Temporal Agents

Jim Cunningham, Imperial College, London SW7 2BZ

rjc@doc.ic.ac.uk, May 10 2002

Abstract

The finite bounds on the rate and duration of physical processes is a natural resource limit for a rational agent, but processes are a not prominent in logical presentations of rationality. Here we consider some of the logical issues in portraying processes in a way which may help explicate linguistic and computational aspects of human agents, assist the design of artificial agents, and perhaps too, explain some mistakes of reason.

1. Appreciating Psychologism

By embracing psychologism in their "New Logic", Gabbay and Woods[1] admit inference which is fallacious in traditional logic. This is justified by arguments such as the usefulness of short cuts for achieving effective response from an agent with limited computational time and space. The cited paper stands as an informal prelude to a logic of wider ambit, including abduction and discovery, an unconventional, ground-breaking logic whose consequences are to be induced from inference steps which include the seemingly fallacious.

Although inspired by the new logic, this paper is more concerned with processes and their logical presentation; the computational processes of a bounded agent in an unbounded environment, with processes to be composed from primitive steps, and where the way processes are composed ultimately constitutes the mental architecture of the agent. If these processes are to operationalise the inferences of a more conventional logic, its consequences will perhaps be perceived as epistemic states of the agent. We take the view that in modelling an agent the processes are no more ephemoral than the states they may produce, that the observational behaviour of an agent depends on rich, dynamic, and communicative ingredients that stative concepts alone do not capture, and that the reasoning processes of agent, through being temporally bounded, can be distinguished from the atemporal stative inference which may be proper for an institutional agent. From this perspective, psychologism may be the misplaced adoption of temporal reasoning processes when not appropriate.

Despite the encapsulated rationality of traditional mathematics and logic, and the hunger of computational theorists for guidance, these disciplines have not been overly successful in characterising real computation. One can argue that a faithful description of computational engines, and of operational agents, depends on a temporal sense of causality and consequence which can be hidden by atemporal interpretation when rationality is expressed in traditional logics. The discipline of symbolic logic may come as inherited background in the drive to represent computational processes, and it has been elegantly exploited by causal rules in logic programming styles, but all too often the underlying logic of a computer program is obscured by cumbersome notation with explicit temporal order, whether Turing machine,

[1] Dov Gabbay and John Woods, The New Logic, Logic Journal of the IGPL, vol.9 No. 2, pp157-190, 2001

pi-calculus, or Java program. At least for reasoning agents we wish to re-examine some of the links between logic and process.

Implicit in our discussion, an agent is assumed finite in its initial knowledge, in its processing speed, in the duration of its previous lifetime, and thus in its current knowledge and beliefs (as we may perceive them). When we allow the agent processing capacity for learning, abstraction, and introspection, we would argue that a primary resource limitation is its temporal nature, not just processing time, but time for gaining experience of its environs, and temporal reasoning itself, which we wish to address. An agent's reaction, and its planning, are limited by its incomplete, finite knowledge. Its best decisions may be incorrect, when viewed from an atemporal global perspective. But rather than being an imperfect reasoner, it may even be locally optimal by some measure, yet prone to error because of imperfect knowledge. Of course, part of an agent's survival strategy, or its inherited strategy for survival of its group, may be for the individual to acquire new beliefs through introspection and abstraction when recovering from mistakes, and to validate these beliefs or seek new knowledge to enhance its control of the environment.

At issue is the coherent design of such an agent. The beliefs, desires and intents of Bratman's agents[2] have proved insightful and durable, although a structurally coarse model for the states of a computational agent's "mind". But they are a weak explanation for the computational processes themselves. Deliberation without learning and enquiry ensures a witless agent. But learning and enquiry are processes too, albiet abstract ones, hidden but salient and dynamic traits of rational agency. Whether enquiry is driven by introspection, and can be an observational basis for a conscious agent is more contentious, but a rational agent must act and react on imperfect knowledge, decide how to recover from apparent inconsistency, and learn from experience. A coherent process architecture for such behaviour should be part of the explication for *mistakes of reason.*

2. Windows on the Mind

It is our privilege to live in an age where formal reasoning can not only be described, but simulated by machines at speeds so surpassing our human capacity that it also becomes evident that algorithmic formalisations of arithmetic, logical, and even chess playing processes are mere trinkets in comparison with the richness of human reasoning. However, if we are indeed reasoning agents we cannot be smug, for there are already more challenging forms of computational agent, including realisations of the BDI agents alluded to above, and their models are quite general in their domain of applicability. Nor should the scientific advance which computer technology provides be underestimated, even if it may seem merely smart technology. Just as the adoption of the toy magnifying glass in Kepler's day provided new sights on the heavens, we too have a toy for the first time; a toy for investigating models of reasoning. The subjectivity of introspective metaphysics can become at last the basis of observation and experiment. Kant saw that Copernicus had distinguished movement of the observer from movement of the heavenly bodies, and sought similar separation between the mind and logical reasoning itself[3]. If we are, in some sense, to re-admit psychologism to formal reasoning, we can still require their integration to be objectively assessed.

For mathematical logicians it is perhaps a fear of non-objectivity which can devalue evidence from the arguably ephemoral and synthetic phenomena of linguistics, even relatively benign taxonomical data. But to ignore linguistic evidence as contaminated by psychologism unless re-interpreted by a more orthodox dogma introduces an unnecessary barrier to mental

[2] see Bratman 1987, Rao and Georgeff 1995

[3] Kant 1781, Critique of Pure Reason, Preface to the Second Edition, Meiklejohn translation.

insight, for language too can be the basis of experiment with artificial agents. For example, although tense logics are inspired by natural language, in a monadic tense logic on a real number time frame we cannot readily distinguish distinct past times when a proposition is true. An intended temporal sequence of propositions can be confused with a sequence which includes replicated, or spurious occurrences. We do not know whether this is a mistake which occurs with human communication, or human memory, but it is an expressive limitation for simple tense logics[4]. In adopting some notions of tense and aspect we do not presume that our suggestions for an interpretation of temporal language are without need of validation, but we believe the approach provides insight. For a start, we adopt an interval model of time; not only because it is well attested as more faithful in the semantic literature[5], enabling aspect as well as tense to be represented, but more contentiously, because by doing so we can represent a progressive sense of process and thereby get a more expressive handle on temporality, and perhaps consciousness, as an inner perception of activity and belief.

2.1 An interval tense logic

For the purposes of exposition we adopt and extend a propositional variant of an interval tense logic introduced by Halpern and Shoham (1986) which we call HS. The expressions of HS are constructed from syntactic atoms using Boolean and modal connectives. The latter incorporate Allen's binary relations on intervals, first introduced as a basis for qualitative reasoning (Allen 1983). These relations are illustrated in figure 1. Each relation R is incorporated in the HS logic as a modality with the dual forms for the possible, $\langle R \rangle$, and the necessary, $[R]$ modality. These are evaluated with respect to a current interval which may loosely be regarded as the period we call "now". The defining relations are not all independent. Thus the *later* relation L can be defined by the modal statement $\langle L \rangle p \leftrightarrow \langle A \rangle \langle A \rangle p$, while a definition for the beginning point of a durative condition p is $[[BP]]p \leftrightarrow \langle B \rangle (p \wedge [B]\bot)$. The operators $\langle \underline{L} \rangle$ and $\langle L \rangle$ and can be used in place of more conventional *Past* and *Future* tense operators to provide a rendering of simple past tense *John ran* and future tense *John will run* on the interval time domain, although this does not remove the expressiveness problem mentioned above. However the conventional rendering of tense is also inadequate if we hold to Reichenbach's use of an auxiliary reference point to distinguish the perfectives (Reichenbach 1947). A fragment of HS was extended with a focus operator in Leith and Cunningham (2001) to provide a computational representation of perfective tense and aspect, with supporting decision procedures for linguistic inferences.

Although more expressive than traditional point based tense logics, the HS logic was also shown by Venema (1990) to have both obscure expressive limitatations and computational complexity, so its illustrative uses for qualitative reasoning appear not to have been pursued. In figure 2 we present a variant of the HS logic that is actually a more expressive, hybrid logic, which we call HSN. HSN extends HS by including nominals for intervals as atoms disjoint from propositional variables[6]. While still being a tense style logic, this enables us to overcome the expressive limitation of simpler logics, to also provide a more direct presentation of the perfective aspect, and a scalable representation of knowledge through named intervals, similar to the use of events (Davidson 1980, Parsons 1990). In this way we bridge the gap between the use of less expressive, but linguistically faithful tense logics, and the thematic relations of event based sentential semantics. HSN includes a satisfaction operator @ binding a nominal to a proposition. This is an internalisation of Allen's metalogical *Holds* predicate (Allen 1994). Relative truth conditions for the propositional forms of HSN are given in

[4]For a more formal treatment of expressive limitations see, e.g. Gabbay et al. 1994
[5]See Vendler 1967, Allen, 1983, 1991, Dowty 1991, Kent 1993, Blackburn 1994.
[6]For an introduction to such Hybrid Logics (from Prior 1968), see Blackburn 2000

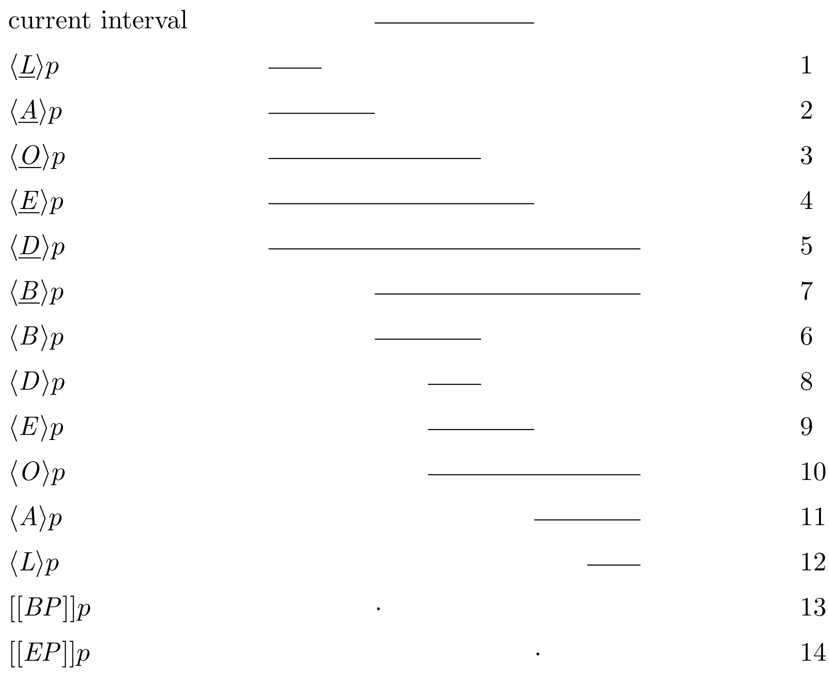

Figure 1: *
Illustration of relative interval positions and their HS formulae

figure 2, with respect a valuation for atomic propositions on closed intervals of a linear time frame, allowing for real or discrete models as is convenient for the application. Temporal linearity can also be ensured by the elaboration of more conventional axioms.

As illustration of the additional expressiveness of HSN we find that a tense operator may apply not just to a sentence, but to a sentence and its temporal reference. Thus the proposition *John ran* may be better represented as $\langle \underline{L} \rangle(\{John\ runs\}' \wedge r)$ where $\{John\ runs\}'$ is a sub-formula representing the sentence *John runs* and r is the nominal for the interval where the sub-formula holds. This nominal can serve as a reference for further discourse. A sense of the English perfective is difficult to capture without interval analysis because the present perfect *John has run* has a sense of immediate past, where the reference for discourse remains the current interval. Thus it is at least tempting to consider the present perfective as explained by use of the operator $\underline{A}$, without modifying the temporal reference, so that the past perfect *John had run* is indeed the composition of a past and perfective operators $\langle \underline{L} \rangle(\langle \underline{A} \rangle\{John\ runs\}' \wedge r)$. This appears to be compatible with Reichenbach's analysis.

2.2 Process Oriented Language

Aspect is as important as tense for the semantic analysis of the English verbs. Indeed, the English verbs are divided by grammarians such as Quirk *et al.* (1985) into the dynamic and the stative depending on whether the progressive aspect, as in *is making*, *is running*, is naturally admitted. Normal verbs like *make* and *run* are dynamic, but a minority, such as the verbs to *be*, *know*, and *desire* are classified as stative, indicative that the expressions **is being*, **is knowing*, and **is desiring* are not progressives in the same sense. The progressive epitomises ongoing process, or activity. Informal analyses of verb aspect by Vendler (1967),

$[s,t] \models x$ for propositional atom x *iff* x is true on the interval $[s,t]$

$[s,t] \models i$ for nominal i *iff* i denotes $[s,t]$

$[s,t] \not\models \bot$

$[s,t] \models (p \rightarrow q)$ *iff* $[s,t] \models p$ implies $[s,t] \models q$

$[s,t] \models @_i p$ *iff* i is a nominal denoting $[u,v]$, and $[u,v] \models p$

$[s,t] \models \langle B \rangle p$ *iff* there is a point u such that $s \leq u < t$ and $[s,u] \models p$

$[s,t] \models \langle \underline{B} \rangle p$ *iff* there is a point u such that $t < u$ and $[s,u] \models p$

$[s,t] \models \langle E \rangle p$ *iff* there is a point u such that $s < u \leq t$ and $[u,t] \models p$

$[s,t] \models \langle \underline{E} \rangle p$ *iff* there is a point u such that $u < s$ and $[u,t] \models p$

$[s,t] \models \langle A \rangle p$ *iff* there is a point u such that $t < u$ and $[t,u] \models p$

$[s,t] \models \langle \underline{A} \rangle p$ *iff* there is a point u such that $u < s$ and $[u,s] \models p$

Figure 2: Truth conditions for formulae of the HSN Logic on a Linear Temporal Frame .

and also by Moens and Steedman (1988), consider the durative character of events they describe, distinguishing processes (*to run*), culminated processes (*run home*), points and culmination (*win*). These distinctions are difficult to capture by atemporal expressions of formal events such as those by Davidson, by similarly motivated event based calculi, or by point based temporal or action logics. This is not to deny a place for events and action in the conceptual frame of an agent, but to emphasise the gulf between process and event models. Again this gulf can be bridged by using the interval nominals of HSN instead of the reified events of Allen (see Blackburn 1994).

In the HS logic a formula may be true on an interval without being homogeneously true on subintervals. To enforce homogeneity on stative conditions we can use the modal operator $[D]$ or define a **hom** operator by **hom** $p \leftrightarrow [D]p$. Indeed, this operator is a candidate for the underlying representation of the verb to *be*, once we have decided our representation for the condition indicated by other sentential components like *Mary* and *happy.* This use of a modal operator to provide the homogeneity associated with the most basic stative verb has an interesting parallel in the well established use of distinguishing modalities to represent statives verbs of propositional attitude, *knows, believes, desires*, and less clearly stative verbs of perception *sees, feels, hears*, where insistence on homogeneity is more contentious. Indeed, the sense of duration associated with the latter, as opposed to the sense of recoverable state associated with the former, may merit more careful consideration of the appropriateness of a homogeneity condition.

Although in general an activity may be interruptible and (de-)composable, the progressive aspect is an important case where an activity is perceived as proceeding or ongoing. For example, *is running*, the present progressive of *run* can be said to require not only an interval model, but *either* that the activity of *running* is also true on each sub-interval of some period embracing the current period (in which case the activity is temporally homogeneous), or else that running can be composed of subactivities such as a sequence of running, resting or walking. So the progressive form of an activity can be treated by a suitably defined operator, for *illustration*, **prog**, where:

$$\textbf{prog}\ p \ \leftrightarrow\ \langle \underline{D} \rangle ([D]p \vee p = r \odot s \wedge (\langle B \rangle \textbf{prog}\ r \wedge \langle E \rangle \textbf{prog}\ s).$$

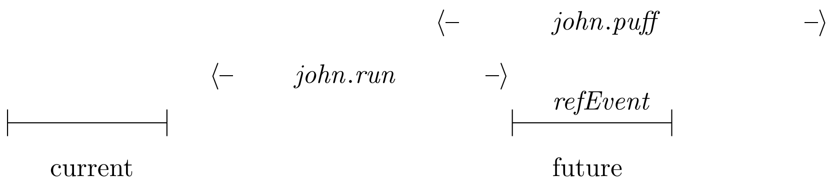

Figure 3: Temporal representation of *John will be puffing after he has been running*

Here $\odot$ is an associative operator denoting sequential, alternative or parallel composition (and where the last two modes of composition are commutative). Thus the **prog** operator forces a process to be ongoing and homogeneous, or composed of subprocesses, at least one of which is progressing in a beginning and an ending sub-interval.

As the dynamic state of an agent must change at the start and stop points of a homogeneous activity, these become the control points for execution of an agent. For example, an activity start point operator is defined by Leith (1997) as:

$$[[SP]]p \leftrightarrow (\langle \underline{A} \rangle [D] \neg p \wedge \langle A \rangle [D] p)$$

More sophistication for linguistic and operational purposes can be introduced by defining points where an activity break, resumption or finish occurs. It thus becomes evident that the use of interval models of time provides a very rich basis for the analysis of processes, since not only can familiar temporal connectives, like *after, before, while etc.* be given plausible representation, but rather than the simple decomposition of an activity allowed above, non-interval based models of process composition can be given an interval interpretation.

A hazard of such a linguistic perspective of process is the difficulty of generality in language itself, because our own fluid interpretation will depend on context[7]. So we have emphasised what Quirk et al. (1985) call the dominant, temporal use of tense, etc. rather than the many literary devices for fictional, presumptive, anaphoric use, etc. By so doing we get a glimpse of a language sufficiently rich in temporal descriptive power to offer possibilities for multi-agent coordination with controlled semantics. Less obviously, but pertinent to the logical role of process is that, when treating a declarative sentence as true on an interval, for instance a sentence like *John will be puffing after he has been running*, we also find that a verifying model can be computed by attention to the way the bounding pairs of points for each interval interleave and associate with logical atoms (also Leith 1997). The sentence may be represented in an HSN logic as $\langle L \rangle (\mathit{refEvent} \wedge \mathbf{prog}\ \mathit{john.puff} \wedge \langle \underline{A} \rangle \mathbf{prog}\ \mathit{john.run})$ where the progressive modality is defined as above. This is illustrated in figure 3.

2.3 Ontological Issues

In introducing an interval model of time for tense and aspect, we have used a propositional tense logic rather than a first order or predicative one. This is not only for presentational reasons, but also because it seems that a multi-modal propositional logic, at least a hybrid one, can represent much more of conventional language than is commonly realised, thereby both simplifying the reasoning processes by avoiding potentially undecidable first order

[7] The work of Moens and Steedman, and of Verkuyl (1993) systematize much of this fluidity.

issues, and also introducing mistakes of conventional reasoning. A case for hybrid multi-modal logics, and correspondences with description logics has been presented elsewhere by Blackburn (2000). Here we sketch some of the issues more specific to our concerns.

It can be taken as self evident that a major issue for agent reasoning is the structure of an agent's knowledge, in particular, the distinctions between the self, the environment, and other agents. However, it is only in very recent years that computational agents have been used experimentally, and the problems of agent communication addressed in practice. In doing so, the organisations concerned, including DARPA, and FIPA have fostered the adoption and extension of many concepts from philosophical discourse, including speech act models of communication, and ontological classifications of knowledge. These endeavours have not been without setbacks, and in particular the most appropriate ways of classifying knowledge appear to require some revision of Parson's analysis, notably to distinguish persons, or agents with intelligence, from objective and instrumental bodies (Schneider 2001). This in turn has impact on the logical representation, because theories of an agent's state of knowledge, perception of the environment, capacity for action and communication, and its beliefs about other agents can depend on an underlying ontological classification. Thus our way of extending a logical representation like HSN to deal with inter-agent communication and environmental perception is to some extent a matter of taste.

By way of illustration, we can suppose that a suitable extension of a logic like HSN to deal with a multi-agent environment will include appropriate modalities of propositional attitude, and of perception for each agent. Furthermore, since the possible states associated with each agent's reasoning, and the identities of each agent are just as extensional as the intervals of our time domain, we may expect to require nominals of each sort for expressiveness. In a classification of an agent's knowledge, these are instances of special kinds of thing, of which there will be many in an environment: *man*, *message*, etc. Somewhere lower down the ontological chain of knowledge we will expect attributes like primitive propositional variables of state and activity associated with each agent.

In a multi-modal representation a kind like *man* can be treated as another form of relation between the holder of knowledge and a percept of the holder, while a propositional variable like *happy*, or *run* can be associated with an instance of a class of a certain kind. The consequences for logical representation are that we can choose to represent the mortality of all men by the modal form $[man]mortal$ and the happiness of the uniquely identified Mary by a suitable hybrid modal form such as $Mary.happy$ rather than the predicative form. While the latter is a trivial although very natural variation on applicative form, the former is indicative that quantification is relative to the relation between the holder of knowledge and the accessible instances of the class. In an executable model one may test this locally, with unsound global results!

3. Mentalistic Processes

3.1 How temporal are processes?

Although a basic working concept of Computer Science, the idea of a process seems to be less of concern for metaphysical discourse, although there are similarities with the supposedly fundamental but unreal temporal A series of McTaggart, at least as discussed in the relatively recent treatment of Mellor (1981). Statistical processes have been proposed by as a weak basis for causation, for example by Salmon (1993), but neither the analysis of time nor of causation hints at the widespread use of process concepts in Engineering and other Applied Science. A cynical explanation would be that while many philosophers are aware of fundamental discoveries in pure science, the more sophisticated constructions of applied science have not been considered of explanatory significance unless couched as metamathematical discovery. While it is fair to say that the notion of process may not have been illuminating for nineteenth century mechanical engines, it is central today in fields like Control Systems and Chemical Engineering, as well as in Computer Systems. Process notions are also of growing importance for analysing the metabolic pathways of biological systems. In each case process is inextricably linked with the passage of time, to the extent that linear systems theory is largely built on the use of differential equations where the rate of change of measurable state properties is central.

Standard notation for the symbolic processes of Computer Science suppresses the explicit representation of time through the use of an operator for sequentiality (;), originating in 1960's programming notation. This then becomes the implicit temporal interpretation of process logics like dynamic logic (Harel 1984), and the μ-calculus (Kozen 1983). These are not interval based logics. When an interval based analysis has become appropriate in Artifical Intelligence, the dominant framework has been metalogical use of standard first order logic with quantified temporal variables, ignoring process logics. An alternative which is possible in HSN is to embed a sequentiality operator:

$$p; q \iff \langle B\rangle(p \wedge pn \wedge \langle A\rangle qn) \wedge \langle E\rangle(q \wedge qn \wedge \langle \underline{A}\rangle pn)$$

where pn and qn are nominals. Without nominals the ambiguity of repeated occurrences arises.

We are not aware of a notion of a process which is not intrinsically linked either with the passage of time, or with physical causality, which itself seems to be an alternative metaphysical foundation for the passage of time. For both human and computational agents the temporal interpretation of a process provides for physical causality, both for internal reasoning processes and for more obviously perceived mechanical actions. Even some details of the interpretation are similar since both brain and computer require energy to run and use electrical signals which provide models for the process logics we can use for representation. We do not yet know the details of internal mental processes, and they are certainly not naive sequential processes, but nor, any longer, are computer systems, particularly not firmware for specialised processing. But at a primitive level one signal can initiate another process, and through the motor cortex and human muscles, or through integrating amplifiers the low energy signals that initiate a process can be converted to physical action. So interval based process logics provide a fair description, subject to the reality constraint that all processes are of finite positive duration and bounded energy.

3.2 Tensed Knowledge Representation and Reasoning

The use of logics for knowledge representation is well established and computationally effective where logic programming representations can be exploited, but it is not without issues of contention. Where representations more faithful to epistemic criteria for introspection and communication with belief are desired, modal style logics can be preferred, although as mentioned above, more standard first order logics have been used metalogically for temporal interval representations, and in order to incorporate ontological structure. None of these considerations address the non-ideal recall and memory of the human agent, but its episodic character is well attested and can be represented through the use of reified events or interval nominals for temporal occurrences.

The linguistic evidence for the episodic character of human memory and recall is implicit in the validity of thematic roles in case grammars (Filmore 1968), as exhibited by an reified event style of representation: $puffingEvent(e) \wedge agent(e, John) \wedge cause(e, e_0) \wedge \cdots$, where the dots indicate a need for more detail on e_0 and its temporal relation to e, which is made more complicated by an interval analysis. While the case style analysis and thematic roles seem to be essential for a scalable form of knowledge representation, it is not provided by conventional tense logics. The HSN logic, when syntactically enhanced to represent ontological roles, appears to provide scalability though named intervals rather than events. This is supported by the intuition that an event is an occurrence when something happens, a dynamic manifestation of change, for which an interval reference can provide as much temporal information as is needed.

There is much more to say on this topic, and particularly on the distinctions between representations of episodic and non-episodic memory in a human agent. Perhaps the key point is that the representation of episodic recall needs to be dynamic, a resource limited process. That this may appear in the guise of a belief state about the past is yet another manifestation of the known capacity of dynamic logics to emulate more normal (and non-normal) modal logics of propositional attitude. But this is not just a formal phenomenon. In the case of human and computational agents, recovering information does require processing, with concommitant delays.

Unlike episodic knowledge, the relatively unproblematic feature of non-episodic knowledge is that it is stative, even knowledge about the effects of action. More temporally acute variants of familiar knowledge representation techniques are facilitated by logics like HSN. For example, conventional knowledge about an action a in terms of its preconditions P and postconditions Q can be expressed temporally by an axiom such as: $a \leftrightarrow (\langle \underline{O} \rangle P \leftrightarrow \langle O \rangle Q)$. Whether such a representation enables adequate treatment of the traditional frame problem, whether known reasoning techniques for intervals can support an effective calculus, and how such a calculus would compare with human reasoning about actions are topics for further study. But an interval calculus of events will not be restricted to sequential reasoning; the concurrent activities and communications of many temporal agents should lie within its ambit.

3.3 Higher Level Mental Processes

One of the motivations for this paper was my desire to improve a rather limited but perhaps insightful account of sentient consciousness as a progressive form of introspective awareness (Cunningham 2001). This can be given logical form in the sort of interval logic we have been discussing. In essence through axioms such as:

$$aware_j\ p \leftrightarrow perceives_j\ p \wedge perceives_j\ perceives_j\ p$$

$$conscious_j \ \leftrightarrow\ \exists p. \mathbf{prog}\ aware_j\ p$$

Sentient consciousness is only a fragment of the range of consciousness issues, but it seems to be compatible with both Baars' (1997) global workspace theory, and some readings of Dennett's (1991) views.

For the present my desire to improve this earlier account remains unfulfilled, but it can serve as a reminder that we can expect mental processes to have tense and aspect too. We can remember previous thought processes, and recover from interruption. Perhaps significantly, interruption is more difficult for those processes which we consider to be logical, as when we are following patterns of calculation associated with formal schooling skills. Hitherto unique to the human mind, but now being investigated for artificial agents too, is interaction with the irrational. Nevertheless, as Baars and others have recognised, the mind has features which superficially correspond to the operating systems processes and supervisory workspaces of a computer system; even though its behaviour can be changed to accord with environmental pressures.

Today one might claim that the greatest challenges for knowledge processing are not knowledge representation, nor reasoning with it (where there has been much computational progress), but the acquisition, discovery, and communication of knowledge. Each gives rise to mental processes. We briefly consider acquisition here, for to the extent that acquisition is learning, it continues to be an active research area of its own. But learning is relevant to temporality in agents because learning is so expensive, both in processing capacity, and in the time required. Explicitly logical approaches to learning, like inductive logic programming (Muggleton 1991) are exciting because they offer the possibility of acquiring knowledge in a logical form, which enables reasoning with it. But learning itself is a process, where the form of the goal ability – the explication of knowledge or the capability for action – can often be described in advance. The learning process can then becomes a form of successive approximation, improving detail. Thus it seems that learning as a mental process can be described like a least fixed point recursion in Kozen's logic, or in a more explicitly temporal extension, where the training process will vary with the goal. One particular facet of learning is the improvement of existing mental processes. We will have achieved something in the logic of temporal processes when such learning can be better explicated.

4. Integrating Temporal Processes

A simple, imperfect and reactive behaviour in a temporal world may well assist survival if it is quick to compute. Rationality for a human agent appears to need the coherent and potentially costly integration of many processes. The simple sequential integration of a few processes seems insufficient. Yet the BDI agents mentioned in section 1 are normally portrayed in just such a manner, a cyclic sequence: generate plan options for achieving goals, deliberate and filter to choose the intended plan for the current subgoal, partially execute through action, then update beliefs and goals for the next cycle, as illustrated in Figure 4.

One can explain the BDI model of rationality as a two stage process to achieve persistent goals, or *desires*, by finite deliberation. In particular, by treating previous intentions as partial plans that can more economically be revised than generated afresh when beliefs are updated. Whenever a goal has not been already been attained, and persists through time, reasoning about options and plans can determine whether the agent believes the goal can be attained in the future, and by which actions. A now well-known modal logic of branching time has been evolved to provide a theory for descrete execution (Rao and Georgeff 1995). In application the simple sequential BDI architecture has been been shown to be a robust, although not trivially extendable scheme. Goals may be meta goals like survival, or updated or communicated subgoals. Even integration with knowledge acquisition, discovery and communication is less radical in principle than it might seem to be in practice. For example, learning and discovery can be accommodated by introducing an abductive goal of explicating the environment, with rational introspection to account for unexpected behaviour, and exploratory learning a consequence of a subgoal to remove uncertainty. In an environment with other agents, communicative acts enhance capacity to acquire knowledge and achieve goals.

Although no explicit mention of time or interval is needed in this sequential architecture, when we allow durative actions, or indirect ramifications of actions, or concurrent processes in general, it becomes more difficult to distinguish cause and effect because effects can overlap in time with sensor updates to beliefs. Thus a need for temporal reference, coordination, and calculation creeps in. Only when we admit other agents does it become evident that an agent needs to accommodate concurrent and potentially durative actions to perceive the temporal world, and in order to coordinate through communication. Then it seems we may need interval logics, and names or dates for the periods we address, and make mistakes in our need for haste. But some of the unresolved issues in realising the BDI model already have a more principled resolution with an interval based model. For a start, there is a underlying dilemma in the treatment of intention. Although recognising intention in action, Bratman (1987, *ibid.*) avoids issues such as volition, and the impredicative treatment of action in conventional logics, preferring to consider intention as a revocable commitment to a future state. But as a result, the status of any commitment to a plan of action becomes problematic in realising a BDI style agent. Once processes become first class predicative entities of our metalogic for agent design we can avoid the underlying dilemma, indeed, re-think the BDI model to properly reflect a mental model of concurrent temporal activity.

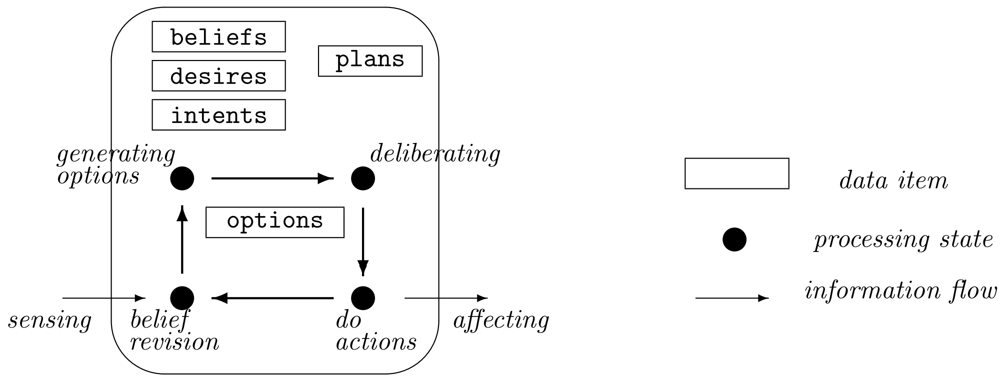

Figure 4: Simple BDI states and processes illustrated

5. Conclusion

By treating language as a window on the mind, and not just as a communication mechanism, and by acknowledging other mentalistic behaviour as processes subject to architectural limitations, such as our episodic and error prone memory, we provide a challenge for logic and for agent models. Some confusing aspects of human rationality are explicated by considering different logical systems. But can a logic plausibly describe an integrated collection of non-ideal reasoning and non-reasoning processes? Can logic help explain evolutionary phenomena, and design artifical agents, and if so, will it still be logic? The tradition of simplification to an elegant essence, perhaps some classical or non-classical logic, or to a single minded view of agenthood in terms of beliefs, desires and intents, has analytic attraction that we may wish to retain, but an unreality that we cannot ignore.

Not only the human mind, but also modern systems of information technology exhibit a complex rationality which co-exists with a real, temporal world. And yet in virtually no case do we see any representation by any logical system. Perhaps the brutal truth is that it is difficult to incorporate pyschologistic elements elegantly into a logic, and we have settled for less psychology rather than emasculated or ineffective logics and models. But if the traditions of Logic and of Philosophy are to have more impact on the technological world, this may need to change so that scalable logical systems of great complexity can explain and guide a future world inhabited by rational and temporally aware agents.

6. Acknowledgement

The author wishes to thank the organisers of the John Woods commemorative event for the invitation to participate, and Lloyd Kamara for his critical comments.

7. Bibliography

References

[1] Allen, J.F. 'Maintaining Knowledge about Temporal Intervals.' Communications of the ACM 26, 11, 832–843, November 1983.

[2] J.F. Allen, 'Time and time again: The many ways to represent time', *Int'l. Jnl. of Intelligent Systems* 6, 4, 341–356, July 1991.

[3] J.F. Allen, and G. Ferguson, 'Actions and Events in Interval Temporal Logic', *J. Logic and Computation* 4, 5, 1994.

[4] A. Artale and E. Franconi, 'A Temporal Description logic for Reasoning about Actions and Plans', *J. of AI research*, 9:462–506, 1998

[5] B.J. Baars, *In the Theater of Consciousness*, OUP, 1997

[6] P. Blackburn, 'Tense, Temporal Reference and Tense Logic', *J. of Semantics*, 11, 83-101, 1994.

[7] P. Blackburn, Representation, Reasoning, and Relational Structures: A Hybrid Logic Manifesto. Logic Journal of the IGPL. Vol 8:3, 339–625, 2000.

[8] M. Bratman, *Intention, Claims and Practical Reason*, Harvard U. Press, 1987.

[9] J. Cunningham, *Towards and Axiomatic Theory of Consciousness*, Logic Journal of the IGPL, 9(2):341-347 2001

[10] D. Davidson, *Essays on Actions and Events* Oxford University Press 1980.

[11] D.Dennett, *Consciousness Explained*, Penguin 1991

[12] D.R Dowty, *Word Meaning and Montague Grammer*, Kluwer, 1979, 1991.

[13] C.J. Fillmore, 'The Case for Case' in E. Bach and R.T. Harms (eds), *Universals in Linguistic Theory*, London 1-88, 1968.

[14] D. Gabbay, I. Hodkinson, and M. Reynolds 'Temporal Logic: Mathematicl Foundations and Computational Aspects', *Oxford University Press*, 1994.

[15] D. Gabbay and J. Woods, 'The New Logic', *Logic Journal of the IGPL*, vol.9 No. 2, pp157-190, 2001.

[16] J. Halpern & Y. Shoham, 'A Propositional Modal Logic of Time Intervals'. In: *Proc. Symposium on Logic in Computer Science*. Cambridge, Massachusetts, pp. 279–292. 1986.

[17] D. Harel, 'Dynamic Logic', *Handbook of Philosophical Logic*, Volume 2, Reidel 1984.

[18] S. Kent, *Modelling Events from Natural Language*, Ph.D. Thesis, Imperial College, University of London, 1993.

[19] D. Kozen, 'Results on the Propositional μ-calculus', *Theoretical Computer Sience*, 27, 333–354. 1983.

[20] M.F. Leith, 'Modelling Linguistic Events', Ph.D. Thesis, Imperial College, University of London, 1997.

[21] M. Leith and J. Cunningham 'Aspect and Interval Tense Logic', *Linguistics and Philosophy*, 24(3):331–381, June 2001.

[22] M. Moens and M. Steedman 'Temporal Ontology and Temporal Reference', *Computational Linguistics*, 14:15–28, June 2001.

[23] D.H. Mellor, *Real Time*, Cambridge University Press, 1981, 1985

[24] S. Muggleton, *Inductive Logic Programming*, New Generation Computing, 8(4):295-318, 1991

[25] B. Nebel. *Reasoning & Revision in Hybrid Representation Systems*, LNAI 422 Springer 1990,95

[26] T. Parsons, *Events in the Semantics of English: A Study in Sub-atomic Semantics*, MIT press, Cambridge Mass. 1990

[27] A. Prior, *Past, Present and Future*, University of Oxford Press 1967

[28] A. Prior, *Papers on Time and Tense*, University of Oxford Press 1968

[29] R. Quirk, S. Greenbaum, G. Leech, J. Svartvik, *A Comprehensive Grammar of the English Language*, Longman Group 1985

[30] A.S. Rao and M.P. Georgeff, 'BDI agents:from theory to practice", Proceedings of ICMAS-95, San Francisco, CA pp312-319, 1995

[31] H. Reichenbach, *Elements of Symbolic Logic*, Macmillan, London, 1967

[32] W.C. Salmon, 'Causality: production and Propogation', in E.Sosa and M.Tooley (eds), *Causation*, OUP 1993

[33] K. Schild, 'On the Relationship between BDI logics and standard Logics of Concurrency' Autonomous Agents and Multi-Agent Systems 3(3):259-283 2000

[34] L.N. Schneider, *Naive Metaphysics* MSc Thesis, Imperial College, University of London 2001

[35] Z. Vendler, 'Verbs and Times', *Linguistics in Philosophy*, 1967

[36] Y. Venema, 'Expressiveness and Completeness of an Interval Temporal Logic', *Notre Dame Journal of Formal Logic*, 31(4),529–547, 1990

[37] H. Verkuyl, *A Theory of Aspectuality*, CUP, 1993