Partial Order Reasoning for a Nonmonotonic Theory of Action

Matthew Stone
University of Pennsylvania
matthew@linc.cis.upenn.edu

Abstract

This paper gives a new, proof-theoretic explanation of partial-order reasoning about time in a nonmonotonic theory of action. The explanation relies on the technique of *lifting* ground proof systems to compute results using variables and unification. The ground theory uses argumentation in modal logic for sound and complete reasoning about specifications whose semantics follows Gelfond and Lifschitz's language $\mathcal{A}$. The proof theory of modal logic represents inertia by rules that can be instantiated by sequences of time steps or events. Lifting such rules introduces string variables and associates each proof with a set of string equations; these equations are equivalent to a set of partial-order tree-constraints that can be solved efficiently. The defeasible occlusion of inertia likewise imposes partial-order constraints in the lifted system. By deriving an auxiliary partial order representation of action from the underlying logic, not the input formulas or proofs found, this paper strengthens the connection between practical planners and formal theories of action. Moreover, the general correctness of the theory of action justifies partial-order representations not only for forward reasoning from a completely specified start state, but also for explanatory reasoning and for reasoning by cases.

Recent advances in planning [McAllister and Rosenblitt, 1991; Pednault, 1991; Penberthy and Weld, 1992] raise the possibility of applying streamlined planning techniques to actions in domains with rich characterizations. At the same time, formal theories of action based on nonmonotonic logic are maturing, achieving validations for a rich range of specifications [Lin and Shoham, 1991; Kartha, 1993; Sandewall, 1994; Baral and Gelfond, to appear]. Ideally, planning researchers could draw on the formal tradition to guide their extensions. In practice, there is a gap, because planners for simple domains already use representations that formal theories cannot account for. Chief among these is the use of a *partial-ordering* to reason about the temporal relationship of planned actions. The partial-order representation is particularly challenging to describe formally because it figures neither in the domain specification given to planners nor in the final output plan.

This paper bridges the gap with a new, proof-theoretic explanation of partial-order reasoning about time in a provably correct nonmonotonic theory of action. Like [McAllister and Rosenblitt, 1991], this account depends on the technique of *lifting* proof systems. Lifting converts a *ground* system, in which instantiated terms are explicit and complete, into an equivalent *lifted* system, in which instantiated terms are represented by variables and assigned values using equality constraints. McAllester and Rosenblitt's construction applies to the terms that represent actions; there, lifting allows variables to abstract the objects to which actions apply. Here, the construction applies to terms that represent possible states, and lifting allows variables to abstract sequences of steps during which a fluent remains unchanged. Unlike terms representing actions, terms representing states are governed by an equational theory. Such equational theories provide the best known treatment of introspection axioms, such as the axiom of inertia, in modal logic [Wallen, 1990; Ohlbach, 1991]; the equational theory gives terms a correct structure as sequences, so that rules like inertia can apply to sequences in a single step of instantiation. Thus, in the lifted system, *string equalities* describe the conditions on variables under which a proof is correct. Using results derived for more general problems in modal deduction, we can show that these equational unification problems are equivalent to a set of partial-order tree-constraints that can be solved efficiently. The defeasible occlusion of inertia likewise imposes partial-order constraints in the lifted system.

By deriving an auxiliary partial order representation of action from the underlying logic, not the input formulas or proofs found, this paper strengthens the connection between implemented planners and formal theories of action. Moreover, the general correctness of the theory of action justifies partial-order representations not only for forward reasoning from a completely specified start state, but also for explanatory reasoning and for reasoning by cases.

The presentation given here involves the use of a rich logic as an intermediate representation in which reasoning is performed. The use of this intermediate logic may seem somewhat distracting, given the simplicity of the inferences needed here. The intention, however, is to lay the groundwork for range of exciting extensions to this approach by explicitly presenting the basic results in their full generality. Modal logic is expressive enough to capture the inferences needed for theories with concurrent actions [Baral and Gelfond, to appear], with taxonomic knowledge (cf. [Baral, to appear]), and even with explicit representations of an agent's knowledge at different points in a plan (as in [Moore, 1985]). In each case, once the correspondence of a ground modal proof system and its intended models has been established, the lifting argument described here will explain how efficient constraint algorithms can be deployed in proof search.

The results also help clarify the methodology of validating nonmonotonic theories of action, in the following sense. A key step of validation is to devise an intuitive system for calculating consequences and to describe it precisely as a set of intended models. This set of intended models has a straightforward, classical characterization. Specifying these intended models might seem to obviate the nonmonotonic theory entirely: Once you are precise about what your intended models are, why not just reason with your classical description of them? Indeed, this is just what is proposed in compiling ADL frame axioms, or adopting explanation closure axioms [Pednault, 1994; Schubert, 1990; Reiter, 1991].

What the caricature leaves out is the advantage of the nonmonotonic formalism in proof-theoretic properties. The nonmonotonic system combines logical proof with some simple metatheoretic notion for 'combining' proofs to take into account the best evidence. Separating these steps can help clarify the grounds for optimizations and the sources of complexity in planning. For example, in our nonmonotic system we find that partial-order reasoning is actually a sound, complete and efficient optimization at the level of constructing arguments, so that we need only invoke heuristics or search to order actions when interacting arguments are compared.

1 A New Validation

This section describes and validates a simple formalism for reasoning about action that uses a framework of argumentation to combine competing deductions in modal logic. The definitions and proofs of this validation owe an obvious debt to previous validations, particularly the program of *logic programming theory of action* represented by [Gelfond and Lifschitz, 1993; Kartha, 1993; Baral and Gelfond, to appear].

Why the combination? Modal logic represents the inferences we need explicitly but compactly. We use modal operators to describe both change and persistence; we use modal introspection axioms to capture inertia. ([Ginsberg 1995] also suggests an analogy between frame and modal operators.) Argumentation leaves the leap to planners small; [Ferguson, 1995] has observed a parallel between argumentation and the SNLP planner. More importantly, it allows us to reason with standard proofs. This design allows the proof theory of modal logic, already well-studied, to give interesting insights into the system.

The particular formalism presented here uses prefix semantic translations for modal logic proof [Wallen, 1990; Ohlbach, 1991] and Dung's presentation of argumentation frameworks for defeasible reasoning [1993]. However, modal logic is a general logic of possible states (see [Halpern and Moses, 1985]), while argumentation provides a general framework for defeasible reasoning (see [Lin and Shoham, 1989; Pollock, 1992; Simari and Loui, 1992]). Thus, we expect the techniques presented here to continue to apply as richer nonmonotonic theories of actions are constructed and validated.

1.1 A Modal Semantics and a Modal Translation

This section describes a set of intended models for a theory of action; it is based roughly on the semantics of $\mathcal{A}$ from [Gelfond and Lifschitz, 1993]. We introduce a set F of fluent names and a set A of action names; a fluent literal has the form f or $\neg f$, with $f \in F$. (The opposite of a literal $\sim f$ is g if f is $\neg g$, $\neg f$ otherwise.) A domain theory is specified by a set R of causal rules and a set O of observations. A causal rule takes the form:

$$a \textbf{ causes } f \textbf{ if } P_1, \ldots, P_n$$

where a is an action name, and f and all P_i are fluent literals. An observation takes either of the forms:

$$a \textbf{ happens at } t \quad f \textbf{ holds at } t$$

where a is an action name; f is a fluent literal; and t is a natural number.

The models we work with are Kripke models $\langle W, AF \rangle$ where W is a set of worlds and AF is a nonempty set of functions from worlds to worlds encoding accessibility. A world is represented as a pair $\langle S, E \rangle$ where S (the state) is a set of fluent names and E (the event) is a set of action names; a fluent f holds at $\langle S, E \rangle$ iff $f \in S$; $\neg f$ holds at $\langle S, E \rangle$ iff $f \notin S$; and an action a happens at $\langle S, E \rangle$ iff $a \in E$.

Definition 1. A model $\langle W, AF \rangle$ represents an *inertial model* of causal rules R if each function $\alpha \in AF$ respects the following constraints for any state $\langle S, E \rangle$:

- For any rule in R of the form a **causes** f **if** $P_1, \ldots, P_n$: if a happens at $\langle S, E \rangle$ and P_1 through P_n hold in $\langle S, E \rangle$, then f holds at $\alpha(\langle S, E \rangle)$.
- Otherwise, f holds at $\alpha(\langle S, E \rangle)$ iff f holds at $\langle S, E \rangle$.

Two actions may occur concurrently if their effects do not interfere with each other; otherwise accessibility functions respect the inertial meaning of causal rules just as transition functions do in [Gelfond and Lifschitz, 1993].

Definition 2. A model $\langle W, AF \rangle$ is a model of an observation at a world w according to the following criteria:

- f **holds at** *1* at w iff f is true at w; otherwise f **holds at** t iff for all $\alpha \in AF$, f **holds at** $t - 1$ is true at $\alpha(w)$.

$$
\begin{aligned}
(a\ \textbf{happens at}\ t)^T = &\ [\mathrm{N}]^{t-1}\ \mathbf{h}a \\
(f\ \textbf{holds at}\ t)^T = &\ [\mathrm{N}]^{t-1}\ [\mathrm{H}]\ f \\
(a\ \textbf{causes}\ f\ \textbf{if}\ P_1, \ldots, P_n)^T\ (l) = & \\
&\ [\mathrm{H}]\ (P_1 \wedge \ldots \wedge P_n \wedge \mathbf{h}a \supset [\mathrm{N}]\ [\mathrm{H}]\ f) \\
&\ \wedge [\mathrm{H}]\ (\mathbf{h}a \supset ab(\sim f, l))
\end{aligned}
$$

Figure 1: Translation $\cdot^T$ to modal logic

- a **happens at** *1* at w iff a happens at w; otherwise a **happens at** t iff for all $\alpha \in AF$, a **happens at** $t-1$.

Definition 3. A model $\langle W, AF\rangle$ is an *exact model* of O at w iff it is a model of O at w, and for every finite sequence $\alpha_1, \ldots, \alpha_m$ of elements of AF (possibly empty), whenever any action a happens at $\alpha_1(\ldots(\alpha_m(w))\ldots)$, there is an observation a **happens at** $m+1$.

Definition 4. A model $\langle W, AF\rangle$ is an *intended model* of R and O at w, iff it is an inertial model of R and an exact model of O at w.

Definition 5. R and O *entails* an observation o iff every intended model of R and O at any w is a model of o at w.

The truth-conditions above mirror those of modal formulas. For example, if $[\mathrm{N}]$ (next) is a modal operator interpreted with reference to accessibility functions AF, and the formula $\mathbf{h}a$ is true of an action a in world w iff a happens there, then a **happens at** t corresponds to the formula $[\mathrm{N}]^{t-1}\ \mathbf{h}a$.

The role of holds in inertia can also be modeled as a modal operator. If f holds at a certain world, then f is true there, and f will continue to hold in accessible worlds until further notice. By this analogy, we introduce a modality $[\mathrm{H}]$ to represent *holding until further notice*; it is subject to the formal axioms $[\mathrm{H}]\ f \supset f$, $[\mathrm{H}]\ f \supset [\mathrm{H}]\ [\mathrm{H}]\ f$ and $[\mathrm{H}]\ f \supset [\mathrm{N}]\ f$. Since $[\mathrm{H}]$ only has these properties "by default", we won't assign models an explicit set of accessibility functions to interpret it. Instead, we simply use this intuition in designing the representations of the proof system.

The proof system represents assumed facts by a translation $\cdot^T$. Under the translation, f **holds at** t becomes $[\mathrm{N}]^{t-1}\ [\mathrm{H}]\ f$. Causal rules are interpreted (1) by a clause establishing effects: $[\mathrm{H}]\ (P_1 \wedge \ldots \wedge P_n \wedge \mathbf{h}a \supset [\mathrm{N}]\ [\mathrm{H}]\ f)$; and (2) by a clause triggering abnormality formulas for occlusion: $[\mathrm{H}]\ (\mathbf{h}a \supset ab(\sim f, l))$. (Each causal law is given a distinct symbol l to index the source of occlusion.) These translations are summarized in Figure 1. Meanwhile, to *prove* an observation, the weaker translation $(f\ \textbf{holds at}\ t)^Q = [\mathrm{N}]^{t-1}\ f$ suffices.

This approach is similar to other translations of $\mathcal{A}$. The differences are the distinctive use of modal operators to encode inertia, the use of facts to encode the occurrence of actions, and the elimination of occurrences of negation-as-failure that appear in logic programming theories of action. The effect of negation-as-failure will be restored by an explicit operation of comparing deductions. This allows an ordinary modal logic proof system to apply in the base case.

$$\frac{\mu = \nu}{\Gamma, f^{\mu} \longrightarrow f^{\nu}, \Delta}\,\text{axiom} \qquad \frac{\mu\nu = \sigma}{\Gamma, f^{\mu}, \neg f^{\mu} \longrightarrow \bot^{\sigma}, \Delta} \rightarrow \bot$$

$$\frac{\Gamma \longrightarrow A \wedge B^{\mu}, A^{\mu}, \Delta \qquad \Gamma \longrightarrow A \wedge B^{\mu}, B^{\mu}, \Delta}{\Gamma \longrightarrow A \wedge B^{\mu}, \Delta} \rightarrow \wedge$$

$$\frac{\Gamma, A \wedge B^{\mu}, A^{\mu}, B^{\mu} \longrightarrow \Delta}{\Gamma, A \wedge B^{\mu} \longrightarrow \Delta} \wedge \rightarrow \frac{\Gamma \longrightarrow A \vee B^{\mu}, A^{\mu}, B^{\mu}, \Delta}{\Gamma \longrightarrow A \vee B^{\mu}, \Delta} \rightarrow \vee$$

$$\frac{\Gamma, [\mathrm{H}]\, f \longrightarrow \Delta \qquad \Gamma, [\mathrm{H}]\, \neg f \longrightarrow \Delta}{\Gamma \longrightarrow \Delta}\,\text{cut}$$

$$\frac{\Gamma, A \supset B^{\mu} \longrightarrow A^{\mu}, \Delta \qquad \Gamma, A \supset B^{\mu}, B^{\mu} \longrightarrow \Delta}{\Gamma, A \supset B^{\mu} \longrightarrow \Delta} \supset\rightarrow$$

$$\frac{\Gamma, [\mathrm{N}]\, A^{\mu}, A^{\mu\alpha} \longrightarrow \Delta}{\Gamma, [\mathrm{N}]\, A^{\mu} \longrightarrow \Delta}[\mathrm{N}] \rightarrow \frac{\Gamma, [\mathrm{H}]\, A^{\mu}, A^{\mu\sigma} \longrightarrow \Delta}{\Gamma, [\mathrm{H}]\, A^{\mu} \longrightarrow \Delta}[\mathrm{H}] \rightarrow$$

$$\frac{\Gamma \longrightarrow [\mathrm{N}]\, A^{\mu}, A^{\mu\alpha}, \Delta}{\Gamma \longrightarrow [\mathrm{N}]\, A^{\mu}, \Delta} \rightarrow [\mathrm{N}]^{\dagger}$$

Figure 2: Path-based, explicitly-scoped sequent calculus for modal logic. † For $(\rightarrow [\mathrm{N}])$, α must not appear in the conclusion.

1.2 Proof Theory

Figure 2 shows an explicitly-scoped proof system that we will use to construct arguments about entailment in intended models. It is based on [Wallen, 1990; Ohlbach, 1991; Auffray and Enjalbert, 1992]. As in these systems each formula is labeled with a string that represents the path of accessibility to the possible world where the formula must be shown true. Figure 2 is specialized to describe the modal fragment for inertia more precisely. The axiom rule applies to any fluent literals or atomic action occurrence statement. The $(\rightarrow \bot)$ rule encodes how contradictory fluents derive a contradiction. Since the contradiction records the paths of the terms that introduce it, we must test for contradiction "when all is said and done". The cut rule formalizes the fact that at the initial state, either f or $\neg f$ is true, and whichever is true will tend to stay that way. The cut is tractable because it applies only to fluent names and can be restricted to "pseudo"-analytic uses—cases where either f or $\neg f$ is a subformula of the sequent already [D'Agostino and Mondadori, 1994]. Finally, the left modal rules build in the axioms relating them: [N] matches any constant; [H] matches any string.

Given a set of causal rules R and observations O, translated to logic via the rules in Figure 1 as R^T and O^T, we can analyze the structure of deductions ending with a sequent $R^T, O^T \longrightarrow \Delta$ to restrict the kinds of proofs that need to be constructed and compared, to a space that can be searched easily, cf. [Gallier, 1986; Kleene, 1951; Miller *et al.*, 1991]. For any deduction has an equivalent form where cuts occur at the root and only one formula appears on the right in sequents. In such deductions, which we shall term *direct*, we can distinguish as an *attack site* each subproof that is an input to a (cut) rule but which is not itself obtained by applying a (cut) rule. Without loss of generality, the end-sequent $\Gamma \longrightarrow A$ of an attack site gives a representative set of assumptions Γ under which any conclusion obtained anywhere in the subproof may be challenged.

1.3 Argumentation

Following [Dung, 1993], an *argumentation framework* as a pair $\mathcal{F} = \langle AR, attacks \rangle$ where *AR* is a set (of arguments) and *attacks* is a binary relation on elements of *AR*. For two arguments *D* and *E*, $attacks(D, E)$ means that *D* argues against the acceptability of *E*.

We take *AR* to be the set of direct modal proofs as given in the previous section. *attacks* is defined as the union of two relations, $attacks_{\mathrm{I}}$ for inertia, and $attacks_{\mathrm{O}}$ for occlusion.

Definition 6. $attacks_{\mathrm{I}}(D, E)$ if *D* has end-sequent $\Gamma \longrightarrow ab(f, l)^{\nu}$, *E* has an attack site with end-sequent $\Gamma \longrightarrow \Delta$, and the attack site contains a rule application $([\mathrm{H}] \rightarrow)$ deriving fluent literal $f^{\lambda\mu}$ from $[\mathrm{H}]\, f^{\lambda}$, where λ is a prefix of ν and ν is a proper prefix of $\lambda\mu$.

Because inertia is handled by an introspection axiom, we can propagate a fluent forward inertially from the result-situation of the action that establishes the fluent to the result of an arbitrary sequence of subsequent actions in a single step of instantiation. However, propagation into the result state of each action is subject to occlusion by abnormality; an abnormality at a given step challenges both application of inertia at that step and subsequent inertial propagation of the fluent. Definition 6 encodes this.

Definition 7. $attacks_{\mathrm{O}}(D, E)$ if *D* has end-sequent $\Gamma \longrightarrow \sim P_i^{\mu}$ and *E* has an attack site with end-sequent $\Gamma \longrightarrow ab(f, l)^{\mu}$ for a rule *l* with precondition P_i.

An argument that an action occludes a fluent is challenged by showing that the conditions where the action occludes the fluent are not actually met, because some fluent P_i is sure not to hold.

To describe the consequences of an argumentation framework requires the following definitions [Dung, 1993].

Definition 8. An argument *D* is *acceptable* wrt a set *S* of arguments iff for each argument *E* in *AR*: if $attacks(E, D)$ then there is a D' in *S* with $attacks(D', E)$.

Definition 9. A set *S* of arguments is *admissible* if there are no arguments *D* and *E* in *S* with $attacks(D, E)$, and every argument in *S* is acceptable wrt *S*.

Definition 10. A *preferred extension* of an argumentation framework *AF* is a maximal (wrt set inclusion) admissible set of arguments of *AF*.

Dung proves that every argumentation framework in which no infinite sequence of attacks is possible has a unique preferred extension, which can be obtained by a least fixed-point construction. This extension, denoted GE_{AF}, represents the natural consequences of the framework. Dung's theorem applies to attacks defined by definitions 6 and 7, because each argument can be assigned a finite *grade*, based on the prefixes that appear in it, such that only arguments of lower grade attack it. For, $ab(f, l)^{\mu}$ arguments can be attacked by arguments for fluents true at μ, but such arguments can only be attacked in turn by other $ab(f', l')^{\nu}$ arguments with ν a proper prefix of μ.

Argumentation is closely connected to logic programming. Search for an acceptable argument in GE_{AF} can be captured as a logic program by the meta-interpreter:

(1) $acc(D) \leftarrow not\ defeated(D).$
$defeated(D) \leftarrow attacks(D, E), acc(E).$

1.4 Validation

We can now prove the following theorem:

Theorem 1. (Correctness) Let O be a set of observations in which the latest time mentioned is t, and let R be a set of rules. Then R, O entails every observation from a finite set O' iff GE_{AF} contains an argument with end-sequent $R^T, O^T \longrightarrow (\bigwedge_{o \in O'} o^Q) \vee [\mathrm{N}]^t \perp$.

The formal statement and proof is omitted for space constraints. Soundness is proved directly, by double induction on the grade of arguments and the structure of proofs. Completeness appeals to a lemma that the models of R and O can be partitioned into a finite set of types, where each type is a finite description of the initial state that determines what formulas change when during the observations. Thus we can apply the cut rule repeatedly until each attack site in the derivation specifies the type of all of the models of the attack site. We then show that by induction that if all of the models have the same type, then cut-free inference is complete: the observations can be used to determine which causal rules are triggered by each action.

2 Lifting the approach

This section applies results from modal proof theory to derive a partial-order representation for actions in arguments. We couple this with a formal analogue to lifting in the argumentation framework as a whole. In the lifted framework, argument *construction* requires imposing partially-ordered constraints on actions to *enable inertial propagation* of fluents; argument *comparison* requires imposing constraints on the partial order to *disable occlusion attacks* on inertial arguments. By considering how these constraints arise, we discover that the steps required to construct and check an acceptable argument in a preferred extension correspond to the causal-linking and threat-resolution strategies of partial-order planners like UCPOP [Penberthy and Weld, 1992].

2.1 Lifting proofs

Proofs in modal logic can be lifted to use unification using standard proof-theoretic techniques [Gallier, 1986; Ohlbach, 1991; Auffray and Enjalbert, 1992]. Skolem terms are used to represent arbitrary [N] transitions on the right in sequents. On the left, string variables whose values must have length 1 are used to represent [N] transitions, while unconstrained string variables are used to represent [H] transitions. Equalities checked at the leaves of the proof in ground derivations become equality constraints that the values of variables must satisfy. In the case of observations $[\mathrm{N}]^t\varphi$, the path to a state at which φ is true can be represented by a single variable whose value must have length t. Thus, we can schematize over different values of t by introducing an arbitrary path x_t and ensuring that all paths to φ have the same length as x_t.

Further proof-theoretic analysis shows that these equalities and constraints are equivalent to a simpler set of partial-order constraints that describe a tree of paths of accessibility [Stone, 1997]. This analysis, which applies more generally to fragments of modal logic, involves three principles. (1) By constructing proofs in an order that applies right modal rules before left ones, we can ensure that modal Skolem terms are simply constants (even if we have first-order quantifiers). (2) Each individual equation can be correctly represented by equality, precedence and length constraints, because each variable and constant is preceded by a unique prefix. Constraints of *agreement* are needed for full paths only, not for variables generally. (3) In any solution, each variable must be bound to a sequence of constants introduced earlier in the proof—this is because all modal constants are distinct and introduced on the right. Hence, the individual solutions can be reconciled by constraining modal constants introduced later never to precede constants introduced earlier. In fact, since the constraints are simple and local, the simplest structure solving all of them can be maintained using an efficient, relaxation-style algorithm modeled after the tree-construction algorithm of [Aho *et al.*, 1981].

Arguments explicitly record causal links. Each such link records that the effect f of one action a establishes the precondition of another action, b. In the argument, the link is represented, in part, by an $(\supset\rightarrow)$ rule whose antecedent has subformula $\mathbf{h}a$ and whose consequent establishes [N] [H] f^x; in part, by ([N] $\rightarrow$) and ([H] $\rightarrow$) rules instantiating [N] [H] f^x to f^{xuy}; and finally by an leaf equating f^{xuy} with a formula f^z, which a lower rule combines with $\mathbf{h}b^z$. As a partial-order constraint, $xuy = z$ corresponds to $x < z$.

2.2 Lifting argumentation

We can complement the lifting of deduction with a lifting of argumentation in general. Technically, this is achieved by relativizing notions of attack and acceptability for lifted arguments to particular substitutions of values to variables. A longer version of this paper will present the definitions and properties of the lifted system in full detail. The behavior of this system still corresponds to the logic program of (1). However, the system now explores sets of ground argument simultaneously, by using lifted arguments and accumulating constraints on the values of variables in them. This requires operationalizing negation by *constraining* a potentially defeated lifted argument *away from* substitutions where it would be defeated. Thus, at each step we wish to find sets of constraints under which we can show $acc(D)$, given that constraints C already apply to D. To do so, we look for any proof of $defeated(D)$ with extended constraints C', and then nondeterministically add appropriate constraints to C to obtain a new set of constraints C'' that are inconsistent with C'. C'' is determined by where the new constraints in C' are imposed in showing $defeated(D)$.

(1) Additional constraints may be added to C' to ensure $attacks(D, E)$. (1a) Suppose that $E\sigma$ can attack inertia in $D\sigma$ at some substitution σ satisfying the constraints C', according to definition 6. Then D derives $f^{\lambda x}$ from [H] f^{λ} while E derives $ab(f, l)^{\nu}$, and $\lambda\sigma < \nu\sigma < (\lambda x)\sigma$. Now, without loss of generality, we may assume that the variable representing ν does not appear in D; so we cannot constrain it directly. However, by the structure of arguments, we know that ν matches the occurrence of some action $\mathbf{h}a^{\nu}$. Since attacking arguments agree on the occurrence of actions, there *is* a way to define a variable $\hat{\nu}$ in D with the same *length* as ν must have. Thus, we can rule out σ by adding one of $\hat{\nu} \le \lambda$; or $\lambda x \le \hat{\nu}$. (Letting $|x|$ denote the length of x, the alternatives are $|\nu| \le |\lambda|$ or $|\lambda x| \le |\nu|$.) These strategies describe the only alternatives to avoid the attack by constraining paths; these strategies correspond to promotion and demotion in planners.

(1b) Suppose that $E\sigma$ can attack occlusion in $D\sigma$ at some substitution σ satisfying the constraints C', according to definition 7. Then D establishes $ab(f, l)^{\nu}$ and E counters by ruling out the application of l, by instantiating the effect [N] [H] P^{λ} of some action a to $(\lambda ux)\sigma = \nu\sigma$. To defuse this, the only option is to add the constraint $|\nu| \le |\lambda|$.

(2) Additional constraints may be added to C' to ensure $acc(E)$. Then we must have some acceptable argument F which attacks E, given some set of constraints C'' extending C. It suffices to impose this set of constraints C'' on the overall argument D. In planning, this captures one of UCPOP's separation strategies. In separation, F and E fall into case (1b), so we restore the equality $\lambda ux = \nu$—or, the constraint $\lambda < \nu$.

Rules:

(1) $puton(X,Y)$ **causes** $on(X,Y)$
if $clear(X) \wedge clear(Y)$

(2) $puton(X,Y)$ **causes** $\neg clear(Y)$
if $clear(X) \wedge clear(Y) \wedge \neg table(Y)$

(3) $puton(X,Y)$ **causes** $clear(Z)$
if $clear(X) \wedge clear(Y) \wedge on(X,Z)$

(4) $puton(X,Y)$ **causes** $\neg on(X,Z)$
if $clear(X) \wedge clear(Y) \wedge on(X,Z)$

Figure 3: The blocks world. Causal rules are obtained by substituting distinct values in $\{a, b, c, t\}$ for X, Y and Z.

3 Examples of Reasoning

The new lifted theory combines three features: partial-order representation of time, (cut) rules to introduce proofs by cases, and the ability to reason by contradiction.

3.1 Partial-order reasoning

Suppose G is a conjunction of fluent literals $g_1 \wedge \ldots \wedge g_m$, and we have rules R and observations O. To construct a plan to achieve G after n actions, we find an acceptable argument in the preferred extension with end-sequent $R^T, O^T, H \longrightarrow [\mathrm{N}]^n\ G$, with H a set of at most n formulas of the form $[\mathrm{N}]^k\ \mathbf{h}a$ with $k < n$. Consider how this characterization of plans describes the Sussman anomaly. We have an initial state description O in which blocks a and b are on a table t, with block c on block b. The problem is to use the theory of action consisting of the rules R in Figure 3 to bring about a state where a is on b, and b is on c. In fact, this can be achieved in three steps, by finding an acceptable argument D with end-sequent:

$$R^T, O^T, M, N, P \longrightarrow [\mathrm{N}]\,[\mathrm{N}]\,[\mathrm{N}]\,(on(a,b) \wedge on(b,c))$$

The actions of the plan are $M = [\mathrm{N}]^i\,\mathbf{h}puton(c,t), N = [\mathrm{N}]^j\,\mathbf{h}puton(a,b)$, and $P = [\mathrm{N}]^k\,\mathbf{h}puton(b,c)$. The basic structure of D is as follows. We must prove the goal at a path $\mu = \alpha\beta\gamma$. By propagating the initial state inertially along a path x of length i, we establish $[\mathrm{N}]\ [\mathrm{H}]\ clear(b)^x$ as a result of putting c on t. By applying inertia along a path y of length j to this result and the initial state, we show that putting a on b results in $[\mathrm{N}]\ [\mathrm{H}]\ on(a,b)^y$. This introduces the constraint $x < y$; using the result to help establish the the goal adds the constraint $y < \mu$. Meanwhile, by propagating the initial state inertially along a path z of length k, and combining this with $[\mathrm{N}]\ [\mathrm{H}]\ clear(b)^x$ (from putting c on t), we get $[\mathrm{N}]\ [\mathrm{H}]\ on(b,c)^z$ as the result of putting b on c. This finishes the proof of the goal, with the constraints $x < z < \mu$.

This proof is subject to attack, because we can use the success of putting a on b to prove $ab(clear(b), 2)^y$. This attacks the plan's inertial propagation of $[\mathrm{N}]\ [\mathrm{H}]\ clear(b)^x$ to z, if $x < y < z$. Since z cannot precede x, we must add the constraint $|z| \leq |y|$ to the plan. This ensures the linear order x, z, y, μ.

3.2 Cases, contradiction and explanation

In the domain scenario defined by the rules R and observation O of Figure 4, a bomb is hidden in one of two packages. We must make a plan to disable this bomb (achieve $\neg enabled$) with some

Rules:

(0)	*hide* **causes** *enabled*	(3)	*dunkA* **causes** $\neg$*enabled* **if** *bombA*
$(1a)$	*hide* **causes** *bombA* **if** *hidingPlaceA*	(4)	*dunkB* **causes** $\neg$*enabled* **if** *bombB*
$(1b)$	*hide* **causes** $\neg$*bombB* **if** *hidingPlaceA*		
$(2a)$	*hide* **causes** *bombB* **if** $\neg$*hidingPlaceA*	**Observation:**	
$(2b)$	*hide* **causes** $\neg$*bombA* **if** $\neg$*hidingPlaceA*	(5)	*hide* **happens at** 1

Figure 4: The bomb in the toilet problem.

dunking. The solution to this planning problem is an acceptable argument D with end-sequent:

$$R^T, O^T, [\mathrm{N}]^i\, \mathbf{h}\mathit{dunkA}, [\mathrm{N}]^j\, \mathbf{h}\mathit{dunkB} \longrightarrow [\mathrm{N}]^3\, \neg \mathit{enabled}$$

The structure of D is as follows. The lowest inference is a cut, to consider separately the case where *hidingPlaceA* is true and that where $\neg$*hidingPlaceA* is true. Either case begins by introducing a path μ of length 3, where we show $\neg \mathit{enabled}^{\mu}$. In the first case, rule (1a) establishes *bombA* after step 1. Instantiating the occurrence of *dunkA* to a path x of length i and including step 1, $\mathbf{h}\mathit{dunkA}^x$ establishes $[\mathrm{N}]\,[\mathrm{H}]\,\neg \mathit{enabled}^x$. Assuming $x < \mu$, inertial instantiation then establishes $\neg \mathit{enabled}^{\mu}$. In the second case, rule (2a) establishes $[\mathrm{H}]\,\mathit{bombB}$ at step 1; inertial instantiation propagates this to the occurrence of *dunkB*; that achieves the goal by rule (4)—with analogous constraints.

No constraints need be imposed to show that this argument is acceptable. The only potential attackers are the use of rule (0) to occlude inertial propagation for $[\mathrm{N}]\,[\mathrm{H}]\,\neg \mathit{enabled}^x$ or $[\mathrm{N}]\,[\mathrm{H}]\,\neg \mathit{enabled}^x$. These attacks are ruled out by the constraint that the dunkings follow the hiding.

Suppose the scenario is as before, with a hidden bomb; we observe that A is dunked and that nothing blows up. We conclude that A was the bomb by finding an acceptable argument D with end-sequent:

$$R^T, O^T, [\mathrm{N}]\, \mathbf{h}\mathit{dunkA}, [\mathrm{N}]^2 \neg \mathit{enabled} \longrightarrow [\mathrm{N}]\, \mathit{bombA} \vee [\mathrm{N}]^2 \perp$$

Again, D begins with a cut to consider the case where *hidingPlaceA* is true and that where $\neg$*hidingPlaceA* is true. In the first case, rule (1) establishes *bombA* after step 1, which suffices. In the second case, rule (2) establishes $[\mathrm{H}]$ *enabled* after step 1, which can propagate by inertia to step 2 and, with the assumption $[\mathrm{N}]^2\, \neg \mathit{enabled}$, derive a contradiction. The possible counterargument of occlusion of *enabled* by *dunkA* is not acceptable: the effect $\neg$*bombA* of the *hide* action negates a precondition of rule (3).

4 Conclusion

In this paper, we have taken a principled nonmonotonic theory of action, introduced a series of general results for streamlining proof search, and applied them to this system. The resulting system permits reasoning about the relationships between actions using the same representations and combinatory operations as partial-order planners. The system thus establishes a new, tight relationship between abstract theories of action and the behavior of implemented planners. In fact, the actual operation of a planner like UCPOP [Penberthy and Weld, 1992] could be seen as a search engine for acceptable arguments *in this system* that simply *interleaves* the construction of an argument with the demonstration that this argument is acceptable.

References

[Aho *et al.*, 1981] A. V. Aho, Y. Sagiv, T. G. Szymanski, and J. D. Ullman. Inferring a tree from lowest common ancestors with an application to the optimization of relational expressions. *SIAM Journal of Computation*, 10(3):405–421, 1981.

[Auffray and Enjalbert, 1992] Yves Auffray and Patrice Enjalbert. Modal theorem proving: an equational viewpoint. *Journal of Logic and Computation*, 2(3):247–295, 1992.

[Baral and Gelfond, to appear] Chitta Baral and Michael Gelfond. Reasoning about effects of concurrent actions. *Journal of Logic Programming*, to appear.

[Baral, to appear] Chitta Baral. Relating logic programming theories of actions and partial order planning. *Annals of Mathematics and Artificial Intelligence*, to appear.

[D'Agostino and Mondadori, 1994] Marcello D'Agostino and Marco Mondadori. The taming of the cut. classical refutations with analytic cut. *Journal of Logic and Computation*, 4(3):285–319, 1994.

[Dung, 1993] Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotic reasoning and logic programming. In *IJCAI*, pages 852–857, 1993.

[Ferguson, 1995] George M. Ferguson. *Knowledge Representation and Reasoning for Mixed-Initiative Planning*. PhD. Thesis, University of Rochester, 1995.

[Gallier, 1986] Jean H. Gallier. *Logic for Computer Science: Foundations of Automated Theorem Proving*. Harper and Row, New York, 1986.

[Gelfond and Lifschitz, 1993] Michael Gelfond and Vladimir Lifschitz. Representing actions and change by logic programs. *Journal of Logic Programming*, 17(2–4):301–321, 1993.

[Ginsberg 1995] Matthew L. Ginsberg. Approximate planning. *Artifical Intelligence Journal*, 76:89–123, 1995.

[Halpern and Moses, 1985] Joseph Y. Halpern and Yoram Moses. A guide to the modal logics of knowledge and belief: preliminary draft. In *9th International Joint Conference on Artificial Intelligence*, pages 480–490, 1985.

[Kartha, 1993] G. Neelakantan Kartha. Soundness and completeness theorems for three formalizations of action. In *IJCAI*, pages 724–729, San Mateo, CA, 1993. Morgan Kaufmann.

[Kleene, 1951] Stephen C. Kleene. Permutation of inferences in Gentzen's calculi LK and LJ. In *Two papers on the predicate calculus*, pages 1–26. American Mathematical Society, Providence, RI, 1951.

[Lin and Shoham, 1989] Fangzhen Lin and Yoav Shoham. Argument systems: a uniform basis for non-monotonic reasoning. In *KR*, pages 245–255, 1989.

[Lin and Shoham, 1991] Fangzhen Lin and Yoav Shoham. Provably correct theories of actions: preliminary report. In *AAAI*, pages 349–354, 1991.

[McAllister and Rosenblitt, 1991] David McAllister and David Rosenblitt. Systematic nonlinear planning. In *Proceedings of AAAI*, pages 634–639, 1991.

[Miller *et al.*, 1991] Dale Miller, Gopalan Nadathur, Frank Pfenning, and Andre Scedrov. Uniform proofs as a foundation for logic programming. *Annals of Pure and Applied Logic*, 51:125–157, 1991.

[Moore, 1985] Robert C. Moore. A formal theory of knowledge and action. In Jerry R. Hobbs and Robert C. Moore, editors, *Formal Theories of the Commonsense World*, pages 319–358. Ablex, Norwood NJ, 1985.

[Ohlbach, 1991] H. J. Ohlbach. Semantics-based translation methods for modal logics. *Journal of Logic and Computation*, 1(5):691–746, 1991.

[Pednault, 1991] Edwin P. D. Pednault. Generalizing nonlinear planning to handle complex goals and actions with context-dependent effects. In *IJCAI*, pages 240–254, 1991.

[Pednault, 1994] Edwin P. D. Pednault. ADL and the state-transition model of action. *Journal of Logic and Computation*, 4(5):467–512, 1994.

[Penberthy and Weld, 1992] J. S. Penberthy and D. S. Weld. UCPOP: a sound, complete partial order planner for ADL. In *KR*, pages 103–114, 1992.

[Pollock, 1992] John L. Pollock. How to reason defeasibly. *Artificial Intelligence*, 57(1):1–42, 1992.

[Reiter, 1991] Raymond Reiter. The frame problem in the situation calculus: a simple solution (sometimes) and a completeness result for goal regression. In Vladimir Lifschitz, editor, *Artificial Intelligence and Mathematical Theory of Computation*, pages 359–380. Academic Press, 1991.

[Sandewall, 1994] Erik Sandewall. *Features and Fluents: Representation of Knowledge about Dynamical Systems*. Oxford University Press, New York, 1994.

[Schubert, 1990] Lenhart K. Schubert. Monotonic solution of the frame problem in the situation calculus: an efficient method for worlds with fully specified actions. In H. E. Kyburg, R. P. Loui, and G. N. Carlson, editors, *Knowledge Representation and Defeasible Reasoning*, pages 23–67. Kluwer, Boston, 1990.

[Simari and Loui, 1992] Guillermo R. Simari and Ronald P. Loui. A mathematical treatment of defeasible reasoning and its implementation. *Artificial Intelligence*, 53:125–157, 1992.

[Stone, 1997] Matthew Stone. Efficient constraints on possible worlds for reasoning about necessity. Submitted, University of Pennsylvania, 1997.

[Wallen, 1990] Lincoln A. Wallen. *Automated Proof Search in Non-Classical Logics: Efficient Matrix Proof Methods for Modal and Intuitionistic Logics*. MIT Press, Cambridge, 1990.