Chapter 18

Temporal Action Logics

Patrick Doherty and Jonas Kvarnström

18.1 Introduction

The study of frameworks and formalisms for reasoning about action and change [67, 58, 61, 65, 70, 3, 57] has been central to the knowledge representation field almost from the inception of Artificial Intelligence as a general field of research [52, 56].

The phrase "Temporal Action Logics" represents a class of logics for reasoning about action and change that evolved from Sandewall's book on *Features and Fluents* [61] and owes much to this ambitious project. There are essentially three major parts to Sandewall's work. He first developed a narrative-based logical framework for specifying agent behavior in terms of action scenarios. The logical framework is state-based and uses explicit time structures. He then developed a formal framework for assessing the correctness (soundness and completeness) of logics for reasoning about action and change relative to a set of well-defined intended conclusions, where reasoning problems were classified according to their ontological or epistemological characteristics. Finally, he proposed a number of logics defined semantically in terms of definitions of preferential entailment[1] and assessed their correctness using his assessment framework.

Several of these logics were intended to correspond directly to existing logics of action and change proposed by others at the time, while the rest were new and were intended to characterize broad classes of reasoning problems which subsumed some of the existing approaches. Each of these definitions of preferential entailment were then analyzed using the assessment framework, giving upper and lower bounds in terms of the classes of reasoning problems for which they produced exactly the intended conclusions. Much insight was gained both in terms of advantages and limitations of previously proposed logics of action and change and in how one might go about proposing new logics of action and change in a principled manner with formal assessments included.

---

[1]Preferential entailment reduces the set of classical models of a theory by only retaining those models that are minimal according to a given preference relation, a strict partial order over logical interpretations [66].

The starting point for Temporal Action Logics was one of the definitions of preferential entailment in Sandewall's book called PMON (Pointwise Minimization of Occlusion with Nochange premises). It was one of the few preferential entailment methods that were assessed correct for the $\mathcal{K}$–**IA** class of action scenario descriptions, where $\mathcal{K}$ is an epistemological characteristic stating approximately that explicit, correct and accurate knowledge is provided (with no requirements on complete knowledge in the initial state and no restrictions on knowledge about other states), and **IA** is an ontological characteristic stating approximately that discrete integer time is used together with plain inertia (without ramification constraints, delayed effects, or other complicating factors).

Thus, PMON solved the frame problem relative to an explicit statement of assumptions ($\mathcal{K}$–**IA**) under which it could be assessed correct. The nature of the definition of preferential entailment was somewhat related to explanation closure [64, 19], although a partitioning of action scenario theories was used where only parts of the theory were minimized and other parts used as filters on the preferred model set for the theory. Though ramifications and qualifications to actions were not allowed in $\mathcal{K}$–**IA**, the class is in fact quite broad, permitting the use of conditional effects, non-deterministic effects, incomplete specification of states and the timing of actions, actions with duration and specification of dynamics within action durations.

18.1.1 PMON and TAL

While the original PMON was characterized semantically in terms of a preferential entailment method, Doherty later developed an equivalent syntactic characterization in classical 2nd-order logic (also called PMON), using a circumscription axiom to formalize the PMON definition of preferential entailment [7, 14]. In these papers, he also showed that the 2nd-order circumscription axiom was equivalent to a 1st-order pointwise circumscription axiom, enabling the use of standard first-order theorem proving techniques to reason about PMON action narratives. In extended versions of PMON which led to TAL, it has also been shown that quantifier elimination techniques or predicate completion techniques (Definition 7.3.6 in Chapter 7 of this Handbook) can be used to reduce TAL circumscribed theories to logically equivalent 1st-order theories under certain assumptions.

Doherty's PMON logic used two languages for representing and reasoning about narratives. The surface language $\mathcal{L}$(SD), Language for Scenario Descriptions, provided a convenient high-level notation for describing narratives, and could be described as a set of macros easily translated into a base language $\mathcal{L}$(FL), which was initially a many-sorted first-order language and was later altered to be an order-sorted[2] first-order language. The $\mathcal{L}$(SD) language was later renamed to $\mathcal{L}$(ND), Language for Narrative Descriptions.

The logic was further extended and generalized in several steps in order to deal with such issues as the ramification and qualification problems, use of concurrent actions, use of structured object-oriented action theories, and use as a specification formalism for TALplanner. Each extension generally implied adding new macros to $\mathcal{L}$(ND), adding additional predicates to $\mathcal{L}$(FL), extending the translation definition to

---

[2]Essentially, an order-sorted language allows the use of sub-sorts; for example, CAR and BICYCLE may be sub-sorts of the VEHICLE sort.

$\mathcal{L}$(FL) and providing slight modifications to the circumscription policies used. It is important to observe that all extensions proposed have been made in a manner which preserves the property of reducibility of the 2nd-order circumscription theory to a 1st-order theory. This is essential for practical reasons.

A number of the main extensions to PMON which led to the TAL family of logics include:

- **PMON-RC** [25], which provides a solution to the ramification problem for a broad, but as yet unassessed class of action scenarios. The main idea is the addition of a new statement type for causal constraints, where changes taking place in the world can automatically trigger new changes at the same time or at a specified delay from the original change. The solution is very fine-grained in the sense that one can easily encode dependencies between individual objects in the domain, work with both boolean and non-boolean fluents and represent both Markovian and non-Markovian dependencies [21]. PMON-RC also correctly handles chains of side effects.

- **TAL 1.0** (**PMON**$^+$) [8], which is an extended version of the original PMON logic incorporating the changes made in PMON-RC together with other useful extensions. This logic was originally called PMON$^+$, but was later renamed TAL 1.0 and provided the first stable kernel for the TAL family of logics.

- **TAL-C** [35], which uses fluent dependency constraints (an extended form of causal constraints) as a basis for representing concurrent actions. A number of phenomena related to action concurrency such as interference between one action's effects and another's execution, bounds on concurrency, and conflicting, synergistic, and cumulative effects of concurrent actions are supported.

- **TAL 2.0** [10], which provides a basic stable kernel of TAL. It is essentially TAL-C with some useful extensions and includes a tutorial on TAL and how it is used.

- **TAL-Q** [11, 42], which introduces the idea of combining an encoding of default values for features using persistence statements together with dependency constraints for representing qualifications to actions.

TAL 2.0 (TAL-C), extended with additions from TAL-Q, has been used as the basis for much of the recent work with Temporal Action Logics and will be described in some detail in this chapter. In the remainder of the chapter, we will use "TAL" as a term to denote the latest stable kernel of this family of logics.

18.1.2 Previous Work

There has been a great deal of previous work in the development of the material described in this chapter. We briefly summarize this work chronologically.

The root node from which TAL originated is the Features and Fluents (F&F) monograph [61]. Later developments with F&F are summarized in Sandewall [63]. Doherty [7] provides a syntactic characterization of PMON using pointwise circumscription and shows how a particular class of narratives can be characterized as first-order theories. Doherty [6] contains a detailed account of PMON circumscription theories

and provides additional characterizations of PMON in terms of predicate circumscription and predicate completion, where syntactic transformations are defined on narratives to provide a definition of the *Occlude* predicate. Doherty and Łukaszewicz [14] provide syntactic characterizations of 7 out of the 9 definitions of preferential entailment considered in F&F, using different forms of circumscription. Doherty and Peppas [18] incorporate the use of primary and secondary fluents in PMON to model a subclass of indirect effects of actions. A framework is also introduced for comparing linear time logics such as PMON with branching time logics such as the situation calculus. Karlsson [29] considers how to formally characterize different modal truth criteria used in planning algorithms such as TWEAK and NONLIN using PMON. Karlsson [31, 30] extends this work. Doherty [8] provides a detailed description of TAL 1.0 used as a basis for early implementations of TAL. Doherty, Łukaszewicz and Szałas [16, 17] develop a quantifier elimination algorithm which constructively generates logically equivalent 1st-order formulas for a certain class of 2nd-order formulas. The intent with the work was to study the possibility of reducing other logics for action and change characterized in terms of circumscription theories, thus making them amenable to classical theorem proving techniques. Gustafsson and Doherty [25] extend TAL to deal with ramifications of actions by introducing causal constraints, which have later been subsumed by the use of dependency constraints in TAL-C. In addition, they show how to represent delayed effects of actions in TAL. Doherty, Łukaszewicz and Szałas [19] consider the relation between the automatic generation of a definition for the *Occlude* predicate using circumscription and quantifier elimination, with the manual generation of Explanation Closure axioms considered in Schubert [64]. Karlsson [32] investigates a number of weaknesses in situation calculus and provides an alternative semantics grounded in intuitions derived from work with TAL. Bjäreland and Karlsson [5] investigates the use of regression operators as a means of doing inference in TAL related formalisms. Bjäreland [4] provides a detailed presentation of the approach in [5] and other approaches using tractable temporal logics. Karlsson, Gustafsson and Doherty [9, 36] examine the use of delayed effects of actions and various problems of interference which arise with their introduction. Doherty and Kvarnström [11] present an initial solution to simple forms of qualification to actions. Kvarnström and Doherty [42] provide a more detailed solution to the qualification problem described in this chapter. Karlsson and Gustafsson [35] consider the problem of modeling concurrent actions in TAL and the variety of interactions that may ensue between actions executing concurrently. Gustafsson [23] provides a detailed study of extensions to TAL involving dependency constraints, concurrency, and delayed effects of actions. Karlsson [33] studies the possibility of introducing narratives as 1st-class objects in the object language of a logic whose semantics is related to that of TAL. Doherty, Łukaszewicz and Madalińska-Bugaj [15] study the relation between TAL and belief update. Karlsson [34] provides detailed accounts of narratives as 1st-class citizens in action logic, concurrent actions and additional extensions to TAL. Gustafsson [24] provides a detailed description of many of the extensions to TAL up to 2001. Gustafsson and Kvarnström [26, 27] provide a novel means of structuring large TAL narratives based on the use of intuitions from object-oriented programming.

Doherty and Kvarnström [12] present a new forward chaining planner which uses TAL as a semantic framework for its development. In Kvarnström and Doherty [43], an early detailed account of TALplanner is provided. Kvarnström, Doherty and

Haslum [44] provide an extension to TALplanner which integrates concurrent actions and resources. Doherty and Kvarnström [13] provide a concise overview of TALplanner. Kvarnström and Magnusson [45] provide a description of some of the control rules used in TALplanner in the Third International Planning Competition (IPC-2002), and the reasoning underlying these rules. Kvarnström [40] discusses application of domain analysis techniques to control rules in TALplanner. Kvarnström [41] provides the most recent and most detailed description of TALplanner.

The thesis work of both Karlsson [34] and Gustafsson [24] provide excellent references to much of the later extensions to TAL. The thesis work of Kvarnström [41] provides an excellent description of TALplanner. A software system VITAL [39] for reasoning about action and change using TAL is available for download and on-line use.

18.1.3 Chapter Structure

In Section 18.2, the main concepts and ideas used in the development of TAL are presented. In Section 18.3, action narratives used in TAL are defined and a complex scenario, the Russian Airplane Hijack (RAH) scenario, is presented. This will be used throughout the chapter to explain the different features provided by TAL. Section 18.4 considers the relation between the high level macro language $\mathcal{L}(\text{ND})$ used to specify action narratives and the base logical language $\mathcal{L}(\text{FL})$ which it is translated to. In Section 18.5, we provide a formal description of the language $\mathcal{L}(\text{ND})$, and in Section 18.6, we provide a formal definition of the base logical language $\mathcal{L}(\text{FL})$. In Section 18.7, the circumscription policy used to specify the definition of preferential entailment used in TAL is presented. In addition, we show how the resulting 2nd-order circumscription theories which characterize action narratives can be reduced to logically equivalent 1st-order theories under certain conditions. Section 18.8 proposes a solution to the ramification problem which is used in TAL. The RAH scenario is modified to incorporate this solution. Section 18.9 proposes a solution to the qualification problem which is used in TAL. The RAH scenario is again modified to incorporate this solution. Section 18.10 provides further examples of the expressivity of TAL actions. Section 18.11 presents an extension to TAL which models the use of concurrent actions where complex types of interaction between such actions may occur. Section 18.12 presents an application of TAL to planning where it is shown how TAL can be used as a semantic framework in the development and implementation of TALplanner, an award winning automated planner. In Section 18.13, we conclude.

18.2 Basic Concepts

When using TAL, we assume there is an *agent* interested in reasoning about a specific *world*. This world might be formally defined, or it might be the "real world", in which case the agent can only reason about a formally defined abstraction of the real world. In either case, it is assumed that the world is dynamic, in the sense that the various properties or *features* of the world can change over time. Conceptually, any feature has a fluent function associated with it representing the stream of values associated with the feature at each state or temporal entity used in the formalism.

The TAL framework also permits the use of multiple *value domains*, which can be used for modeling different types of *objects* that might occur in the world which is

being modeled. For example, the well-known blocks world contains blocks that can be stacked on top of each other. The blocks world can then be modeled using a value domain for blocks, containing values such as A, B and C, together with parameterized boolean-valued features representing relations such as on($block_1$, $block_2$), which holds iff $block_1$ is on top of $block_2$, and clear($block$), which holds iff there is no block on top of $block$. Of course, values can also be used to represent properties of objects rather than the objects themselves. For example, if the color of each block should be modeled, then this could be done using a value domain for colors containing values such as red, green and blue, together with a color-valued (non-boolean) feature color($block$). In summary, instantiated parameterized features take specific values (boolean or non-boolean) at specific times. In this manner, both relations and properties are capable of being represented.

Time itself can be viewed differently depending on the nature of the world being reasoned about and the reasoning abilities of the agent. TAL offers a modular means of choosing the temporal structure to be used. Currently, TAL uses linear time structures, as opposed to branching time structures. Research within the TAL framework has been focused on discrete non-negative integer time structures, and such a structure will be used throughout this chapter, though most concepts should carry over directly or with little modification to a real-valued time structure using ideas from [59, 60, 65].

The development of the world over a (possibly infinite) period of discrete time can be viewed in two different ways. Fig. 18.1 shows what would happen in a simple blocks world scenario where block A is initially on top of B, which is on the table, and where one unstacks A from B, places it on the table, picks up B, and finally stacks this block on top of A. The information about this scenario can be viewed as a sequence of *states*, where each state provides a value to all features (or "state variables") for a single common timepoint, or as a set of *fluents*, where each fluent is a function of time which specifies the development of a single feature. We sometimes use the terms "feature" and "fluent" interchangeably to refer to either a specific property of the world or the function specifying its value over time.

Consequently, a logical model in TAL is a sequence of states indexed by time, where each state contains a value for each feature in the vocabulary at the timepoint

	time 0	time 1	time 2	time 3	time 4	
on(A,A)	false	false	false	false	false	
on(A,B)	true	false	false	false	false	
on(B,A)	false	false	false	false	true	**fluent**
on(B,B)	false	false	false	false	false	
ontable(A)	false	false	true	true	true	
ontable(B)	true	true	true	false	false	
clear(A)	true	false	true	true	false	
clear(B)	false	true	true	false	true	
handempty	true	false	true	false	true	

(annotations: **feature** points to clear(B); **state** points to the time 2 column)

Figure 18.1: Viewing a development as fluents or states.

associated with the state. In the logical language, the assertion that a feature has a value at a specific time is denoted as $[\tau]\ f(\bar{\omega}) \hat{=} \omega$ in the macro language $\mathcal{L}(\text{ND})$ and $Holds(\tau, f(\bar{\omega}), \omega)$ in the logical language $\mathcal{L}(\text{FL})$, where τ is a temporal expression, $f(\bar{\omega})$ is a parameterized feature and ω is a value from the feature's value domain.

Since there is an agent, there is usually also a set of *actions* that the agent can perform. Such actions can only be performed when the requisite *preconditions* are satisfied. Performing an action changes the state of the world according to a set of given rules. Such rules are not necessarily deterministic. For example, the action of tossing a coin can be modeled within the TAL framework, and there will be two possible result states. TAL offers a highly expressive language for specifying actions where non-deterministic, context-dependent, concurrent and durational actions are expressible, among other types of actions.

Background knowledge associated with a reasoning domain can be modeled in a number of ways in TAL. *Observation statements* represent observations made by an agent. *Domain Constraint statements* represent facts true in all scenarios associated with a particular reasoning domain. *Dependency Constraint statements* can be used to represent causal theories or assertions which model intricate dependencies describing how and when features change relative to each other.

All of these concepts are modeled in a narrative specified in the language $\mathcal{L}(\text{ND})$. $\mathcal{L}(\text{ND})$ is a high-level extendable macro language which provides support to the knowledge engineer when constructing narratives and permits specification of narratives at a higher level of abstraction than logical statements. An extendable translation function is provided which translates narratives specified in $\mathcal{L}(\text{ND})$ into 1st- and 2nd-order logical theories.

One of the fundamental problems in developing logics for reasoning about action and change has been in finding both representational and computationally efficient ways to encode the fact that there is a great deal of invariant structure in the world at a particular level of abstraction in which agents often describe and reason about the world. Even though the world is often dynamic and changing, from the perspective of an agent functioning in the world, properties and relations among entities are more often than not inert. On the other hand, there are often reasons for features in the world to change or reasons that provide the possibility for change. Many of these are obvious. For example, if an agent executes a physical action, the intent is usually to change some aspect of the world to the agent's advantage in completing a task. Others are less obvious, for example the subtle ramifications and aftereffects of an action. Developing theories of action and change is very much about identifying and representing normative rules which capture invariant and non-invariant epistemic and physical structure in environments in which agents are embedded and in which they operate.

Many of the representational and computational problems associated with modeling action and change have been given names, such as the frame, ramification and qualification problems, while others have not. Many useful techniques for capturing normative behavior have also been developed such as default reasoning. The principal intuition used in the development of TAL to deal with many of these issues is very simple to state, but quite difficult to make operational in an efficient manner in a logical formalism such as TAL.

In any TAL model, a time series is implicitly associated with any feature in the vocabulary. Whether a feature may change value or not in a transition from one timepoint to another in the time series is specified by occluding or marking that feature as being given the *possibility* of changing value relative to other constraints in the theory. Policies for occluding features at timepoints are both contextually and temporally dependent on a number of factors and done for a number of reasons. The definitions of the frame, ramification and qualification problems specify some of these reasons. For whatever reason this is done, to the greatest extent possible, this labeling process should be achieved in a principled manner and remain more or less hidden from the knowledge engineer via the use of macro mechanisms in the $\mathcal{L}(\text{ND})$ language and the translation into the base logical language $\mathcal{L}(\text{FL})$.

At the level of $\mathcal{L}(\text{ND})$, there are a number of ways to incrementally provide an occlusion policy for a feature, some more explicit than others. At the $\mathcal{L}(\text{FL})$ level, the policies result in a set of labels for each feature represented as $Occlude(\tau, f(\bar{\omega}))$. The generation of such policies provide sufficient conditions for features being given the possibility to change value in state transitions (from $\tau - 1$ to τ). A circumscription policy then provides the necessary conditions and a definition of the occlusion predicate in the logical theory. An additional specification of whether and when a feature is persistent, durational, or dynamic in nature is also provided. These statements provide a means of filtering "bad" models out of the model set for a particular narrative, such as models where persistent features change value without being occluded.

If one uses this technique in a principled manner and restricts the generation of such policies to only include positive occurrences of the predicate *Occlude* in the theory, then a reduction of the 2nd-order circumscription theory to a logically equivalent 1st-order theory is always guaranteed. It is in this manner we provide partial solutions to the frame, ramification and qualification problems in the context of TAL.

18.3 TAL Narratives

A narrative in $\mathcal{L}(\text{ND})$ can be said to consist of two parts: The *narrative background specification* (NBS), which provides background information that is common to all narratives for a particular domain, and the *narrative specification* (NS), which provides information specific to a particular instance of a reasoning problem. Most of this information is represented as a set of labeled narrative statements in the surface language $\mathcal{L}(\text{ND})$.

Before providing a formal definition of the $\mathcal{L}(\text{ND})$ language, we will introduce most of the macros, formula types and statement classes using a rather complex example scenario called the *Russian Airplane Hijack* (RAH) scenario, which in order to be adequately represented in any logical formalism would require robust solutions to the frame, ramification and qualification problems. We say robust because a complete description of the RAH world requires the representation of concurrent actions, incomplete specifications of states, ramification with chaining, the use of non-boolean features, fine-grained dependencies among objects in different feature value domains, actions with duration, two types of qualification (*weak* and *strong*) and the use of explicit time, in addition to other features.

The RAH narrative description will be used as a vehicle for considering different facets of Temporal Action Logics and demonstrating how various aspects of a domain

can be modeled in TAL. This will be done in stages. In this section, we will represent the narrative without the use of side effects and under the assumption that actions always succeed if their basic preconditions are satisfied. In other words, we will omit solutions to the ramification and qualification problems.[3] After having provided formal specifications of the $\mathcal{L}$(ND) and $\mathcal{L}$(FL) languages (Sections 18.5 and 18.6), we will once more return to the RAH scenario in order to consider how ramification constraints (Section 18.8) and qualification constraints (Section 18.9) can be modeled in TAL.

18.3.1 The Russian Airplane Hijack Scenario

The Russian Airplane Hijack scenario[4] can be described as follows.

Example 18.1 *(Russian Airplane Hijack scenario).* A Russian businessman, Boris, travels a lot and is concerned about both his hair and safety. Consequently, when traveling, he places both a comb and a gun in his pocket. A Bulgarian businessman, Dimiter, is less concerned about his hair, but when traveling by air, has a tendency to drink large amounts of vodka before boarding a flight to subdue his fear of flying. A Swedish businessman, Erik, travels a lot, likes combing his hair, but is generally law abiding.

One ramification of moving between locations is that objects in your pocket will follow you from location to location. Similarly, a person on board a plane will follow the plane as it flies between cities.

Generally, when boarding a plane, the only preconditions are that you are at the gate and you have a ticket. However, if you try to board a plane carrying a gun in your pocket, which will be the case for Boris, this should qualify the action. Also, a condition that could sometimes qualify the boarding action is if you arrive at the gate in a sufficiently inebriated condition, as will be the case for Dimiter. When the boarding action is qualified, attempting to board should have no effect.

Boris, Erik and Dimiter already have their tickets. They start (concurrently) from their respective homes, stop by the office, go to the airport, and try to board flight SAS609 to Stockholm. Both Erik and Boris put combs in their pockets at home, and Boris picks up a gun at the office, while Dimiter is already drunk at home and may or may not already have a comb in his pocket. Who will successfully board the plane? What are their final locations? What will be in their pockets after attempting to board the plane and after the plane has arrived at its destination?

Let us assume that the scenario is encoded correctly in TAL and that we agree on our commonsense intuitions regarding what solutions to the frame, ramification and qualification problems would imply. Then the following inferences should be entailed by the logical theory associated with the RAH scenario[5]:

---

[3]This will initially result in a scenario where it is assumed that any attempt to board a plane always succeeds, regardless of whether a person carries a gun or is drunk. In addition, ramifications of action effects will be included in action specifications rather than being specified separately.

[4]This scenario is an elaboration and concretization of a sketch for a scenario proposed by Vladimir Lifschitz in on-line discussions in the Electronic Transactions on Artificial Intelligence (ETAI/ENAI), and was previously published in [11, 42].

[5]Assume that Boris, Erik and Dimiter own the combs comb1, comb2 and comb3, respectively.

1. Erik will board the plane successfully, eventually ending up at his destination.

2. An indirect effect of flying is that a person ends up at the same location as the airplane he is on. In addition, because items in pockets follow a person, a transitive effect results where items in a person's pocket are at the same location as the plane which that person is on. Consequently, Erik's comb, comb2, will also end up at his destination.

3. Boris will get as far as the airport with a gun and comb1 in his pocket. He will be unable to board the plane.

4. Dimiter will get as far as the airport, and may or may not be able to board the plane. If he is able to board the plane, he will eventually end up at his destination. Otherwise, he will remain at the airport. In any case, if he initially carried a comb, it will end up in the same location.

18.3.2 Narrative Background Specification

A narrative background specification contains a collection of statements of the following types:

- *Persistence statements* (labeled **per**[6]) allow each fluent to be specified as being persistent (normally retaining its value from the previous timepoint), durational (normally reverting to a default value), or dynamic (varying freely, subject to other constraints involving this fluent).

- *Domain constraint statements* (labeled **dom**) characterize acausal information which is always true in the world being modeled.

- *Action type specifications* (labeled **acs**) provide generic definitions of action types.

- *Dependency constraint statements* (labeled **dep**) characterize causal and directional dependencies among features.

A narrative background specification also contains a vocabulary for the narrative. In the following subsections, each of the statement types and the vocabulary specification will be described in detail and correlates to the RAH scenario will be listed.

Vocabulary

The vocabulary of an $\mathcal{L}(\text{ND})$ narrative defines the constant symbols, feature symbols, action symbols, and other symbols that are available for use in narrative formulas. Since narrative examples used in the literature have traditionally been quite simple, the vocabulary has usually either been considered to be implicit in the remainder of the narrative specification or has been described informally in the main text of the article. Here, however, vocabularies will be described in terms of labeled narrative

---

[6]A number is often suffixed to each label, as in **per3**. These numbers are used to disambiguate references in the text and have no semantic meaning.

declaration statements using a syntax borrowed from the software tools VITAL [39] and TALplanner [41].

For the Russian Airplane Hijack scenario, we define a domain LOCATION for locations, and a domain THING containing everything that has a location. We also define the subdomains RUNWAY for LOCATIONs that are runways, PLANE for THINGs that are airplanes, PERSON for THINGs that are people, and PTHING for THINGs that people can pick up.

> **domain** LOCATION :elements { home1, home2, home3, office, airport, run609, run609b, air }
>
> **domain** THING :elements { gun, comb1, comb2, comb3, boris, dimiter, erik, sas609 }
> **domain** RUNWAY :parent LOCATION :elements { run609, run609b }
> **domain** PLANE :parent THING :elements { sas609 }
> **domain** PERSON :parent THING :elements { boris, dimiter, erik }
> **domain** PTHING :parent THING :elements { gun, comb1, comb2, comb3 }

We also use the boolean domain, which is present by default in all narratives and behaves as if it had been specified in the following manner:

> **domain** BOOLEAN :elements { true, false }

Note that the domain specification in $\mathcal{L}$(ND) describes a type hierarchy. This will translate into the order-sorted vocabulary in the base logic $\mathcal{L}$(FL).

Finally, four fluents and four actions are used where the arguments to these are typed relative to the domain specification above.

> **fluent** loc(THING) :domain LOCATION
> **fluent** inpocket(PERSON, PTHING) :domain BOOLEAN
> **fluent** onplane(PLANE, PERSON) :domain BOOLEAN
> **fluent** drunk(PERSON) :domain BOOLEAN
>
> **action** pickup(PERSON, PTHING)
> **action** travel(PERSON, LOCATION, LOCATION)
> **action** board(PERSON, PLANE)
> **action** fly(PLANE, RUNWAY, RUNWAY)

Persistence statements

Persistence statements are a novel feature of TAL and offer a very powerful and fine-grained mechanism for specifying inertia and default value assumptions for individual features when used together with the occlusion labeling mechanism mentioned previously. The majority of existing formalisms for action and change build in an assumption that a property or relation is either *always* assumed to be inert and subject to nochange by default or to be dynamic and subject to change by default. Through the use of persistence statements TAL permits the specification of contextually and temporally dependent inertia assumptions and default value assumptions per feature and down to the feature object level. This is an important feature of any action and change formalism since the inertial granularity of physical and other objects differs greatly. For example, a mountain will remain in place much longer than a ball on the ground which under certain weather conditions is in fact not inert at all.

Persistence statements can be used to classify features as being *persistent*, *durational*, or *dynamic*. In fact, a specific instantiated feature or set of features may be classified differently in the same scenario relative to context.

A feature declared as *persistent* at a timepoint is only allowed to change value when an action or dependency constraint in the scenario explicitly allows it to change, by labeling the feature at that timepoint as being occluded using a reassignment macro (Section 18.5.2, subsection *Formulas*). Otherwise, it retains the same value it had at the previous timepoint (the *persistence assumption* or *inertia assumption*). For example, the persistence statement below declares that all instantiated features of the form loc(*thing*) are inert at all timepoints:

per $\forall t, thing\ [Per(t, loc(thing))]$

The translation from $\mathcal{L}(ND)$ to $\mathcal{L}(FL)$ is performed using the *Trans* function defined in Section 18.6.1. If the timeline is infinite in both directions, the following translation could be used for the *Per* predicate, stating that unless a feature is occluded at τ, it will retain its previous value.

$$Trans(Per(\tau, f)) = \neg Occlude(\tau, f) \supset Holds(\tau, f, v) \equiv Holds(\tau - 1, f, v)$$

However, TAL is generally used with a non-negative time structure, where this translation would lead to problems at the boundary where $\tau = 0$, where $\tau - 1$ does not exist. Thus, the translation has to be changed to the following, where for $\tau = 0$ there exists no t such that $\tau = t + 1$ and the antecedent of the implication will be false, which correctly models the intuition that persistence should not affect the possible values of a fluent at the beginning of time:

$$Trans(Per(\tau, f)) = \forall t . \tau = t + 1 \wedge \neg Occlude(t + 1, f) \supset$$
$$\forall v[Holds(t + 1, f, v) \equiv Holds(t, f, v)]$$

A feature declared as *durational* is associated with a default value, and can only take on another value when an action, dependency constraint, or other constraint allows it to (the *default value assumption*). At timepoints when no action or dependency constraint explicitly allows it to take on another value, it will immediately revert back to its default value. Through the use of durational features, TAL can encode simple types of default rules and assumptions. For instance, one may model the presence of noise using a durational fluent with default value false, capturing the intuition that there is no noise unless an action is currently generating it. An action generating noise would then use a suitable reassignment operator to exempt the noise fluent from its default value assumption during the appropriate temporal interval. This is especially important in the presence of concurrent actions, where an alternative solution that sets noise to true at the beginning of the action and explicitly sets it to false at the end does not work properly with partially overlapping actions, potentially generating the conclusions that when one action ends, it cancels the noise generated by another action.

For example, the persistence statement below declares that the fluent noise should have the default value false at all timepoints. Further examples will be given when qualification is discussed (Section 18.9).

per $\forall t\ [Dur(t, noise, false)]$

The translation of a *Dur* declaration in $\mathcal{L}(\text{ND})$ into $\mathcal{L}(\text{FL})$ would be:

$$Trans(Dur(\tau, f, \omega)) = \neg Occlude(\tau, f) \supset Holds(\tau, f, \omega)$$

Unless the feature is occluded at τ, it will take on its default value.

Finally, a TAL feature can also be *dynamic* if it is not declared to be persistent or durational. Since no persistence or default value assumption is applied, dynamic fluents can vary freely over time to satisfy observations and domain constraints.

Note that some earlier TAL logics (including PMON) used a fixed *nochange axiom* instead of persistence statements, forcing all fluents to be persistent. Using persistence statements provides a more flexible and fine-grained approach to controlling the default behavior of fluents and is currently the technique used in TAL to specify inertia and default value assumptions.

Intuitively, the features used in the Russian Airplane Hijack scenario describe properties that do not change unless something changes them. These features are all declared to be persistent. The declarations for the RAH scenario are as follows:

per1 $\forall t, thing \, [Per(t, \mathsf{loc}(thing))]$
per2 $\forall t, person, pthing \, [Per(t, \mathsf{inpocket}(person, pthing))]$
per3 $\forall t, person \, [Per(t, \mathsf{drunk}(person))]$
per4 $\forall t, plane, person \, [Per(t, \mathsf{onplane}(plane, person))]$

Domain constraints

Domain constraints represent knowledge about logical feature dependencies which are not specific to a particular reasoning problem instance but which are known to hold in every possible scenario taking place within a domain. An even stronger assumption often made in other formalisms is that these are formulas true in all states (universally quantified over all timepoints, situations or states) and behave much as a classical logical formula would behave in a standard theory. In domain constraints, as well as other TAL formulas, the fact that a feature f takes on a particular value ω is denoted by the elementary fluent formula $f \,\hat{=}\, \omega$. For the boolean domain, the formula $f \,\hat{=}\, \mathsf{true}$ ($f \,\hat{=}\, \mathsf{false}$) can be abbreviated as f ($\neg f$). Elementary fluent formulas can be combined using boolean connectives and quantification over values to form fluent formulas. The fixed fluent formula $[\tau] \, \phi$ states that the fluent formula ϕ holds at the timepoint τ.

For the Russian Airplane Hijack scenario we will define three domain constraints: No PTHING can be carried by two PERSONs at the same time, no PERSON can be on board two PLANEs at the same time, and any PTHING in a PERSON's pocket must be at the same location as that PERSON.

dom1 $\forall t, pthing, person_1, person_2 \, [person_1 \neq person_2 \wedge [t] \, \mathsf{inpocket}(person_1, pthing) \supset$
      $[t] \, \neg\mathsf{inpocket}(person_2, pthing)]$
dom2 $\forall t, person, plane_1, plane_2$
      $[plane_1 \neq plane_2 \wedge [t] \, \mathsf{onplane}(plane_1, person) \supset [t] \, \neg\mathsf{onplane}(plane_2, person)]$
dom3 $\forall t, person, pthing \, [[t] \, \mathsf{inpocket}(person, pthing) \supset$
      $[t] \, \mathsf{loc}(pthing) \,\hat{=}\, value(t, \mathsf{loc}(person))]$

Action types

Actions can be invoked by the agent in order to change some properties in the world. If *person* picks up a thing *pthing* in the Russian Airplane Hijack scenario, then this

should cause inpocket(*person*, *pthing*) to become true, for example. But since the inpocket feature is persistent, simply stating the fact that inpocket(*person*, *pthing*) will be true at the end of the action invocation is not sufficient. Instead, it is necessary to use a *reassignment macro* to explicitly release this feature from the persistence assumption at the specific point in time where it should change values from false to true.

There are three different reassignment macros: X, R and I. They can all be used with a temporal interval, for example, $R((\tau, \tau'] \, \alpha)$, or a single timepoint, for example $I([\tau] \, \alpha)$. Each of these operators has the effect of releasing the features occurring in α from the persistence and default value assumptions during the given interval or at the given timepoint. However, the operators differ in whether they place further constraints on the values of these features, and if so, at what time.

The X operator is used for *occlusion*. Its purpose is simply to allow the value of the features in the formula α to vary at a timepoint or during an interval, and therefore it does not further constrain the features occurring in α. Intuitively, the X operator occludes (hides) any changes in a feature value from the persistence or default value constraints generated by the persistence statements in the narrative.

The R operator is used for *reassignment*, and ensures that α will hold at the final timepoint in the interval. During the rest of the interval, the features occurring in α are allowed to vary freely, unaffected by the persistence or default value assumption (but still subject to other constraints that may also be present in the narrative).

The I operator is used for *interval reassignment* and ensures that α will hold during the entire interval. Note that if α is a disjunctive formula, features occurring in α may still vary during the interval as long as the formula remains satisfied throughout the interval.

An *action type specification* uses reassignment macros to define what will happen if and when a particular action is invoked. Note that it does not state that an action does occur. This is specified in the narrative specification using action occurrence statements.

In many existing action formalisms, actions do not have duration and are essentially single step. If actions with duration are introduced, it is often the case that during the duration nothing can happen or be specified to happen. TAL offers highly expressive action types. They can be single-step or durational, inert during the duration or highly dynamic. Additional constraints specifying what goes on during the execution of an action can easily be included in the action specification.

In the Russian Airplane Hijack scenario, four actions were declared in the narrative background specification. Here, those actions will be defined without taking qualifications into account and without making use of ramification constraints to specify side effects, resulting in a narrative where guns do not qualify the boarding action and where the fact that people inside an airplane move when the airplane moves must be expressed explicitly in the action definition. These action definitions will later be modified in Section 18.8.

acs1 $[t_1, t_2]$fly(*plane*, *runway*$_1$, *runway*$_2$) $\rightsquigarrow$ ($[t_1]$ loc(*plane*) $\hat{=}$ *runway*$_1$ $\supset$
$I((t_1, t_2)$ loc(*plane*) $\hat{=}$ air) $\wedge$ $R([t_2]$ loc(*plane*) $\hat{=}$ *runway*$_2$) $\wedge$
$\forall$*person*$[[t_1]$ onplane(*plane*, *person*) $\supset$
  $I((t_1, t_2)$ loc(*person*) $\hat{=}$ air) $\wedge$ $R([t_2]$ loc(*person*) $\hat{=}$ *runway*$_2$) $\wedge$
  $\forall$*pthing*$[[t_1]$ inpocket(*person*, *pthing*) $\supset$
    $I((t_1, t_2)$ loc(*pthing*) $\hat{=}$ air) $\wedge$ $R([t_2]$ loc(*pthing*) $\hat{=}$ *runway*$_2$)]]])

acs2 $[t_1, t_2]$pickup(*person*, *pthing*) $\rightsquigarrow [t_1]$ loc(*person*) $\hat{=}$ *value*(t_1, loc(*pthing*)) $\supset$
  $R((t_1, t_2)$ inpocket(*person*, *pthing*))

acs3 $[t_1, t_2]$travel(*person*, loc_1, loc_2) $\rightsquigarrow [t_1]$ loc(*person*) $\hat{=} loc_1 \supset$
  $R([t_2]$ loc(*person*) $\hat{=} loc_2) \wedge$
  $\forall pthing[[t_1]$ inpocket(*person*, *pthing*) $\supset R([t_2]$ loc(*pthing*) $\hat{=} loc_2)]$

acs4 $[t_1, t_2]$board(*person*, *plane*) $\rightsquigarrow [t_1]$ loc(*person*) $\hat{=}$ airport $\supset$
  $R([t_2]$ loc(*person*) $\hat{=}$ *value*(t_2, loc(*plane*)) $\wedge$ onplane(*plane*, *person*))

For reasons of representational efficiency, it is quite clear from observing these action specifications that a solution to the ramification problem is really necessary.

18.3.3 Narrative Specification

In the narrative specification, *observation statements* (labeled **obs**) represent observations of feature values at specific timepoints while *action occurrence statements* (labeled **occ**) specify which instances of the generic action types occur and during which time intervals.

Observation statements

Observation statements are intended to describe specific facts that have been observed to hold in the world, permitting complete or incomplete specifications of the initial state or any other state in the world development corresponding to a narrative. They provide information about a particular reasoning problem instance within a domain, and are therefore part of the narrative specification.[7]

For this scenario, we define the initial locations of all THINGS, as well as who is drunk in the initial state. On the other hand, we do not observe which things are in whose pockets.

obs1 $[0]$ loc(boris) $\hat{=}$ home1 $\wedge$ loc(gun) $\hat{=}$ office $\wedge$ loc(comb1) $\hat{=}$ home1 $\wedge$ ¬drunk(boris)
obs2 $[0]$ loc(erik) $\hat{=}$ home2 $\wedge$ loc(comb2) $\hat{=}$ home2 $\wedge$ ¬drunk(erik)
obs3 $[0]$ loc(dimiter) $\hat{=}$ home3 $\wedge$ loc(comb3) $\hat{=}$ home3 $\wedge$ drunk(dimiter)
obs4 $[0]$ loc(sas609) $\hat{=}$ run609

Action occurrence statements

Action occurrence statements specify which actions actually do take place in a narrative. Like observations, they are part of the narrative specification—the instance-specific part of the narrative.

For the Russian Airplane Hijack scenario, the following action occurrences are also required. The exact timepoints used below were not specified in the RAH scenario, but have been chosen arbitrarily. Alternatively, exact timepoints could have been avoided by using non-numerical temporal constants. Note, however, that many of the actions are concurrent, sometimes with partially overlapping intervals.

---

[7]In some earlier versions of TAL, an explicit *Observe* predicate was introduced in the base logical language $\mathcal{L}$(FL) to which observation statements are translated. Distinguishing sensor-generated facts about the world from other facts is useful when interfacing such logics to robotic systems. One might choose to view observation statements as perception statements, although this is not done in the current version of TAL.

occ1	[1,2] pickup(boris, comb1)	**occ8**	[7,9] travel(erik, office, airport)
occ2	[1,2] pickup(erik, comb2)	**occ9**	[8,10] travel(boris, office, airport)
occ3	[2,4] travel(dimiter, home3, office)	**occ10**	[9,10] board(dimiter, sas609)
occ4	[3,5] travel(boris, home1, office)	**occ11**	[10,11] board(boris, sas609)
occ5	[4,6] travel(erik, home2, office)	**occ12**	[11,12] board(erik, sas609)
occ6	[6,7] pickup(boris, gun)	**occ13**	[13,16] fly(sas609, run609, run609b)
occ7	[5,7] travel(dimiter, office, airport)		

Note that this action scenario has been simplified for expository purposes. A number of additional extensions to the scenarios would in fact make it more realistic. For example, one could add more realistic timing actions, perhaps by explicitly modeling distances between locations and dividing by expected speed. In addition, upon introducing truly concurrent actions, one must be aware that there may be different types of interactions and these would have to be dealt with in an appropriate manner. TAL allows such extensions and we refer the interested reader to Section 18.11 where a summary of concurrent actions in TAL is provided.

18.4   The Relation Between the TAL Languages $\mathcal{L}(\text{ND})$ and $\mathcal{L}(\text{FL})$

In order to reason about a particular narrative, it is first mechanically translated into the base language $\mathcal{L}(\text{FL})$, an order-sorted classical first-order language with equality using a linear discrete time structure (Fig. 18.2). A circumscription policy is applied to the resulting theory, foundational axioms are added, and quantifier elimination techniques are used to reduce the resulting second order theory to first order logic. This is possible only under certain assumptions pertaining to the use of the *Occlude* predicate and the nature of the temporal structure used.
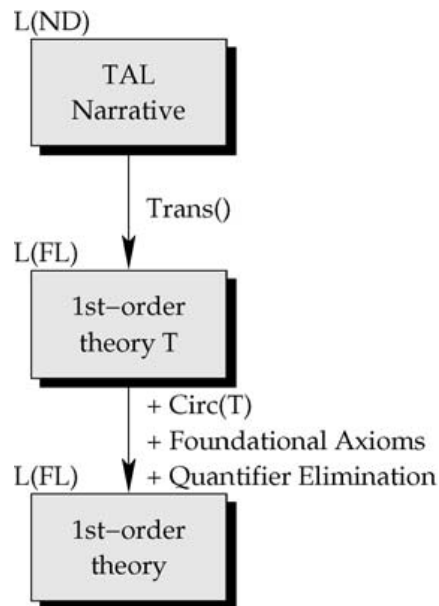

Figure 18.2: The relation between $\mathcal{L}(\text{ND})$ and $\mathcal{L}(\text{FL})$.

In Section 18.5, we will present the TAL surface language $\mathcal{L}$(ND). In Section 18.6, we will present the TAL base language $\mathcal{L}$(FL), and in Section 18.7, we will consider the circumscription policy used in TAL and reducibility results.

18.5 The TAL Surface Language $\mathcal{L}$(ND)

This section defines the surface language $\mathcal{L}$(ND). The translation to the first-order language $\mathcal{L}$(FL) is presented in Section 18.6.1. In the following, the overline is used as an abbreviation for a sequence, when the contents of the sequence is obvious. For example, $f(\bar{x}, \bar{y})$ means $f(x_1, \ldots, x_n, y_1, \ldots, y_m)$.

18.5.1 Sorts, Terms and Variables

Definition 18.1 (Basic sorts). *There are a number of sorts for values $\mathcal{V}_i$, including the boolean sort $\mathcal{B}$ with the constants {true, false}. TAL is order-sorted, and a sort may be specified to be a subsort of another sort. The sort $\mathcal{V}$ is a supersort of all value sorts.*

There are a number of sorts for features $\mathcal{F}_i$, each one associated with a value sort $dom(\mathcal{F}_i) = \mathcal{V}_j$ for some j. The sort $\mathcal{F}$ is a supersort of all fluent sorts.

There is also a sort for actions $\mathcal{A}$ and a temporal sort $\mathcal{T}$.

The sort $\mathcal{T}$ is often assumed to be interpreted and semantic attachment is used in implementations, but it can be axiomatized in various ways, For example, in first-order logic, it can be axiomatized as a subset of Presburger arithmetic [38] (natural numbers with addition), or in second-order logic as Peano arithmetic.

Definition 18.2 (Terms). *A value term, often denoted by ω, is a variable v or a constant v of sort $\mathcal{V}_i$ for some i, an expression value(τ, f) where τ is a temporal term and f is a fluent term, or an expression $g(\omega_1, \ldots, \omega_n)$ where $g : \mathcal{V}_{k_1} \times \cdots \times \mathcal{V}_{k_n} \to \mathcal{V}_i$ is a value function symbol and each ω_j is a value term of sort $\mathcal{V}_{k_j}$.*

A temporal term, often denoted by τ, is a variable t or a constant $0, 1, 2, 3, \ldots$ or $s_1, t_1, \ldots$, or an expression of the form $\tau_1 + \tau_2$, all of sort $\mathcal{T}$.

A fluent term, often denoted by f, is a feature variable or an expression $f(\omega_1, \ldots, \omega_n)$ where $f : \mathcal{V}_{k_1} \times \cdots \times \mathcal{V}_{k_n} \to \mathcal{F}_i$ is a feature symbol and each ω_j is a value term of sort $\mathcal{V}_{k_j}$.

An action term Ψ is an expression $A(\omega_1, \ldots, \omega_n)$ where $A : \mathcal{V}_{k_1} \times \cdots \times \mathcal{V}_{k_n} \to \mathcal{A}$ is an action symbol and each ω_j is a value term of sort $\mathcal{V}_{k_j}$.

Variables are typed and range over the values belonging to a specific sort. Although the sort is sometimes specified explicitly in narratives, it is more common to simply give the variable the same name as the sort but (like all variables) written in italics, possibly with a prime and/or an index. For example, the variables *plane*, *plane'* and *plane*$_3$ would be of the sort PLANE. Similarly, variables named t or τ are normally temporal variables, and variables named n are normally integer-valued variables.

The function *value*(τ, f) returns the value of the fluent f at the timepoint τ, where $[\tau]\ f \hat{=} v$ iff *value*$(\tau, f) = v$. The expression $[\tau]\ f \hat{=} g$, where f and g are fluent terms, is shorthand notation for $[\tau]\ f \hat{=}$ *value*(τ, g).

18.5.2   Formulas

Definition 18.3 (Temporal and value formulas). *If τ and τ' are temporal terms, then $\tau = \tau'$, $\tau < \tau'$ and $\tau \leqslant \tau'$ are temporal formulas. A value formula is of the form $\omega = \omega'$ where ω and ω' are value terms, or $r(\omega_1, \ldots, \omega_n)$ where $r \colon \mathcal{V}_{k_1} \times \cdots \times \mathcal{V}_{k_n}$ is a relation symbol and each ω_j is a value term of sort $\mathcal{V}_{k_j}$.*

We will sometimes write $\tau \leqslant \tau' < \tau''$ to denote the conjunction $\tau \leqslant \tau' \wedge \tau' < \tau''$, and similarly for other combinations of the relation symbols $\leqslant$ and $<$.

Definition 18.4 (Fluent formula). *An elementary fluent formula, sometimes called an isvalue expression, has the form $f \,\hat{=}\, \omega$ where f is a fluent term of sort $\mathcal{F}_i$ and ω is a value term of sort $dom(\mathcal{F}_i)$. A fluent formula is an elementary fluent formula or a combination of fluent formulas formed with the standard logical connectives and quantification over values.*

The elementary fluent formula $f \,\hat{=}\, \mathsf{true}$ ($f \,\hat{=}\, \mathsf{false}$) can be abbreviated f ($\neg f$).

Definition 18.5 (Timed formulas). *Let τ and τ' be temporal terms and α a fluent formula. Then*:

- $[\tau, \tau']\,\alpha$, $(\tau, \tau']\,\alpha$, $[\tau, \tau')\,\alpha$, $(\tau, \tau')\,\alpha$, $[\tau, \infty)\,\alpha$, $(\tau, \infty)\,\alpha$ *and* $[\tau]\,\alpha$ *are* fixed fluent formulas,

- $C_T([\tau]\,\alpha)$, $C_F([\tau]\,\alpha)$ *and* $C([\tau]\,\alpha)$ *are* fluent change formulas,

- $R([\tau, \tau']\,\alpha)$, $R((\tau, \tau']\,\alpha)$, $R([\tau, \tau')\,\alpha)$, $R((\tau, \tau')\,\alpha)$ *and* $R([\tau]\,\alpha)$ *are* reassignment formulas,

- $I([\tau, \tau']\,\alpha)$, $I((\tau, \tau']\,\alpha)$, $I([\tau, \tau')\,\alpha)$, $I((\tau, \tau')\,\alpha)$ *and* $I([\tau]\,\alpha)$ *are* interval reassignment formulas, *and*

- $X([\tau, \tau']\,\alpha)$, $X((\tau, \tau']\,\alpha)$, $X([\tau, \tau')\,\alpha)$, $X((\tau, \tau')\,\alpha)$ *and* $X([\tau]\,\alpha)$ *are* occlusion formulas.

Fixed fluent formulas, fluent change formulas, reassignment formulas, interval reassignment formulas and occlusion formulas are called timed formulas.

Definition 18.6 (Static formula). *A static formula is a temporal formula, a value formula, a fixed fluent formula, a fluent change formula,* `true`*,* `false`*, or a combination of static formulas formed using the standard logical connectives together with quantification over values and time.*

Note that the formulas `true` and `false` are not the same as the boolean values true and false.

Definition 18.7 (Change formula). *A change formula is a formula that is rewritable to the form $Q\bar{v}(\alpha_1 \vee \cdots \vee \alpha_n)$ where $Q\bar{v}$ is a sequence of quantifiers with variables, and each α_i is a conjunction of static, occlusion and reassignment formulas. The change*

formula is called balanced *iff the following two conditions hold.* (a) *Whenever a feature $f(\bar{\omega})$ appears inside a reassignment or occlusion formula in one of the α_i disjuncts, it must also appear in all other α_i's inside a reassignment or occlusion formula with exactly the same temporal argument.* (b) *Any existentially quantified variable v in the formula, whenever appearing inside a reassignment or occlusion formula, only does so in the position $f \,\hat{=}\, v$.*

Definition 18.8 (Application formula). *An* application formula *is any of the following*: (a) *a balanced change formula*; (b) $\Lambda \supset \Delta$, *where Λ is a static formula and Δ is a balanced change formula*; *or* (c) *a combination of elements of types* (a) *and* (b) *formed with conjunction and universal quantification over values and time.*

Application formulas will be used in dependency constraints and action type specifications. The structural constraints on balanced change formulas and application formulas guarantee the proper generation of the occlusion predicate in the translation from $\mathcal{L}(\text{ND})$ to $\mathcal{L}(\text{FL})$. Restricting the structure of these formulas will guarantee first-order reducibility of the circumscription policy applied to the narrative.

Definition 18.9 (Occurrence formula). *An occurrence formula has the form $[\tau, \tau']\, \Psi$, where τ and τ' are temporal terms and Ψ is an action term.*

Definition 18.10 (Persistence formula). *A persistence formula is an expression of the form $Per(\tau, f)$ where τ is a temporal term and f is a fluent term, an expression of the form $Dur(\tau, f, \omega)$ where τ is a temporal term, f is a fluent term and ω is a value term, or a combination of persistence formulas formed with conjunction and universal quantification over values or time.*

18.5.3 Statements

Definition 18.11 (Narrative statements). *The following types of narrative statements are available in the current version of TAL.*

An action type specification or action schema (labeled **acs***) has the form $[t, t']\, \Psi \rightsquigarrow \phi$, where t and t' are temporal variables, Ψ is an action term and ϕ is an application formula.*

A dependency constraint (labeled **dep***) consists of an application formula.*

A domain constraint (labeled **dom***) consists of a static formula.*

A persistence statement (labeled **per***) consists of a persistence formula.*

An observation statement (labeled **obs***) consists of a static formula.*

An action occurrence statement (labeled **occ***) consists of an occurrence formula $[\tau, \tau']\, \Psi$ where τ and τ' are ground temporal terms and Ψ is a ground action term.*

All of these statement types have been provided with intuitive meanings in Section 18.3 except dependency constraints, which will be used to model side effects of actions (Section 18.8) and qualifications to actions (Section 18.9).

18.6   The TAL Base Language $\mathcal{L}$(FL)

This section defines the current base language $\mathcal{L}$(FL) used in TAL. The translation from $\mathcal{L}$(ND), which has already been described, to the first-order language $\mathcal{L}$(FL) is presented in Section 18.6.1. The base language $\mathcal{L}$(FL) is an order-sorted classical first-order language with equality. We assume familiarity with standard ways to define vocabulary and variable types in sorted logics. Additionally, a temporal structure must be chosen for the temporal sort $\mathcal{T}$. This would include a domain such as the natural numbers or integers and associated operators. It was mentioned previously that a number of choices as to temporal structure could be made.

   $\mathcal{L}$(FL) currently uses the following predicates where $\mathcal{T}$ is the temporal sort, $\mathcal{F}$ is a supersort of all fluent sorts and $\mathcal{V}$ is a supersort of all value sorts.

- *Holds*: $\mathcal{T} \times \mathcal{F} \times \mathcal{V}$—The *Holds* predicate expresses what value a feature has at each timepoint, and is used in the translation of fixed fluent formulas; for example, the formula $[0]$ loc(boris) $\hat{=}$ home1 $\wedge$ loc(gun) $\hat{=}$ office can be translated into *Holds*$(0,$ loc(boris), home1$) \wedge$ *Holds*$(0,$ loc(gun), office$)$.

- *Occlude*: $\mathcal{T} \times \mathcal{F}$—The *Occlude* predicate expresses the fact that a persistent or durational feature is exempt from its persistence or default value assumption, respectively, at a given timepoint. It is used in the translation of the R, I and X operators, which are intended to change the values of features.

- *Occurs*: $\mathcal{T} \times \mathcal{T} \times \mathcal{A}$—The *Occurs* predicate expresses that a certain action occurs during a certain time interval, and is used in the translation of action occurrence statements and action type specifications.

18.6.1   Translation from $\mathcal{L}$(ND) to $\mathcal{L}$(FL)

The following translation function is used to translate $\mathcal{L}$(ND) formulas into $\mathcal{L}$(FL).

Definition 18.12 (*Trans* Translation function). *Trans is called the* expansion transformation, *and is defined as follows. All variables occurring only on the right-hand side are assumed to be fresh variables.*

The formulas true *and* false *need no translation*:

$$Trans(\text{true}) = \text{true}$$
$$Trans(\text{false}) = \text{false}$$

Basic macros are translated into $\mathcal{L}$(FL) *predicates*:

$$Trans([\tau] \, f(\bar{\omega})) = Holds(\tau, \, f(\bar{\omega}), \text{true})$$
$$Trans([\tau] \, f(\bar{\omega}) \hat{=} \omega) = Holds(\tau, \, f(\bar{\omega}), \omega)$$
$$Trans(X([\tau] \, f(\bar{\omega}))) = Occlude(\tau, \, f(\bar{\omega}))$$
$$Trans(X([\tau] \, f(\bar{\omega}) \hat{=} \omega)) = Occlude(\tau, \, f(\bar{\omega}))$$
$$Trans([\tau, \tau'] \, \Psi) = Occurs(\tau, \tau', \Psi), \quad \textit{where } \Psi \textit{ is an action term}$$

In some versions of TAL, the $\mathcal{L}$(ND) *functions Per and Dur are also translated into* $\mathcal{L}$(FL) *predicates. Here, they are translated directly into constraints on fluent values*

and occlusion.

$$Trans(Per(\tau, f)) = \forall t.\tau = t + 1 \land \neg Occlude(t + 1, f) \supset$$
$$\forall v[Holds(t + 1, f, v) \equiv Holds(t, f, v)]$$
$$Trans(Dur(\tau, f, \omega)) = \neg Occlude(\tau, f) \supset Holds(\tau, f, \omega)$$

Top-level connectives and quantifiers are left unchanged:

$$Trans(\neg\alpha) = \neg Trans(\alpha)$$
$$Trans(\alpha \mathcal{C} \beta) = Trans(\alpha)\mathcal{C} Trans(\beta), \quad where \; \mathcal{C} \in \{\land, \lor, \supset, \equiv\}$$
$$Trans(\mathcal{Q}v[\alpha]) = \mathcal{Q}v[Trans(\alpha)], \quad where \; \mathcal{Q} \in \{\forall, \exists\}$$

Action type specifications use $\rightsquigarrow$ to denote the definition of an action. For this version of TAL, it is simply translated into an implication.

$$Trans([\tau, \tau'] \; \Psi \rightsquigarrow \phi) = Trans([\tau, \tau'] \; \Psi) \supset Trans(\phi)$$

Fixed fluent formulas can contain nested connectives and quantifiers, which are transferred outside the scope of the temporal context $[\tau]$.

$$Trans([\tau] \; \mathcal{Q}v[\alpha]) = \mathcal{Q}v[Trans([\tau] \; \alpha)], \quad where \; \mathcal{Q} \in \{\forall, \exists\}$$
$$Trans([\tau] \; \neg\alpha) = \neg Trans([\tau] \; \alpha)$$
$$Trans([\tau] \; \alpha \mathcal{C} \beta) = Trans([\tau] \; \alpha)\mathcal{C} Trans([\tau] \; \beta),$$
$$where \; \mathcal{C} \in \{\land, \lor, \supset, \equiv\}$$

Nested connectives and quantifiers can also occur within occlusion formulas. However, the translation of these formulas has to be modified somewhat to take into account the fact that any occlusion formula should occlude all *fluents occurring within the scope of the occlusion operator: Even a disjunctive formula such as $X([\tau] \; \alpha \lor \beta)$ should occlude all fluents in α and* all *fluents in β and is therefore not equivalent to $X([\tau] \; \alpha) \lor X([\tau] \; \beta)$ but to $X([\tau] \; \alpha) \land X([\tau] \; \beta)$. The translation procedure takes this into account by removing negations inside the X operator, translating connectives occurring inside X into conjunctions, and converting all quantifiers inside X into universal quantification.*

$$Trans(X([\tau] \; \neg\alpha)) = Trans(X([\tau] \; \alpha))$$
$$Trans(X([\tau] \; \alpha \mathcal{C} \beta)) = Trans(X([\tau] \; \alpha) \land X([\tau] \; \beta)),$$
$$where \; \mathcal{C} \in \{\land, \lor, \supset, \equiv\}$$
$$Trans(X([\tau]\mathcal{Q}v[\alpha])) = \forall v[Trans(X([\tau]\alpha))], \quad where \; \mathcal{Q} \in \{\forall, \exists\}$$

Fixed fluent formulas can contain infinite temporal intervals. This is a shorthand notation; infinity is not part of the temporal sort and disappears in the translation.

$$Trans([\tau, \infty) \; \alpha) = \forall t[\tau \leqslant t \supset Trans([t]\alpha)]$$
$$Trans((\tau, \infty) \; \alpha) = \forall t[\tau < t \supset Trans([t]\alpha)]$$

Finite temporal intervals are permitted both in fixed fluent formulas and in the occlusion operator. Only one form of interval is shown; the extension to allow open, closed

and semi-closed intervals is trivial.

$$Trans([\tau, \tau']\,\alpha) = \forall t[\tau \leqslant t \leqslant \tau' \supset Trans([t]\alpha)]$$

$$Trans(X((\tau, \tau']\,\alpha)) = \forall t[\tau < t \leqslant \tau' \supset Trans(X([t]\alpha))]$$

The R and I operators are defined as follows. Again, one form of interval is shown.

$$Trans(R((\tau, \tau']\,\alpha)) = Trans(X((\tau, \tau'], \alpha)) \land Trans([\tau']\alpha)$$

$$Trans(R([\tau]\,\alpha)) = Trans(X([\tau]\,\alpha)) \land Trans([\tau]\,\alpha)$$

$$Trans(I((\tau, \tau']\alpha)) = Trans(X((\tau, \tau']\,\alpha)) \land Trans((\tau, \tau']\,\alpha)$$

$$Trans(I([\tau]\,\alpha)) = Trans(X([\tau]\,\alpha)) \land Trans([\tau]\,\alpha)$$

Finally, the C_T "changes to true" operator is defined as follows, with the operators C_F (changes to false) and C (changes) added for symmetry.

$$Trans(C_T([\tau]\,\alpha)) = \forall t[\tau = t + 1 \supset Trans([t]\,\neg\alpha)] \land Trans([\tau]\,\alpha)$$

$$Trans(C_F([\tau]\,\alpha)) = \forall t[\tau = t + 1 \supset Trans([t]\,\alpha)] \land Trans([\tau]\,\neg\alpha)$$

$$Trans(C([\tau]\,\alpha)) = Trans(C_T([\tau]\,\alpha) \lor C_F([\tau]\,\alpha))$$

Example 18.2 *(Narrative translation).* The following is a translation of several of the $\mathcal{L}(ND)$ statements in the Russian Airplane Hijack scenario into $\mathcal{L}(FL)$. For brevity, the translation is limited to one statement of each statement class; the remaining formulas are translated in a similar manner.

per1 $\forall t, thing, t'\ [t = t' + 1 \land \neg Occlude(t' + 1, \mathsf{loc}(thing)) \supset$
$\forall v[Holds(t' + 1, \mathsf{loc}(thing), v) \equiv Holds(t', \mathsf{loc}(thing), v)]]$

dom1 $\forall t, pthing, person_1, person_2\ [\neg(person_1 = person_2) \land$
$Holds(t, \mathsf{inpocket}(person_1, pthing), \mathsf{true}) \supset$
$\neg Holds(t, \mathsf{inpocket}(person_2, pthing), \mathsf{true})]$

acs2 $\forall t_1, t_2, person, pthing\ [Occurs(t_1, t_2, \mathsf{pickup}(person, pthing)) \supset$
$(Holds(t_1, \mathsf{loc}(person), value(t_1, \mathsf{loc}(pthing))) \supset$
$Holds(t_2, \mathsf{inpocket}(person, pthing), \mathsf{true}) \land$
$\forall t\ [t_1 < t \land t \leqslant t_2 \supset Occlude(t, \mathsf{inpocket}(person, pthing))])]$

obs1 $Holds(0, \mathsf{loc}(\mathsf{boris}), \mathsf{home1}) \land Holds(0, \mathsf{loc}(\mathsf{gun}), \mathsf{office}) \land$
$Holds(0, \mathsf{loc}(\mathsf{comb1}), \mathsf{home1}) \land \neg Holds(0, \mathsf{drunk}(\mathsf{boris}), \mathsf{true})$

occ1 $Occurs(1, 2, \mathsf{pickup}(\mathsf{boris}, \mathsf{comb1}))$

18.7 Circumscription and TAL

The commonsense intuition one would like to capture and formally model in TAL is the fact that at a particular level of abstraction, relations between and properties of objects generally have reasons for changing and if not, we can assume, unless observed otherwise, that these are the only *possible* changes we need to be concerned about when reasoning about the specific environment around us and knowledge associated with that environment. So far, we have shown how one can encode in a principled manner sufficient reasons for the possibility of change by using a combination of the *Occlude* predicate and automatic translations from the surface language $\mathcal{L}(ND)$ to the base language $\mathcal{L}(FL)$. When specifying narratives in TAL, all sufficient reasons for the possibility of change are specified using the reassignment macros *R*, *I* and *X* in

dependency constraints and action type definitions. When translated, these statements result in constraints on the *Occlude* predicate. Unfortunately, what has been achieved so far under-constrains our notion of normative change, for we would also like to state that these are the only, or necessary, reasons for *possible* change.

In order to add this additional constraint to our action theories, we will appeal to the use of circumscription (Section 6.4 in Chapter 6 of this Handbook; [54, 55]) with an additional twist. Rather than applying a circumscription policy to the whole action theory, the theory will be partitioned and we will apply circumscription selectively to different partitions. Although this technique, which we call filtered circumscription [14], is now commonly used in other action theories [65, 48], in the context of action and change, it was first proposed in Sandewall [60]. Here, it was called filtered preferential entailment and was used as a basis for several of the definitions of preferential entailment in [61].

The basic idea will be to first circumscribe the predicate *Occlude* in that part of the action theory containing action occurrence statements and dependency constraint statements. This will result in a set of preferred or minimal models for the action theory providing a definition of all timepoints and features where it is possible for them to change value based on the constraints in the theory. Of course, these models will also contain spurious change since we have only provided sufficient and necessary conditions for the definition of *Occlude*. To rule out spurious change, we will then filter the resulting circumscriptive sub-theory with that part of the theory containing persistence statements. Persistence statements specify when features should not change value, assuming a predefined definition of *Occlude* which circumscription provides. In this manner, any model containing feature change not mandated by the implicit occlusion policy in the action theory will be excluded as a model of the action theory. For example, the Yale Shooting Problem [28] involves loading a gun, waiting, and shooting. Since waiting occludes no fluents, interpretations where the gun becomes unloaded while waiting are filtered out, yielding the intended conclusion that the gun remains loaded at the start of the shooting action. A separate circumscription policy will be used for that part of the action theory containing action occurrence statements, where the predicate *Occurs* will be circumscribed. Finally, all partitions will be conjoined.

Due to the structural syntactic constraints built into statement definitions in $\mathcal{L}(\text{ND})$, we can show that the two circumscribed sub-theories which are 2nd-order due to the use of circumscription, can be reduced to logically equivalent first-order theories. In fact, since only positive occurrences of the predicates *Occlude* and *Occurs* occur in the two circumscribed partitions of the action theory, respectively, a standard syntactic transformation on formulas may be used to generate the necessary conditions for both predicates. This in fact is a form of predicate completion, and it is related to Definition 7.3.6 in Chapter 7 of this Handbook.

The formal definition of the circumscription policy used in TAL will use the following terminology:

- Let $\mathcal{N}$ denote the collection of narrative statements contained in a narrative in $\mathcal{L}(\text{ND})$, and let $\mathcal{N}_{\text{per}}$, $\mathcal{N}_{\text{obs}}$, $\mathcal{N}_{\text{occ}}$, $\mathcal{N}_{\text{acs}}$, $\mathcal{N}_{\text{dom}}$, and $\mathcal{N}_{\text{dep}}$ denote the sets of persistence statements, observation statements, action occurrence statements, action type specifications, domain constraint statements, and dependency constraint statements in $\mathcal{N}$, respectively.

- Let Γ denote the translation of $\mathcal{N}$ into $\mathcal{L}(\text{FL})$ using the *Trans* translation function, and let Γ_{per}, Γ_{obs}, Γ_{occ}, Γ_{acs}, Γ_{dom}, and Γ_{dep} denote the persistence formulas, observation formulas, action occurrence formulas, action type specifications, domain constraint formulas, and dependency constraint formulas in Γ, respectively.

- Let Γ_{fnd} denote the set of foundational axioms in $\mathcal{L}(\text{FL})$, containing unique names axioms, unique values axioms, etc.

- Let Γ_{time} denote the axiomatization of the particular temporal structure used in TAL.

In the following, we assume familiarity with circumscription [54, 55] and common notation used to denote circumscription policies [47]. Let

$$\Gamma = \Gamma_{\text{per}} \wedge \Gamma_{\text{obs}} \wedge \Gamma_{\text{dom}} \wedge \Gamma_{\text{occ}} \wedge \Gamma_{\text{dep}} \wedge \Gamma_{\text{acs}}$$

be the translation of an action narrative in $\mathcal{L}(\text{ND})$ into a first-order theory in $\mathcal{L}(\text{FL})$ as described previously. Based on the discussion above, we use circumscription to minimize *Occurs* in Γ_{occ} and *Occlude* in $\Gamma_{\text{dep}} \wedge \Gamma_{\text{acs}}$ as follows:

$$\Gamma_1 = \Gamma_{\text{per}} \wedge \Gamma_{\text{obs}} \wedge \Gamma_{\text{dom}} \wedge CIRC[\Gamma_{\text{occ}}; \textit{Occurs}] \wedge$$
$$\quad CIRC[\Gamma_{\text{dep}} \wedge \Gamma_{\text{acs}}; \textit{Occlude}]$$

In addition, let

$$\Gamma_2 = \Gamma_{\text{fnd}} \wedge \Gamma_{\text{time}}$$

For any narrative $\mathcal{N}$ in TAL, a preferred narrative theory in the base logic $\mathcal{L}(\text{FL})$ is defined as

$$\Delta_{\mathcal{N}} = \Gamma_2 \wedge \Gamma_1$$

We say that a formula α in the base logic $\mathcal{L}(\text{FL})$ is preferentially entailed by the narrative $\mathcal{N}$ whose translation into $\mathcal{L}(\text{FL})$ is Γ iff

$$\Delta_{\mathcal{N}} \models \alpha$$

Observe that there are several equivalent formalizations of $\Delta_{\mathcal{N}}$ due to the following general property of circumscription (p. 311, [47]): for any sentence B not containing P, Z (where P is minimized and Z is varied),

$$CIRC[\Gamma(P, Z) \wedge B; P; Z] \equiv CIRC[\Gamma(P, Z); P; Z] \wedge B \qquad (18.1)$$

and the observation that

$$CIRC[\Gamma_{\text{occ}}; \textit{Occurs}] \wedge CIRC[\Gamma_{\text{dep}} \wedge \Gamma_{\text{acs}}; \textit{Occlude}]$$
$$\quad \equiv CIRC[\Gamma_{\text{occ}} \wedge \Gamma_{\text{dep}} \wedge \Gamma_{\text{acs}}; \textit{Occurs}, \textit{Occlude}]$$

From this, it follows that $\Delta_{\mathcal{N}}$ is equivalent to

$$\Delta' = \Gamma_2 \wedge \Gamma_{\text{per}} \wedge \Gamma_{\text{obs}} \wedge \Gamma_{\text{dom}} \wedge CIRC[\Gamma_{\text{occ}} \wedge \Gamma_{\text{dep}} \wedge \Gamma_{\text{acs}}; \textit{Occurs}, \textit{Occlude}]$$

and since $\Gamma_2 \wedge \Gamma_{\text{obs}} \wedge \Gamma_{\text{dom}}$ does not contain *Occurs* or *Occlude*, this is equivalent to

$$\Delta' = \Gamma_{\text{per}} \wedge CIRC[\Gamma_2 \wedge \Gamma_{\text{obs}} \wedge \Gamma_{\text{dom}} \wedge \Gamma_{\text{occ}} \wedge \Gamma_{\text{dep}} \wedge \Gamma_{\text{acs}}; Occurs, Occlude]$$

Note that it is important that Γ_{per} is outside the circumscriptive theory due to the fact that it contains occurrences of the predicate *Occlude*. Consequently, filtered circumscription is fundamental to the approach used in TAL.

As it stands, $\Delta_{\mathcal{N}}$ is in fact, a second-order theory due to the fact that Γ_1 contains two second-order circumscription formulas, $CIRC[\Gamma_{\text{occ}}; Occurs]$ and $CIRC[\Gamma_{\text{dep}} \wedge \Gamma_{\text{acs}}; Occlude]$. Due to structural syntactic constraints in the narrative $\mathcal{N}$ in $\mathcal{L}(\text{ND})$ which are carried over to its translation Γ in $\mathcal{L}(\text{FL})$, it can be shown that both $CIRC[\Gamma_{\text{occ}}; Occurs]$ and $CIRC[\Gamma_{\text{dep}} \wedge \Gamma_{\text{acs}}; Occlude]$ are reducible to logically equivalent first-order formulas. We now show this. First some preliminaries.

An occurrence of a predicate symbol in a formula is *positive* if it is in the range of an even number of negations (this is assuming that the connectives $\supset$ and $\equiv$ have been eliminated and replaced by other connectives in some equivalent normal form). A formula $A(P)$ is *positive* (relative to P) if all occurrences of P in $A(P)$ are positive.

Based on Definitions 18.7 and 18.8 for change and application formulas in $\mathcal{L}(\text{ND})$ and the definition of the translation function *Trans* from $\mathcal{L}(\text{ND})$ into $\mathcal{L}(\text{FL})$, it is straightforward to show that the predicate *Occurs* can only appear positively in Γ_{occ} and that the predicate *Occlude* can only appear positively in $\Gamma_{\text{dep}} \wedge \Gamma_{\text{acs}}$. The following proposition can then be applied to show that both $CIRC[\Gamma_{\text{occ}}; Occurs]$ and $CIRC[\Gamma_{\text{dep}} \wedge \Gamma_{\text{acs}}; Occlude]$ are reducible to logically equivalent first-order formulas [17]:

Proposition 18.1. *(See p. 316, [47].) If $A(P, Z)$ is positive relative to P, then the circumscription $CIRC[A(P, Z); P; Z]$ is equivalent to*

$$A(P, Z) \wedge \neg \exists x, z[P(x) \wedge A(\lambda y (P(y) \wedge x \neq y), z)]$$

In fact, it can be shown that predicate completion can be applied to Γ_{occ} and $\Gamma_{\text{dep}} \wedge \Gamma_{\text{acs}}$, respectively. The following proposition will be of use. Let $\bar{x}$ be a tuple of variables, and $F(\bar{x})$ be a formula with all parameters explicitly shown.

Proposition 18.2. *(See p. 309, [47].) If $F(\bar{x})$ does not contain P, then the circumscription $CIRC[\forall \bar{x}(F(\bar{x}) \supset P(\bar{x})); P]$ is equivalent to $\forall \bar{x}(F(\bar{x}) \equiv P(\bar{x}))$.*

This proposition generalizes to conjunctions of formulas of the form $\forall \bar{x} F(\bar{x}) \supset P(\bar{x})$. Using a number of syntactic transformations [8, 10], it can be shown that

$$\Gamma_{\text{occ}} = \bigwedge_{i=1}^{n} \forall \bar{x}(F_i(\bar{x}) \supset Occurs(\bar{x})) \tag{18.2}$$

where n is the number of *Occurs* formulas in Γ_{occ}. By the generalization of Proposition 18.2 and (18.2), it follows that

$$CIRC[\Gamma_{\text{occ}}; Occurs] = \forall \bar{x} \left[\left(\bigvee_{i=1}^{n} F_i(\bar{x}) \right) \equiv Occurs(\bar{x}) \right] \tag{18.3}$$

In addition it can be shown [8, 10] that

$$
\Gamma_{\text{dep}} \wedge \Gamma_{\text{acs}}
$$
$$
= \left[\bigwedge_{i=1}^{n} B_i \wedge \bigwedge_{j=1}^{k} C_j \wedge \forall \bar{x} \left[\left(\bigvee_{i=1}^{n} F_i(\bar{x}) \vee \bigvee_{j=1}^{k} G_j(\bar{x}) \right) \supset Occlude(\bar{x}) \right] \right]
$$
(18.4)

where B_i and C_i contain no occurrences of the *Occlude* predicate.

By (18.1), it follows that,

$$
CIRC[\Gamma_{\text{dep}} \wedge \Gamma_{\text{acs}}; Occlude]
$$
$$
= \bigwedge_{i=1}^{n} B_i \wedge \bigwedge_{j=1}^{k} C_j \wedge
$$
$$
CIRC \left[\forall \bar{x} \left[\left(\bigvee_{i=1}^{n} F_i(\bar{x}) \vee \bigvee_{j=1}^{k} G_j(\bar{x}) \right) \supset Occlude(\bar{x}) \right]; Occlude \right] \quad (18.5)
$$

By the generalization of Proposition 18.2 and (18.5), it follows that

$$
CIRC[\Gamma_{\text{dep}} \wedge \Gamma_{\text{acs}}; Occlude]
$$
$$
= \bigwedge_{i=1}^{n} B_i \wedge \bigwedge_{j=1}^{k} C_j \wedge \forall \bar{x} \left[\left(\bigvee_{i=1}^{n} F_i(\bar{x}) \vee \bigvee_{j=1}^{k} G_j(\bar{x}) \right) \equiv Occlude(\bar{x}) \right] \quad (18.6)
$$

Consequently, one can reduce Γ_1 in $\Delta_{\mathcal{N}}$ to a logically equivalent first-order formula. Under the assumption that Γ_2 in $\Delta_{\mathcal{N}}$ is also first-order (the temporal structure has a first-order axiomatization), for any narrative $\mathcal{N}$, its translation into a preferred narrative in $\mathcal{L}(\text{FL})$, $\Delta_{\mathcal{N}}$, is a first-order theory.

Example 18.3 *(Circumscription of the RAH scenario).* Though the circumscription of the *Occlude* predicate can be translated into a single first order formula, we have instead chosen for expository purposes to generate a separate formula for each ground fluent in the narrative, where the conjunction of these formulas is entailed by the original 2nd-order circumscription axiom. Here, we show a subset of these formulas for the Russian Airplane Hijack scenario.

First, the following are the necessary and sufficient conditions for loc(boris) to be occluded at any given point in time. For example, if boris is at home at time 3, which is a precondition for the action occurrence [3,5] travel(boris, home1, office), then his location will be occluded at time 5, when the final effects of the travel action take place.

$\forall t \, [Occlude(t, \text{loc(boris)}) \equiv t = 5 \wedge Holds(3, \text{loc(boris)}, \text{home1}) \vee$
$t = 10 \wedge Holds(8, \text{loc(boris)}, \text{office}) \vee t = 11 \wedge Holds(10, \text{loc(boris)}, \text{airport}) \vee$
$14 \leqslant t \wedge t \leqslant 15 \wedge$
$Holds(13, \text{loc(sas609)}, \text{run609}) \wedge Holds(13, \text{onplane(sas609, boris)}, \text{true}) \vee$
$t = 16 \wedge Holds(13, \text{loc(sas609)}, \text{run609}) \wedge$
$Holds(13, \text{onplane(sas609, boris)}, \text{true})]$

The conditions for occlusion of loc(dimiter) are quite similar, but differ in certain time-points given that boris and dimiter do not travel at the same time.

$\forall t$ [$Occlude(t,$ loc(dimiter)$) \equiv t = 4 \wedge Holds(2,$ loc(dimiter), home3$) \vee$
$t = 7 \wedge Holds(5,$ loc(dimiter), office$) \vee t = 10 \wedge$
$Holds(9,$ loc(dimiter), airport$) \vee$
$14 \leqslant t \wedge t \leqslant 15 \wedge Holds(13,$ loc(sas609), run609$) \wedge$
$Holds(13,$ onplane(sas609, dimiter), true$) \vee$
$t = 16 \wedge Holds(13,$ loc(sas609), run609$) \wedge$
$Holds(13,$ onplane(sas609, dimiter), true$)]$

The fluent inpocket(boris,gun) may be occluded at time 7 if boris is at the required location when attempting to pick it up, but inpocket(dimiter,gun) can never be occluded.

$\forall t$ [$Occlude(t,$ inpocket(boris, gun)$) \equiv t = 7 \wedge Holds(6,$ loc(boris), loc(gun)$)]$
$\forall t$ [$Occlude(t,$ inpocket(dimiter, gun)$) \equiv$ `false`]

18.8 Representing Ramifications in TAL

The ramification problem [46, 20, 37, 48, 51, 21, 62, 25, 68] states that it is unreasonable to explicitly specify all the effects of an action in an action specification itself. One would rather prefer to state the direct effects of actions in the action specification and then use deductive machinery to derive the indirect effects of actions using the direct effects of actions together with general knowledge of *directional* dependencies among features specified in some background theory. The feature dependencies specified do not necessarily have to be based solely on notions of physical or other causality, but often are. A solution to the ramification problem is important from the representational perspective, where one strives for incremental, modular and intuitive characterizations of action and change. When one thinks of actions at a certain level of abstraction, one normally thinks of actions in terms of their direct effects and one would like to represent actions as such. On the other hand, causality plays an important role in any type of reasoning about action and change, therefore modular and incremental theories of causal and other dependencies among features is equally important to represent as is the interaction between actions and causal theories.

Some earlier approaches to solving the ramification problem made use of pure domain constraints (essentially logical implication) in order to infer side effects of actions. For the Russian Airplane Hijack scenario, for example, one might specify that everyone onboard an airplane is always in the same physical location as the airplane. Should one fly the airplane to another location, a direct effect would constrain the airplane to be in the new location, and the locations of everyone onboard would have to change location to the airplane's new location in order to still satisfy the domain constraint. This type of solution is non-directional, which may sometimes be of advantage but may also lead to unexpected or unintended results. For example, in some representations, invoking an action that moves a single person would also cause the airplane and everyone else onboard to move.

The key insight in providing a good solution to the ramification problem is that of finding appropriate and representationally efficient ways of encoding directionality in dependencies among features which cause change in addition to allowing longer

chains of directional feature change. Logical implication, for example, is one way to encode a dependency constraint among features, but it under-constrains the directionality of a dependency due to the fact that the contrapositive to an implication formula is logically equivalent to that formula. Another important point to keep in mind is the way in which dependencies are triggered. This is highly contextual and though it is often the case that change triggers change, it is also the case that state triggers change. Both types of context and combinations of both should be expressible in action theories.

The TAL solution to the ramification problem involves the use of dependency constraints, which were formally specified in Definition 18.11. In this definition, the similarity between action type specifications and dependency constraints may be noted. Whereas an action type specification is an application formula conditionalized by the occurrence of an action ($[t, t']\ \Psi$ where Ψ is an action term) and then a precondition once that action is invoked, a dependency constraint consists of an application formula without such an action occurrence precondition. In a sense, while actions must be explicitly invoked, dependency constraints are constantly active. In both cases, there is an explicit directionality of feature change implicit in the representation. Technically, this is achieved by noting that features are occluded via assignment operators on the right hand side of implications, whereas they are not on the left hand side. This together with the minimization policy for occlusion and persistence statements permits the encoding of directionality of change in a fine-grained manner.

This solution to the ramification problem can be directly applied to the Russian Airplane Hijack scenario. Recall that the definition of the travel and fly actions included formulas explicitly causing anything a person was carrying to move to the same destination (Section 18.3.2). Clearly it would be better if such a formula could be factored out and modeled as a side effect of a person moving between two locations in any manner, rather than having to be specified for every action that causes a person to move. This can be represented using the following feature dependency constraint, stating that if the fact that *person* is at *loc* becomes true (changes to true)—in other words, if the person has just moved to *loc*—then anything the person carries in his pockets will also move to the same location. The use of explicit reassignment with the R operator ensures that such changes are permitted despite the general persistence assumption for the loc fluent.

dep1 $\forall t, person, pthing, loc\ [[t]\ \text{inpocket}(person, pthing) \wedge C_T([t]\ \text{loc}(person) \hat{=} loc) \supset$
$R([t]\ \text{loc}(pthing) \hat{=} loc)]$

With this change, the travel action can be simplified as follows:

acs3 $[t_1, t_2]\text{travel}(person, loc_1, loc_2) \rightsquigarrow [t_1]\ \text{loc}(person) \hat{=} loc_1 \supset$
$R([t_2]\ \text{loc}(person) \hat{=} loc_2)$

The fly action can be simplified in a similar manner. Before showing the new definition of this action though, we will consider one more indirect effect: people on board an airplane move when the airplane moves.

dep2 $\forall t, plane, person, loc\ [[t]\ \text{onplane}(plane, person) \wedge C_T([t]\ \text{loc}(plane) \hat{=} loc) \supset$
$R([t]\ \text{loc}(person) \hat{=} loc)]$

Note that the context for the dependency constraint in **dep2** has both a triggering condition (C_T) and a standard state condition. This is useful for encoding chaining of indirect effects.

Though this is quite similar to the previous indirect effect, it serves to illustrate an important property of fluent dependency constraints: It is possible to trigger not only a single indirect effect but a *chain* of indirect effects, which can be utilized to further modularize the specification of a narrative. In this particular scenario, causing an airplane to move will cause all people on board the airplane to move, which in turn will cause anything they are carrying to move, allowing the fly action to be modeled as follows:

acs1 $[t_1, t_2]$fly$(plane, runway_1, runway_2) \rightsquigarrow [t_1]$ loc$(plane) \mathbin{\hat{=}} runway_1 \supset$
    $I((t_1, t_2)$ loc$(plane) \mathbin{\hat{=}}$ air$) \wedge R([t_2]$ loc$(plane) \mathbin{\hat{=}} runway_2)$

18.9   Representing Qualifications in TAL

The qualification problem [67, 46, 20, 49, 51, 65, 11, 42, 69] was identified by McCarthy [53, 54] while developing systems for representing general commonsense knowledge. McCarthy showed a way to deal with the representational problem by using circumscription. In his own words,

The "qualification problem" immediately arose in representing general commonsense knowledge. It seemed that in order to fully represent the conditions for the successful performance of an action, an impractical and implausible number of qualifications would have to be included in sentences expressing them. [54]

A solution to the qualification problem would involve a normative representation of an action which would model the fact that an action can be invoked unless *something* prevents it from being invoked, where that *something* is assumed by default not to exist unless explicitly represented in an action theory. Additionally, when qualifications to actions are learned, the representation should permit an incremental and elaboration tolerant means of adding such qualifications to the action theory.

We have now modeled most of the Russian Airplane Hijack scenario in TAL, but we have not provided a means for modeling qualifications to actions in a representationally efficient, incremental and elaboration tolerant manner. Some examples of qualifications to actions in the RAH scenario would be: someone who carries a gun cannot board a plane, or someone who is drunk may or may not be able to board a plane. In fact, it may be the case that there are qualifications to qualifications. For example, security personnel should be able to board a plane with a gun.

There are already a number of solutions to various aspects of the qualification problem in the literature, some of which would be applicable in TAL. However, many of these solutions are dependent on the assumption of highly constrained action types, where (for example) actions must correspond to simple state transition with a precondition state and an effect state with no description of what happens in the duration of an action. As we have shown, actions in TAL go far beyond this limited form of representation. We would like to provide a solution that retains at least the following features of TAL:

- Any state, including the initial state, can be completely or incompletely specified using observations and domain constraints.

- Actions can be context-dependent and non-deterministic. They can have duration and internal state, and the duration may be different for different executions of the action. There may be concurrent actions with partially overlapping execution intervals.

- There can be dynamic processes continuously taking place independently of any actions that may occur.

- Domain constraints can be used for specifying logical dependencies between fluents generally true in every state or across states. They may vary over time.

- Actions can have side effects, which may be delayed and may affect the world at multiple points in time. They may in turn trigger other delayed or non-delayed side effects.

We would also like to retain the first-order reducibility of the circumscription axiom in any solution to the qualification problem in TAL. In order to do this, the following restrictions and assumptions will apply. First, we will be satisfied with a solution where invoking a qualified action either has no effect or has some well-defined effect. Secondly, we will restrict the solution to the off-line planning and prediction problems and not claim a complete solution for the post-diction problem, which would require being able to conclude that an action was qualified because its successful execution would have contradicted an observation of some feature value after that action was invoked.

18.9.1   Enabling Fluents

To handle the qualification problem, we use a solution based on defaults where each action type in a narrative is associated with an *enabling fluent*, a boolean durational fluent with default value true and with the same number and type of arguments as the action type. This fluent will be used in the precondition of the action and will usually be named by prefixing "poss-" to the name of the action. For example, the boarding action in the RAH scenario will be associated with an enabling fluent poss-board(*person*, *plane*). We also add a persistence statement for this fluent stating that it is a durational fluent. Recall that a durational feature retains its default value unless an additional constraint specifies that there is an exception to that value at a particular point or points in time. **acs4** is then modified as follows:

per5 $\forall t, person, plane \, [Dur(t, \text{poss-board}(person, plane), \text{true})]$
acs4$'$ $[t_1, t_2] \, \text{board}(person, plane) \rightsquigarrow$
$[t_1] \, \text{poss-board}(person, plane) \wedge \text{loc}(person) \doteq \text{airport} \supset$
$R([t_2] \, \text{loc}(person) \doteq value(t_2, \text{loc}(plane)) \wedge \text{onplane}(plane, person))$

The other action types are modified in a similar way. Note that the existing precondition that loc(*person*) $\doteq$ airport remains in the action definition and will not be moved to the definition of poss-board. This is a modeling issue, where some conditions are identified as "ordinary" preconditions whereas others are identified as qualifications which are moved outside the action type specification. A similar modeling issue already arose in the case of ramifications, where some effects are considered "ordinary"

action effects whereas others are considered to be indirect effects modeled using dependency constraints.

Now, suppose that board(*person, plane*) is executed between timepoints t_1 and t_2. If poss-board (*person, plane*) is false at t_1 for some reason, the action is qualified, or *disabled*. On the other hand, if the fluent is true at t_1, the action is *enabled*. Of course, it can still be the case that the action has no effects, if other parts of its precondition are false.

To generalize this, a context-independent action that should have no effect at all when qualified can be defined using a simple action definition of the form[8]

acs$m$ $[t_1, t_2]$ *action* $\rightsquigarrow$ $[t_1]$ *poss-action* $\land \alpha \supset R([t_2] \beta)$

where α is the precondition and β specifies the direct effects of the action (context-dependent actions are defined analogously). However, we also wanted to be able to define actions that do have some effects when they are qualified. This can be done by defining a context-dependent action that defines what happens when the enabling fluent is false:

acs$n$ $[t_1, t_2]$ *action* $\rightsquigarrow$ $([t_1]$ *poss-action* $\land \alpha_1 \supset R([t_2] \beta_1)) \land$
        $([t_1] \neg$ *poss-action* $\land \alpha_2 \supset R([t_2] \beta_2))$

For example, suppose that whenever anyone tries to board a plane but the action is qualified, they should try to find new transportation. In order to model this, we would add a new persistent fluent find-new-transportation(PERSON) : BOOLEAN and modify the boarding action from Section 18.3.2 as follows:

acs4$''$ $[t_1, t_2]$ board(*person, plane*) $\rightsquigarrow$
        $([t_1]$ poss-board(*person, plane*) $\land$ loc(*person*) $\hat{=}$ airport $\supset$
            $R([t_2]$ loc(*person*) $\hat{=}$ *value*$(t_2,$ loc(*plane*)) $\land$ onplane(*plane, person*))) $\land$
        $([t_1] \neg$poss-board(*person, plane*) $\land$ loc(*person*) $\hat{=}$ airport $\supset$
            $R([t_2]$ find-new-transportation(*person*)))

In this alternative scenario, if anyone is at the airport and tries to board a plane, and the action is qualified, they will have a goal of finding new transportation. If they are at the airport but the action is not qualified, they will board the plane. If they are not at the airport, none of the preconditions will be true, and invoking the action will have no effect. Note that it may very well be the case that they can not board for a more serious reason such as carrying a gun. This is a case where the original qualification might have to be qualified.

Regardless of whether a qualified action has an effect or not, its enabling fluent is a durational fluent with default value true. Therefore, the fluent will normally be true, and the action will normally be enabled. In the remainder of this section, we will examine some of the ways in which we can disable an action using strong and weak qualification.

---

[8]Note that due to the regularity of the solution, such extensions could be implicit in an action macro, thus avoiding unneeded clutter in the representation and delegating representation responsibility to the system rather than the knowledge engineer.

18.9.2 Strong Qualification

When there is sufficient information to conclude that an action will definitely not succeed, it is *strongly qualified*. This can be modeled by forcing its enabling fluent to be false at the timepoint at which the action is invoked.

For example, suppose that when a person has a gun in his pocket, it should be impossible for that person to board a plane. Then, whenever inpocket(*person*, gun) holds, poss-board(*person*) necessarily becomes false. This can be represented using a dependency constraint:

dep3 $\forall t, person, plane$ [[t] inpocket(*person*, gun) $\supset I$ ([t] ¬poss-board(*person*, *plane*))]

At any timepoint t when a person has a gun in his pocket, we use the I macro both to occlude poss-board(*person*, *plane*) for all airplanes, thereby releasing it from the default value axiom, and to make it false. This implies that as long as a person has a gun in his pocket, poss-board will be false for that person on all airplanes. If the gun is later removed from the pocket, this dependency constraint will no longer be triggered. At that time, assuming no other qualifications affect the enabling fluent, it will automatically revert to its default value, true.

18.9.3 Weak Qualification

Although strong qualification can often be useful, we may sometimes have enough information to determine that an action *may* fail, even though we cannot conclusively prove that it will. We call this *weak* qualification.

For example, we may want to model the fact that when a person is drunk, he *may or may not* be able to board an airplane, depending on whether airport security discovers this or not. We may not be able to determine within our model of the RAH scenario whether airport security does discover that any given person is drunk. In this case, whenever drunk(*person*) holds, we must release poss-board from the default value assumption, which would otherwise have forced poss-board to be true:

dep4 $\forall t, person$ [[t] drunk(*person*) $\supset \forall plane$ [X ([t] poss-board(*person*, *plane*))]]

At any timepoint t when a person is drunk, we occlude poss-board(*person*, *plane*) for all airplanes, but since we do not state anything about the *value* of the enabling fluent, it is allowed to be either true or false.

Although being able to state that an action *may* fail is useful in its own right, it is naturally also possible to restrict the set of models further by adding more statements to the scenario which could make it possible to infer whether poss-board(dimiter, sas609) is true or false at some or all timepoints. For example, we may know that people boarding sas609 are always checked more carefully, so that it is impossible for anyone who is drunk to be on board that airplane, which could be expressed using an additional domain constraint. In the context of postdiction, observation statements could be used in a similar manner. For example, adding the observation statement **obs5** [13] onplane(sas609, boris) to the narrative would allow us to infer that Boris did in fact board the plane and that poss-board(boris, sas609) was in fact true. He would then end up at his intended destination. If instead we added

the observation statement **obs6** [13] ¬onplane(sas609, boris), we could infer that he was unable to board the plane and he did not end up at his destination.

It should be noted that this approach to modeling qualification has similarities to a standard default solution to the qualification problem, but with some subtle differences. For example, it permits more control of the enabling precondition, even allowing it to change during the execution of an action. More importantly, it involves no changes to the minimization policy already used in TAL to deal with the frame and ramification problems, consequently the circumscription theory is still first-order reducible.

18.9.4 Qualification: Not Only For Actions

As we have shown, this approach to qualification is based on general concepts such as durational features and fluent dependency constraints, instead of introducing new predicates, entailment relations or circumscription policies specifically designed for dealing with the qualification problem. This is appealing not only because we avoid introducing new complexity into the logic, but also because reusing these more general concepts adds to the flexibility of the approach. In fact, exactly the same approach can be used for specifying qualifications to any rule or constraint. Most notably, one can provide qualifications for ramification constraints, thereby introducing defeasible side effects—or one can even qualify qualification constraints themselves.

As an example, when we initially considered the boarding action, the "natural" preconditions were that one had to be at the airport; this is the precondition encoded in the definition of board (**acs4**). Later, we found another condition that should qualify the action: No one should be able to board a plane carrying a gun. Now, however, we may discover that this qualification does not always hold: Airport security *should* be able to board a plane carrying a gun.

Assuming that there is a fluent is-security(person) : BOOLEAN, this exception to the general qualification rule could of course be modeled by changing the dependency constraint **dep3** in the following way:

dep3′ $\forall t, person, plane$ [[t] inpocket($person$, gun) ∧ ¬is-security($person$) ⊃
I ([t] ¬poss-board($person$, $plane$))]

However, we may later discover additional conditions under which it should be possible for a person to board a plane with a gun, and we do not want to modify **dep3** each time. Instead, the qualification itself should be qualified. This can easily be done using the same approach as for actions. A new enabling fluent guns-forbidden(PERSON, PLANE) : BOOLEAN is added for the qualification constraint, and **dep3** is modified as follows:

dep3″ $\forall t, person, plane$ [[t] inpocket($person$, gun) ∧ guns-forbidden($person$, $plane$) ⊃
I ([t] ¬poss-board($person$, $plane$))]

Now, we can qualify the qualification **dep3** simply by making guns-forbidden false for some person and airplane. In order to do this, we add a new dependency constraint:

dep5 $\forall t, person, plane$ [[t] is-security($person$) ⊃
I ([t] ¬guns-forbidden($person$, $plane$))]

18.9.5 Ramifications as Qualifications

A problem related to the qualification problem occurs in formalisms where ramification constraints and qualification constraints are expressed as domain constraints [20, 49]. Assume, for example, that we are reasoning about the blocks world, and that we have the following domain constraint (expressed using TAL syntax), stating that no two blocks can be on top of the same block:

$$\textbf{dom} \quad \forall t, x, y, z \ [[t] \ \mathsf{on}(x, z) \wedge \mathsf{on}(y, z) \supset x = y]$$

Now, suppose that the direct effect of the action $\mathsf{put}(\mathsf{A}, \mathsf{C})$ is $\mathsf{on}(\mathsf{A}, \mathsf{C})$, and the action is executed in a state where $\mathsf{on}(\mathsf{B}, \mathsf{C})$ is true. Then, we cannot determine syntactically whether the domain constraint should be interpreted as a ramification constraint (since no two blocks can be on top of C, B must be removed) or as a qualification constraint (since no two blocks can be on top of C, the action should fail).

In TAL, however, all indirect effects of an action must be expressed as *directed* dependency constraints. Therefore, this problem simply does not arise. For example, if a ramification constraint is required, the following dependency constraint can be used:

$$\textbf{dep} \quad \forall t, x, y, z \ [[t] \ \mathsf{on}(x, z) \wedge C_T([t + 1] \ \mathsf{on}(y, z)) \wedge x \neq y \supset R([t + 1] \ \neg\mathsf{on}(x, z))]$$

If x is on z, and we then place y on z, then an indirect effect is that x is removed from z. On the other hand, if a qualification constraint is required, an enabling fluent $\mathsf{poss\text{-}put}(\textsc{block}, \textsc{block})$ can be used and the following qualification condition would then be added:

$$\textbf{dep} \quad \forall t, x, y, z \ [[t] \ \mathsf{on}(x, z) \wedge x \neq y \supset I([t] \ \neg\mathsf{poss\text{-}put}(y, z))]$$

Clearly, the problem of determining whether a constraint should be implicitly interpreted as a qualification or a ramification does not arise in this approach. One could criticize such a solution as over-constraining the action theory model, but then again, use of domain constraints could equally well be criticized for under-constraining the model.

A description of the TAL representation of the Russian Airplane Hijack scenario is now complete and the general methods used to resolve the frame, ramification and qualification problems have been described. The partial translations into $\mathcal{L}(\mathrm{FL})$ were done using VITAL [39], a research tool that can be used to study problems involving action and change within TAL and generate visualizations of action scenarios and preferred entailments.

18.10 Action Expressivity in TAL

For the sake of brevity, narratives used as examples in the literature are generally modeled at a rather high level of abstraction. This is especially true when a narrative is used for the purpose of demonstrating the properties of a solution to a specific problem; for example, the Russian Airplane Hijack scenario was explicitly designed for the demonstration of qualification constraints. This, however, should not be taken

to mean that this is the *only* level of abstraction possible in TAL. We briefly illustrate this point by adding more realistic timing to several action types in the RAH scenario and by introducing effects at inner timepoints during the execution of the fly action.

As shown in previous examples, the timing of an action occurrence has often been completely specified in the corresponding action occurrence statement:

occ $[1, 2]$ pickup(boris, comb1)

In many cases, the duration of the action is better specified in the action itself. An action specification can contain arbitrary constraints on its parameters, which can be used to constrain the time required for boarding as well as the amount of time required when boarding fails:

acs $[t_1, t_2]$ board(*person*, *plane*) $\rightsquigarrow$
  $([t_1]$ poss-board(*person*, *plane*) $\wedge$ loc(*person*) $\hat{=}$ airport $\supset$
    $R([t_2]$ loc(*person*) $\hat{=}$ *value*(t_2, loc(*plane*)) $\wedge$ onplane(*plane*, *person*) $\wedge$
    $t_2 = t_1 + 100)) \wedge$
  $([t_1] \neg($poss-board(*person*, *plane*) $\wedge$ loc(*person*) $\hat{=}$ airport$) \supset$
    $t_2 = t_1 + 10)$

This also illustrates the use of contextually dependent effects, where the exact outcome of the action is determined by the state of the world when it is invoked (though non-deterministic and incompletely specified effects are also possible). An arbitrary number of conditions (mutually exclusive or not) can be used to specify the effects of an action.

The timing in the action occurrence statements is then relaxed by introducing a number of temporal constants. Here, boris begins picking up comb1 at time 1. He does not know when he will finish, but at the next timepoint (boris1 + 1), he will begin traveling to the office.

occ $[1,$ boris1$]$ pickup(boris, comb1)
occ $[$boris1 $+ 1,$ boris2$]$ travel(boris, home1, office)
occ $\ldots$

Of course, action durations do not have to be defined using a constant. If distances between locations are modeled using a dist fluent, one can specify the duration of a fly action as follows:

acs $[t_1, t_2]$ fly(*plane*, *runway*$_1$, *runway*$_2$) $\rightsquigarrow$
  $[t_1]$ loc(*plane*) $\hat{=}$ *runway*$_1$ $\supset$
  $I((t_1 + 200, t_2)$ loc(*plane*) $\hat{=}$ air$) \wedge R([t_2]$ loc(*plane*) $\hat{=}$ *runway*$_2$) $\wedge$
  $t_2 = t_1 + 200 +$ dist(*runway*$_1$, *runway*$_2$)$/200$

Here, flying between two locations takes an initial 200 timepoints for taxiing, plus time proportional to the distance between the two locations. This more accurate model of the fly action can be further extended by modeling the remaining distance at any given time when the plane is in the air, by conjoining the following formula to the effects given above:

$\forall t$ $[t_1 + 200 < t \wedge t \leqslant t_2 \supset I([t]$ remaining-distance(*plane*) $\hat{=}$
  dist(*runway*$_1$, *runway*$_2$) $\cdot (t - t_1 - 200)/(t_2 - t_1 - 200))]$

At any timepoint within the interval $(t_1 + 200, t_2]$, the remaining distance is assigned a new value. Further elaborations to effects and timing can be added as required by the task to which the model will be applied.

18.11   Concurrent Actions in TAL

Much work in reasoning about action and change has been done under the (sometimes implicit) assumption that there is a single agent performing sequences of non-overlapping actions. The use of explicit metric time in TAL clearly enables the specification of narratives where action execution intervals are partly or completely overlapping, whether those actions are performed by a single agent or by multiple cooperating or adversarial agents. Similarly, the fact that actions can have non-unit duration and that one can specify in detail what happens during the execution interval enables richer domain models where a larger class of phenomena related to concurrency can be modeled. However, a complete treatment of concurrency also requires the ability to model interactions between concurrent effects of multiple actions. Such interactions can be *synergistic*, where two actions must be executed concurrently in order to achieve the desired effect. For example, moving a table requires lifting both sides of the table simultaneously in order to avoid the undesired side effects of everything on the table sliding off onto the floor. Interactions may also be *accumulative*, as when a number of agents are placing packages in a vehicle for transportation, or *harmful*, where one action provides the desired effect unless certain other actions are executed concurrently.

In each of these cases, the composite effect of executing several actions is not equivalent to the logical conjunction of the individual effects. For example, lifting the left side of the table causes the table to tilt, as does lifting the right side of the table, but lifting both sides at once cancels this effect. Though this could in theory be handled by modeling all possible interactions within each action definition, this would clearly be an extremely non-modular solution and would suffer from a combinatorial explosion in the number of conditional effects required in each action definition. This is especially true when dealing with actions with duration, where the number of combinations is determined not only by the number of actions but also by the number of ways two or more actions can overlap in time. The use of ramification constraints also complicates the issue by introducing interactions between actions and chains of (potentially delayed) ramification effects.

For these reasons, a more principled and indirect solution was proposed by Karlsson and Gustafsson [35], where actions do not directly change the state of the world but instead produce a set of influences. Fluent dependency constraints can then be used to model how the world is affected by a combination of influences.

18.11.1   Independent Concurrent Actions

The use of independent concurrent actions involving disjoint sets of features is unproblematic in TAL. This is illustrated in the following narrative, describing a world with two types of actions (LightFire and PourWater), and a number of agents (bill and bob) and other objects (wood1 and wood2). All variables appearing free are implicitly universally quantified.

acs1 $[s, t]$ LightFire$(a, wood) \supset ([s]$ dry$(wood) \supset R((s, t]$ fire$(wood)))$
acs2 $[s, t]$ PourWater$(a, wood) \supset R((s, t]$ ¬dry$(wood) \wedge$ ¬fire$(wood))$
obs1 $[0]$ dry(wood1) $\wedge$ ¬fire(wood1) $\wedge$ wood(wood1)
obs2 $[0]$ dry(wood2) $\wedge$ ¬fire(wood2) $\wedge$ wood(wood2)
occ1 $[2, 7]$ LightFire(bill, wood1)
occ2 $[2, 7]$ LightFire(bob, wood2)
occ3 $[9, 12]$ PourWater(bob, wood1)

The first action law states that if an agent *a* lights a fire using a piece of wood, and the wood is dry, then the wood will be on fire. The second action law states that if somebody pours water on an object, then the object will no longer be dry, and will cease being on fire. There are two pieces of wood (wood1 and wood2) which are initially dry and not burning. Two fires are lit by bill and bob during the temporal interval [2, 7], and then bob pours water on bill's fire at [9, 12]. Since no concurrency is involved, the expected effects will take place: Both pieces of wood will be on fire at 7, and wood1 will no longer be burning at 12.

18.11.2   Interacting Concurrent Actions

Now consider the case where actions affecting the same fluents occur concurrently. For example, suppose bob pours water on wood1 while bill is still lighting the fire. Intuitively, the wood should not be on fire at 7. We formalize this in TAL by modifying **occ3**.

occ3 $[3, 5]$ PourWater(bob, wood1)

From the modified narrative one can still infer that wood1 is on fire at time 7, because the effects of LightFire(bill, wood1) are only determined by whether the piece of wood is dry at time 2, whereas in reality the effects of any action may also be altered by the direct and indirect effects of other concurrent actions. A slight modification of the narrative above illustrates another problem. Assume that **occ3** is replaced with the following:

occ3 $[3, 7]$ PourWater(bob, wood1)

Now, the lighting and pouring actions end at the same time. From **acs1** and **occ1** one can infer the effect $[7]$ fire(wood1) and from **acs2** and **occ3** one can infer $[7]$ ¬fire(wood1). Both effects are asserted to be direct and indefeasible. Thus, the narrative becomes inconsistent. The conclusion one would like to obtain is again that the wood is not on fire.

18.11.3   Laws of Interaction

Karlsson and Gustafsson [35] considers two solutions to these problems.

In the first solution, action laws are extended to allow references to other action occurrences and the effects of LightFire are made conditional on the fact that there is no interfering PourWater action. As noted in the introduction, this solution makes action descriptions less modular and there may be a combinatorial explosion in the number of conditional effects for each action. Other problems include the fact that

a concurrent action might only interfere with part of an action's effects, leading to further complexity in action laws.

The second solution is based on the assumption that interactions resulting from concurrency are best modeled not on the level of actions but on the level of features. Action laws encode the influences that an action has upon the environment of the agent; in the fire example, $[s, t]$ LightFire$(a, wood)$ would have the effect $I((s, t]$ fire*$(wood, \text{true}))$ where fire*$(wood, \text{true})$ is a fluent representing an influence to make the feature fire$(wood)$ true. This example follows the convention of representing the influences on an actual fluent $f(\bar{\omega})$ with $f^*(\bar{\omega}, v)$, where v is a value in the domain of f. Similarly, dependency constraints are modified to result in influences rather than actual fluent changes. The actual effects that these influences have on the environment are then specified in a special type of dependency laws called influence laws. Applying this solution to the fire example yields the following narrative:

dom1	$Per(\text{fire}(wood)) \wedge Dur(\text{fire}^*(wood, v), \text{false})$
dom2	$Per(\text{dry}(wood)) \wedge Dur(\text{dry}^*(wood, v), \text{false})$
acs1	$[s, t]$ LightFire$(a, wood) \supset I((s, t]$ fire$^*(wood, \text{true}))$
acs2	$[s, t]$ PourWater$(a, wood) \supset I((s, t]$ dry$^*(wood, \text{false}))$
dep1	$[s] \neg\text{dry}(wood) \supset I([s]$ fire$^*(wood, \text{false}))$
inf1	$[s, s + 3]$ fire$^*(wood, \text{true}) \wedge \neg\text{fire}^*(wood, \text{false}) \supset R([s + 3]$ fire$(wood))$
inf2	$[s]$ fire$^*(wood, \text{false}) \supset R([s] \neg\text{fire}(wood))$
inf3	$[s, s + 3]$ dry$^*(wood, \text{true}) \wedge \neg\text{dry}^*(wood, \text{false}) \supset R([s + 3]$ dry$(wood))$
inf4	$[s]$ dry$^*(wood, \text{false}) \supset R([s] \neg\text{dry}(wood))$
obs1	$[0] \neg\text{fire}(wood1) \wedge \text{dry}(wood1)$
occ1	$[2, 6]$ LightFire$(\text{bill}, wood1)$
occ2	$[3, 5]$ PourWater$(\text{bob}, wood1)$

The action laws **acs1** and **acs2** and dependency law **dep1** produce influences; for example, **dep1** states that the fact that the wood is not dry produces an influence fire*$(wood, \text{false})$ to extinguish the fire (if there is one). The effects of these influences, alone and in combination, are specified in **inf**x; for example, in order to affect the feature fire$(wood)$, the influence fire*$(wood, \text{true})$ for starting the fire has to be applied without interference from fire*$(wood, \text{false})$ for an extended period of time. In the preferred models of this narrative, wood1 will be wet at $[4, \infty)$, implying that fire*$(\text{wood1}, \text{false})$ will hold at $[4, \infty)$; consequently there is no interval $[s, s + 3]$ where fire*$(\text{wood1}, \text{true}) \wedge \neg\text{fire}^*(\text{wood1}, \text{false})$, and fire$(\text{wood1})$ will never become true.

The case when an effect of one action enables the effect of another action can also be handled with conditional influence laws. For instance, the following influence law states that opening a door requires initially keeping the latch open (the example is originally due to Allen [1]):

inf1	$[t]$ latch-open $\wedge [t, t + 5]$ open$^*(\text{true}) \supset R([t + 5]$ open$)$

Though not explicitly shown here, it is possible to use separate modular influence laws to specify the result of arbitrary combinations of influences, including combinations that lead to no effect at all. Influences can naturally also be combined with the TAL approach to ramification, both in the sense that ramifications may lead to influences and

in the sense that influences may cause chains of ramifications. One can also use influence laws to model resource conflicts, with either deterministic, non-deterministic or prioritized outcomes when two agents attempt to use the same resource [35, 24]. This results in a highly flexible and modular solution to many problems associated with concurrency, regardless of whether that concurrency is due to actions, ramifications or delayed effects.

18.12   An Application of TAL: TALplanner

The flexibility of TAL as a language for describing and modeling actions with concurrent effects, dependencies between fluents and other commonly occurring aspects of dynamic domains also makes it eminently suitable for modeling planning domains. This is especially true for planners that make extensive use of domain knowledge in various forms. For this reason, TAL has been used as the semantic basis for a planner called TALplanner [12, 13, 41], where TAL is used for modeling not only actions, initial states and standard state-based goals but also a set of control formulas acting as constraints on the set of valid plans. This latter use of logical formulas was initially inspired by the planner TLPLAN [2].

One of the intended uses of TAL in TALplanner is as a specification language providing a declarative semantics for planning domains and plans. This is an important difference from TLPLAN where only control formulas are based on the use of logic and actions are instead modeled using an operational semantics. But unlike Green's approach [22], which involved not only representing planning domains in logic but also *generating* plans using a resolution theorem prover, the declarative semantics of TAL currently serves mainly as a specification for the proper behavior of the planning algorithm. The TALplanner implementation generates plans using a procedural forward-chaining search method together with a search tree which is pruned with the help of temporal control formulas.

Given that performance is of paramount importance in a planner, the full expressivity of TAL is intended to be introduced into the planner implementation in stages; the full power of the language, including non-deterministic actions, chains of ramifications and arbitrary interactions between concurrent actions, must be approached carefully. Having the specification of the proper semantics of such constructs available from the beginning is useful even in the initial phase, providing a better view of what extensions will be desired in the future, which sometimes affects the basic framework of an implementation. Nonetheless, the language currently used for domain specifications in TALplanner is a subset of the full language for TAL described previously in this chapter.

Planning domains and planning problems also require the specification of certain types of information that were not originally supported in TAL or its predecessors. This required a set of new additions to the language which will be described.

Thus, both extensions and limitations relative to TAL are in order. This falls neatly within the TAL policy of providing macro languages adapted to specific tasks together with a translation into a single unified first-order base language $\mathcal{L}(\text{FL})$ with a well-defined semantics and circumscription policy. While the details of the new macro language $\mathcal{L}(\text{ND})^*$ are beyond the scope of this chapter (see Kvarnström [41] for a
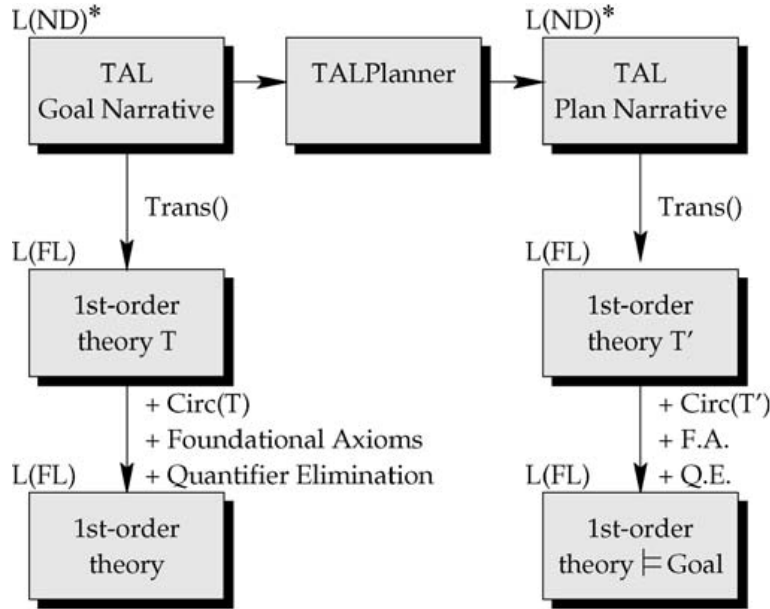
Figure 18.3: The relation between TAL and TALplanner.

complete definition), most of the extensions are used in the example planning domain discussed below.

Planning as narrative generation

TAL is based on the use of narratives, and automated planning can be viewed as a form of narrative generation where an initial narrative, specifying an initial state as well as various forms of domain knowledge, is incrementally extended by adding new action occurrences—in other words, new steps in the plan. The intention, then, is to generate a suitable set of action occurrences such that the desired goals are satisfied in the resulting complete narrative.

Fig. 18.3 contains an extended version of the diagram previously shown in Fig. 18.2. As seen in the top row of this figure, the input to TALplanner is a narrative in the extended macro language $\mathcal{L}(ND)^*$. This narrative is sometimes called a *goal narrative*, emphasizing the fact that it specifies a planning problem instance, and is usually denoted by $\mathcal{N}$. The goal narrative consists of two parts: A domain description, defining among other things the operators that are available to the planner, and a problem instance description, defining the initial state and the goal. TALplanner uses this high-level description of a planning problem to search for a set of TAL action occurrences (plan steps) that can be added to this narrative so that in the corresponding logical model, a goal state is reached. If this succeeds, the output is a new TAL narrative in $\mathcal{L}(ND)^*$ where the appropriate set of TAL action occurrences has been added. This narrative is sometimes called a *plan narrative*, emphasizing the fact that it represents a solution to a planning problem. Both goal narratives and plan narratives can be translated into $\mathcal{L}(FL)$ (the second row in the figure). As in pure TAL, a number of foundational axioms are required, and a standard TAL circumscription policy is applied, yielding complete definitions of the *Occlude* and *Occurs* predicates (the third row). Further details are available in Kvarnström [41].

Adding action occurrences to a standard TAL narrative is a non-monotonic operation, in the sense that conclusions entailed by the original narrative may have to be retracted once a new action occurrence is added. However, at each step in the planning process, one would also prefer to be able to determine whether a certain conclusion will remain valid regardless of what new actions may be added to a plan. This is especially important in the context of temporal control formulas, where a candidate plan should not necessarily be discarded for violating a control formula if this violation might be "repaired" by adding new actions.

The key to solving this problem lies in the flexibility of the TAL solution to the frame problem. By selecting a search space where new action occurrences are constrained not to begin before any of the actions already present in the plan—that is, if there are actions beginning at times 0, 10 and 273, one cannot add a new action beginning at 272—one can guarantee that along any branch of the forward-chaining search tree, there is a monotonically increasing temporal horizon such that any new effects introduced by future actions will take place strictly after this horizon.[9] Then, the standard definition of inertia can be altered to ensure that persistence is applied up to and including this temporal horizon, while leaving fluents unconstrained at all later timepoints. This is easily done by changing the TAL translation function while retaining the same circumscription policy.

It should be noted that this approach is not equivalent to assuming a complete lack of knowledge after the temporal horizon. On the contrary, any fluent constraints resulting from action effects or (in a future implementation) domain constraints are still equally valid after the temporal horizon; only the persistence assumption has been relaxed at those timepoints where the complete set of effects is unknown. Thus, depending on the effects that have been applied so far, it can still be possible to prove that a control formula has been definitely violated after the temporal horizon, which is essential for the performance of the concurrent version of TALplanner.

An example planning domain

We will now show some examples of the use of $\mathcal{L}(\text{ND})^*$ in modeling the timed version of the ZenoTravel domain, originally used in the AIPS 2002 International Planning Competition [45, 50]. Due to space limitations, the complete domain description will not be provided. Nevertheless, the most pertinent aspects of the modeling language will be presented in sufficient detail.

The ZenoTravel domain contains a number of aircraft that can fly people between cities. There are five planning operators available: Persons may board and debark from aircraft, and aircraft may fly, zoom (fly quickly, using more fuel), and refuel. There are no restrictions on how many people an aircraft can carry. Flying and zooming are equivalent except that zooming is generally faster and uses more fuel. Fig. 18.4 shows a tiny example problem, with arrows pointing out goal locations.

Objects in a planning problem are modeled using standard TAL values, and state variables are modeled using TAL fluents.

---

[9]Note that this does not rule out the generation of plans with concurrent actions and one version of TALplanner does generate actions concurrently.

Figure 18.4: A ZenoTravel problem instance.

domain   THING :elements {... }
domain   AIRCRAFT :parent THING :elements {... }
domain   PERSON :parent THING :elements {... }
domain   CITY :elements {... }

feature   at(THING, CITY), in(PERSON, AIRCRAFT) :domain BOOLEAN
feature   fuel(AIRCRAFT) :domain INTEGER
. . .

Operators are modeled using a new form of operator statement, which uses a new syntax with explicit preconditions, prevail conditions, durations and effects. As specified by the competition organizers, the time required to board a plane is specified using the boarding-time fluent, which is here multiplied by 1000 in order to provide higher precision timing. Note also that the plane is required to remain at its location during boarding. The time required to fly between two cities is proportional to the distance and inversely proportional to the speed of the aircraft.[10]

operator   board($person$, $aircraft$, $city$) :at s
 :duration $value(t, 1000 * \text{boarding-time})$:as dur
 :precond $[s]$ at($person$, $city$) $\wedge$ at($aircraft$, $city$)
 :prevail $[s + 1, s + dur]$at($aircraft$, $city$)
 :effects $[s + 1]$at($person$, $city$) := false, $[s + dur]$in($person$, $aircraft$) := true

operator   fly($aircraft$, $city_1$, $city_2$) :at s
 :duration $value(t, 1000 * \text{distance}(city_1, city_2) / \text{slow-speed}(aircraft))$:as dur
 :precond $[s]$ at($aircraft$, $city_1$) $\wedge city_1 \neq city_2 \wedge$
            $[s]$ fuel($aircraft$) $\geqslant$ distance($city_1$, $city_2$) $*$ slow-burn($aircraft$)
 :effects $[s + 1]$ at($aircraft$, $city_1$) := false,
            $[s + dur]$ at($aircraft$, $city_2$) := true,
            $[s + dur]$ fuel($aircraft$) := $value(s, \text{fuel}(aircraft)) - 1$

---

[10]We appeal to the use of *semantic attachment* [71] techniques in the implementation of TAL and TALplanner by liberal use and invocation of built in mathematical and other functions associated with value domains for features.

Control formulas specify constraints that must be satisfied in the logical model corresponding to a solution plan. In some respects, the central use of explicit control formulas is really what stands out from other automated planning paradigms. Control formulas are intended to represent the high-level heuristics or commonsense smarts that one usually assumes a human might use when faced with specific planning problems in well-defined domains. Initially, a person may not have sufficient competence about a domain. Consequently, the plans generated may not be the best and will certainly take longer to generate. As a person acquires a feel for a domain, certain constraints are then applied when generating plans which in turn minimize the search space. It is this intuition which is behind the use of control formulas as a domain dependent means of limiting the huge search space of action combinations one is faced with when using a forward chaining planner. The technique is also incremental in nature. Control formulas may be added incrementally as one learns more about the domain in question thus improving the efficiency of the planner.

The following two control formulas used in the ZenoTravel domain state that passengers should only board an aircraft when they desire to be in another city, and that they should only debark when they have reached their destination. Free variables are assumed to be universally quantified.

control :name "only-board-when-necessary"
$[t] \neg \text{in}(person, aircraft) \wedge [t+1] \text{in}(person, aircraft) \supset$
$\exists city, city_2 \, [\, [t] \, \text{at}(person, city) \wedge \text{goal}(\text{at}(person, city_2)) \wedge city \neq city_2 \,]$

control :name "only-debark-when-in-goal-city"
$[t] \text{in}(person, aircraft) \wedge [t+1] \neg \text{in}(person, aircraft) \supset$
$\exists city \, [\, [t] \, \text{at}(aircraft, city) \wedge \text{goal}(\text{at}(person, city)) \,]$

In addition to these statements, which are valid in an entire planning domain, the planner also needs a complete specification of the initial state (using TAL observation statements) and a specification of the state-based goals that should be achieved. The latter is specified using goal statements, consisting of TAL fluent formulas that must hold in the final state resulting from executing a solution plan.

The following are possible goal and initial state statements for the example in Fig. 18.4:

goal at(person1, city3) $\wedge$ at(person2, city1) $\wedge$ at(person3, city3) $\wedge$ at(person5, city1)
obs $\forall city \, [\, [0] \, \text{at}(person1, city) \equiv city = \text{city0} \,]$
obs $\forall city \, [\, [0] \, \text{at}(person2, city) \equiv city = \text{city0} \,]$
obs $\forall city \, [\, [0] \, \text{at}(plane1, city) \equiv city = \text{city2} \,]$
obs $[0] \, \text{fuel}(plane1) \,\hat{=}\, 5$

The main statement types for goal and plan narratives have now been introduced. A goal narrative is input to the forward-chaining TALplanner system and if possible, a TAL plan narrative is generated by the planner which contains action occurrence statements and timings for such statements which entail the goal and control statements originally included in the goal narrative.

It was stated that the strategy used in TALplanner is not "planning as theorem-proving", but using TAL as a specification language for developing planners. Perhaps

a better way to describe TAL and its relation to TALplanner is not only as a specification framework, but as "theorem-proving as an aid to plan generation". One can clearly see from Fig. 18.3 that in the plan generation process, one can use TAL to reason about partial plans being generated. In fact, during plan generation, a simple form of inference is currently used to verify that control formulas are satisfied in theories associated with partial plans. In the plan execution process, one can use TAL to verify and monitor the plan execution process by querying the current state of a robotic system with TAL formulas. This is a form of on-line model-checking similar to the progression algorithms used in TALplanner and TLPLAN. TALplanner also uses a limited form of resolution to reason about control formulas and operators during the initial (preprocessing) phase of the planning process, inferring a set of facts that must necessarily be true during the invocation of an operator and thereby improving the performance of checking control formulas during the planning phase. The flexible framework described in this section offers great potential for leveraging the use of logic with planning in a pragmatic and efficient manner.

18.13 Summary

This chapter provides a presentation of the latest stable version of TAL, a temporal action logic for reasoning about action and change. In the article, we present the basic narrative framework for specifying action scenarios using two languages $\mathcal{L}(\text{ND})$ and $\mathcal{L}(\text{FL})$. A definition of the circumscription policy used for TAL is provided in addition to proposals for partial solutions to the frame, ramification and qualification problems. Solutions are obviously dependent on the nature of the application domains to which they are applied. We say the solutions are partial because it is unclear whether such solutions would hold up practically unless one had specifications of the environments in which TAL would be used and a means of assessing whether the formalism would cover the spectrum of reasoning problems associated with a particular domain. Such a qualification would apply to any action and change formalism and assessments should be done using either formal assessment frameworks such as that described in the introduction to this chapter or empirical testing. TAL has been partially assessed for a particular type of application domain but much remains to be done in terms of assessing many of the newer extensions to TAL. That being said, TAL is one of the most expressive logical formalisms for reasoning about action and change, the underlying semantic framework is highly intuitive and TAL has been shown to correctly model the majority of benchmark problems proposed in the action and change research community. In the chapter, we have also provided a description as to how one could deal with the very complex problem of true concurrent actions and their interactions. We have concluded with an application of TAL to an award winning automated planner, TALplanner.

Acknowledgements

and Information Science, Linköping University: Marcus Bjäreland, Joakim Gustafsson, Lars Karlsson, Martin Magnusson, Andrzej Szałas and Witold Łukaszewicz.

Bibliography

[1] J. Allen. Temporal reasoning and planning. In Allen, Kautz, Pelavin, and Tenenberg, editors. *Reasoning about Plans*. Morgan Kaufmann, 1991.

[2] F. Bacchus and F. Kabanza. Using temporal logics to express search control knowledge for planning. *Artificial Intelligence*, 116 (1–2):123–191, 2000.

[3] C. Baral and M. Gelfond. Logic programming and reasoning about actions. In M. Fisher, D.M. Gabbay, and L. Vila, editors. *Handbook of Temporal Reasoning in Artificial Intelligence*. Elsevier Publications, 2005.

[4] M. Bjäreland. Two aspects of automating logics of action and change: Regression and tractability. Master's thesis, Linköping University, 1998. Thesis No 674. LiU-Tek-Lic 1998:09.

[5] M. Bjäreland and L. Karlsson. Reasoning by regression: Pre- and postdiction procedures for logics of action and change with nondeterminism. In *Proceedings of the 15th International Joint Conference on Artificial Intelligence (IJCAI'97)*, Nagoya, Japan, August 1997. Morgan Kaufmann, 1997.

[6] P. Doherty. Notes on PMON circumscription. Technical Report LITH-IDA-94-43, Dept. of Computer and Information Science, Linköping University, Linköping, Sweden, December 1994.

[7] P. Doherty. Reasoning about action and change using occlusion. In A.G. Cohn, editor. *Proceedings of the Eleventh European Conference on Artificial Intelligence (ECAI-1994)*, Amsterdam, The Netherlands, 1994, pages 401–405. John Wiley and Sons, Chichester, England, 1994.

[8] P. Doherty. PMON$^+$: A fluent logic for action and change, formal specification, version 1.0. Technical Report LITH-IDA-96-33, Department of Computer and Information Science, Linköping University, Linköping, Sweden, December 1996.

[9] P. Doherty and J. Gustafsson. Delayed effects of actions = direct effects + causal rules. *Linköping Electronic Articles in Computer and Information Science*, 3, 1998.

[10] P. Doherty, J. Gustafsson, L. Karlsson, and J. Kvarnström. TAL: Temporal Action Logics—language specification and tutorial. *Electronic Transactions on Artificial Intelligence*, 2(3–4):273–306, September 1998.

[11] P. Doherty and J. Kvarnström. Tackling the qualification problem using fluent dependency constraints: Preliminary report. In L. Khatib and R. Morris, editors. *Proceedings of the Fifth International Workshop on Temporal Representation and Reasoning (TIME-1998)*, Los Alamitos, CA, USA, May 1998, pages 97–104. IEEE Computer Society Press, 1998.

[12] P. Doherty and J. Kvarnström. TALplanner: An empirical investigation of a temporal logic-based forward chaining planner. In C. Dixon and M. Fisher, editors. *Proceedings of the Sixth International Workshop on Temporal Representation and*

Reasoning (TIME-1999), Orlando, Florida, USA, May 1999, pages 47–54. IEEE Computer Society Press, 1999.

[13] P. Doherty and J. Kvarnström. TALplanner: A temporal logic-based planner. *AI Magazine*, 22(3):95–102, 2001.

[14] P. Doherty and W. Łukaszewicz. Circumscribing features and fluents: A fluent logic for reasoning about action and change. In D.M. Gabbay and H.J. Ohlbach, editors. *Proceedings of the First International Conference on Temporal Logic (ICTL-1994)*, *Lecture Notes in Artificial Intelligence*, vol. 827, pages 82–100. Springer-Verlag, London, 1994.

[15] P. Doherty, W. Łukaszewicz, and E. Madalińska-Bugaj. The PMA and relativizing change for action update. In *Proceedings of the 6th International Conference on Principles of Knowledge Representation and Reasoning (KR'98)*, 1998.

[16] P. Doherty, W. Łukaszewicz, and A. Szałas. Computing circumscription revisited: Preliminary report. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI'95)*, vol. 2, pages 1502–1508, 1995.

[17] P. Doherty, W. Łukaszewicz, and A. Szałas. Computing circumscription revisited: A reduction algorithm. *Journal of Automated Reasoning*, 18:297–336, 1997.

[18] P. Doherty and P. Peppas. A comparison between two approaches to ramification: PMON(R) and $\mathcal{AR}_0$. In X. Yao, editor. *Proceedings of the 8th Australian Joint Conference on Artificial Intelligence*, pages 267–274. World Scientific, 1995.

[19] P. Doherty, W. Łukaszewicz, and A. Szałas. Explaining explanation closure. In *Proceedings of the 9th International Symposium on Methodologies for Intelligent Systems (ISMIS'96)*, 1996.

[20] M.L. Ginsberg and D.E. Smith. Reasoning about action II: The qualification problem. *Artificial Intelligence*, 35(3):311–342, 1988.

[21] E. Giunchiglia and V. Lifschitz. Dependent fluents. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI-1995)*, Montréal, Québec, Canada, pages 1964–1969. Morgan Kaufmann Publishers, San Mateo, CA, USA, 1995.

[22] C. Green. Applications of theorem proving to problem solving. In *Proceedings of the First International Joint Conference on Artificial Intelligence (IJCAI-1969)*. Morgan Kaufmann, 1969.

[23] J. Gustafsson. Extending temporal action logic for ramification and concurrency. Master's thesis, Linköping University, 1998. Thesis No 719. LiU-Tek-Lic 1998:54.

[24] J. Gustafsson. Extending temporal action logic. PhD thesis, Linköping Studies in Science and Technology, Dissertation No. 689, 2001.

[25] J. Gustafsson and P. Doherty. Embracing occlusion in specifying the indirect effects of actions. In L.C. Aiello, J. Doyle, and S.C. Shapiro, editors. *Proceedings of the Fifth International Conference on Principles of Knowledge Representation and Reasoning (KR-1996)*, pages 87–98. Morgan Kaufmann Publishers, San Francisco, CA, USA, 1996.

[26] J. Gustafsson and J. Kvarnström. Elaboration tolerance through object-orientation. In *Proceedings of the Fifth Symposium on Logical Formalizations of Commonsense Reasoning (Common Sense-2001)*, 2001.

[27] J. Gustafsson and J. Kvarnström. Elaboration tolerance through object-orientation. *Artificial Intelligence*, 153:239–285, March 2004.

[28] S. Hanks and D.V. McDermott. Default reasoning, nonmonotonic logics, and the frame problem. In *Proceedings of the Fifth National Conference on Artificial Intelligence (AAAI-1986)*, Philadelphia, Pennsylvania, USA, August 1986, pages 328–333. Morgan Kaufmann Publishers, Los Altos, CA, USA, 1986.

[29] L. Karlsson. Specification and synthesis of plans using the Features and Fluents framework. Master's thesis, Linköping University, 1995. Thesis No 469. LiU-Tek-Lic 1995:01.

[30] L. Karlsson. Causal links planning and the systematic approach to action and change. In *Proceedings of the AAAI 96 Workshop on Reasoning about Actions, Planning and Control: Bridging the Gap*, Portland, OR, August 1996. AAAI Press, 1996.

[31] L. Karlsson. Planning, truth criteria and the systematic approach to action and change. In *Proceedings of the 9th International Symposium on Methodologies for Intelligent Systems (ISMIS'96), Lecture Notes for Artificial Intelligence*. Springer-Verlag, 1996.

[32] L. Karlsson. Reasoning with incomplete initial information and nondeterminism in situation calculus. In *Proceedings of the 15th International Joint Conference on Artificial Intelligence (IJCAI'97)*, 1997.

[33] L. Karlsson. Anything can happen: on narratives and hypothetical reasoning. In *Proceedings of the 6th International Conference on Principles of Knowledge Representation and Reasoning (KR'98)*. Morgan Kaufmann, 1998.

[34] L. Karlsson. Actions, interactions and narratives. PhD thesis, Linköping Studies in Science and Technology, Dissertation No. 593, 1999.

[35] L. Karlsson and J. Gustafsson. Reasoning about concurrent interaction. *Journal of Logic and Computation*, 9(5):623–650, October 1999.

[36] L. Karlsson, J. Gustafsson, and P. Doherty. Delayed effects of actions. In H. Prade, editor. *Proceedings of the Thirteenth European Conference on Artificial Intelligence (ECAI-1998)*, Brighton, UK, August 1998, pages 542–546. John Wiley and Sons, Chichester, England, 1998.

[37] G.N. Kartha and V. Lifschitz. Actions with indirect effects. In *Proceedings of the International Conference on Knowledge Representation and Reasoning (KR)*, pages 341–350. Morgan Kaufmann, 1994.

[38] M. Koubarakis. Complexity results for first-order theories of temporal constraints. In J. Doyle, E. Sandewall, and P. Torasso, editors. *Proceedings of the Fourth International Conference on Principles of Knowledge Representation and Reasoning (KR-1994)*, pages 379–390. Morgan Kaufmann Publishers, San Francisco, CA, USA, 1994.

[39] J. Kvarnström. VITAL. An on-line system for reasoning about action and change using TAL. 1997–2006.

[40] J. Kvarnström. Applying domain analysis techniques for domain-dependent control in TALplanner. In M. Ghallab, J. Hertzberg, and P. Traverso, editors. *Proceedings of the Sixth International Conference on Artificial Intelligence Planning and Scheduling (AIPS-2002)*, Toulouse, France, April 2002, pages 101–110. AAAI Press, Menlo Park, CA, USA, 2002.

[41] J. Kvarnström. TALplanner and other extensions to temporal action logic. PhD thesis, Linköpings universitet, April 2005. Linköping Studies in Science and Technology, Dissertation no. 937.

[42] J. Kvarnström and P. Doherty. Tackling the qualification problem using fluent dependency constraints. *Computational Intelligence*, 16(2):169–209, May 2000.

[43] J. Kvarnström and P. Doherty. TALplanner: A temporal logic based forward chaining planner. *Annals of Mathematics and Artificial Intelligence*, 30:119–169, 2000.

[44] J. Kvarnström, P. Doherty, and P. Haslum. Extending TALplanner with concurrency and resources. In W. Horn, editor. *Proceedings of the Fourteenth European Conference on Artificial Intelligence (ECAI-2000)*, *Frontiers in Artificial Intelligence and Applications*. Berlin, Germany, August 2000, pages 501–505. IOS Press, Amsterdam, The Netherlands, 2000.

[45] J. Kvarnström and M. Magnusson. TALplanner in the Third International Planning Competition: Extensions and control rules. *Journal of Artificial Intelligence Research*, 20:343–377, December 2003.

[46] V. Lifschitz. Formal theories of action. In F.M. Brown, editor. *The Frame Problem in Artificial Intelligence: Proceedings of the 1987 Workshop*, Lawrence, Kansas, USA, April 1987, pages 35–58. Morgan Kaufmann Publishers, Los Altos, CA, USA, 1987.

[47] V. Lifschitz. Circumscription. In D.M. Gabbay, C.J. Hogger, and J.A. Robinson, editors. *Handbook of Artificial Intelligence and Logic Programming*, vol. 3, pages 297–352. Oxford University Press, 1991.

[48] F. Lin. Embracing causality in specifying the indirect effects of actions. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI-1995)*, Montréal, Québec, Canada, August 1995. Morgan Kaufmann Publishers, San Francisco, CA, USA, 1995.

[49] F. Lin and R. Reiter. State constraints revisited. *Journal of Logic and Computation*, 4(5):655–678, 1994.

[50] D. Long and M. Fox. The third international planning competition: Results and analysis. *Journal of Artificial Intelligence Research*, 20:1–59, December 2003.

[51] N. McCain and H. Turner. A causal theory of ramifications and qualifications. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1978–1984, 1995.

[52] J. McCarthy. Programs with common sense. In *Proceedings of the Teddington Conference on the Mechanization of Thought Processes*, London, pages 75–91. Her Majesty's Stationary Office, 1959.

[53] J. McCarthy. Epistemological problems of artificial intelligence. In R. Reddy, editor. *Proceedings of the Fifth International Joint Conference on Artificial Intelligence (IJCAI-1977)*, Cambridge, MA, USA, pages 1038–1044. William Kaufmann, 1977.

[54] J. McCarthy. Circumscription—a form of non-monotonic reasoning. *Artificial Intelligence*, 13(1–2):27–39, 1980. Reprinted in [56].

[55] J. McCarthy. Applications of circumscription to formalizing common sense knowledge. *Artificial Intelligence*, 28(1):89–116, 1986. Reprinted in [56].

[56] J. McCarthy. *Formalization of Common Sense, papers by John McCarthy edited by Vladimir Lifschitz*. Ablex, 1990.

[57] E.T. Mueller. *Commonsense Reasoning*. Morgan Kaufmann, 2006.

[58] R. Reiter. *Knowledge in Action*. MIT Press, 2001.

[59] E. Sandewall. Combining logic and differential equations for describing real-world systems. In R.J. Brachman, H.J. Levesque, and R. Reiter, editors. *Proceedings of the International Conference on Knowledge Representation and Reasoning (KR)*, pages 412–420. Morgan Kaufmann, 1989.

[60] E. Sandewall. Filter preferential entailment for the logic of action and change. In N.S. Sridharan, editor. *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence (IJCAI-1989)*, San Francisco, August 1989. Morgan Kaufmann Publishers, San Mateo, CA, USA, 1989.

[61] E. Sandewall. *Features and Fluents: A Systematic Approach to the Representation of Knowledge about Dynamical Systems*, vol. 1. Oxford University Press, 1994.

[62] E. Sandewall. Assessments of ramification methods that use static constraints. In *Proceedings of the International Conference on Knowledge Representation and Reasoning (KR)*, pages 99–110. Morgan Kaufmann, 1996.

[63] E. Sandewall. Cognitive robotics and its metatheory: Features and fluents revisited. *Electronic Transactions on Artificial Intelligence*, 2(3–4), 1998.

[64] L.K. Schubert. Monotonic solution of the frame problem in situation calculus. In H.E. Kyburg, R.P. Loui, and G.N. Carlson, editors. *Knowledge Representation and Defeasible Reasoning*, pages 23–67. Kluwer, 1990.

[65] M. Shanahan. *Solving the Frame Problem: A Mathematical Investigation of the Common Sense Law of Inertia*. The MIT Press, Cambridge, MA, USA, 1997.

[66] Y. Shoham. Nonmonotonic logics: Meaning and utility. In J.P. McDermott, editor. *Proceedings of the Tenth International Joint Conference on Artificial Intelligence (IJCAI-1987)*, Milan, Italy, August 1987, pages 388–393. Morgan Kaufmann Publishers, Los Altos, CA, USA, 1987.

[67] Y. Shoham. *Reasoning about Action and Change*. MIT Press, 1987.

[68] M. Thielscher. Ramification and causality. *Artificial Intelligence*, 89(1–2):317–364, 1997.

[69] M. Thielscher. The qualification problem: A solution to the problem of anomalous models. *Artificial Intelligence*, 131(1–2):1–37, 2001.

[70] M. Thielscher. *Reasoning Robots—The Art and Science of Programming Robotic Agents*. Springer, 2005.

[71] R.W. Weyhrauch. Prolegomena to a theory of mechanized formal reasoning. *Artificial Intelligence*, 13(1–2), 1980.