

ICCRTS 2006 Adapting C2 to the 21st Century
Track: Cognitive and Social Issues
Chair: Valerie Summers
Co-chair: Anne-Marie Grisogono
TITLE: A Philosophical Approach to Knowledge Management for a Networked Army
AUTHORS: Philp W., Aidman E., Gaidow S., Hughes D., Kirby B., Latham C., Lui,F.,
 Martin C., Passmore S., Stothard C., Turner T., Yen L. and Yue Y.
POC: Wayne PHILP
EMAIL: wayne.philp@defence.gov.au
PAPER NUMBER: I-036

Name of Organization: Land Operations Division, Defence Science and Technology Organisation (DSTO)
Complete Address: PO Box 1500, Edinburgh SA 5016, Australia.
Telephone: +61-8-82595625
E-mail Address: wayne.philp@defence.gov.au

ABSTRACT

This paper explains the concept of temporal convergence in the context of the time-value of knowledge. We explore epistemological aspects of knowledge/information management through notions of human cognitive awareness and mental models by which we may share situational knowledge and common intentions. We explore the temporal relationships of information (lag, reiteration, timeliness, lifetime and obsolescence) to decision-making based on shared intentions and desired end-states.

Temporal convergence is associated with a human construct of retrogressive linear time and sequential actions taken leading to a planned event; much like the timely execution of a launch sequence of necessary actions, activities and checks during the count-down to a rocket launch. An individual, or a team of like-minded individuals, may deem a desired future state or goal to be set firmly at a future time (as if it has already happened) while working in the present to make it so; correcting any misalignment in the immediate chain-of-events that may skew the planned way forward to the desired end state. Generally, temporal convergence describes the process of realising a human cognitive goal in the more “distant” future. Moreover, temporal convergence needs to be considered in parallel with temporal divergence where temporal divergence is a representation of the near future as an expanding set of possibilities originating from one fixed reality, the present moment.

We developed the concept of temporal convergence as a model to help us information- and knowledge-management theory to some of the complex military processes addressed under Network-Centric Warfare, such as: command direction and intent; team self-synchronization; endgame focused “Course of Action”; taking an effects-based [combined] approach; shaping the environment, and shared situational awareness.

INTRODUCTION

This paper describes the concept of *temporal convergence*¹, in the context of the time-value of knowledge. In so doing, we may better understand the base requirements for shared situational awareness and distributed decision-making.

Temporal convergence describes the process of real-time creativity, whereby the timely and sequential exercise of intentional action is applied to achieve a predetermined objective outcome or end-state; essentially, to “make it so”! This is in accord with Karl Popper’s notion of *Oedipus Effect*² in the sense of a prediction that, in being made, actually causes itself to become true. In our case however, it is the predictor’s timely actions (by a person, or a collective of people who believe in the prediction) that, after the prediction has been stated, actually cause the prediction to come true. If one is true to one’s word (absolutely reliable and determined), with the power and opportunity to make one’s claim happen (in real-time), then the realisation of a future prediction or intention may be deemed to be inevitable. The term *temporal convergence* embraces this concept, wherein the future becomes the central reality and the present merely the instrument of affecting that reality. The French phrase “fait accompli” (accomplished [future] fact) also well describes the concept. In general, *temporal convergence* applies to the process of realising a human cognitive goal in the more “distant” future.

However, if by frustration or contingency one cannot do as intended (e.g. a change in the environmental conditions, or a loss of power, or the opportunity), one simply fails. The desired end-state was not so when it was meant to be.

But, the mind-set of self-realised future intentions or goals cannot stand alone: there is also the present to deal with. Whereas *temporal convergence* readily applies to the “big picture”, it is not helpful in dealing with the issues and conditions that surround our present and near-future state or situation. We need a more linear cause-effect (deterministic) mind-set to help inform our decision-making in the present while we concurrently hold in mind our intended future situation. That is, we work our actions in the present to both: (a) effect/shape ourselves and the future environment to converge on a “more distant” desired end-state (*temporal convergence*), and (b) manage the unfolding moment-by-moment situation, based on the immediate past. This cause-effect sense of events, expanding before us in time as we gauge and project what is most likely to follow from the immediate past and present, is herein described as *temporal divergence*. *Temporal divergence* best applies to the “near” future, scoped by an event horizon based on our limited mental ability to process and understand the microcosm of our present circumstances and our ability to deterministically project the same into the future.

¹ 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”, Proc. 3rd Asia-Pacific International Conference on Knowledge Management (KMAP06), Hong Kong, 11-13 Dec 2006.

² “One of the ideas I had discussed in *The Poverty [of Historicism]* was the influence of a prediction upon the event predicted. I had called this the “Oedipus effect”, because the oracle played a most important role in the sequence of events which led to the fulfilment of its prophecy. . . . For a time I thought that the existence of the Oedipus effect distinguished the social from the natural sciences. But in biology too -- even in molecular biology -- expectations often play a role in bringing about what has been expected” Karl Popper (1976), “Unended Quest: An Intellectual Autobiography”, Open Court Publishing Company, La Salle, Ill., 1982.

An interesting condition exists when the event horizon for our mindset of *temporal divergence* approaches (as time progresses) the proximity of our intended future end-state. In this case hope reveals its substance in reality. It is at this stage that deliberate planning is replaced by immediate planning and we (the collective of people who believe in the prediction) wrestle with the probability of success or failure. The closer in time we get to our desired future state, the less degrees of freedom we have, the more important the “time to take effect” becomes and the greater the likelihood of failure if we over/under estimate the spatial or temporal effect of our actions. It may be that the rate or magnitude of environmental variation exceeds our capacity to respond such that we can no longer reach our desired future state in time... in which case we have failed. Although interesting, the meeting of divergent and convergent mental models will not be addressed further in this paper, although it is the subject of ongoing research by the authors.

The development of these concepts has been found necessary in the formalisation of Information Exchange Requirements (IERS) for military decision-making within the paradigm of Network-Centric Warfare (NCW). It is intuitively apparent that two different knowledge-bases are needed for NCW distributed decision-making; (a) detailed knowledge and change information about a particular tactical situation in a warfighting environment (for which *temporal divergence* well applies), and (b) relevant knowledge and change update about the cooperative and combined progress of different force-elements toward an operational-level military objective (applied *temporal convergence*). Both knowledge sets are spatially consistent, but rarely the same: the time (temporal) parameters and the scale/resolution of the two are quite different. The tactical situation requires detailed knowledge about the microcosm of local cause and effect, where that knowledge has a finite (mostly short) lifetime or relevance in a dynamically changing environment. The temporal nature of information reiteration/update, reliability, lifetime, lag/latency, sequence, timeliness and obsolescence are critical as immediate events unfold in a deterministic, but dynamic, way. On the other hand, the operational-level picture of “how goes the battle?” has quite different time-dependent parameters. If naively one were to impose an operational-level mindset upon soldiers engaged in a tactical mission through a “common operating picture” or an automated “shared situational awareness” system, or conversely, a tactical mindset was imposed upon the mental image of an operational-level commander, confusion would inevitably result.

Temporal convergence, as a mindset shared by a team in support of a commander’s intentions, is necessary to support NCW military concepts such as:

- Self Synchronisation,
- Effects based approach,
- Combined effects,
- Shared situation awareness,
- Focus of wills.

Prior to commencing an operational-level campaign plan, the military staff team, would initially have a shared future reality of the culmination point³ and the operational end-state. There would be an agreed course of action that would reflect how the assigned task force would behave to achieve that end state. Once forces were committed, the initiated plan has momentum and orientation toward the goal(s). The job of the operational commander is now to modify, in real-time, the trajectory of combined battlespace effects so as to maintain course and focus in the presence of enemy interference/deflection/obstruction. The instrument to actualise the real-time corrective bias effect is the tactical commander. The tactical commander is *temporally*

³ The point in a battle/war following which an enemy is no longer able to mount or maintain a successful offensive

divergent in his immediate actions and *temporally convergent* in effecting his commander's intent.

This paper attempts to formalise the concepts of *temporal convergence* and *temporal divergence* for ongoing application to the domain of NCW; especially, in the focus areas of Command, Control, Communications and Intelligence (C3I) and Intelligence, Surveillance, Target Acquisition and Reconnaissance (ISTAR).

CONCEPT

The authors follow the position of Piaget (1977), in *The Development of Thought Equilibrium of Cognitive Structure*, in the view that knowledge held by an individual is subjective and constructed from personal experiences and interpretations, resulting in personal (and unique) schemata or mental models. An individual's *situation image* 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 were engaged in first^{4, 5}. Further, in the style of Thomas Kuhn (1970)⁶ both individuals and groups can simultaneously maintain and selectively employ seemingly disparate mental models in tandem, such as, for the physicist, classical and quantum theories of physics.

We here describe two basic ways of viewing the world around us: the *causal situation* (deterministic projection of the past onto the present and possible future states) and the *intentional situation* (belief in one's opportunity and ability to take action in shaping the environment to actualise a desired outcome as a future end-state⁷). When held in tandem, these complimentary *Weltanschauung*⁸ (worldviews, see Sire, 2004) may support many real-life decision-making processes.

Causal Situation. In the *causal situation*, as time progresses, possible near-future states may be considered to diverge from the present and recent past (*temporal divergence*). In order to appreciate our present state (now) and to gauge what action we should take (now), we follow a series of deliberations:

1. OBSERVE: What is our current state with respect to the environment? What are the causes & effects that have brought us {where/when} to {here/now}?
2. ORIENT: Given our state {here/now}, in what state are we likely to be {there/then}, all variables remaining the same in our present environment?
3. DECIDE: Rationalise the projected change of state required, if any.
4. ACT: Initiate change action {here/now}, which by cause and effect is likely to bring about the required state else {where/when}.
5. Repeat

The near-future may be predicted, based on known causes and effects, with the assumption that environmental conditions will not substantially change. The future is viewed to be the most probable state of affairs after a step in time, projected from the present. This well represents the physical dimension, akin to the concepts of energy conservation and motion. The acronym of steps 1-4 above spells OODA and its reiteration is the OODA Loop as described by the late

⁴ Stojanov G. and Kulakov A. (2003) *Interactivist approach to representation in epigenetic agents*. In Prince, Berthouze et al., Eds. Proceedings Third International Workshop on Epigenetic Robotics: Modelling Cognitive Development in Robotic Systems 101, pages pp. 123-130, Boston, MA, USA.

⁵ Bickhard, M. H. (1998), *Levels of representationality*. Journal of Experimental & Theoretical Artificial Intelligence, 10(2), 179–215.

⁶ Thomas Kuhn (1970), *The Structure of Scientific Revolutions*, 2nd ed., University of Chicago Press, Chicago.

⁷ Internal locus of control.

⁸ *Weltanschauung* (German) meaning “look into the world.” Sire J. *The Universe Next Door: A Basic Worldview Catalog* 4th ed., InterVarsity Press, Downers Grove, IL. 2004.

USAF Col. John Boyd (1996)⁹ to describe generically how individuals or groups observe, understand and interact with the world to shape their action.

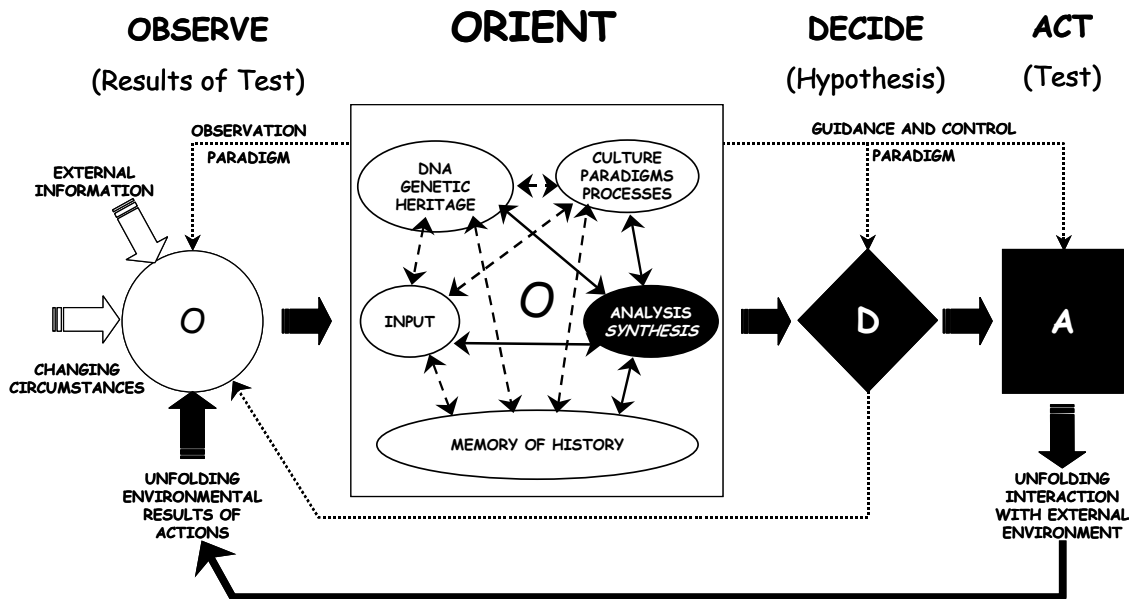


Figure 1: John Boyd's OODA Loop concept

Observation is the means by which one collects/registers information about the state of the external world. Orientation comprises the internal processes by which observations are compared with prior knowledge and experience to update an understanding of the world. Decision is the internal process by which various tentative solutions are assessed and one selected for action. Action is the process by which the internally constructed solution is applied to the world. Observations test the constructed worldview at each step in the process.

By way of illustration, consider a hypothetical scenario based on the notion of appropriate social behaviour and arranged marriages, as may have been typical between conservative western families at the turn of the last century¹⁰. Loving parents may have deemed it desirable that their 18 year old son marry a caring and loving young wife from a socially appropriate family upon the completion of his military tour of duty, at about the age of 30 years. Should at some time between now and then their young soldier son fall in love with a young maiden from the local working-class village, the parents may actively seek to dissuade the relationship and directly initiate effects to cause their son to abandon the undesirable (from the parent's perspective) liaison. Examples of such direct effects could be to have the young man posted to another unit, apply parental pressure to the son, discourage the young woman in some way, or to encourage his friends to intervene. In any case, these immediate plans would be made to deliver near-term outcomes based on deterministic actions and their consequent effects. This is a *causal situation*, the parent's causally calculate the relative effect of their actions on the son's situation to achieve a predictable outcome, from set of possibilities. Several solutions to the dilemma are possible but one is considered best. Subsequent intervention may still be

⁹ Boyd (1996), "The essence of winning and losing."

<http://www.belisarius.com/modern_business_strategy/boyd/essence/eowl_frameset.htm>.

¹⁰ This hypothetical scenario is fictitious but makes clear social classes, biases and prejudices that were relevant in the 1800s that are unacceptable in 2007, under non-discriminatory norms and legislation. The author's have no intent to offer offence to any reader in the selection of this example, but use the illustration only because of its direct application to the concept of temporal divergence and convergence.

necessary, but on a day-by-day basis, as the immediate situation changes: all is based on known cause and probable effect. The cognitive frame of reference is in the “present” where the recent past bears predictably on the near-future.

In taking this scenario further, the loving parents may enhance their simple plan to have their son marry upon the completion of his military service (in 12 years time). It would be ideal if the wife-to-be was about 20 years of age at that time, was savvy in business affairs, and was fit and healthy to support childbearing. The future couple could then settle in the family estate and raise healthy children so that they may share in the family inheritance and businesses to carry forward the family traditions through future generations. This fictitious scenario represents the parent’s *situation image*, consisting of the ongoing interactive properties of the environment, likely future states and interactions, and states and interactions which are now available, or which might become available, if certain other interactions were first engaged. It is most likely that these parents would/will make/take every opportunity (as they present in the evolving future) to make introductions for their son to significantly younger daughters from socially appropriate families to facilitate their desired future. In the parent’s *Lebenswelt* (lifeworld, see Husserl, 1936)¹¹, they are shaping a future reality as they work in the present to confirm that reality. The cognitive frame of reference is in the “more distant” future where the present is merely a preliminary state. We thus describe a *intentional situation* as the parents attempt to shape reality (by whatever means at their disposal) as an instance within their *situation image*.

The point of the illustration is to make clear how the parent’s must manage the day-to day contingencies surrounding their son’s present circumstances, independence and experiences. The frame of reference is the “present” where the recent past bears predictably on the near-future (*causal situation*). At the same time they must also maintain their goals and intentions for the future (in the belief that their aspirations will come to pass) and take every opportunity, as it presents, to shape and form their desired end-state (*intentional situation*). They just work-it-out along the way as they can, seeding preliminary trends which should in combination and longitudinal effect collectively help realise their ambitions. They seek to “make it so” and may, through the agreement of a like family with similar ambitions, negotiate a marriage with, or without, the knowledge of the husband and wife-to be. The contrived union may just seem to fall naturally into place, at the ideal time and place in accord with human nature, without the son ever being aware of the behind-the-scenes manipulations of his parents and the parents of his beloved.

Of course, unforeseen events may make the intended future impossible. In the above illustration, should the son die or go insane whilst on military service, the parent’s *situation image* is defunct.

Intentional situation. A future end-state (in time and/or space) is locked-in as a goal and present actions (deliberate or *ad hoc*) are deemed to either help or hinder reaching that end-state (based on past experience of relational causes and effects). It does not matter much if environmental conditions change in the process, as management is by perceived “delta”, rather than by absolute reference; however, the set of deliberations is still seen to follow:¹²

¹¹ *Lebenswelt* (German) meaning “as lived”. . . . Husserl E. (1936) *Die Krisis der europäischen Wissenschaften und die transzendentale Phänomenologie: Eine Einleitung in die phänomenologische Philosophie*. [The Crisis of European Sciences and Transcendental Phenomenology, Northwestern University Press, Evanston, 1970.]

¹² Grant T. (2005), Unifying Planning and Control using OODA-based Architecture, Conf. Proc. SAICSIT 2005 South African Institute of Computer Scientists and Information Technologists, Mpumalanga, SA, 20-22 Sep, 111-113.

Grant and Kooter point out that “OODA exhibits shortcomings in psychological validity, in that it lacks concepts of memory and attention, and cognitive representations of inter alia world states and domain knowledge. This should be rectified by explicitly including a world model in OODA, with world states and models represented appropriately and store/recall functions for updating the world model. In addition, OODA needs a mechanism that determines the focus of attention appropriate to the current situation.”

Further to Grant’s revised formulation, the OODA construct does not seem readily applicable to collaborative decision-making by teams, as in military command staffs and combat teams. For example, no processes are shown representing negotiation or collaboration between team members or [initially] shared *Situation Awareness*: there remains a need for information distribution, the development of shared situation awareness, task re-allocation, confirmation and authorisation of decisions, and team maintenance.

DISCUSSION

In the context of Kuhn’s paradigms, we here consider two seemingly contrary worldviews – the *causal situation* and the *intentional situation* – which we somehow manage to hold together in making decisions about future events in achieving our intentions. The first is “now”-centric, based on our native sense that worldly events happen around us. In such a model, at our “now” point in time, we consider the world to exist as we perceive it to be; “as it is”. The future then unfolds through time from this point towards an “uncertain future” by actualising particular outcomes of probabilistic events. As such, the “certain present” is fixed and real, while the “uncertain future” deviates more and more from the “certain”, as the one-time “present” recedes further and further into the past. This is a forward projection in time.

To make this clearer, consider “now” as a fixed benchmark or milestone corresponding to this instant in time; my “certain present”. The “now” is only ever an individual cognitive state, or mental “situation image”, where the best we can do to share this “situation image” is to establish an idealised point of reference (t_0) and record what we thought the world was like at that instant, which then forms our benchmark. We can then discuss how the dynamic world evolved away from that benchmark as various possibilities became actualised to some specific value that then determines what possibilities are next available for actualisation. Most scientists are comfortable with this world view because it is based on probable cause and effect; this is the world of *temporal divergence*, as illustrated in Figure 3.

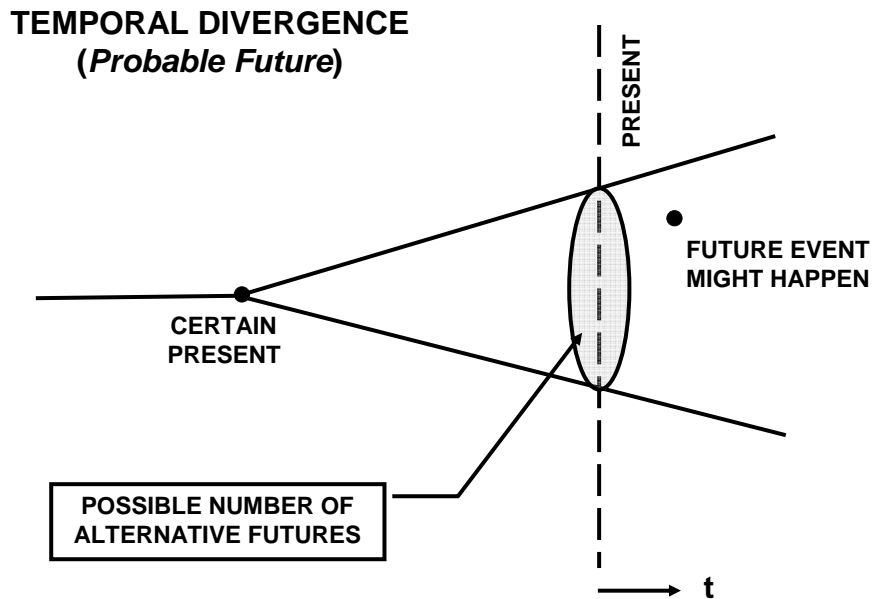


Figure 3: Temporal Divergence

The past and present have only one state – “as it was” and “as it is” – whereas the future is unbounded. Likewise, one has no ability to change what is or has been (the present or past) only the future is malleable, dependent on how much future “freedom of action” one has and the number of likely opportunities for possible future change implementation.

The alternative worldview is that the world around us is malleable and fashioned by our will, our power to affect our environment and our freedom to act, in time, to dynamically make our desired end-state the future (“as it will be”). In this model, the desired future is the fixed reality, where we, as creative beings, actualise our determinations and intentions via our individual and collective opportunity to effect change. There will be only one “future” end-state (because we will “make it so”), with a decreasing number of “now” possibilities, depending on our temporal separation from our “future” destination. We have limited freedom of action to effect the present (because we can do only as opportunity presents) until eventually there is “nothing more we can do” since the planned future deadline is imminent: where the predicted “future” has become the “now”. In this worldview, success or failure is very tangible. In this mindset, the “future reality” becomes the clearly defined milestone, where through intent and determination we take every opportunity to effect change, such that we eventually achieve (or not) the desired state of the world in accord with the pre-established milestone. This is the world of *Temporal Convergence*, as illustrated in Figure 4.

TEMPORAL CONVERGENCE
(*Fait Accompli*)

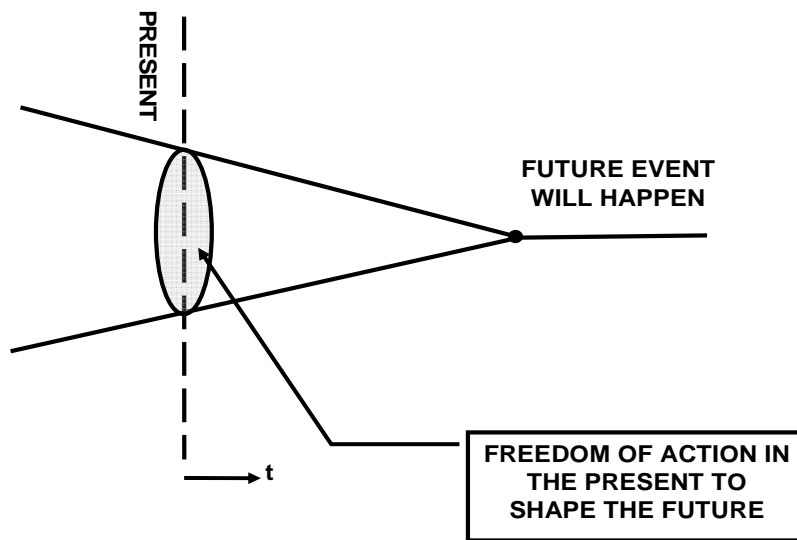


Figure 4: Temporal Convergence

Whereas for *temporal divergence* the past and present had only one state for *temporal convergence* the future has just one state – “as it will be” – whereas the past is unbounded according to “what might have been”. One has no ability to change what is predestined to happen, nor what might have been done in the past. Rather only the “now” is malleable, dependent on how close the present lies to the selected future time and how much “freedom of action” one has to effect change. As time converges, our freedom to act becomes more constrained, since any effect from our actions will take time to be evident. As the available time to cause an effect decreases, the fewer choices of action one has to create the necessary correction.

The causal and intentional situations can also be considered using Pascal’s Triangles¹⁴, as shown in Figure 5. With this tool we can see that for any combined set of binomial decisions (right/left, better/worse, profit/loss), the number of possible end-states increases divergently from the initial decision, but the central option in the final row $C(6,3)$, remains the most likely, with 20 decision paths to its realisation. Probability defines the most likely future.

¹⁴ Pascal's Triangle – Combinations: The number of paths by which you can arrive at a particular point in a binomial tree, say, the 6th row is the number of ways that you can string together Left and Right paths such that you end up at that point. If you want the total routes for all points in the 6th row we have two choices (right or left) for each step on the way giving $2^6 = 64$ as the total of all routes. The number of paths leading to any given fork along the way is given by $C(n,r) = n!/(r!(n-r)!)$, representing the number of branches you can take in n steps. Then, with $n = 6$ and $r = \{1,2,3,4,5,6\}$ we get the 6th row of the triangle as:
 $C(6,0) + C(6,1) + C(6,2) + C(6,3) + C(6,4) + C(6,5) + C(6,6) = 1 + 6 + 15 + 20 + 15 + 6 + 1 = 64$.
 1, 6, 15, 20, 15, 6, 1 are therefore the entries for the 6th row of Pascal's triangle.

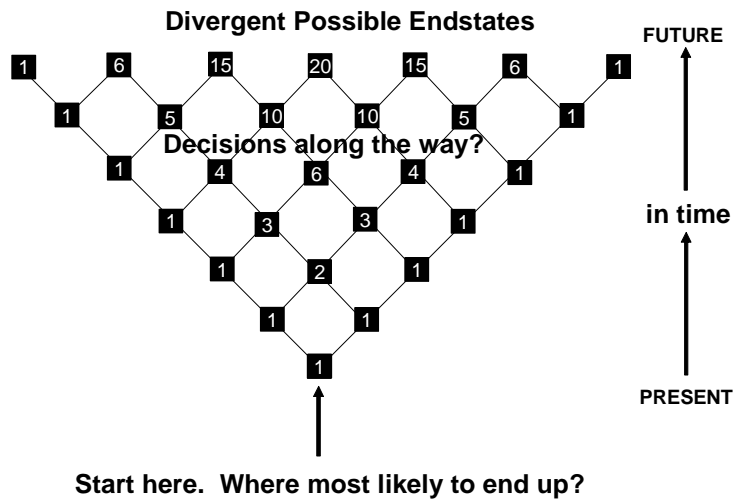


Figure 5: Pascal's Triangle with temporal divergence

But, Pascal's Triangle can also be inverted, as shown in Figure 6, in which we see that given the reality of one end-state, there exist many sets (combinations) of binomial decisions that will lead to that desired end-state. With this mindset, decisions converge in the future. The question here is no longer "Where best to end up?" but, "Where best to start?" Figure 5 illustrates six starting points, but the central option in the first row $C(6,3)$ would be the best, because it provides the greatest likelihood of realising the desired end-state. Here the number of future opportunities to make decisions, so to effect change, is very much dependent on the start point.

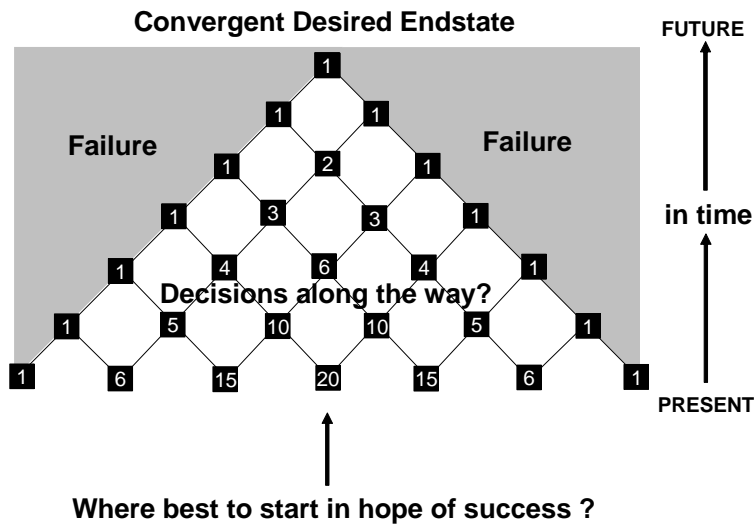


Figure 6: Pascal's Triangle with temporal convergence

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. 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, for example, 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 - 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 7 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).




Figure 7: The number of paths to each reachable node in a binomial state-space.

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:

$$P(\text{Finish} | \text{Start}) = \frac{20}{64} = 0.3125$$

By inverting the triangle as shown in Figure 8, it can be seen that of all nodes on the *Start* row, the *Start* node offers the most paths — twenty (20) — to the *Finish* node. There is a 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.




Figure 8: The number of paths to a given node from all possible starting nodes.

Now consider Figure 9 which shows the token part-way through a journey. It can be seen that there is only one (1) path which 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 $(\frac{1}{8} = 0.125)$.

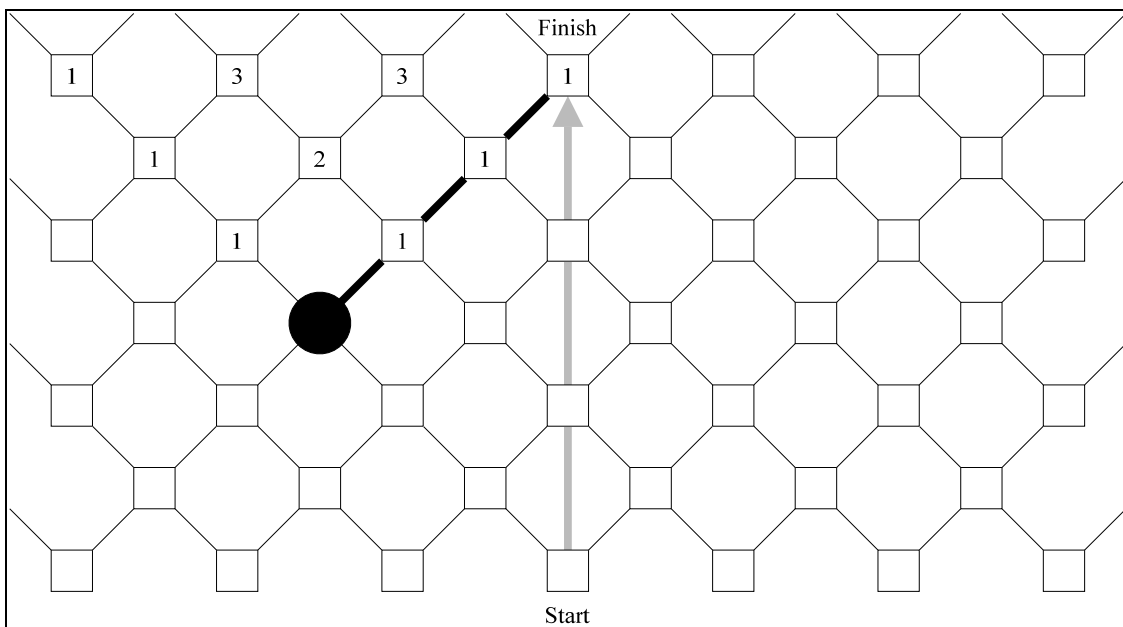


Figure 9: Only one path out of eight will reach the desired goal.

But now consider Figure 10 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, and one (1) of them will lead to the *Finish* node. That is, the probability of reaching the desired goal has now increased to $(\frac{1}{2} = 0.5)$.




Figure 10: Only one degree of freedom but a greater chance 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 in Figure 7 from 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 $(\frac{10}{32} = 0.3125)$; 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 9 or Figure 10 to arise. Although in Figure 10 the token has a 50% chance *a priori* 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.

KNOWLEDGE FOR GOAL-DRIVEN AGENTS

We can formalise our common-sense 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 outcome¹⁵. If some actions lead to desirable outcomes, then by implication, others are not, and 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 as 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. More broadly, action-outcome pairs serve to differentiate situational classes. We here describe the process by which an agent uses an action-outcome pair to differentiate a situational class as “situation image bifurcation”. Consider a simple case, as illustrated in Figure 11.

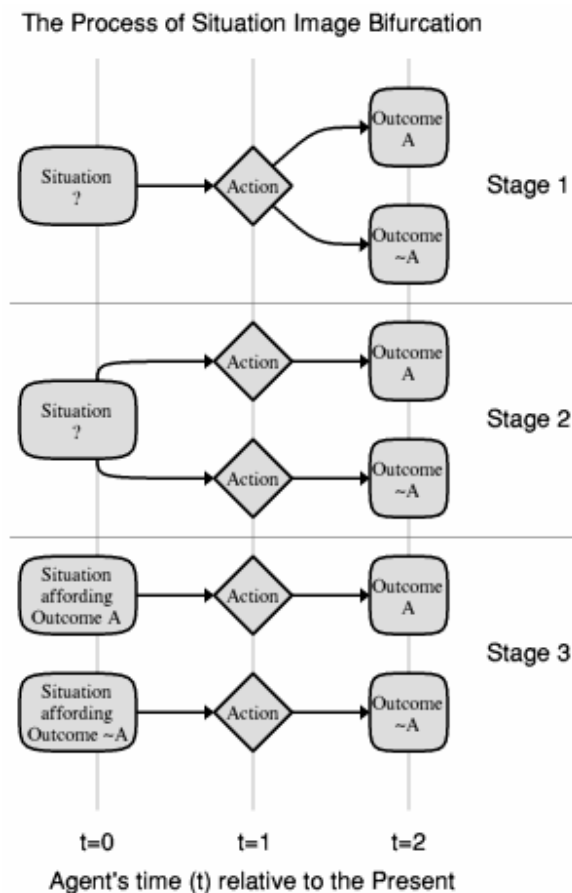


Figure 11

¹⁵ 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.
- The agent is passive and merely waits for the desired outcome to arrive via environmental flux.

At time $t=0$ representing the agent's present, the agent projects the future performance of an action at time $t=1$. However, the current situation is undifferentiated in that the agent cannot predict which outcome will be realized at time $t=2$. When the agent's clock is advanced to time $t=2$, then there is no longer any question. One of the outcomes has been realized, and so the agent can “unzip” its *situation image* to differentiate the class of situation at time $t=0$.

If there is no possibility of an undesirable outcome, then the agent has nothing to lose by trying an action. However, it does the agent no good to determine that it is unsafe to perform an action by performing that action! So we should focus on cases for which: (i) an action has either a distinctly desirable or distinctly undesirable outcome; and (ii) the agent intends to attain the desirable outcome; and (iii) it is vital for an agent to determine that the current situation is a member of the desirable-outcome class, as opposed to the undesirable-outcome alternative, before performing the action.

Suppose that Outcome A in Figure 11 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 image. A stage in this process of “active sensing” is shown in Figure 12.

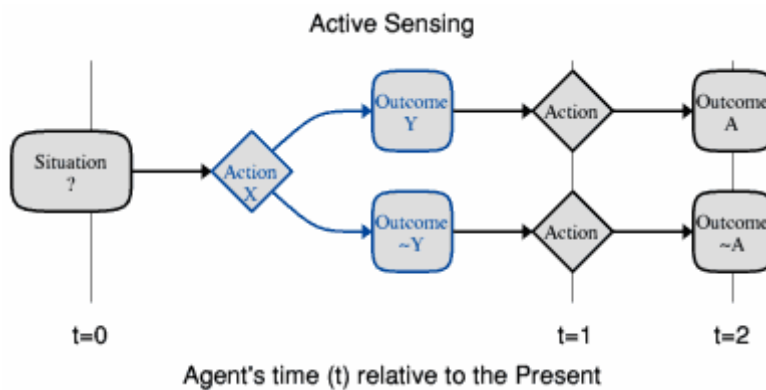


Figure 12

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 image* 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.

Traditional	Situation Image
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 - Rule Expressions

The anticipation of outcomes from actions allows the agent to detect when its representations are in error. If a *situation image* 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.¹⁶ 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.¹⁷ The point here is that whether one is managing knowledge in a theatre of temporal convergence or temporal divergence - success is highly dependent on the multiplicity of connections to reality, the degree to which they are criticised to detect anomalies, and the speed with which these connections can be made and criticised relative to the competition. The timeliness of knowledge is critical in either case.

It follows that in any large organisation we will find parts in which there is rarely any doubt about the current situation, how it was reached and what to do next. We may also find parts where situations are fluid and often novel. In the former case, the *situation image* is mostly static, long service is devalued, and knowledge management tools dominate. But in the latter case, the *situation image* will be undergoing frequent bifurcation and expansion. Team members will greatly value those who have a long history in that part of the organisation. And knowledge management tools (if tried at all) will have little (if any) positive impact. Long-serving staff are valued not just for the complexity of their situation images built up in the course of much experience, but also in the skilled processes they have developed for *situation image* construction and bifurcation. That is, a *situation image* can change, but so too can the process by which the *situation image* is changed. We say that people who develop the latter capability have been learning how to learn and their fluency in constructing new knowledge is at least as valuable as the knowledge they have already constructed.

If knowledge management systems are to be employed/applicable/useful/effective beyond well proscribed domains, then beyond focusing on the storage and retrieval of fully-developed *situation images*, they must begin to explicitly support *situation image construction and bifurcation*. Ideally, they will even come to support meta-processes by which the processes of *situation image construction and bifurcation* are enhanced.

¹⁶ With one condition: the outcome must be specified only in terms from the domain of the agent's internal organisation.

¹⁷ The ability to detect that the world is not as expected is a fundamental requirement for any adaptive system.

To be continued by 02 April 2007.....

In this regard, we present a conceptual framework as a model to help us understand some of the complex military processes addressed under Network-Centric Warfare, such as: command direction and intent; team self-synchronization; endgame focused “Course of Action”; effects-based approach and combined effect; shaping a malleable environment; and shared situational awareness.

We explore the temporal relationships of information (lag, reiteration, timeliness, lifetime and obsolescence) to decision-making based on shared intentions and desired end-states.

We contrast our conceptual model to others.