it with automated decision making, when it is likely to be hazardous. Again, overriding human decision-making is a highly contentious suggestion. But in cases where a machine detects a violation of rules of engagement, such as the train illustration at the close of Section 3.3.1, then it should be able to at least question the order before compliance. From the authors' perspective, the appropriate balance between the exercise of intent by people and machines is something best determined empirically, rather than on the basis of *a priori* belief. A simple illustration of machine decision equality to human decision is the negotiation of traffic lights at a motor vehicle intersection. The machine takes the initiative, commanding human drivers to stop – which one would generally do, for good reason!

When automated components substitute functionality that is currently provided by people in hierarchic structures, including social coordination functionality, then those automated agents must accept the authority, responsibility and competency [19] associated with that functionality. For automated agents, this should be ordered by competency, then responsibility, and then authority.

- An automated agent's *competency* will depend on the expertise embedded within it, and the agreements it forms should primarily derive from its competencies.
- An automated agent's *responsibility* will follow from the social agreements it forms, given available competencies.
- An automated agent's *authority* is not determined by *a priori* rank, but depends upon the role it assumes in social agreements, given available competencies. In the end, authority is a matter of agreement.

3.4.2. Improved interaction

Following on from Section 3.3.2, the integration tenet also contends that as the machines acquire an ability to reason about their environment, *id est* comprehend and project, they will also require a means of presenting information to people that goes beyond simple "dots on maps" displays and the desktop metaphor. In essence, the machines need to have a storytelling capability.

In our everyday lives, television news often provides our situation awareness about the world. It does this by assembling presenters, maps, diagrams and video footage to convey stories about daily events of interest. A Future Operations Centre Analysis Laboratory (FOCAL) embraced this approach by developing software virtual advisers, virtual battlespaces, virtual interaction mechanisms and environments, and virtual videos, as respective software counterparts to the presenters, maps, diagrams and video footage featured in news services [15]. Figure 6 illustrates the correspondence. As software, it allows the machine to generate stories from its accessible information. As software it is portable, being easily replicated, adapted and distributed throughout a network. And unlike television news services, as software it is interactive[2], allowing the user to access the information of interest to them.

3.5. Distributed, decentralised and non-linear

This principle advocates diversity by endorsing C^2 that is distributed, decentralised and non-linear.

---

[2]In FOCAL the interaction mechanisms include speech recognition, natural language processing, speech synthesis, stereoscopic displays, open hand gesture recognition and gaze tracking.
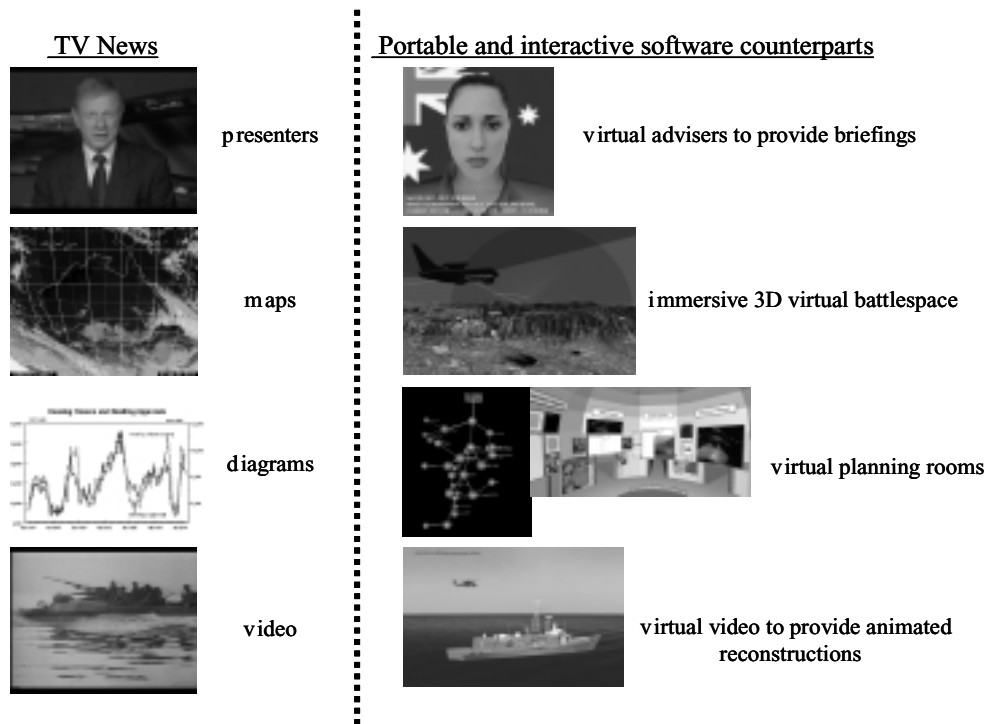
TV News	Portable and interactive software counterparts
presenters	virtual advisers to provide briefings
maps	immersive 3D virtual battlespace
diagrams	virtual planning rooms
video	virtual video to provide animated reconstructions

Fig. 6. Storytelling Technology.

3.5.1. Distributed systems

Distributed systems postulates that C^2 should be distributed across location. This reflects the adaptation of location theme presented in Section 2.1. Current web portal technology characterises the desirable attributes.

- Awareness should be retained by members of the collective in accordance with their roles.
- Related pieces of the distributed awareness should be accessible to members of the collective from a single point of access.
- The single point of access should be location independent.
- Access to awareness should occur without members of the collective having to know the distributed locations from which it derives.

Distributed systems affords location independent access (unity) while the physical distribution of information (diversity) offers protection from spatio-temporally constrained strike capabilities like missiles.

3.5.2. Decentralised systems

Decentralised systems postulates that C^2 should support the decentralisation of intent. Each member of the collective should have the capacity to ask (pull awareness), tell (push awareness), command (push intent)

and obey (accept intent). The decentralisation of intent therefore allows for agreements about intent as well as awareness. Decentralised systems afford protection from strike capabilities that target centralised will (the origin and ownership of intent) like assassination and blackmail. The framework combines distributed systems with decentralised systems. It accommodates a diversity of intent situated at a diversity of locations.

Decentralised systems give rise to what we might term "mission agreement". Mission agreement allows for agreements that are not restricted to a hierarchical top down cascading of intent. Consequently, mission agreement supersedes the conventional notion of mission command because it allows for intent network structures of which intent hierarchical structures are but one type. Intent can be introduced at "the edge" of an organisation and propagate inwards if it garners sufficient support. This introduces a "command fusion" (intent) issue to complement the "information fusion" (awareness) issue already present under mission command.

The generalisation of hierarchies to networks allows for the use of hierarchies when they are appropriate, and non-hierarchical networks when they are inappropriate. Intent hierarchies assume that tasks are decomposable into simpler subtasks, such that each subtask
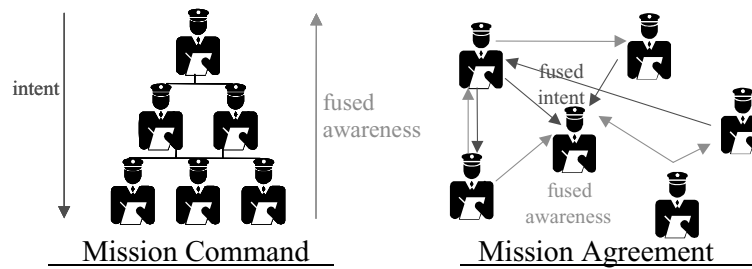
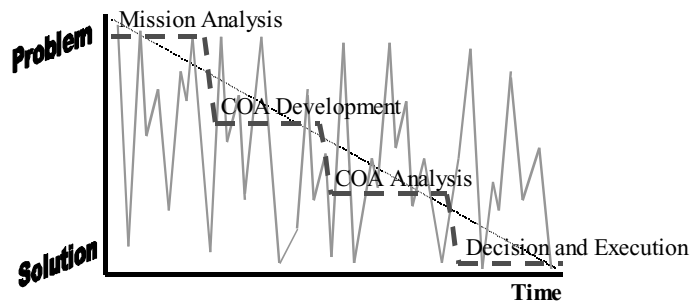Fig. 7. Mission Command and the more generalised form Mission Agreement.

Fig. 8. How groups of human experts faced with a novel problem, reason in arriving at a solution (solid line) as opposed to doctrinal guidance of JMAP (dashed line).

may be completed independently, and thus in parallel, with others. Decomposition applied to unordered domain situations and problems, usually fails to represent the relationships between objects that make behaviours 'more than the sum of their parts'. The reason for this is that a decomposition into M tasks also requires representation of 2^M combinations of relationship groupings if a complete representation of the original problem is to be retained. It is not surprising then to find for even small values of M, that most of these relationships are discarded or trivialised to maintain tractability. A focus on the M tasks at hand is much easier than the 2^M relationships.

Intent hierarchies may not be robust. When assigned tasks have inherent relationship overlap, message passing will occur between actors crossing the hierarchy in an attempt to recover that which has been lost through task decomposition. High rates or volume of such traffic will tend to overload key nodes causing congestion. Furthermore, a hierarchy with N nodes has a mere $N - 1$ links, so there is no reserve ready if links are lost – they lack adequate connectivity. A breakdown near the top of a hierarchy has the effect of isolating all lower nodes from the rest of the network. If a link is lost between physically distributed commanders, the two isolated structures will have reduced mutual unity in the wake of any subsequent environmental changes.

Study of some of these effects has been conducted using a variant of the game of checkers that includes network structural overlay, where actions of pieces are frozen in the event of becoming isolated [5].

The traditional view of the organisation as a vertically integrated hierarchy may be critically incomplete, yet it remains a vital legacy consideration in C^2. Fortunately, there is a form of network topology that can provide dynamic robustness by augmenting hierarchies, yet remains close to optimum. Multi-scale networks [9], appear to uniquely combine both congestion robustness and connectivity robustness to a wide range of environmental conditions, and retain good scalability. Multi-scale networks combine local unity (highly ordered clusters) with long-reaching diversity (random shortcut links across the network) with the long-reaching links grading from a large number at the top to a small number at the base. This means that teams function at various scales with the information burden distributed across multiple scales to avoid overloading individuals.

3.5.3. Non-linear systems

Non-linear systems postulates that C^2 should support the non-linear achievement of intent. Complexity is inherent in human decision-making processes. Figure 8, adapted from Conklin [6], illustrates the linear

"Waterfall model" process of progressing from problem to solution, compared with an example from an experiment with a team of designers (solid line). We have overlaid the four levels of a doctrinal planning process called the Joint Military Appreciation Process (JMAP) [33] as a dotted line.

The following describes their experiment and key result:

A number of designers participated in an experiment in which the exercise was to design an elevator control system for an office building. All of the participants in the study were experienced and expert integrated circuit designers, but they had never worked on elevator systems before. Indeed, their only experience with elevator systems came from riding in elevators. Each participant was asked to think out loud while they worked on the problem. The sessions were videotaped and analyzed in great detail.

The analysis showed, not surprisingly, that these designers worked simultaneously on *understanding the problem* and *formulating a solution.* They exhibited two ways of trying to *understand the problem:*

– efforts to understand the requirements for the system (from a one page problem statement they were given at the beginning of the session); and

– mental simulations (e.g. "Let's see, I'm on the second floor and the elevator is on the third floor and I push the 'Up' button. That's going to create this situation . . . ").

They would start by trying to understand the problem, but they would immediately jump into formulating potential solutions. Then they would jump back up to refining their understanding of the problem. Rather than being orderly and linear, the line plotting the course of their thinking looks more like a seismograph for a major earthquake, as illustrated.

In particular, the experiment showed that, faced with a novel problem, human beings do not simply start by gathering and analyzing data about the problem.

It is also striking, from (the) figure, that problem understanding continues to evolve until the very end of the experiment. Our experience in observing individuals and groups working on design and planning problems is that, indeed, their understanding of the problem continues to evolve – forever! Even well into the implementation of the

design or plan, the understanding of the problem, the "real issue," is changing and growing.

This non-linear process is not a defect, but rather the mark of an intelligent and creative learning process.

As chaotic as this pattern of activity appears, it reflects a deeper order in the cognitive process. It shows that people formulate possible solutions and try them out in order to better understand the problem. The new insights into the problem gave them fresh ideas about the shape of the solution [6, pp. 2–3].

A linear process may be used to justify the rationality of a solution post-hoc, however, faced with a novel problem it would appear to fail to support the non-linear process natural to human individuals and their social cognition, when engaged on novel problems.

3.6. Social coordination

In general systems will have a demand pool of human and machine agents offering intent, and a supply pool of human and machine agents offering capability. Moreover, the two pools will generally overlap, as any member of the collective can be a member of both pools concurrently. The challenge is to manage this level of flexibility without anarchy.

Systems can achieve social coordination by instituting social agreement protocols that coordinate collectives composed of both people and machines. The social coordination can be instituted through software, id est, as more sophisticated variants of existing workflow systems. In essence, eBay is a social agreement protocol implemented through software. The cost of finding information and expertise in this system is low and the agreement and monitoring mechanisms provide feedback for self-regulation. Examples of social contracting protocols, include:

1. Contract net protocol (CNP [26]);
2. Extended contract net protocol (ECNP [11]);
3. Provisional agreement protocol (PAP [24]); and
4. Legal agreement protocol (LAP).

The above protocol ordering reflects an increase in computational complexity, and an increase in rights for the proletariat. The legal agreement protocol introduced in [17], offers a facility for full contract law agreements between agents, be they human or machine. Figure 9 outlines the LAP. Anytime something is bought in everyday life, a contract exists, and so it is a familiar practice for humans. LAP formalises the

SPEECH ACT	SEND	RECEIVE
Request for Proposal (requests X α)	Contractor X uses domain knowledge to determine potential proposers $\{Y_1, \ldots, Y_n\}$ for RFP (request for proposal) goal α and sends X requests that α to each Y_i. A request can be sent at anytime.	Potential proposer Y_i receives X's requests for α. In response, Y_i determines its best proposal β_i, and if Y_i can undertake the proposal β_i, then Y_i sends a propose speech act.
Retraction of Request (retracts X α)	Contractor X decides not to proceed with a sent RFP for a and sends this to the individuals $\{Y_1, \ldots, Y_n\}$ that the RFP goal α was initially send to. A retraction of an RFP can occur, without penalty, any time before acceptance of an offer contributing to the RFP.	Potential proposer Y_i receives a retraction for RFP goal α and deletes all proposals and offers associated with it. Damages can be sought by Y_i against X for each of Y_i's offers accepted by X.
Proposal (proposes $Y_i \beta_i$ X α)	Proposer Y_i sends a proposal for β_i to X to achieve α in response to a request from X. The proposal can be sent any time before the request is retracted or a contracts for a has been let.	Contractor X receives a proposal bi to achieve a from Yi. In response, X determines whether bi is a better proposal for a than any currently received. If it is, then X sends an invitation for Y_i to offer β_i formally.
Withdrawal of Proposal (withdraw $Y_i \beta_i$ X α)	Proposer Y_i sends a withdrawal to X for proposal β_i to achieve α. Y_i will do this if it becomes clear that Y_i cannot perform the proposal. A withdrawal for β_i can be sent any time before an offer for β_i is sent.	Contractor X receives a withdrawal of proposal β_i from Y_i. X removes proposal β_i and all of its dependencies.
Invitation to Offer (invites X $Y_i \beta_i \alpha$)	Contractor X sends an invitation to proposer Y_i to formally offer proposal β_i for request α. X can send such an invitation whenever X believes Y_i's proposal β_i is the best proposal for α.	Proposer Y_i receives an invitation from X for proposal β_i to achieve α. In response, Y_i can send an offer θ_i to X for β_i, but this will be contractually binding on Y_i if X accepts that offer before Y_i revokes it.
Decline of Invitation (decline $Y_i \beta_i$ X α)	Proposer Y_i sends to X a decline to offer for proposal β_i to achieve α. Y_i can do this at any time.	Contractor X receives from Y_i a decline to offer for proposal β_i to achieve α. In response, X must delete β_i and all dependent states.
Offer (offers $Y_i \theta_i \beta_i$ X α)	Proposer Y_i sends an offer θ_i to X for proposal β_i to achieve α. Y_i will do this if Y_i believes it is prepared to be obligated, as θ_i will be contractually binding on Y_i if X accepts that offer before Y_i revokes it.	Contractor X receives an offer θ_i for proposal β_i to achieve α. If X accepts the offer then X is contractually obligated to it. X can instead delay acceptance, but risks Y_i revoking the offer before X can accept it.
Revocation of Offer (revokes $Y_i \theta_i \beta_i$ X α)	Proposer Y_i sends a revocation of offer θ_i to X for proposal β_i to achieve α. Y_i will do this if Y_i believes it has a better alternative. This succeeds provided X has not already accepted the offer, in which case Y_i will be liable for damages. The revocation of the offer occurs at the time at which it is sent, not received.	Contractor X receives a revocation of offer θ_i from Y_i for proposal β_i to achieve . In response, X must delete θ_i and all dependent states and is not entitled to damages unless θ_i has been accepted.
Acceptance of Offer (accepts X $Y_i \theta_i \beta_i \alpha$)	Contractor X sends acceptance of offer θ_i from Y_i for proposal β_i to achieve α. The effect is that both X and Yi are contractually obligated to achieve θ_i. If either party renigs, then they are liable for damages under breach of contract.	Proposer Y_i receives an acceptance of offer θ_i for proposal β_i to achieve α. At this point Y_i becomes aware of the mutual obligation.
Rejection of Offer (rejects X $Y_i \theta_i \beta_i \alpha$)	Contractor X sends a rejection of offer θ_i from Y_i for proposal β_i to achieve α. The effect is that no contractual obligation exists between X and Y_i. X will do this if X believe it has a better offer or potential offer.	Proposer Y_i receives a rejection of offer θ_i for proposal β_i to achieve α. At this point Y_i deletes offer θ_i and proposal β_i and all of their dependencies.

Fig. 9. Legal Agreement Protocol.

agreement process, and can be embedded in software so as to hide legal complexities unless required.

Social protocols, such as LAP, facilitate adaptive co-operative alliances of the sort canvassed in Section 2.2, through the formation of contractual agreements between members of the collective. They can also generate adaptive competitive factions, as members of the collective compete for capability resource to satisfy their intent. In their primitive form, the aforementioned protocols admit a *laissez-faire* management style.

– Agents are free to contract non-linearly rather than having to adhere to a linear waterfall model, such as the JMAP of Section 3.5.3. Agents are also
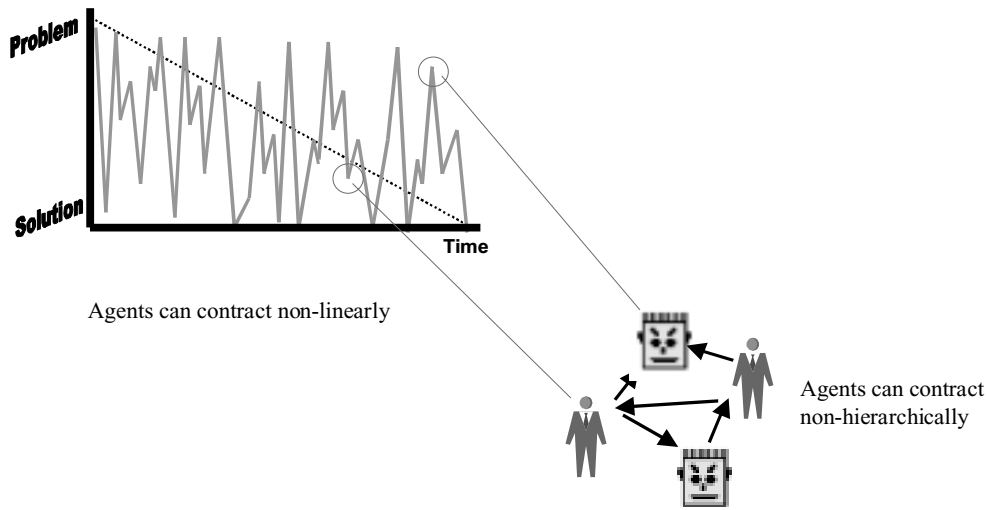
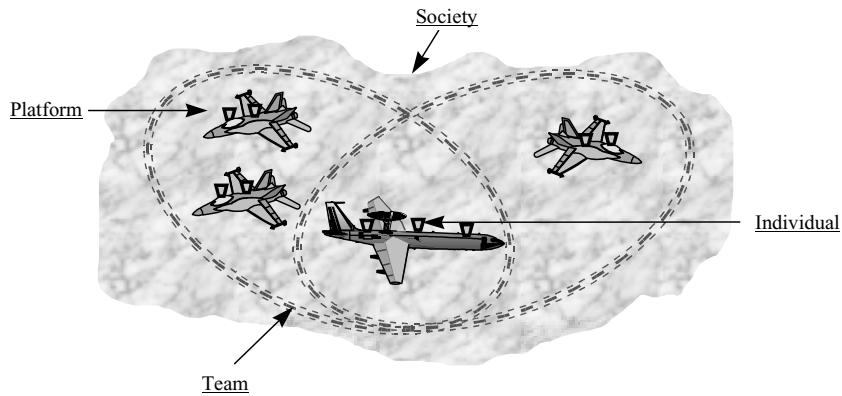Fig. 10. "Diverse Appreciation" – Agreement in UC^2.

Fig. 11. Levels of Management.

free to contract non-hierarchically rather than be constrained to the hierarchy of mission command. We term this alternative to JMAP, "diverse appreciation". This is illustrated in Fig. 10.

– Under the more advanced protocols, intent may be adapted to match the awareness of the contracted capability available. This allows behaviour to occur to match the level of awareness of the situation, where the adapted intent can precipitate action.

– Finally, in the spirit of Section 3.2.2, under the more advanced protocols, agents are also free to maintain public and private intent, so that competing intents within a UC^2 system can be managed rather than discarded, just in case they subsequently prove to be valuable.

Additional social policies will be required if individuals are to trade self-benefit for the benefit of a col-

lective. This could involve, for example, a multi-scale networking constraint to ensure efficiency within a collective. It could involve a prioritising of intent according to the rank of the members of the collective. It could have social constraints that govern how to trade membership and the quality of decisions with the time available. These issues will vary with the nature of the collective.

3.7. Management

3.7.1. Management levels

The framework identifies at least four management levels, characterised by diminishing proximity and increasingly flexible options for social coordination. Figure 11 identifies an instance of the four levels of management pictorially.

The four levels of management identify natural and social constraints that will necessarily be imposed on what might otherwise be the *laissez-faire* management style alluded to in Section 3.6.

– Individuals are the smallest unit of management. Whether human or machine, the individual practises self-management by relying on cognitive capabilities.
– Platforms provide the second unit of management. Despite advances in virtual presence noted in Section 2.1, some individuals will be collocated on platforms that must be socially coordinated.
– Teams constitute the third unit of management. Teams are formed on the basis of a commonality of intent, rather than a commonality of location. This allows for a more flexible approach to social coordination.
– Societies form the fourth unit of management. Societies form on the basis of interaction, be it physical or virtual. Societies accommodate the mix of collaborative and competitive ingredients noted in Section 2.2.

The system is perhaps best understood as a society of societies. The social agreement protocols and constraints have to contend with both dynamic intra and inter social group consequences. Individuals will generally belong to multiple social groups concurrently. Societies are dynamic, often with membership changes according to the mission.

3.7.2. Social regulation

Any coalition to achieve a specific mission intent will comprise agents with varying degrees of synergy (unity) and antagonism (diversity) towards that intent. Although unity is proposed to be achieved primarily through agreements, not everyone in a coalition need agree privately with the publicly accepted view. Some individuals may choose to accept the public position and retain their private position without further action so as to maintain public unity. Others may be less willing to do so. Given varying levels of dissent, it is necessary to institute regulation to manage conflicts.

We identify four progressive levels of regulation:

– Negotiation generally includes multiple offers and rejections before acceptance, and is the primary level at which conflicts may be resolved. The rejection of an offer is a formal communication of a conflict.

– Mediation may apply to agents that cannot reach agreement through negotiation. This service would be delivered by impartial agent(s) with the aim to improve the dialogue and awareness between conflicting agents.
– Arbitration may be employed in the event of failure of mediation. Impartial agent(s) provide a ruling for the conflicting agents to abide by. Of course, even arbitration is a matter of agreement!
– Probation may be enforced on agents who fail to abide by an arbitration decision. This may result in a loss of privileges or other form of social ostracism.

Of course all of these regulatory roles may be performed by the collective itself without a ruling class as per Section 3.1, rather than by a separate ruling class as is the case with Wikipedia.

3.8. System design

Guidance for the human engineering of conventional military systems [8] does not consider the machine capabilities described in this paper and their commensurate human integration (Section 3.4), nor does it provide for contracting systems development based on new and emerging methods of software engineering.

The design of UC2 systems might be considered in several phases. An initial phase would be implementation of human networking support, including distributed mission agreement and social coordination protocol technologies. A second phase would be progressive implementation of agent capabilities, followed by a third phase of human-agent integration, which would necessarily revisit the first phase, and so forth.

3.8.1. Requirements – Preferential and critical

Software agents are a key capability of UC^2 systems. Agents are a relatively new paradigm for adaptable information systems. They are: a means for reducing the communicative gap between human and computer systems [14]; key processors to realise distributed (high-level) data fusion [15]; social coordinators and at the very least a conceptual structure for the design of complex information systems. However, they will introduce new challenges to security and assurance. The design of UC^2 systems may be based on the notion of preferential and critical requirements [13].

Preferential requirements are addressed differently at each management level of Section 3.7.1, according to the increasing degree of autonomy.

– Individual: routines and real-time scheduling.
– Platform: a temporal shortest path problem in which the commander must satisfy intent, given routines and beliefs.
– Team: multi-agent collaboration based upon joint mission intentions.
– Society: negotiation to share resources for collective benefit.

Preferential requirements may be achieved by coding the appropriate strategies into agent designs.

The satisfaction of critical requirements, that specify behavioral boundary conditions of the system, typically by citing failsafe conditions, requires verification.

> The critical requirements for a UC^2 system identify constraints which the system *must comply with at all times*. The critical requirements specify the behavioural boundaries of the UC^2 system, typically by citing fail-safe conditions. As all of the critical requirements must be met, a measure of guarantee is inappropriate [13, p. 39].

If verification is necessary for the deployment of a specific system configuration, then a formal proof of system design would be required. Our approach to design is to define adaptable capability that can be adaptively combined, while ensuring that certain boundary conditions must be met.

The lack of quantitative formal design elements in architecture frameworks approach (Department of Defense Architecture Framework Working Group, 1997), is addressed by augmenting this descriptive approach with an agreed set of formal semantic primitives (Fig. 4). These primitives allow for a mathematical logic base [18,23] to underpin the behaviour of aspects of a system. This will mean there is no "translation" between the descriptive architectural products and the design of detailed behaviour models of networked operations (where compromises and misinterpretations can sneak in). These primitives would facilitate the design and monitoring of a system. Moreover, in the spirit of Section 3.3.2, these primitives will need to be semantically represented within the machines, so that their meaning is accessible to others, and so that alternative theoretical frameworks are comparable with it.

3.8.2. Capability development – Partial design contracting

In describing the *use* of a system, the authors have proposed an adaptive, devolved, automation integrated, distributed, decentralised, non-linear approach. In this section the authors argue that these same sentiments equally apply to the *design* of the system. In the current military mindset, *design* and *use* are different things. Under this framework, the distinction between design and use becomes blurred. In seeking to *use* capability, it may prove necessary to adapt intent to generate new capability, and so capability design becomes an ingredient of capability use.

The development of a system would not follow the conventional linear waterfall model with complete, detailed (and consequently out-of-date) system requirements, but involve a non-linear development based on a model of "partial design" contracting. It is non-linear because deliverables span the spectrum of specification through to solution, without starting at specification and ending with solution, but by combining components determined by the competencies of capability developers, and the unique characteristics of the capability to be addressed. This shift in mindset brings with it a stronger ambition for adaptive capability, composable capability and shorter development times. Techniques of user-centred design [22] and agile programming [1] are illustrative of progress towards this.

It is inevitable that the user will become the requirements developer, capability composer, deployer and maintainer of systems. Web services provide an extant example of how capability development should function in the future.

4. Conclusion

In brief form, we have presented the following principles.

Adaptations in transport and telecommunications are adapting the influence distance has on presence, which is adapting organisational function toward strategic alliances, which is adapting organisational structure toward networks, which is adapting both the impact on localised events and the pace of change.

Decision Devolution enables the social collective to decide, rather than governing individuals, in order to benefit from the diversity of expertise.

Ubiquity of C^2 offers extreme robustness through agreements between similar, rather than identical, C^2 capabilities on every platform.

Automation provides the basis for ubiquity by extending intrinsic human capabilities with automated semantic and cognitive decision makers and aids.

Integration between people and machines is managed through mixed initiative strategies and by equipping cognitive machines with storytelling technologies.

Distributed locations allow seamless virtual integration with the robustness of physical diversity **and Decentralised** intent provides unity through mission agreements with robustness through a diversity of underlying intent.

Social Coordination among people and machines in a collective can be flexibly achieved through automated social agreement protocols and social policies.

Management provides bounds on social coordination, through use of levels which naturally arise from commonalities of location and intent, and through forms of social regulation.

Design attends to preferential and critical requirements while embracing capability design as an integrated part of capability use.

References

[1]   k. Beck, *Extreme Programming Explained: Embrace Change*, Boston, MA: Addison-Wesley.

[2]   M.E. Bratman, D.J. Israel and M.E. Pollack, Plans and Resource-Bounded Practical Reasoning, *Computational Intelligence* **4**(4) (1988), 349–355.

[3]   D. Browning, ed., *Philosophers Of Process*, Random House Inc., New York, 1965.

[4]   W.P. Burgers, C.E.L. Hill and W.C. Kim, A Theory of Global Strategic Alliances: The Case of the Global Auto Industry, *Strategic Management Journal* **14** (1993), 419–432.

[5]   G. Calbert, P. Smet, J.B. Scholz and H.-W. Kwok, *Dynamic Games to Assess Network Value and Performance*, Proceedings of the Australian Joint Conference on Artificial Intelligence, Perth, 2003.

[6]   J. Conklin, *Dialog Mapping: Making Sense of Project Fragmentation*, Chapter 2. forthcoming, 2004. See: http://www.cognexus.org.

[7]   Department of Defense Architecture Framework Working Group, *DoD Architecture Framework Version 1.0.* Washington, DC: Department of Defense, 2003.

[8]   Department of Defense, *Human Engineering Program Process and Procedures*, MIL-HDBK-46855A, 17 May, 1999.

[9]   P.S. Dodds, D.J. Watts and C.F. Sabel, Information Exchange and the Robustness of Organisational Networks, *PNAS* **100**(21) (October 14, 2003), 12516–12521.

[10] M.R. Endsley, *Design and evaluation for situation awareness enhancement*, in Proceedings of the Human Factors Society 32nd Annual Meeting, Santa Monica, CA: Human Factors Society, 1988, 97–101.

[11] K. Fischer, J.P. Muller and M. Pischel, Cooperative transportation scheduling: an application domain for DAI, *Applied Artificial Intelligence* **10**(1) (1996), 1–33.

[12] M. Hammer and J. Champy, *Reeengeering the Corporation*, Nicholas Brealey, 1993.

[13] D.A. Lambert, *Ubiquitous Command and Control*, Proceedings of the 1999 Information, Decision and Control Conference, IEEE, Adelaide, Australia, 1999, 35–40.

[14] D.A. Lambert, *Advisers with Attitude for Situation Awareness*, Proceedings of the 1999 Workshop on Defense Applications of Signal Processing, edited by A. Lindsay, B. Moran,

J. Schroeder, M. Smith and L. White, La Salle, Illinois, USA, 1999, 113–118.

[15] D.A. Lambert, *Grand Challenges of Information Fusion*, Proceedings of the 6^{th} International Conference on Information Fusion. Cairns, Australia, 2003, 213–219.

[16] D.A. Lambert, *Automating Cognitive Routines*, Proceedings of the Sixth International Symposium on Information Fusion, Cairns, Australia, 2003, 986–993.

[17] D.A. Lambert and J.B. Scholz, *A Dialectic for Network Centric Warfare*, Proceedings of the 10^{th} International Command and Control Research and Technology Symposium (ICCRTS), MacLean, VA, June 13–16, 2005.

[18] D.A. Lambert, *Formal Theories for Semantic Fusion*, Proceedings of the 9^{th} International Conference on Information Fusion, Florence, Italy, 2006.

[19] C. McCann and R. Pigeau, (eds), The Human in Command: Exploring the Modern Military Experience, Kluwer Academic Publishers, 2000.

[20] H. Mintzberg, Crafting Strategy, reprinted in *The Strategy Process: Concepts, Contexts, and Cases*, 3^{rd} ed., H. Mintzberg and J.B. Quinn, eds, Prentice-Hall, New Jersey, 1996, pp. 101–110.

[21] H. Mintzberg, The Design School: Reconsidering the Basic Premises of Strategic Management, *Strategic Management Journal* **11** (1990), 171–195.

[22] D. Norman and S. Draper, eds, *User-Centered System Design: New Perspectives on Human-Computer Interaction*, Lawrence Erlbaum Associates, 1986.

[23] C. Nowak and D.A. Lambert, *The Semantic Challenge for Situation Assessments*, submitted to the 8^{th} International Conference on Information Fusion, 2005.

[24] D. Perugini, D.A. Lambert, L. Sterling and A. Pearce, *Distributed Information Fusion Agents*, Proceedings of the Sixth International Symposium on Information Fusion "Fusion 2003", Cairns, Australia, 2003, 86–93.

[25] P.M. Senge, *The Fifth Discipline: The Art and Practice of the Learning Organisation*, Doubleday Currency, New York, 1990.

[26] R.G. Smith, The contract net protocol: high level communication and control in a distributed problem solver, *IEEE Transactions on Computers* **C-29**(12) (1980), 1104–1113.

[27] G. Tesauro, Temporal Difference Learning and TD-Gammon, *Communications of the ACM* **38**(3) (March 1995).

[28] A.A. Thompson and A.J. Strickland, *Strategic Management*, Irwin Press, 1995.

[29] United States Department of Defense, 2001, *Department of Defense Dictionary of Military and Associated Terms*, Joint Publication 1-02, 12 April (as amended through 22 March 2007), p. 101.

[30] US Congress, 2000, *Unmanned advanced capability aircraft and ground combat vehicles*, Department of Defense Authorizations: Title II Research, Development, Test and Evaluation – Subtitle B, Section 217. 106^{th} Congress, 2D Edition, Calendar No. 543, S.2549, Report 101–292.

[31] J.W. Walker, *Human Resource Strategy*, McGraw-Hill Inc, New York, 1992.

[32] D.J. Watts, *Six Degrees – The Science of a Connected Age*, Vintage, 2004.

[33] L. Zhang, L. Falzon, M. Davies and I. Fuss, *On Relationships between Key Concepts of Operational Level Planning*, Proceedings of the 5^{th} International Command and Control Research & Technology Symposium (ICCRTS), Canberra, ACT, 2000.