

TEMPORAL CONVERGENCE FOR KNOWLEDGE MANAGEMENT

Christopher Phillip Martin

Land Operations Division,
Defence Science and Technology Organisation,
Department of Defence,
Australian Government.

Wayne Philp

Land Operations Division,
Defence Science and Technology Organisation,
Department of Defence,
Australian Government.

William P Hall

National Fellow,
Australian Centre for Science, Innovation and Society,
University of Melbourne.

ABSTRACT

Time and knowledge have tended to be conceptualised in conventional knowledge management systems as either 'timeless' recordings of procedures, or time-stamped records of past events and states. The concept of temporal convergence was previously developed to help apply knowledge-management theory to complex military processes such as commander's intent, shared situation awareness, and self-synchronisation. This paper clarifies the concept and introduces several others in forming a framework to assist discussion and exploration of the types of knowledge required for complex endeavours, such as warfighting, characterised by opposition and uncertainty. The approach is grounded in a pragmatist philosophy and constructivist epistemology. Argument proceeds along mathematical lines from a basis that the types of knowledge most valuable to goal-directed agents in uncertain environments can be modelled as directed graph topologies. The framework is shown to be useful in describing and reasoning about the knowledge requirements and prerequisites for distributed decision-making through the sharing of situational knowledge and common intentions, with practical application to the planning and execution of operations. To the designers of knowledge management systems seeking to address this space, it presents a challenge that cannot be addressed merely by construction, storage, search and retrieval of documents and records pertaining to the past.

INTRODUCTION

Business process improvement focuses on the recording and analysis of processes in hopes of making them more efficient. This paper extends into the realm of business processes characterised by uncertainty and opposition to address the question, “What is the nature of knowledge specific to such processes?” In so doing, we propose an epistemological framework which clarifies *temporal convergence* and the *time-value of knowledge* (Dalmaris, Hall and Philp, 2006) and introduces several related concepts. We also show how this framework illuminates the knowledge requirements and prerequisites for distributed decision-making through the sharing of situational knowledge and common intentions.

The paper begins by establishing the psychological and mathematical foundations for the theory to follow. The theory proper commences by presenting two contrasting perspectives an agent may hold regarding how to relate to its own possible futures, and concludes by describing processes internal to the agent that show complex interrelations between these two perspectives. An application of the theory follows, giving insight to some concepts of knowledge management in military affairs. Further implications of the theory are drawn for processes of learning and decision-making. The paper concludes with some general suggestions for knowledge management practitioners.

MATHEMATICAL FOUNDATIONS

We take it as given that a business faced with a problem will engage in some process in an attempt to solve the problem. A well designed knowledge management system should play an active role in the problem-solving process through assisting the assimilation, filtering, alerting, accessing and building of knowledge. For many business problems, solution processes are already well-known to someone and so the primary purpose of a knowledge management system in such cases is to record that process knowledge and enable others to reproduce it, correctly and efficiently.

In this paper, we consider processes which can be described in terms of a system undergoing transitions from one discrete internal state to another. Some state transitions occur due to external forces impinging upon the system, perturbing it. Other state transitions occur as the result of the activity of components within the system. Components may in turn be seen as systems in their own right with their own internal states.

Not considered in this paper are other classes of process which can be better described by other formalisms such as linear or dynamic equations. Some problems can be addressed through optimisation techniques in which a fully known, closed set of parameters represent the dimensions of a coordinate space and the objective of the problem solving process is to search the coordinate space as efficiently as possible to find a set of coordinates—a combination of parameter values—that best satisfies some criteria; that is, an optimal solution. It is important to recognise that these problems have a ‘timeless’ quality; unless time is a parameter, solutions are otherwise independent of the passage of time. (This is not to say that we do not want them solved as quickly as possible.) If an optimal solution’s coordinates were known in advance, then that optimal solution could be arrived at immediately. That is, there is no order to the points in the coordinate space and so any point in that space can be immediately accessed, independently of any other point in that space, simply by setting the parameters to appropriate values. Furthermore, the coordinate space is ‘timeless’ in that the parameters themselves never change; their values do, but the space itself will not suddenly acquire a new parameter, an extra dimension.

This paper addresses a class of problems that *require* transitioning through a sequence of intermediate states in search of one that satisfies some solution criteria. That is, *even if* all the relevant parameters and the solution state were known in advance, one could not necessarily attain that state immediately by ‘leaping’ directly from the initial state to the desired goal state. These problems have an inherent notion of time as a product of two factors: state transitions not being instantaneous; and changes in the set of relevant parameters concomitant with changes in state. In the worst cases, future states are indeterminate and hence fundamentally unknowable. In this paper, we address this type of problem space which can be represented as a directed graph topology.

In the following discussion, in the interest of clarity some examples will be based on a problem space in which the goal state is defined as a specific point within a directed graph topology. However, we realise that this cannot adequately exemplify all problem spaces. Rather, goals are more generally definable in functional terms. A problem space may contain many points which satisfy a given function. A functionally defined goal state may include a location parameter, as will some examples in the following discussion. Alternatively, the function that defines a goal state may include a time parameter such as a deadline. It is a functionally defined goal that allows one to say, “I don’t know what it looks like, but I’ll know it when I see it.” The parameters of the function may be interrelated in complex ways; the more complex the interrelation, the more difficult it becomes to predict the end-state of a sequence of parameter adjustments. Nevertheless, we presume that the use of simpler examples will not detract from the basic argument.

PSYCHOLOGICAL FOUNDATIONS

We follow the position of Jean Piaget (1977) in the view that knowledge held by an individual is subjective and constructed from personal experiences and interpretations, resulting in personal (and unique) *schemata*. Furthermore, in the style of Thomas Kuhn (1970), both individuals and groups can simultaneously maintain and selectively employ seemingly disparate—even conflicting—*schemata* in parallel, such as classical and quantum theories in the mind of a physicist.

The following discussion also presumes the existence of *goal-directed intentional agents*. That is, an agent’s behaviour need not always be goal-directed but the agent has at least a presupposed desire to continue existing and a presupposed belief in its own ability, given opportunities, to take action in reshaping itself and its environment to satisfy its desires. Such an agent may be attributed with an *internal locus of control orientation* with a *high sense of self-efficacy* (Rotter, 1966; Bandura, 2001).

An agent’s *situation knowledge web* (Bickhard, 1993) consists of information about the ongoing interactive properties of the environment, which interactions have been engaged in, which sorts of interactions are now available, and which might become available if certain other interactions are engaged in first.

“In more complex organisms, such indications may branch and iterate into vast webs of organizations of conditional indications of interactive potentiality. This web is the basic resource for complex representation. Such a web constitutes the organism’s situation knowledge: knowledge of the potentialities constituting the current situation. A web of any complexity will of necessity be ongoingly changing, due to processes in the environment, new interactions and differentiations, and just from the passage of time: the environment is not, in general, static. The processes of creating, maintaining, and updating situation knowledge are those of apperception.” (Bickhard, 2007, p. 4)

PERSPECTIVES ON THE FUTURE

We contrast two perspectives on the future which could be adopted for decision-making: a view of the future as *stochastic*, and a view of the future as *intentional*.

The *stochastic future* is anticipated by a divergent projection of the immediate past and present onto possible future states. The near future may be predicted, based on known causes and effects, with the assumption that environmental conditions will not substantially change. Of course, the stochastic future tends to be quite near-term because environmental conditions *do* change continually over a continuum of time scales. Some changes are highly predictable; some are imperceptible (or at least nearly so) but highly significant, such as an increase in ionising radiation; and some are unexpected and catastrophic. An agent's relation to change must be somewhat opportunistic. It is important for the following discussion to note that the term *stochastic future* relates to an agent's situation knowledge web. That is, it is about the ability of an agent to anticipate potentialities in the stochastic unfolding of events in its world (including itself). In a less formal manner of speaking, an agent's stochastic future is its knowledge of what lies ahead. This knowledge will usually be incomplete and error-prone, requiring continual reconstruction.

In contrast, an *intentional future* is anticipated by a belief in one's ability and opportunity to take action in shaping oneself and the environment to realise a predetermined goal-state. The timing of the goal could be in the next second, or the next decade, or unspecified; what matters is that at the present moment, the goal has yet to be realised but at least some preconditions can be defined.

Taken together, these complementary perspectives may support many real-life decision-making processes categorised by choosing actions in the present to both: (a) manage the unfolding moment-by-moment situation; and (b) effect change in oneself and one's environment to shape future possibilities to converge on a goal-state. The following section presents a detailed example that, though simple and idealised, illustrates how relations between the *stochastic* and *intentional* perspectives on the future can be subtle and complicated.

Stochastic and Intentional Futures

Consider a very simple state-space of a type commonly found in games such as checkers. The simplest such space is described by the binomial tree, also known as Pascal's Triangle. In such a state-space, states are represented by nodes in a directed graph. From any node, one may transition to only one of two other nodes; one chooses heads-or-tails, left-or-right.

Figure 1 shows such a state-space with a circular token on a node at the bottom of the figure representing the start state and a node at the top representing the goal state. The objective is to weave the token at *Start* upwards from node to node, choosing either the left or right branch at each step, until either the *Finish* node is reached or it becomes impossible to reach without backtracking (at which time one has failed).

If one considers all possible six-step journeys from *Start*, it can be calculated that there are exactly sixty-four (64) distinct paths and all of them will take the token to a node on the *Finish* row. Each node in the diagram is labelled by the number of paths that reach it from the *Start* node. However, only twenty (20) of those paths will arrive at the *Finish* node; the relevant branches are highlighted in the figure. If the choice to go left or right is based on the toss of an unbiased coin, then one can calculate the probability of reaching *Finish* from *Start* as:

$$P(\textit{Finish} \mid \textit{Start}) = \frac{20}{64} = 0.3125$$

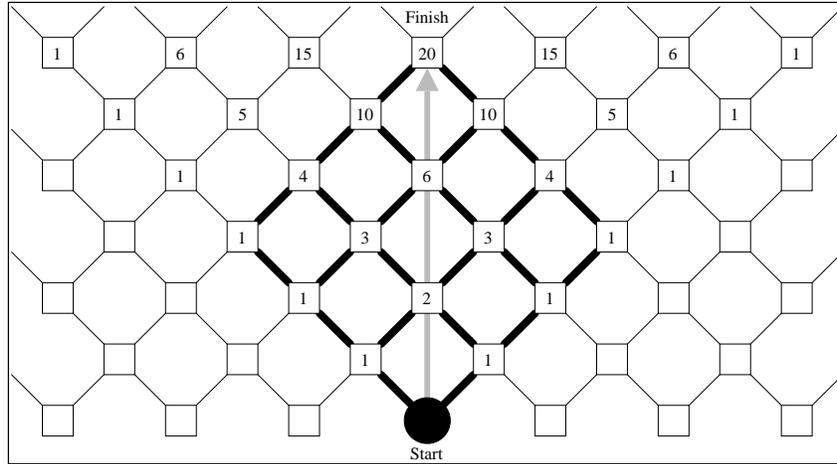


Figure 1: The number of paths to each reachable node in a binomial state-space.

By inverting the triangle as shown in Figure 2, it can be seen that of all nodes on the *Start* row, the *Start* node itself offers the most paths — twenty (20) — to the *Finish* node. (Readers may have noticed symmetry here in that a randomly chosen path from the *Start* node is most likely to reach the *Finish* node, and the *Finish* node is most likely to be reached from the *Start* node.) All possibilities within the stochastic future, from the *Start* node as ‘now’, are revealed in. All possibilities of realising an intentional future at *Finish* are revealed in Figure 2. The highlighted nodes and branches in

and Figure 2 represent the **intersection** between stochastic and intentional futures respectively. This is a key observation and will return to focus later in the discussion.

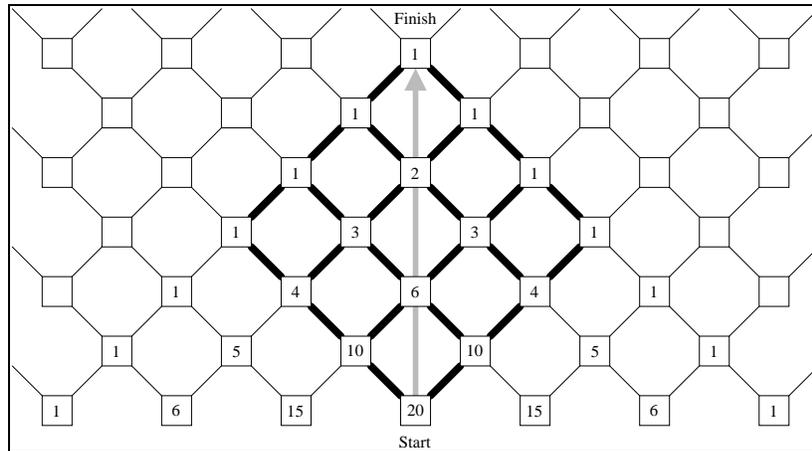


Figure 2: The number of paths to a given node from all possible starting nodes.

It may be intuitive to think that as one approaches a goal state, one's degrees of freedom will inevitably decrease and thus increasing vigilance will be required to ensure that the goal state is reached. Consider Figure 3, which shows the token part-way through a journey. It can be seen that there is only one (1) path that will lead to the *Finish* node; the other seven (7) paths will inevitably lead away from the *Finish* node. In this case, the probability of reaching the *Finish* node has indeed decreased to $(\frac{1}{8} = 0.125)$. It is also worth noting that there are now two most equally likely end-states, neither of which is the *Finish* node. That is, outcome probabilities are changing as the token moves through the state-space.

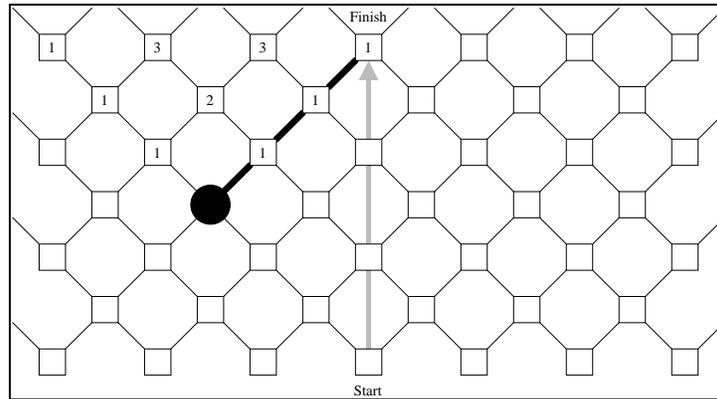


Figure 3: Only one path out of eight will reach the desired goal.

However, even though degrees of freedom may decrease with increasing temporal or spatial proximity to the goal, it does not follow that the risk of missing the goal will necessarily increase. Rather, a decrease in degrees of freedom may counter-intuitively afford a decreased vigilance. Consequently, it follows that there can be cases in which the goal state can be considered attained long before it is reached—a foregone conclusion—if, in fact, there are no degrees of freedom between the current and goal states. Consider Figure 4 in which the token has been fortunate to advance to within one step of the *Finish* node. In this case, it can be seen that there are only two (2) possible actions, step-left or step-right, one (1) of which leads to the *Finish* node. That is, as the token has approached the goal, the probability of reaching it has **increased** to $(\frac{1}{2} = 0.5)$.

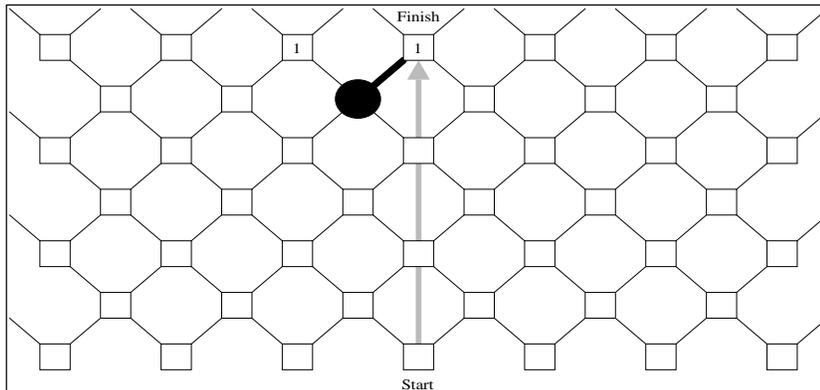


Figure 4: Only one degree of freedom but a greater likelihood of reaching the goal.

This simple thought experiment provides a counter-example to the intuitive notion that as one approaches a goal state, decreasing degrees of freedom will necessarily demand increasing vigilance. That is, the mere decrease in degrees of freedom associated with an impending goal does not *per se* cause a decrease in the probability of reaching that goal.

It was said above that a goal state can be considered attained when there are no degrees of freedom between the current and goal states. However, this may not be true when there is opposition to attaining the goal. Consider a hypothetical game in which an opponent is seeking to prevent the token reaching the *Finish* node by blocking the token's movement options. If the opponent blocks the token's movement down the right branch off the *Start* node, then it can be calculated that the probability of the token reaching the *Finish* node will remain at $\left(\frac{10}{32} = 0.3125\right)$; that is, counter-intuitively, it will have no effect on the probable outcome. However, the opponent could reduce the probability to zero with a single action simply by waiting for a situation as in Figure 3 or Figure 4 to arise. Although in Figure 4 the token has a 50% *a priori* probability of reaching the *Finish* node, blocking the right branch leaves the token with no way to do so under the rules of the game. Under such conditions, it matters not how favourably the token's weaving carries it toward the *Finish* node; it is simply impossible for the token to reach it from the very outset. Probabilistic decision-making is futile when the probabilities can be forced to zero or one.

To conclude: on approaching a goal state with decreasing degrees of freedom, the probability of reaching it increases under no opposition; but should opposition be encountered, the effectiveness of that opposition also increases.

The Intersection of Future Perspectives

Figure 5 depicts the common experience of not being able to predict sufficiently the stochastic and intentional futures. In such cases, the stochastic and intentional futures are being continually (re)constructed. The agent has advanced its token through a few state transitions. At each state in the sequence, it has been able to see only a step or two ahead. Furthermore, it can only foresee the last step or two from which it might reach the goal state. However, at the moment depicted in the figure, the stochastic and intentional futures have not intersected. We can now more precisely define what was called the 'point' of *temporal convergence* (Dalmaris, Hall and Philp, 2006) as the moment when an agent identifies an intersection region of the stochastic and intentional futures. In contrast, *temporal divergence* applies to the moment when it becomes apparent that the stochastic and intentional futures cannot intersect. (However, all is not lost if the agent can survive the failure to reach the desired goal. Instead, surviving a case of temporal divergence provides the agent with an opportunity to learn. This will be addressed in a later section.)

The phrase, "it becomes apparent" was used above; but apparent to whom? This is especially problematic for the definition of temporal divergence. It may be apparent to an omniscient observer, but not an involved agent, when the moment of temporal divergence has passed and failure is inevitable. In relation to Figure 5, if the token was to take two more left branches, we may imagine filling in the graph to show that temporal divergence will have occurred at the moment of taking the second left. An agent may continue to engage in activity for quite some time before they realise that temporal divergence has occurred and failure is thus inevitable. In many real-world cases, the question of whether stochastic and intentional futures temporally diverge may be undecidable. Recognising temporal convergence requires, at a minimum, finding a single point of intersection between the two futures. However, recognising temporal divergence requires either an exhaustive search of both futures, or a mathematical proof, to ensure that there can be no point of intersection. How can one be sure that an intersection is not "just around the next corner?" (In such cases, time limits become important parameters of functionally defined goal states.) In case of confusion, one

should be prepared to specify whether one is speaking of a moment of temporal convergence or divergence in `absolute' terms from the perspective of some detached, omniscient observer; or relative to the perspective of the agent(s) involved.

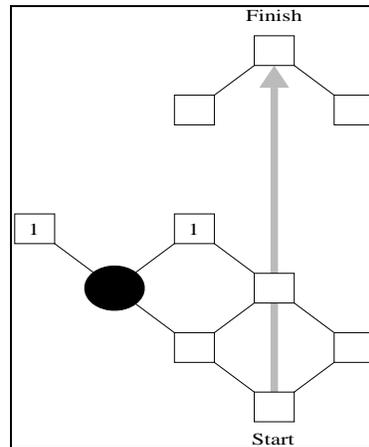


Figure 5: Stochastic and intentional futures which have not yet intersected.

Finally, note that it is possible for temporal divergence to occur *after* temporal convergence. The “second left” mentioned above could occur as in Figure 4 where clearly the two futures have intersected but, by chance, a single action leads to instant temporal divergence. Temporal convergence is about finding the intersection between stochastic and intentional futures. Substantiating the intentional future is another matter because the intersection changes as moment-by-moment the anticipated potentialities of the stochastic future collapse into the experienced present. Temporal divergence can occur at any time.

PROCESSES OF TEMPORAL CONVERGENCE AND DIVERGENCE

As well as moments of temporal convergence and divergence, one may also speak of processes of temporal convergence and divergence. A process of temporal convergence is one which seeks to construct a non-empty intersection region between stochastic and intentional futures; that is, it seeks a moment of temporal convergence. A process of temporal divergence is one which seeks to ensure that a non-empty intersection region between stochastic and intentional futures cannot be constructed; that is, it seeks to preclude moments of temporal divergence.

A single agent may have the computational resources required to be simultaneously constructing both stochastic and intentional futures. However, in some contexts there may be two agents involved where each agent's computational resources are sufficiently consumed by the process of constructing just one of the futures. In such cases, temporal convergence will occur through both agents engaging in a dialogue as they search for a point or region of intersection between their two respective futures, or until it becomes clear that temporal divergence has occurred and the two futures cannot intersect.

Temporal divergence applies to both failure to attain one's desired goal, and prevention of an opposing force's attainment of its goals. By opposing forces, we mean that each force has its own intentional future which, at least in part, cannot intersect with the other's. On the other hand, there is a 'commonality' between the two forces' stochastic futures in that each one's stochastic future includes a representation of the other force and the other's stochastic future. Both forces are embroiled in the conflict and each is somewhat aware of the other. Note that the two stochastic futures may not intersect at all if either force is completely mistaken about the other's situation. Nonetheless, in times of conflict each force will seek to construct a functional stochastic future that includes the other force, and an intentional future purporting to be of the other force. Furthermore, each force from its own perspective is trying to temporally converge what it believes to be a common stochastic future with its own intentional future while simultaneously trying to temporally diverge the common stochastic future with the other's intentional future. In this there is much potential for error because each force is trying to construct five things: a common stochastic future (SF); their own intentional future (IFself); the other's intentional future (IFother); the temporal convergence of SF and IFself; and the temporal divergence of SF and IFother. Any of these constructions may be insufficiently complete or in error.

APPLICATION TO MILITARY KNOWLEDGE MANAGEMENT

The conceptual framework described above has general application to any strategic business operation, but perhaps the best source of apposite examples is the military domain. Military operations are characterised by danger, high-tempo and uncertainty, and military practitioners have responded with command hierarchies, formal documentation and intensive training in standardised procedures. Publicly available military policy documents contain terms whose meanings and implications can be clarified by recasting into our proposed framework. This in turn may illustrate the applicability of such concepts to other strategic business operations. Three public-domain military concepts will be considered here:

- Commander's intent;
- Shared situation awareness;
- Self-synchronisation.

Commander's intent

Commander's intent, in our terms, is the declaration of an intentional future by the Commander. It is to be expected that the intentional future of a military campaign will be more difficult to define than that of our simple Pascal's Triangle example above. However, the most difficult aspect of this is revealed when one considers that it is not possible for a commander to directly imprint his intentional future in the minds of his subordinates. Instead, each subordinate must engage in a private process of *situation knowledge web* construction that will, hopefully, result in a *situation convention* (Bickhard, 1980, 1992, 2003) that includes the same intentional future as the commander's. Military tradition and training practices have long addressed this problem through uniform training, and by raising up commanders through the ranks, in the hope that uniformity of behaviour and experience will produce uniformity in situation conventions and processes of situation knowledge web construction (in our terms)... but this is by no means guaranteed.

Consider a situation in which a Commander gives a verbal briefing to his command staff who then prepare a series of written briefs for field commanders. The Commander's verbal briefing is the result of a highly complex constructive process drawing upon his private situation knowledge web and his personal perspective on the interaction with his command staff. Consequently, an intentional

future in each staffer's private situation knowledge web is constructed by a similar but uniquely personal process. Thus with a Commander and N staff, we have $N+1$ processes all engaged in the construction of private situation knowledge webs and public artefacts (verbal utterances, hand-written notes, sketches, etc.) through which the processes interact. Another set of N processes is then engaged to construct the written briefs for the field commanders. Strictly speaking, a written brief is not actually the Commander's Intent *per se*. Rather, it is an artefact constructed by a command staffer in the hope that it can be used by a field commander as a resource in constructing an intentional future that intersects the intentional future of the Commander.

We should be amazed that so many distinct processes can produce intentional futures that sufficiently intersect. This is far more than a question of mere data-replication accuracy. Analogies are often drawn with computers exchanging data over networks but such analogies are misleading because computers are far more uniform than organic entities such as people. Furthermore, the capabilities of computers to sense and appreciate their environment are far more limited. That is, the universe of a computer is very simple. Even so, data that was perfectly replicated can still be inappropriately processed. For example, a classic class of errors arises from confusion over what is called *endian-ness*. If a sending computer uses a numeric format that issues the *most* significant bit first, but the receiving computer uses a numeric format that expects the *least* significant bit first, then it doesn't matter that the bits are perfectly replicated over the transmission medium; the overall result of the entire process is confusion. Thus, a focusing of effort on network construction is misguided; the focus of effort must be on the constructive processes within each node.

Shared situation awareness

In the definition of situation awareness as "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" (Endsley, 1988) we see a stochastic future. Shared situational awareness then exists only in as much as two or more agents can engage in similarly constructive processes with similar resources such that their stochastic futures are likely to have a high degree of intersection. As already described with regard to intentional futures, this is no mean feat. Within a command headquarters, shared situational awareness is the intended product of the continual back-briefing that occurs after the Commander's staff take his statement of intent and generate plans by which it may be achieved. As already mentioned, the Chief of Staff acts as the nexus through which the stochastic future is gradually constructed by the command staff shops.

When the staff believe that they have identified a temporal convergence between the Commander's intentional future and their shared image of the stochastic future, this is presented to the Commander as a plan to achieve his intent. In as much as the command staff have training and experience in common, and they share an immediate environment and resources, and there is an individual—the Chief of Staff—as a nexus, it is reasonable to expect that that there may be shared situational awareness. However, all this means is that the individuals have engaged in similarly constructive processes with similar resources such that their stochastic futures have a high degree of intersection. It does not mean that they are 'complete' or 'correct' in their understanding of the situation from an omniscient observer's point of view. Nor does it mean that they have not fallen victim to phenomena such as groupthink or polarisation; in fact, it is likely that they have (Whyte, 1989).

Finally, a field commander or a soldier who does not have the same constructive resources on which to draw in constructing their own stochastic future is not likely to share the situation awareness of the command staff.

Self-synchronisation

Self-synchronisation “involves subordinate commanders re-tasking their units so that the commander’s intent is maintained... [and] helps warfighters adapt to changing circumstances and allows them to apply a warfighting advantage more effectively” (Director General Capability Development and Plans, 2006, pp. 10–11). In this definition we see references to both stochastic and intentional futures; stochastic with regard to the warfighter’s changing circumstances, and intentional with regard to the commander’s intent. The effective application of warfighting advantage is about identifying and substantiating temporal convergence between the two futures. The term *self-synchronisation* is perhaps a misnomer. The concept is more about subordinate units interacting directly to achieve their commander’s intent without needing their commander’s control or mediation.

It is claimed that “shared situational awareness will enable self-synchronisation” (Director General Capability Development and Plans, 2006, p. 11). However, it is clear that the conditions necessary for shared situational awareness are a *proper subset* of the conditions necessary for self-synchronisation. Therefore, shared situational awareness is necessary for self-synchronisation but far from sufficient. Self-synchronisation between subordinates A and B, with Commander C, requires that:

- A and B each hold an intentional future that intersects with C’s;
- The stochastic futures held by A and B substantially intersect;
- A and B have each identified a temporal convergence between their respective stochastic and intentional futures;
- The temporal convergences of A and B substantially intersect;
- In the temporal convergence held by A, there is a representation of B. Ideally, B would also hold in his temporal convergence a representation of A.

All the above conditions are required in order for A to reasonably select an action to perform in anticipation of (i) B’s actions or needs, and (ii) satisfaction of C’s desired end-state. It should not be surprising if self-synchronisation is difficult to achieve given the highly demanding conditions required for it to occur.

TEMPORAL DIVERGENCE AND LEARNING

It was said above that surviving temporal divergence provides an agent with an opportunity to learn. In fact, a learning process could be engaged at any time that temporal divergence is even suspected although doing so may be costly in terms of processing resources. The following discussion outlines the nature of knowledge required for agents to learn from temporal divergence.

Knowledge for goal-driven agents

We can formalise our approach thus far by considering an agent (or decision maker) to be goal-driven when, in the simplest case, the agent performs some action to achieve a desired outcome. (Note, outcomes are always specified in terms that are internal to the agent. We take as trivial the case where the agent has already realised all desired outcomes, with no further action, or the agent is passive and merely waits for the desired outcome to occur through environmental flux.) If some actions lead to desirable outcomes, then by implication others do not. For some actions, there may be multiple possible outcomes of which only some are desirable. Whichever outcome is realised

may depend upon factors in the current situation defined as the totality of state with respect to both agent and environment. For our discussion, it does not matter where the influencing factors lie, in the agent or the environment, or both: what matters is that a desirable outcome will be realised from only some, not all, situations.

For an action-outcome pair, any given situation may be classified in simplest terms as leading to either success or failure. In order to perform this classification, the agent performs the action and notes whether or not the anticipated outcome is realised. In other words, action-outcome pairs serve to differentiate situational classes. This paper introduces the concept of bifurcation of the situation knowledge web as the outcome of a process by which an agent uses an action-outcome pair to differentiate a situational class. Bifurcations of the situation knowledge web are behind sentences that begin, "If I had known then what I know now..."

Consider the simplest possible case, as illustrated in Figure 6. At time ($t=0$) representing the agent's present, the agent projects the future performance of a particular action at time ($t=1$). Proceeding from that action are two parallel arcs to two different possible outcomes at time ($t=2$). However, which outcome will be realised? The agent's present situation at time ($t=0$) is undifferentiated with respect to the two outcomes in that the agent cannot predict which outcome will be realised at time ($t=2$). At time ($t=1$) the agent will perform the action. By time ($t=2$) there will no longer be any question. One of the outcomes, either A or $\sim A$, will have been realised and the agent will be able to bifurcate its situation knowledge web to determine which class of situation held at time ($t=0$). That is, at time ($t=2$) the agent will have the hindsight to know which situation it was in at time ($t=0$): either a situation leading to particular outcome A, given a particular action; or a situation leading to some outcome other than A (*i.e.* $\sim A$) given that same action. Figure 6 shows the stages of the bifurcation process from the outcome realisation back to the identification of the actual situation.

If there is no possibility of any undesirable outcome, then the agent has nothing to lose by trying the action. However, it does the agent no good to determine that it was unsafe to perform an action by performing that action! A risky situation for an agent can be characterised by three properties: (i) an action has either a distinctly desirable or distinctly undesirable outcome; (ii) the agent intends to attain the desirable outcome; and (iii) the desirable outcome cannot (to the agent's knowledge at that moment) be attained through alternative, safer actions. It becomes vital for an agent to somehow determine that the current situation is a member of the desirable-outcome class, as opposed to the undesirable-outcome alternative, before performing the action. (In practice, there could be a number of actions that the agent may perform. However, whether there is one action or many, the principle remains that a wise agent would determine the class of situation before acting.)

Suppose that Outcome A in Figure 6 is a desirable, safe outcome while Outcome $\sim A$ is an undesirable, unsafe alternative. The agent needs to determine at time ($t=0$) if it is safe to perform the Action at time ($t=1$). This may be achieved if the agent can identify some other action—with multiple possible outcomes, all of which are safe—to perform between times ($t=0$) and ($t=1$) such that the outcome realised will bifurcate the situation knowledge web. A stage in this process of 'active sensing' is shown in Figure 7.

The agent is choosing to perform Action X, not because either Outcome Y or $\sim Y$ is a desirable end in itself, but because of what the realised outcome will indicate for future action-outcomes. This is different to either passively waiting for whatever sensory input arrives, or applying some filter to all sensory input received. Rather, the agent engages in activity that is intended to elicit or invite a particular class of sensory input/information from the environment, where the absence of that resultant stimulation may be just as telling as its presence.

Note that what we have been calling the “current situation” is actually the ultimate outcome of the penultimate action in a (potentially very long) sequence of overlapping outcome-action-outcome triads. Bifurcation could be continued into the agent’s history of actions and outcomes, but this is not necessary. For the agent to predict the realisation of either Outcome A or $\sim A$, it is sufficient merely to stop bifurcation at Action X. Note that at time ($t=0$), Action X is yet to be performed. That is, the agent is able to make decisions for action entirely on the basis of future events without any reference to, or even awareness of, past or current situations.

Readers who are familiar with traditional rule-based expert systems should note that to encode a situation knowledge web would require more complex rule expressions than is traditional. The difference is illustrated in Table 1. In short, each traditional condition-action rule is bracketed by a preceding action and an expected post-action condition.

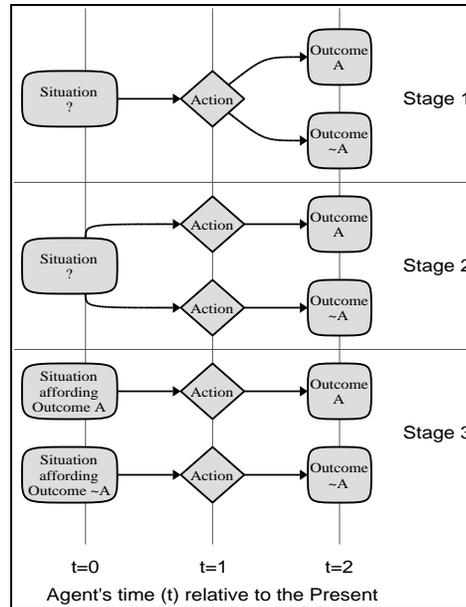


Figure 6: The process of situation knowledge web bifurcation.

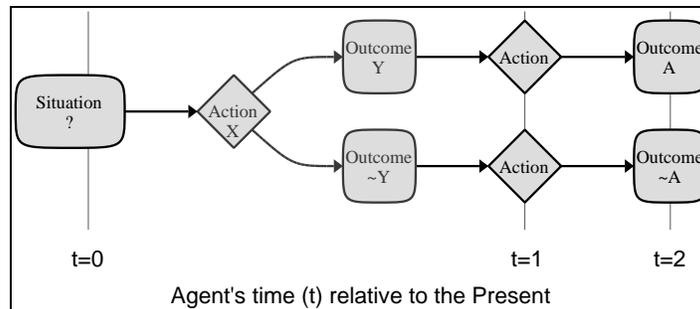


Figure 7: Active sensing.

Traditional	Situation Knowledge Web
If condition C holds, Then perform action A.	If action X was performed Resulting in condition C holding, Then perform action A Expecting condition Y to hold.

Table 1: Example rule expressions.

The anticipation of outcomes from actions allows the agent to detect when its representations are in error. If a situation knowledge web composed of linked action-outcomes is a representation of the world, in that it can differentiate classes of situation, then it is possible for an agent to detect for itself when its representation is in error. The agent does not need a homunculus, nor an external advisor, to know when an anticipated outcome was not realised. (Note, with one condition: the outcome must be specified only in terms from the domain of the agent's internal organisation.) The agent may not know what to do about the fact that an action produced an unexpected outcome, but it must at least recognise that there is an error to address, since the ability to detect that the world is not as expected is a fundamental requirement for any adaptive system.

IMPLICATIONS FOR KNOWLEDGE MANAGEMENT SYSTEMS

In any large organisation, we may find some parts in which there is rarely any doubt about the current situation, how it transpired and what to do next. We may also find other parts that frequently deal with novel or fluid situations. In the former case, situations are static or change in small degrees at slow time scales relative to the rate at which organisation members can adjust their situation knowledge webs. The organisation has the opportunity to create and disseminate procedural documentation via content management and portal systems. Long-serving members seem to know without thinking how to react to particular situations because their stochastic and intentional futures have been well-formed—they have ‘seen it all before’—possibly to the point of being rather fixed. But in the latter case, the threat of temporal divergence is common and members’ situation knowledge webs undergo frequent (though costly) bifurcation and expansion. Long-serving staff are valued not just for the complexity of their situation knowledge webs built up in the course of much experience, but also in the skilled and relatively efficient processes they have developed for situation knowledge web construction and bifurcation, the ability to create a situation knowledge web being more valuable than the image itself. People who develop the latter capability have been ‘learning how to learn’ (Bateson, 1973) and their fluency in constructing new knowledge is at least as valuable as the knowledge they have already constructed. Success is highly dependent on the speed and efficiency with which temporal divergence can be detected and addressed through situation knowledge web bifurcation processes.

If knowledge management systems are to be employed effectively beyond well proscribed domains, then beyond focusing on the storage, search and retrieval of documents and records with a focus on the past, they must begin to expressly support stochastic and intentional future concepts. Ideally, they will even come to support meta-processes by which the processes of situation knowledge web construction and bifurcation are enhanced. We recognise that it may be difficult for some readers to accept ‘knowledge’ as having a future aspect. Some philosophers hold that the future is epistemologically unknowable because: it cannot be experienced (e.g. empiricists); and induction from past events (or beliefs about past events) is prone to error, thus uncertain, thus unworthy of the verb *to know* (e.g. falsificationists). However, the kinds of knowledge in question here relate to one’s own and others’ future goals, and the likelihood of future events and how preceding actions and events may shift that likelihood. That is, we are not talking about knowledge of *the* future but knowledge of possible futures and of agents’ intentions for the future. Because this ‘future’

knowledge is never certain, knowledge management systems must accept continual reconstruction. The time-stamped record of a past event or state may later turn out to be fictional, but that the record was made and was believed to be factual at the time of articulation is another kind of valuable knowledge that has a lifetime distinct from the record itself as long as it features in the situation knowledge web of some agent. The promulgation of false truths and intentions is a commonly used method for inducing temporal divergence in an opponent's situation knowledge web, and that such falsehoods are available and amenable to that purpose is also a kind of knowledge. Data aggregation systems and alerting systems may allow the anticipation of events. Good search and retrieval on content management systems may allow those faced with unfamiliar situations to filter a large volume of prior experience and records for ideas and possibly matching situations. Better still for organisational knowledge management systems to support the continual construction, extension and sharing of members' situation knowledge webs (cf. Bosse *et al.*, 2005).

CONCLUSION

Dalmaris, Hall and Philp (2006) attempted to extend the Knowledge-Intensive Business Processes (KBPI) to incorporate the idea of the time-value of knowledge, defined as a qualification helpful in assessing the applicability of knowledge to solving problems with a variety of time horizons. This examination extended from the simple observation that knowledge requirements—knowledge objects, relationships and supporting infrastructure that make up business processes—for the typical knowledge worker such as a help desk customer support agent, were very different to those of a military commander or CEO, who selects actions in the present on the basis of future effects. This paper has presented a framework of terms and concepts to assist exploration and discussion of the types of knowledge required for complex endeavours, such as warfighting, characterised by opposition and uncertainty. We are not alone in arguing for the central importance of planning and human cognition as goal-directed constructive processes (Bryant, 2003). Knowledge management systems may play a very important role in that a commander needs to know, above all else, when a temporal convergence has been found between his intentional future and the stochastic future, or when temporal divergence has occurred and it is time for a strategic withdrawal and change of intent. In between is the grey area in which a commander does not know whether success is either possible or impossible, and he will draw on information from all sources and ask his knowledge management systems to help him substantiate temporal convergence or divergence. Finally, he will want to use knowledge management systems in the production and dissemination of artefacts that his subordinates may use to (re)construct their own situation knowledge webs which primarily represent beliefs about relationships between possible actions and events that may or may not happen.

REFERENCES

- Bandura, A. (2001), "Social cognitive theory: An agentic perspective", *Annual Review of Psychology*, Vol. 52, pp. 1–26.
- Bateson, G. (1973), *Steps to an Ecology of Mind*, Paladin, St. Albans, UK.
- Bickhard, M. H. (1980), *Cognition, Convention and Communication*, Praeger, New York.
- Bickhard, M. H. (1992), "How Does the Environment Affect the Person?", in Winegar, L. T. and Valsiner, J. (Eds.), *Children's Development within Social Contexts: Metatheory and Theory*, Lawrence Erlbaum Associates, Mahwah, NJ, pp. 63–92.

- Bickhard, M. H. (1993), "Representational content in humans and machines", *Journal of Experimental & Theoretical Artificial Intelligence*, Vol. 5 No. 4, pp. 285–333.
- Bickhard, M. H. (2003), "The Social Ontology of Persons", in Carpendale, J. I. M. and Muller, U. (Eds.), *Social Interaction and the Development of Knowledge*, Lawrence Erlbaum Associates, Mahwah, NJ, pp. 111–132. Online as <<http://www.lehigh.edu/~mhb0/SocOntPersons.pdf>>, last modified 06 February 2003 17:08:00 GMT, last accessed 13 March 2008 23:02:25 GMT.
- Bickhard, M. H. (2007), "Interactivism", online as <<http://www.lehigh.edu/~mhb0/InteractivismRout2Jan07.pdf>>, last modified 23 January 2007 20:03:50 GMT, last accessed 13 March 2008 22:29:50 GMT.
- Bosse, T., Jonker, C. M., Schut, M. C. and Treur, J. (2005), "Simulation and Analysis of a Shared Extended Mind", *SIMULATION*, Vol. 81 No. 10, pp. 719–732.
- Bryant, D. J. (2003), *Critique, Explore, Compare, and Adapt (CECA): A new model for command decision making*, Defence R&D Canada, Toronto.
- Dalmaris, P., Hall, W. P. and Philp, W. R. (2006), "The time-value of knowledge: a temporal qualification of knowledge, its issues, and role in the improvement of knowledge intense business processes", in *Proceedings of the 3rd Asia-Pacific International Conference on Knowledge Management (KMAP06)*, Hong Kong, 11–13 Dec 2006.
- Director General Capability Development and Plans (2006), *Explaining NCW*, Defence Publishing Service, Department of Defence, Commonwealth of Australia, Canberra, ACT 2600.
- Endsley, M. R. (1988), "Design and evaluation for situation awareness enhancement", in *Proceedings of the Human Factors Society 32nd Annual Meeting*, Human Factors Society, Anaheim, CA, pp. 97–101.
- Kuhn, T. (1970), *The Structure of Scientific Revolutions*, 2nd edition, University of Chicago Press, Chicago.
- Piaget, J. (1977), *The Development of Thought: Equilibrium of Cognitive Structures*, Viking, NY.
- Rotter, J. B. (1966), "Generalized expectancies for internal versus external control of reinforcement", *Psychological Monographs*, Vol. 80.
- Whyte, G. (1989), "Groupthink Reconsidered", *The Academy of Management Review*, Vol. 14 No. 1, pp. 40–56.