Action-State Semantics for Practical Reasoning

Trevor Bench-Capon and Katie Atkinson
Department of Computer Science, University of Liverpool, UK
{tbc, katie}@liverpool.ac.uk

Abstract

There are two aspects of practical reasoning which present particular difficulties for current approaches to modelling practical reasoning through argumentation: temporal aspects, and the intrinsic worth of actions. Time is important because actions change the state of the world, we need to consider future states as well as past and present ones. Equally, it is often not what we do but the way that we do it that matters: the same future state may be reachable either through desirable or undesirable actions, and often also actions are done for their own sake rather than for the sake of their consequences. In this paper we will present a semantics for practical reasoning, based on a formalisation developed originally for reasoning about commands, in which actions and states are treated as of equal status. We will show how using these semantics facilitates the handling of the temporal aspects of practical reasoning, and enables, where appropriate, justification of actions without reference to their consequences.

Introduction

Argumentation can be used to justify both beliefs and actions. There are, however, great differences between theoretical reasoning, arguing about what should be believed, and practical reasoning, arguing about what should be done. Many of these differences arise from the direction of fit: we attempt to make our beliefs fit what is the case, but we use our actions to make the world fit our desires. Three elements are important: first, the subjectivity of practical reasoning, in that different people may rationally make different choices because they have different aspirations and values. Second, time is important, because actions change the state of the world, we need to consider future states as well as past and present ones. Third it is often not what we do but the way that we do it that matters: the same future state may be reachable through a desirable and an undesirable action, and often also actions are done for their own sake rather than for the sake of their consequences. It is important that any reasonable account of practical reasoning deals with these three aspects.

In the next section we will discuss three leading approaches to modelling arguments for practical reasoning.While argumentation effectively captures the subjective nature of practical reasoning, the other two aspects present some difficulties for these approaches. In section 3 we will briefly recall the definition of an Action-Based Alternating Transition System (AATS). In section 4 this formalism will be extended to accommodate the action-state semantics of (Reed and Norman 2007). In section 5 we will apply their semantics to the argument scheme introduced in (Atkinson and Bench-Capon 2007), and introduce new argument schemes to allow for arguments justified by actions and by plans as well as immediate goals. In section 6 we present a worked example to demonstrate the approach. Section 7 will give concluding remarks.

Current Approaches

Current work on modelling argumentation for practical reasoning can be broadly divided into those based on the Belief-Desire-Intention model of multi-agent systems (see, e.g., (Wooldridge 2000)), and those based on action transition systems. We will take (Rahwan and Amgoud 2006) and (Amgoud, Devred, and Lagasquie-Schiex 2008) as representative of the first approach, and (Atkinson and Bench-Capon 2007) as representative of the second. In BDI approaches, the underlying knowledge representation is a set of epistemic rules, allowing beliefs to be deduced on the basis of current beliefs; a set of desire rules, allowing desires to be deduced on the basis of current beliefs and desires; and a set of plans allowing desires to be realised in situation satisfying their preconditions. The alternative approach in (Atkinson and Bench-Capon 2007), represents knowledge of actions, their pre- and post-conditions, and the values promoted by actions in the form of an AATS (Wooldridge and van der Hoek 2005), originally developed to reason about norms in multi-agent systems and based on the Alternating-time Temporal Logic of (Alur, Henzinger, and Kupferman 2002).

In (Rahwan and Amgoud 2006) an instantiation of Dung's abstract argumentation framework (Dung 1995) is used first to generate a consistent set of desires, and consistent plans for achieving these desires, with strengths of arguments being based on the worth of desires and the cost of the resources required to achieve them. In (Amgoud, Devred, and Lagasquie-Schiex 2008) argumentation is used to generate a feasible set of justified desires, that is a set of desires which hold in the current state and which have a plan for achieving them, but does not consider how preferences can be used to choose between different sets of feasible and justified desires. In the approach offered by (Atkinson and Bench-Capon 2007), a specific argument scheme, based on the Sufficient Condition Scheme for Practical Reasoning (Walton 1996) is used to allow reasoning about what action should be selected. Preferences between arguments are based on the social values promoted by realising the feasible goals, and the ordering of these values subscribed to by the agent concerned.

In all of these treatments the emphasis is on states: actions are simply components of plans in (Wooldridge 2000), this is also true of (Rahwan and Amgoud 2006) and (Amgoud, Devred, and Lagasquie-Schiex 2008), although they do not always clearly distinguish actions from states: e.g. in an example of (Rahwan and Amgoud 2006), both 'interesting keynote speech' and 'attend keynote speech' are represented by literals from their underlying propositional language. In (Amgoud, Devred, and Lagasquie-Schiex 2008), plans are represented as triples of the start state, the end state and the desire that is 'reached' by the plan, without any specification of the actions that are performed to move between the states. But practical reasoning is not simply about deciding which state to reach: *how* that state is reached is also important, and often actions are performed for their own sake, such as walking in the park, or even attending a speech. Further, practical reasoning is intimately concerned with time – the agent is attempting to bring about one of a variety of possible futures. These temporal aspects are also not comfortably handled in (Rahwan and Amgoud 2006) and (Amgoud, Devred, and Lagasquie-Schiex 2008). In (Atkinson and Bench-Capon 2007), which is based on an AATS there is an explicit representation of actions as transitions between states, but states retain their primacy. The agent continues to justify actions only in terms of the state the action will achieve. Here too there is a restricted notion of time, but it is limited to a single step, and not capable of being explicitly expanded into a history.

Inspiration to extend the work of (Atkinson and Bench-Capon 2007) to properly handle actions as of equal status to states, and to introduce a proper notion of time, can be found in work on imperatives, namely (Hamblin 1987), which has recently been given a formal characterisation in (Reed and Norman 2007). Like practical reasoning, imperatives are concerned with both actions and states: indeed practical reasoning can be seen as issuing an imperative to oneself (cf (Atkinson et al. 2008)). Both practical reasoning and imperatives refer even-handedly to states and actions: one may order/desire a state of affairs, without concern for how it is brought about; one may order/desire an action without reference to its consequences; or one may order/desire a state of affairs to be brought about by a particular action. Hamblin's model gives states and actions equal status, and although it might be considered that this model is more lavish than is required for many purposes (as was recognised by Hamblin), we believe that it is necessary for practical reasoning. In particular, the fact that a given state may be reached from some other particular state by a variety of actions means that it is hard to reduce actions to simple transitions between states in any remotely natural way. Yet, choosing *how* various goals are realised, and how and when particular states are reached, is the very essence of practical reasoning. This will become clear in the example given in section 5.

Our idea is to connect the AATS of (Atkinson and Bench-Capon 2007) which has provided a foundation for their limited account of practical reasoning to the formal characterisation of Hamblin's action-state semantics given in (Reed and Norman 2007). This will retain the advantages of the approach of (Atkinson and Bench-Capon 2007), which include the use of argumentation schemes and critical questions to conduct defeasible reasoning, and the use of social values to represent individual motivations, but will also enable us also to handle time and the intrinsic worth of actions in a more natural way. Moreover, we can then use the logic of imperatives axiomatised in (Reed and Norman 2007) to provide a logic for practical reasoning, with distinct modalities for performing an action and achieving a state of affairs.

Action Based Alternating Transition Systems

AATS were introduced in (Wooldridge and van der Hoek 2005), based on the Alternating-time Temporal Logic of (Alur, Henzinger, and Kupferman 2002). We begin with a finite set Q of possible *states*, with $q_0 \in Q$ designated as the *initial state*. Systems are populated by a set Ag of *agents*. Each agent $i \in Ag$ is associated with a set Ac_i of possible actions, and it is assumed that these sets of actions are pairwise disjoint (i.e., actions are unique to agents). The set of actions associated with the set of agents Ag is denoted by Ac_{Ag}, so $Ac_{Ag} = \bigcup_{i \in Ag} Ac_i$.

A joint action j for a set of agents Ag is a tuple $\langle \alpha_1,...,\alpha_k \rangle$, where for each α_j (where $j \leq k$) there is some $i \in Ag$ such that $\alpha_j \in Ac_i$. Moreover, there are no two different actions α_j and $\alpha_{j'}$ in j_{Ag} that belong to the same Ac_i. The set of all joint actions for a set of agents Ag is denoted by J_{Ag}, so $J_{Ag} = \prod_{i \in Ag} Ac_i$. Given an element j of J_{Ag} and an agent $i \in Ag$, agent i's action in j is denoted by j_i.

As given in (Atkinson and Bench-Capon 2007), this allows an AATS to be defined as follows:

Definition 1: *An* Action-based Alternating Transition System *(AATS) is an (*$n + 7$*)-tuple* $S = \langle Q, q_0, Ag, Ac_1, \ldots, Ac_n, \rho, \tau, \Phi, \pi \rangle$, *where:*

- Q *is a finite, non-empty set of* states*;*
- $q_0 \in Q$ *is the* initial state*;*
- $Ag = \{1,...,n\}$ *is a finite, non-empty set of* agents*;*
- Ac_i *is a finite, non-empty set of actions, for each* $i \in Ag$ *where* $Ac_i \cap Ac_j = \emptyset$ *for all* $i \neq j \in Ag$*;*
- $\rho : Ac_{Ag} \rightarrow 2^Q$ *is an* action precondition function, *which for each action* $\alpha \in Ac_{Ag}$ *defines the set of states* $\rho(\alpha)$ *from which* α *may be executed;*
- $\tau : Q \times J_{Ag} \rightarrow Q$ *is a partial* system transition function, *which defines the state* $\tau(q, j)$ *that would result by the performance of* j *from state* q*. Note that, as this function is partial, not all joint actions are possible in all states (cf. the precondition function above);*
- Φ *is a finite, non-empty set of* atomic propositions*; and*

- π : Q $\rightarrow 2^{\Phi}$ *is an* interpretation function, *which gives the set of primitive propositions satisfied in each state: if* p $\in \pi$(q), *then this means that the propositional variable* p *is satisfied (equivalently, true) in state* q.

In order to express preferences between states, the AATS was extended in (Atkinson and Bench-Capon 2007) to include a notion of values, whereby a value may be promoted or demoted (or neither) by a transition between two states.

- Av_i *is a finite, non-empty set of values* $Av_i \subseteq V$*, for each* $i \in Ag$*. The set of all values for a set of agents* Ag *is denoted by* Av_{Ag}.
- δ*:* $Q \times Q \times Av_{Ag} \rightarrow \{+, -, =\}$ *is a* valuation function *which defines the status (respectively, promoted, demoted or neutral) of a value* $v_z \in Av_{Ag}$ *ascribed to the transition between two states:* $\delta(q_i, q_j, v)$ *labels the transition between* q_i *and* q_j *with one of* $\{+, -, =\}$ *with respect to the value* $v \in Av_{Ag}$.

Note in particular that this means that the transition between two given states will affect a given value in the same way, whichever joint action gives effect to the transition.

Action-State Semantics

Hamblin's idea in (Hamblin 1987) is that at every time point $t \in T$ there is a state, a collection of 'happenings' (events not attributable to any agent) and a set of deeds (actions of agents, one for each agent present at t). We can map these concepts to elements of the AATS. The state is simply some $q \in Q$: we will refer to pairs of states and times as *state worlds*. Thus the set of state worlds U is a set of pairs $\langle q, t\rangle$ such that $q \in Q$ and $t \in T$. Similarly the set of deeds executed at a time point are simply some joint action $j \in J_{Ag}$, but what should we do with happenings? Like Reed and Norman, we will conflate deeds and happenings into *event worlds*, E, so that happenings can be seen as the action of a special 'world' agent, $ag_0 \in Ag$. Now E is a set of pairs $\langle j, t\rangle$ such that $j \in J_{Ag}$ and $t \in T$. Reed and Norman then have a ternary accessibility relation $\mathcal{R_H}$ which links two state worlds by an event world. Thus an element of $\mathcal{R_H}$ will be of the form $\langle\langle q_i, t_n\rangle, \langle j, t_n\rangle, \langle q_k, t_{n+d_j}\rangle\rangle$, where $q_i, q_k \in Q$, $j \in J_{Ag}$ and $t_n, t_{n+1} \in T$, where d_j is the duration of j. Where convenient we will write elements of $\mathcal{R_H}$ as $\langle u_i, e, u_k\rangle$, to be read as "state-world u_k is accessed from state-world u_i through event-world e.".

So, this gives us the following definition, based on (Reed and Norman 2007):

Definition 2:

- $W = U \cup E$*: the set of possible worlds is the set of* state-worlds U *and the set of* event-worlds E. $U \cap E = \emptyset$.
- $\mathcal{R_H}$ *is a ternary relation* $\langle u_i, e, u_k\rangle$ *where* $u_i, u_k \in U$*, and* $e \in E$.

Just as in the AATS where it is convenient to distinguish the initial state, we will distinguish the state world and the event world at the current time point, t_0, as 'now'. Thus 'now' = $\langle\langle q, t_0\rangle, \langle j, t_0\rangle\rangle$ such that $q_0 = q$ and j is the joint action performed at t_0.

The relation $\mathcal{R_H}$ is also constrained by the AATS. Since $\mathcal{R_H}$ links two state worlds, there must be a way to get between them. In a given state, only those joint actions whose preconditions are satisfied are possible. Thus the set of joint actions possible in some q_i, is $J_{q_i} \subseteq J_{Ag}$ such that $j \in J_{q_i}$ if and only if for all $ag_i \in Ag, q_i \in \rho(j_{ag_i})$. Thus $\langle\langle q_i, t_n\rangle\langle j, t_n\rangle, \langle q_j, t_{n+d_j}\rangle\rangle \in \mathcal{R_H}$ if and only if $j \in J_{q_i}$ and $\tau(q_i, j) = q_j$, and the duration of j is d_j.

We now define a history and a future:

Definition 3: *A history,* h_t^{t+n}*, is a sequence* $\{r_t, r_{t+d}, ..., r_{t+n-1}\}$ *where every* $r_i \in \mathcal{R_H}$ *such that if the* n^{th} *element =* $\langle\langle q_k, t_j\rangle, \langle j_1, t_j\rangle, \langle q_l, t_{j+d_{j_1}}\rangle\rangle$*, the* $n+1^{th}$ *element will be* $\langle\langle q_l, t_{j+d_{j1}}\rangle, \langle j_2, t_{j+d_{j1}}\rangle, \langle q_m, t_{j+d_{j_1}+d_{j_2}}\rangle\rangle$*. Histories starting now will be* futures *and histories ending now will be* pasts. *The set of futures* F *are all* $h_{t_0}^{t_2} \in H$ *such that* $t_2 > 0$*. The set of pasts* P *are all* $h_{t_1}^{t_0} \in H$ *such that* $t1 < 0$.

A future f_i is a sub-future of a future f_j if $f_i = h_0^n$ and $f_j = h_0^{n+m}$ and the first *n* terms of f_i are identical to the first *n* terms of f_j. We now extend the functions π and δ so that they can be applied to worlds.

Definition 4: $\pi*$ *:* U $\rightarrow 2^{\Phi}$ *is an* interpretation function, *which gives the set of primitive propositions satisfied in each state world: if p* $\in \pi*$*(u), then this means that the propositional variable p is satisfied (equivalently, true) in state world u. This is true if and only if* $u = \langle q, t\rangle$ *and* $p \in \pi(q)$.

Definition 5: $\delta*$*:* $U \times W \times Av_{Ag} \rightarrow \{+, -, =\}$ *is a* valuation function *which defines the status (respectively promoted, demoted, or neutral) of a value* $v \in Av_{Ag}$ *ascribed to the transition between two worlds:* $\delta * (u, w, v)$ *labels the transition between* u *and* w *with one of* $\{+, -, =\}$ *with respect to the value* $v \in Av_{Ag}$*. Note that the transition is from a state world to* either *a state world or an event world, reflecting the fact that actions as well as goals can promote and demote values.*

We are particularly interested in what is true at the end of a future. First we are interested in whether some proposition, *p*, is satisfied:

Definition 6: *We say that a future* $h_0^n \in F$ *is a* p-future *if* $p \in \pi * (u_n)$ *where* $\langle u_{n-1}, e_{n-1}, u_n\rangle$ *is the the last element of* h_0^n.

Goals are expressed as a conjunction of propositions. This gives the notion of a *g-future*, a future in which the goal is satisfied:

Definition 7: *For G* $\equiv p_1 \wedge,...,\wedge\, p_n$*, f* $\in$ *F is a* g-future *iff f is a* p-future *for all* $p_i \in G$.

Similarly we are interested in the values promoted by the future. Thus:

Definition 8: *A future* $h_0^n \in F$ *is a* v-future *if* $\delta * (u_{n-1}, e_{n-1}, v) = +$ *or* $\delta * (u_{n-1}, u_n, v) = +$.

Finally we are interested in the actions executed at the last transition and so we have:

Definition 9: *A future* $h_0^n \in F$ *is an* ac_i-future *if* $ac_i = j_i$ *where* $e_{n-1} = \langle j, t_{n-1}\rangle$.

Note that in the case of *p-futures*, *v-futures* and ac_i-*futures*, it is the last element of the sequence which we consider. The desired state or event represents the horizon of interest.

Application to Practical Reasoning

We now apply this new machinery to practical reasoning. As mentioned, we will use the argumentation scheme approach proposed in (Atkinson and Bench-Capon 2007). This approach is to provide a *prima facie* justification for the choice of an action using an argument scheme which is an extension of the *sufficient condition scheme for practical reasoning* (Walton 1996). Just as we distinguished ag_0 as the world agent responsible for all happenings, we now distinguish ag_1 as the agent towards which the reasoning is directed.

AS1 *In the current circumstances R, Agent 1 should perform action ac_1, which will result in new circumstances S, which will realise goal G, which will promote value V.*

This scheme distinguishes between three aspects of the effects of an action: the new state of affairs achieved (S); the desirable features of the new state of affairs, the goal (G); and the reason why those features are desirable, the value they promote (V). Following the conception of argumentation schemes in (Walton 1996), the instantiation of AS1 provides a presumptive justification for performing action ac_1, but this must be able to withstand the critical questions which can be posed against it, questioning the various elements of the scheme. For example one might deny that the current circumstances were as designed, that the action would have the claimed consequences, that these consequences promoted the value, that they demoted some other value, and so on. Seventeen critical questions were given in (Atkinson and Bench-Capon 2007), where both AS1 and its critical questions were formalised in terms of an AATS.

With our new machinery we can define the conditions for the instantiation of AS1 to justify an action ac_1:

Definition 10: *There is an AS1 justification for ac_1 if there is some r = $\langle u_i, e, u_k\rangle \in \mathcal{R}_{\mathcal{H}}$ such that $u_i = \langle q_i, t_0\rangle, e = \langle j, t_0\rangle$ and $u_k = \langle q_k, t_1\rangle$ such that for all $p \in R, p \in \pi * (u_i)$, $ac_1 = j_1$, for all $p \in S, p \in \pi * (u_k)$, $G \subseteq S$ and $\delta * (u_i, u_k, v) = +$.*

Note that on this definition there may be many such *r*, reflecting the fact that there may be several possible state worlds satisfying the *relevant* facts of the initial situation, and several resulting state worlds satisfying both S and G reached by ac_1, since the other agents may choose a range of actions which do not affect these aspects of the state. This already improves on (Atkinson and Bench-Capon 2007), which required the reasoner to identify a particular state for the initial circumstances and the resulting circumstances. The first imposed an unrealistic, and unnecessary, epistemic burden, while the second gave rise to irrelevant critical questions relating to the consequences of the action, pointing to differences which were not relevant to the justification.

We can now suggest a variant scheme which justifies the action in terms of its intrinsic merits, rather than its consequences. For example I might decide to go for a walk around the block to enjoy the sunshine: I am not trying to achieve any goal, the activity itself promotes the value of enjoyment.

AS1a *In the current circumstances R, Agent 1 should perform action ac_1, which will promote value V.*

The conditions for an instantiation of AS1a to justify ac_1 are:

Definition 11: *There is an AS1a justification for ac_1 if there is some $r = \langle u_i, e, u_k\rangle \in \mathcal{R}_{\mathcal{H}}$ such that $u_i = \langle q_i, t_0\rangle, e = \langle j, t_0\rangle$ and $u_k = \langle q_j, t_1\rangle$ such that for all $p \in R, p \in \pi{*}(u_i)$, $ac_1 = j_1$, and $\delta * (u_i, e, v) = +$.*

A second limitation of the original scheme is that it considered only the next state. We can now extend the time horizon. Here we will want to argue that an action performed now will at some time in the future help to realise a goal, or enable an action. When we look beyond the next time point, however, it may be that we could perform actions with no relevance to the ultimate goal, without jeopardising its eventual attainment. The justification, however, should only apply to relevant actions. For this reason we will initially consider not the sufficient condition scheme for practical reasoning, but rather the other scheme for practical reasoning given in (Walton 1996), namely the *necessary condition scheme for practical reasoning*:

G is a goal for agent, Doing A is necessary for agent to carry out G, Therefore, agent should do A.

We will make the same distinctions with respect to separating the goal and the value as was made by (Atkinson and Bench-Capon 2007) with respect to the sufficient condition scheme. We therefore state the scheme as:

AS2 *In the current circumstances R, Agent 1 should perform action ac_1, since otherwise goal G will not be realised, and realising G would promote value V.*

In terms of the action-state semantics this can be properly instantiated if for all G-futures which are also V-futures, the first term is some $r \in \mathcal{R}_{\mathcal{H}}$ such that $r = \langle u_i, e, u_k\rangle$ and $ac_1 = j_1$.

A similar scheme can be produced relating to some future action, as in ‘I should book a ticket so that I can go to the theatre tomorrow’:

AS2a *In the current circumstances R, Agent 1 should perform action ac_1, since otherwise it will not be possible to perform ac_2, which would promote value V.*

This can be instantiated if for all ac_2-futures (i.e. futures in which ac_2 is the final action) which are also V-futures, the first term is some $r \in \mathcal{R}_{\mathcal{H}}$ such that $r = \langle u_i, e, u_k\rangle$ and $ac_1 = j_1$.

Although AS2 and AS2a are legitimate arguments to use in practical reasoning, the necessary condition scheme interpreted in this way is rather restrictive. It may well be essential to book a ticket if I am to attend the theatre tomorrow, but it need not be done *now*. We therefore allow an action to be justified if it needs to be performed *at some time* in order to promote the value. Thus if I need to book a ticket to go to the theatre, I need not do it now, but I must do it at some time, and now may be as good as any. Thus AS2 may be instantiated if for all G-futures which are V-futures, there is a *sub-future* which is an ac_1-future, and AS2a can be instantiated if for all ac_2-futures which are V-futures there is a *sub-future* which is an ac_1-future.

Example

In this section we present an example to demonstrate the use of our action-state semantics. The particular problem scenario is discussed in terms of an AATS, then represented in terms of our action-state semantics, from which we make clear the benefits that the proposed new representation brings.

Our example is of a person waiting for a train. Currently he is on the platform, but there is no sign of the train. He would like some refreshment, and could go to the buffet but it is on a distant platform, and if the train came he would miss it. In his bag he has a novel, which he is looking forward to reading, and a draft thesis from a student, which he is not. He thus has a choice of three actions, go to buffet, read novel or read thesis. The train may arrive or not, giving six joint actions. The states of interest are whether he is on the platform, and whether the train is on the platform. Notice that in formulating the problem we try to keep the number of states to the minimum necessary, and so we do not need to discriminate between being in the buffet or anywhere else: all that matters is whether or not he is on the platform. We thus have four states, both off platform (q_2 = 00), both on the platform (q_1 = 11), person on platform and train not, (the initial state q_0 = 10) and train on platform and person not (q_3 = 01). The transitions are straightforward: thesis and no train (j_0) and novel and no train (j_1) move from 10 to 10, thesis and train (j_2) and novel and train (j_3) move from 10 to 11, buffet and train (j_4) moves from 10 to 01 and buffet and no train (j_5) moves from 10 to 00. In the buffet there is the additional action of returning to the platform: if the train arrives (j_6) this will move from 00 to 11, and if it does not (j_7), this will move from 00 to 10. We will relate duration to the time taken to reach the buffet and have a drink. We will say that having a drink takes the same time as the journey between platforms, so that j_4 and j_5 have duration 2 and all other actions duration 1, since we can divide periods of reading into any length we choose.

Values are also straightforward: reading the thesis promotes Duty (D), reading the novel promotes Enjoyment (E), going to the buffet promotes Refreshment (R), catching the train promotes Punctuality (P) and missing the train demotes P. So moving from q_0 to q_2 promotes R and moving from q_0 to q_3 promotes R but demotes P. Moving from q_0 to q_1 promotes P, but whether it promotes D or E also depends on the choice of action. Similarly if the person stays in q_0 the value promoted will depend on what was done. In order to represent this in an AATS we would need to be able to discriminate between the state where the novel had been read and the the state where the thesis had been read. So we might include propositions for 'chapter 1 of novel read' etc. This would lead to a major proliferation of states, and is not very natural: it is the *reading* itself that gives the enjoyment, not having read a chapter. In contrast consider the action-state representation shown in Figure 1. This does give a very natural representation. The transitions are labelled with the values promoted and demoted: note that where the action promotes the value the labelled transition is from a state world to an event world, and where the state promotes the value, it is the transition from event world to state world

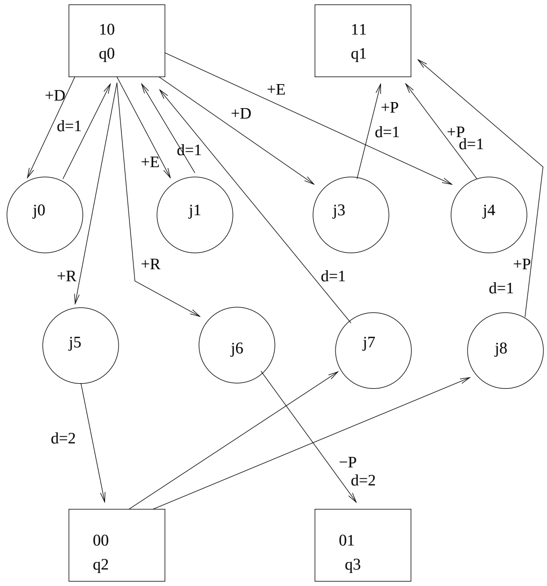

Figure 1: Action-state semantics representation of the example scenario

that is labelled. The transitions from event worlds to state worlds are also labelled with the duration appropriate to the event world. Note that each state corresponds to a family of possible worlds, the time being t_0 initially, and the subscript increasing each time an edge labelled with a duration is traversed.

We can now consider the arguments that the agent can make in the scenario shown in Figure 1. Space precludes a full enumeration of arguments and critical questions, but we will sketch how our four schemes can be instantiated. The initial situation, *now*, is $\langle q_0, t_0 \rangle$. AS1 is based on a value promoted in the next state world. Since P is promoted by either j_3 or j_4, we will have two arguments based on this scheme suggesting that the agent should perform its component of j_3 and j_4 respectively:

A1: *In* q_0 *we should read thesis to reach* $q1$ *which promotes P*

A2: *In* q_0 *we should read novel to reach* $q1$ *which promotes P*

But we also have several transitions to event wolds promoting values, allowing the instantiation of AS1a:

A3: *In* q_0 *we should read thesis which promotes D*

A4: *In* q_0 *we should read novel which promotes E*

A5: *In* q_0 *we should go to buffet which promotes R*

There are no instantiations of AS2 in q_0, but there is an argument based on AS2a, since if the train arrives at t_1, it will not be possible to go to the buffet unless we do so immediately.

A6: *In* q_0 *we should go to buffet since otherwise it might not be possible to go the buffet which would promote R*

For an example of AS2, suppose we go to the buffet. Now we must return to the platform, or we will never catch the train:

A7: *In* q_3 *we should return, since otherwise* q_2 *may not be* *reached, and q_2 promotes P.*

The arguments will be subject to critical questioning: for example we may reject A1 and A2 on the grounds that the train is unlikely to arrive at t_1. A3, A4 and A5 have the objection that they exclude one another. Additionally, A5 has the potential problem that the train may arrive at t_1 or t_2, and so demote P, but we may choose to reject this argument. For example, if the train is not expected until t_3: that is we believe we have time to go to the buffet and return before the train arrives. Similarly, A6 turns on when we expect the train and how cautious we wish to be.

Once we have discounted those arguments which do not survive their critical questions, we can use the preferences of the particular agent over values to resolve the resulting argumentation framework. As one might expect, going to the buffet will only be chosen if R is preferred to P, and otherwise whether the thesis or the novel is read will depend on the relative valuation of E and D. Note that when we return to q_0 we will be in a new situation. Arguments against A5 and A6 will be stronger as the train is increasingly likely to arrive as time moves forward. Also preferences may change: once the buffet has been visited the importance of R will decrease, and if the novel has been read for a while the call of duty may grow in strength.

This small example illustrates the benefits of the proposed representation. It captures important temporal aspects. If, following (Amgoud, Devred, and Lagasquie-Schiex 2008) we define plans in terms only of their starting and end points, we can reach q_1 by reading either the thesis or the novel, or by going to the buffet and returning. Abstracting from the individual actions, the effect of the actions of the other agents and the duration of the actions in the plans makes it impossible to discriminate between them, since the important differences lie in the values promoted by the actions themselves, not the state they reach. Of course, the underlying logic could be extended to a temporal logic, but this would still not accommodate the intrinsic worth of actions. Moreover, the ability to vary the actions if the train does not arrive immediately is obscured.

Allowing the use of argument scheme AS1a also improves on (Atkinson and Bench-Capon 2007), since there is no need to say whether reading will reach q_1 or q_0: since the state reached has no influence on the value promoted this is as it should be. Further the availability of AS2 allows us to produce an argument such as A7, even where the action promotes no value by reaching the successor state.

Concluding Remarks

In this paper we have presented an action-state semantics to act as the basis for practical reasoning using argument schemes. This new formalism solves problems with current approaches to practical reasoning which use only states, both with regard to actions which are performed for their own sake, and with regard to temporal considerations. In previous work these issues could be handled only (if at all) with a considerable degree of contrivance and an unnecessary proliferation of states and actions. Using this formalism, and exploiting the greater expressiveness afforded by the additional three argument schemes it supports will allow the modelling of larger and more realistic problems than has been possible so far. In particular this formalism will help with situations where there are temporary windows of opportunity, where coordination is required, and where the likelihood of an event varies with time.

The mapping from AATS to action-state semantics also has application beyond practical reasoning. The semantics were designed in (Hamblin 1987) and (Reed and Norman 2007) to support reasoning about imperatives. A recent paper (Atkinson et al. 2008) proposed a means of critiquing commands using an argumentation scheme and critical questions approach. The semantics described here will provide a sound basis for further exploration of this topic also, allowing not only states to be ordered, but also actions, and particular ways of bringing states about. For future work we would like to investigate the relation of our approach to *stit* (seeing to it that) semantics (Belnap and Perloff 1988).

References

Alur, R.; Henzinger, T. A.; and Kupferman, O. 2002. Alternating-time temporal logic. *ACM Journal* 49 (5):672–713.

Amgoud, L.; Devred, C.; and Lagasquie-Schiex, M.-C. 2008. A constrained argumentation system for practical reasoning. In *Proceedings of AAMAS-08*, 429–436.

Atkinson, K., and Bench-Capon, T. J. M. 2007. Practical reasoning as presumptive argumentation using action based alternating transition systems. *Artificial Intelligence* 171(10–15):855–874.

Atkinson, K.; Girle, R.; McBurney, P.; and Parsons, S. 2008. Command dialogues. In *Proceedings of ArgMAS 2008*, 9–23. AAMAS 2008.

Belnap, N., and Perloff, M. 1988. Seeing to it that: A canonical form for agentives. *Theoria* 54(3):175–199.

Dung, P. M. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence* 77:321–357.

Hamblin, C. L. 1987. *Imperatives*. Oxford, England: Blackwell.

Rahwan, I., and Amgoud, L. 2006. An argumentation based approach for practical reasoning. In *Proceedings of AAMAS-06*, 347–354.

Reed, C., and Norman, T. J. 2007. A formal characterisation of Hamblin's action-state semantics. *Journal of Philosophical Logic* 36:415–448.

Walton, D. N. 1996. *Argumentation Schemes for Presumptive Reasoning*. Mahwah, NJ, USA: Lawrence Erlbaum Associates.

Wooldridge, M., and van der Hoek, W. 2005. On obligations and normative ability: Towards a logical analysis of the social contract. *Journal of Applied Logic* 3:396–420.

Wooldridge, M. 2000. *Reasoning About Rational Agents*. Cambridge, Mass: MIT Press.