


Formal Semantics of Asbru – An Overview

Michael Balser, Christoph Duelli, Wolfgang Reif
 Lehrstuhl Softwaretechnik und Programmiersprachen
 Universität Augsburg
 86135 Augsburg, Germany
 {balser,duelli,reif}@informatik.uni-augsburg.de

ABSTRACT: This paper gives an overview of the formal semantics of a planning language called Asbru which has been specifically designed for the medical framework. A formal semantics is an important step within the Protocure project which is concerned with the quality assurance of medical guidelines and protocols. We have constructed a formal semantics in the style of Structural Operational Semantics (SOS). However, this paper is addressed to Asbru users. Therefore, we present sets of SOS rules as semiformal statecharts, which leads to a compact, graphical overview of the operational behaviour. In this style, the semantics documents the language best.

I. INTRODUCTION

This work is part of a European project called Protocure [11], which is concerned with the quality assurance of medical protocols. The idea is to model existing informal medical guidelines and protocols in the planning language Asbru [7] [12] and to verify certain properties. Already, Asbru has been used to formalize a variety of examples from different fields of medicine: diabetes mellitus, jaundice in new born babies, artificial ventilation of prematured babies, and others. Other approaches to model medical protocols are e.g. [4] [8] [6]. One of our major goals is to further utilize formal methods in the medical domain by verifying properties of protocols with mathematical rigour in the interactive theorem prover KIV [3], leading to the following overall picture.



Defining a formal semantics of Asbru is an important step within this project. It is the basis for the KIV representation of Asbru plans and the calculus rules. On the other hand it should also help to understand the Asbru language with all its details.

Our semantics for Asbru is operational and mainly consists of a number of Structural Operational Semantics

```

plan Regular_Treatments_3
intentions
  intermediate-state maintain bilirubin ≠ transfusion
  overall-state achieve bilirubin = observation
conditions
  filter-precondition bilirubin ≠ transfusion
  abort-condition  bilirubin = transfusion
                  ∨  bilirubin = pt-intensive
                  ∧  bilirubin_decrease < 1
                  [[4 h, -], [-, 6 h], [-, -], now]
plan-body any-order
  wait-for Observation
  plan-activation Phototherapy_Intensive
  plan-activation Phototherapy_Normal_Prescription
  plan-activation Phototherapy_Normal_Recommended
  plan-activation Observation
    
```

Fig. 1. Example Asbru plan from Jaundice case study

(SOS) rules [9]. However, the complete set of rules turned out to be lengthy and difficult to understand for Asbru users. Therefore, we present sets of SOS rules in semiformal a statechart notation. The notation does not contain all the technical details of the original rules, but serves as a very good overview of the operational behaviour. The notation has been very useful for further discussions within our heterogeneous group of people from formal methods, medicine, planning, knowledge bases and language design.

The main part of this paper gives statecharts for the most important features of Asbru. In order to show that the behaviour is indeed formally defined, we will also sketch the mapping of the graphical notation to SOS rules.

The paper is organized as follows. In Sect. II we will give a short overview of Asbru followed by notational issues in Sect. III. Section IV gives an overview of the semantics. The hierarchy of plans is explained in Sect. V, which is followed by the basic plan state model of Asbru in Sect. VI. This model is enriched with further important concepts of Asbru in Sections VII to X. Section XI gives an outlook on how KIV is used to formally verify properties of Asbru plans and Sect. XII concludes.

II. ASBRU IN A NUTSHELL

As an example, Fig. 1 displays a simplified version of one of the plans in the jaundice case study. Treating jaundice in newborn babies requires monitoring the level of bilirubin in the blood. Quantitative bilirubin levels are abstracted to qualitative values observation, pt_normal, pt_recommended, pt_intensive, and transfusion. The intention of this treatment plan is to maintain a bilirubin level lower than transfusion and to finally achieve a very low level of bilirubin called observation. The filter condition states that this plan is only applicable, if the bilirubin level is not too high in the beginning. The plan will be aborted, if bilirubin is too high or if it is very high and the decrease within 4 to 6 hours is not large enough. Four different alternative treatments are available. The applicability of these alternatives is determined by their own filter conditions (which are not contained in Fig. 1). For example, plan Phototherapy_Intensive can only be used, if bilirubin level is pt_intensive. Because of the "wait-for" construct, the successful completion of plan Observation is mandatory, other plans are optional.

Asbru is a plan oriented language. Several plans are organized in a hierarchy of plans. A parent plan can refer to other sub plans in its plan body. Conditions are used to control the applicability of a plan and to monitor its execution. Conditions can be monitored over time according to so called time annotations. The sub plans in the plan body can be organized using different body types (e.g. any-order). The current state of a plan – especially if a plan has been rejected, aborted, or completed – is propagated according to the plan hierarchy to its parent and sub plans. If a plan is mandatory, it must be completed, otherwise it may also be rejected or aborted.

III. NOTATION


A. EBNF

We will use an EBNF-like notation to describe the syntax of constructs of Asbru. Terminal symbols are written in normal style, names of the grammar rules are typeset in *italic*. Square brackets $[\cdot]$ denote optional parts, and alternatives are written as $(\cdot | \cdot)$. Zero or more repetitions are denoted as $\cdot^*$.

B. Statecharts

A statechart is a directed graph representing a state machine (a nondeterministic automaton). It is used to specify a system's dynamic behaviour. In this paper we will adopt the syntax of STATEMATE [10]. We will explain its basic features and semantics on the basis of Figure 2.

States are depicted as rounded rectangles. *Superstate* contains two substates. As the system can be only in one of these at a given time, *Superstate* is called an OR-state. When the system enters *Superstate*, it's initially in both *Superstate* itself and its substate *Default_state*. The default



$$T : event[condition]/action$$

Fig. 2. Statechart notation

substate is marked with an arc pointing from a black bullet to the state.

Possible state transitions are represented as directed arcs that may be labelled with guards of the form $event[condition]/action$. A guarded transition can only be taken when *event* occurs and *condition* holds at the same time. Note that an enabled transition *must* be taken (non-deterministic choice if several transitions from a state are enabled).

Parallel is an AND-State. Its substates *Thread1* and *Thread2* – separated by a dashed line – are executed synchronously in parallel. Once transition *T* is enabled, *action* will be executed and then *Parallel* will become active, changing the system's active states to *Superstate*, *Parallel*, *Thread1*, *Thread2*, *Thread1.D* and *Thread2.D*. By means of composite AND- and OR-states, we can create a *state hierarchy*, thus facilitating the readability of the statechart.

C. SOS rules

The formal semantics is given in SOS rules. Our rules will be of the following form.

$$\frac{\llbracket \varphi \rrbracket_{\sigma} \quad g(\sigma) \rightarrow g(\sigma') \quad h(\sigma) \nrightarrow \dots}{f(\sigma) \rightarrow f(\sigma')}$$

A configuration $f(\sigma)$ may step to a configuration $f(\sigma')$, if formula φ holds in valuation σ (a positive premise of style $\llbracket \cdot \rrbracket_{\sigma}$), configuration $g(\sigma)$ is able to step to $g(\sigma')$ (a positive premise of style $c \rightarrow c'$) and $h(\sigma)$ is currently not able to take a step (a negative premise of style $c \nrightarrow$). Valuation σ assigns values to variables.

Details on this notation can e.g. be found in [1].

IV. SEMANTICS OVERVIEW

Plans may refer to sub plans in their plan body leading to a hierarchy of plans as described in Sect. V. The behaviour of a single plan is defined in the so called plan state model: conditions are used to control selection and execution of plans. This is explained in Sect. VI. The relationship between parent and sub plans is encoded in events which



Fig. 3. Plan hierarchy

synchronize the execution of sub plans (see Sect. VII), and the concept of propagation (see Sect. VIII). Timeouts can be defined for the execution of plans (see Sect. IX). Conditions are evaluated by an underlying data abstraction unit (see Sect. X). The abstraction unit also takes care of monitoring data over a longer period of time as defined by time annotations in conditions. In our semantics intentions describe properties of plans and can be used as proof obligations for verification (see Sect. XI).

For reasons of space, we only consider part of Asbru version 7.2 (as described in [12]) within this paper. However, we claim that the major concepts of Asbru are covered. Concepts which are neglected, include

- local variables and return values,
- retrieval of aborted plans,
- complex "wait-for" constructs, and
- iterative plan execution.


V. PLAN HIERARCHY

In Asbru, plans are organized in a hierarchy as shown on the left of Figure 3: a parent plan A refers to a number of sub plans B , C , and D in its plan body. Sub plans may refer to further plans resulting in a tree hierarchy. The plan name is used to reference a plan.

One and the same plan may occur several times within this hierarchy (e.g. plan B). Therefore we distinguish between plan references and plan instances. Each reference corresponds to a unique instance. On the right of Fig. 3, plan references have been numbered to give unique instances. The first occurrence of plan B is instance B_1 , the second is instance B_2 .

A. Semantics overview

In our semantics, all existing instances of plans are executed in parallel. The hierarchy of instances is flattened, leading to one top level control. For the situation in Fig. 3, we denote this top level control with the following state-chart.



In this paper, we consider the list of plan instances fixed, i.e. no instances are created or discarded during execution. We will refer to the list of instances as $[P_1, \dots, P_n]$.

B. SOS rules

We use the configuration $\text{tlc}(\sigma)$ to define the semantics of the top level control. A single step of the top level control must adhere to the relation

$$\text{tlc}(\sigma) \rightarrow \text{tlc}(\sigma')$$

with σ and σ' being the states before and after execution.

The top level control executes all plan instances $P_1, \dots, P_n$ synchronously. In one top level step, a step of every plan is executed. The semantics of a step of plan P_i is described by the following relation (also see Sect. VI).

$$\text{psm}_{P_i}(\sigma) \rightarrow \text{psm}_{P_i}(\sigma')$$

Not all plans may be able to take a step in the current state, i.e.

$$\text{psm}_{P_i}(\sigma) \nrightarrow$$

holds. Therefore we divide the list of plans into two lists $P_{m_1}, \dots, P_{m_l}$ with plans which are able to progress, and $P_{m_{l+1}}, \dots, P_{m_n}$ which are currently blocked. If $1 \leq l$ at least some of the plans are currently able to progress. In this case a single step of the top level control is defined by the following SOS rule.

$$\frac{\begin{array}{l} \text{psm}_{P_{m_1}}(\sigma) \rightarrow \text{psm}_{P_{m_1}}(\sigma'_1) \\ \dots \\ \text{psm}_{P_{m_l}}(\sigma) \rightarrow \text{psm}_{P_{m_l}}(\sigma'_l) \\ \text{psm}_{P_{m_{l+1}}}(\sigma) \nrightarrow \\ \dots \\ \text{psm}_{P_{m_n}}(\sigma) \nrightarrow \end{array}}{\text{tlc}(\sigma) \rightarrow \text{tlc}(\sigma'_1 \cup \dots \cup \sigma'_l)}$$

(We omit details on how result states $\sigma'_1, \dots, \sigma'_n$ are united. It is sufficient to know, that each state variable is written by one plan instance only. Other plans may only read variables.)


If all plans are blocked, an environmental step is taken which is described by the following rule

$$\frac{\text{psm}_{P_1}(\sigma) \nrightarrow \dots \text{psm}_{P_n}(\sigma) \nrightarrow \text{env}(\sigma) \rightarrow \text{env}(\sigma')}{\text{tlc}(\sigma) \rightarrow \text{tlc}(\text{reset}(\sigma'))}$$

with relation $\text{env}(\sigma) \rightarrow \text{env}(\sigma')$ describing the nondeterministic update of external variables. Function `reset` resets all boolean variables corresponding to internal events to `false`.

VI. PLAN STATE MODEL

The overall plan state model defines the semantics of the different conditions of a plan. Conditions are used to decide if the plan body is applicable (selection phase) and while executing the body, if execution should be interrupted (execution phase).



- F : $[\text{Elter_tp} = \text{true}]$
 R : $[\text{Elter_tp} = \text{false}]$
 $E.1$: $[\text{activate_mode} \equiv \text{automatic}]$
 $E.2$: $\text{select}[\text{activate_mode} \equiv \text{manual}]$
 A : $[\text{abort_tp} = \text{true}]$
 C : $[\text{complete_tp} = \text{true}]$

Fig. 4. Semantics of plan state model

A. Syntax

The syntax of a plan is as follows.

$\text{plan} = \text{plan name}$
 $\quad [\text{intentions}]$
 $\quad [\text{conditions}]$
 $\quad \quad [\text{Elter-precondition } \textit{temporal-pattern}]$
 $\quad \quad [\text{activate-mode } (\text{automatic}|\text{manual})]$
 $\quad \quad [\text{complete-condition } \textit{temporal-pattern}]$
 $\quad \quad [\text{abort-condition } \textit{temporal-pattern}]$
 $\quad [\text{plan-body}]$

A plan consists of intentions (see Sect. XI), the definition of conditions (see below), and the plan body (see Sect. VII).

B. Semantics overview

A variation of the standard plan state model described in [7] is given in Fig. 4 to define the semantics of conditions. The *Plan_Control* is divided into the selection phase *Selection* and the execution phase *Execution*. Initially a plan is *Considered*. In this state, the Elter condition Elter_tp is checked. If the condition evaluates to *true*, control advances to state *Ready* (transition F). If activation is automatic, the plan is activated at once (transition E.1). If it is manual, an external event *select* is a prerequisite for activation (transition E.2). This event is controlled by the environment. If, during the selection phase, the Elter condition evaluates to *false*, the plan is *Rejected* (transition R). In state *Activated*, the sub plans of the current plan are executed. This is described in Sect. VII. The execution of sub plans can be either completed successfully (transition C) or aborted in the case of emergency patient readings (transition A). We can refer to *Terminated*, if the reason for termination – rejection, completion, or abortion – is irrelevant.

C. SOS rules

In principle, each transition of Fig. 4 corresponds to one SOS rule. The current state of a plan P is stored in variables $P.state$ in σ , and also the external event *select* corresponds to variables $P.select$.

As an example, transition F originates from the following SOS rule:

$$\frac{\llbracket P.state = \text{Considered} \rrbracket_{\sigma} \quad \text{da}(\Phi, \sigma) \rightarrow^* \text{da}(\text{true}, \sigma')}{\text{psm}_P(\sigma) \rightarrow \text{psm}_P(\sigma' [P.state / \text{Ready}])}$$

if $\text{Elter_precondition}(P) \equiv \Phi$. If P is in state *Considered* and the Elter precondition Φ is evaluated to *true* by the data abstraction unit (see Sect. X), plan P may progress to state *Ready*. Similar rules for the other transitions are defined.

Additionally, if P is currently active and the complete and abort conditions Φ and Ψ do not hold, we will execute the body of P (see Sect. VII). This is described by the following rule.

$$\frac{\llbracket P.state \in \text{Activated} \rrbracket_{\sigma} \quad \begin{array}{l} \text{da}(\neg \Phi, \sigma) \rightarrow^* \text{da}(\text{true}, \sigma') \\ \text{da}(\neg \Psi, \sigma') \rightarrow^* \text{da}(\text{true}, \sigma'') \\ \text{bdy}_P(\sigma'') \rightarrow \text{bdy}_P(\sigma''') \end{array}}{\text{psm}_P(\sigma) \rightarrow \text{psm}_P(\sigma''')}$$

VII. PLAN BODY

The hierarchy of plan instances has been flattened. How the parent plan controls the plan instances in its plan body will be explained next.

A. Syntax

The syntax of the plan body is as follows.


plan-body
 $= \text{plan-body } (\text{sequential}|\text{parallel}|\text{any-order}|\text{unordered})$
 $\quad \text{wait-for name}^*$
 $\quad (\text{plan-activation name } [\text{time-annotation}])^*$

The type of the body is either sequential, parallel, any-order or unordered. The "wait-for" construct defines mandatory and optional plans (see below), and the names of the sub plans are listed as "plan-activations" with an optional time annotation (see Sect. IX).

B. Semantics overview

Sub plans $C_1, \dots, C_n$ are controlled in the body of a plan P . Their execution can be organized differently: they can be executed *sequentially* starting with C_1 , they can be executed in parallel either with synchronization (*parallel*) or without synchronization (*unordered*) of the selection and execution phases, and finally they can be executed sequentially, but *any order*, i.e. only one sub plan is executed at once, but the sequence is not fixed.

Some of the sub plans are mandatory for the successful execution of the parent plan, others are optional. Whenever




SC : *consider*
 $E.1$: $[activate_cnd = automatic]$
 $E.2$: $select[activate_cnd = manual]$
 E : *activate*
 $[ess + ref \leq time \leq lss + ref]$
 $/$ $start_time := time$

Fig. 5. Synchronization states in plan state model

we need to distinguish between this, we will divide the list of sub plans $C_1, \dots, C_n$ into two lists $C_1^m, \dots, C_k^m$ (for the mandatory plans) and $C_1^o, \dots, C_l^o$ (for the optional ones).

In order to allow synchronization of the selection and execution phases of the sub plans, the plan state model has to be enriched with intermediate states *Inactive* and *Selected*, resulting in the adapted statechart of Fig. 5. The additional events *consider* and *activate* are used to externally control progress of a plan. A parent plan can thus synchronize the sub plans in its plan body. For this, the *Activated* state of the parent is re $\mathcal{F}$ ined with a controlling statechart.

If no restriction on the progress of sub plan C_i is required, the following controlling statechart $Sub_i\ C_i$ can be used.




$i.1$: $/C_i.consider$
 $i.2$: $[in(C_i.Selected)]/C_i.activate$

Sub plan C_i is considered immediately (transition $i.1$) and is activated as soon as it reaches state *Selected* (transition $i.2$).

The different body types may oppose restrictions on the execution of sub plans. This is done by deferring the generation of the newly added events. Controlling statecharts for the different types are explained next.


B.1 Sequential execution



$i.3$: $[in(C_i.Terminated)]$

The first sub plan is considered. As soon as it terminates, we continue with the second plan (transition 1.3). During execution of one plan, no synchronization is required. Thus, we use *Sub_Control* C_i to execute each sub plan.


B.2 Unordered execution



The sub plans are executed in parallel and no further synchronization is necessary.

Implicitly, if the "activate-mode" of sub plans is manual, arbitrary execution orders (sequential, parallel, etc.) are possible, because the *select* event is controlled by the environment (see Sect. VI).


B.3 Parallel execution



1 : $/C_1.consider; \dots; C_n.consider$
 2 : $[\bigwedge_{i=1}^k in(C_i^m.Selected)]$
 $/$ $C_1.activate; \dots; C_n.activate$

The parallel operator synchronizes selection and execution phases of all sub plans. The sub plans are considered immediately (transition 1). They may only proceed to *Activated* state, if all mandatory sub plans are *Selected* (transition 2).

B.4 Any order execution



Only the selection phases are executed in parallel. The execution phases of the sub plans are synchronized such that at most one sub plan is active at the same time. For this the plans are considered immediately (transition 1). The first plan to become selectable is activated (transition $i.2$). Only if this plan terminates (transition $i.3$) another one can be activated. If several sub plans reach state *Selected* simultaneously, the choice is nondeterministic.

C. SOS rules

Again, the transitions correspond to SOS rules. Examples are given below. The additional events are stored for each plan P as variables $P.consider$ and $P.activate$ in σ .

Transition $i.2$ in *Sub_Control* C_i :

$$\frac{\llbracket C_i.state = Selected \rrbracket_{\sigma}}{sub_{C_i}(\sigma) \rightarrow sub_{C_i}(\sigma[C_i.activate/true])}$$

Transition $i.3$ in *Sequential_Control P*:

$$\frac{\llbracket P.state \in \text{Sub_Control}_i \rrbracket_\sigma \quad \llbracket C_i.state \in \text{Terminated} \rrbracket_\sigma}{\text{seq}_P(\sigma) \rightarrow \text{seq}_P(\sigma[P.state/\text{Sub_Exec}_{i+1}])}$$

for $1 \leq i < n$.

Sequential execution uses standard control to execute each sub plan. Therefore an additional rule is necessary to embed standard control:

$$\frac{\begin{array}{l} \llbracket P.state = \text{Sub_Control}_i \rrbracket_\sigma \\ \llbracket C_i.state \notin \text{Terminated} \rrbracket_\sigma \\ \text{std}_{C_i}(\sigma) \rightarrow \text{std}_{C_i}(\sigma') \end{array}}{\text{seq}_P(\sigma) \rightarrow \text{seq}_P(\sigma')}$$

for $1 \leq i \leq n$.

Transition 1 of *Parallel_Control P*:

$$\frac{\llbracket P.state = \text{Sub_Init} \rrbracket_\sigma}{\text{par}_P(\sigma) \rightarrow \text{par}_P(\sigma[P.state/\text{Sub_Select}] \quad [C_i.consider/true] \quad \dots \quad [C_n.consider/true])}$$

The semantics of the overall body of P depends on its type. If $\text{bodytype}(P) \equiv \text{sequential}$, we receive the additional rule

$$\frac{\text{seq}_P(\sigma) \rightarrow \text{seq}_P(\sigma')}{\text{bdy}_P(\sigma) \rightarrow \text{bdy}_P(\sigma')}$$

and similar for the other types.

VIII. PROPAGATION


The hierarchy of plans has been flattened. The parent is able to control and synchronize progress of its sub plans as described in the previous section. Nevertheless additional control to propagate execution states of a sub plan to its parent and vice versa is necessary. For example, if a mandatory sub plan C_i^m aborts, then also the parent aborts. This is known as propagation in Asbru. There are a number of dependencies between sub plans and parent similar to this example. All of them are displayed as additional or refined transitions in Fig. 6. SOS rules for the added and refined transitions are straightforward.

IX. TIMEOUTS

In the example of Fig. 1, a time annotation has been used to describe the monitoring of conditions over time. This is taken care of in the data abstraction unit (see Sect. X). Additionally, time annotations can be used to define timeouts for plan execution. For example

Regular_Treatments_3 $[[_, 3 \text{ h}], [4 \text{ h}, 6 \text{ h}], [2 \text{ h}, _], \text{now}]$

would state that the plan needs to be activated within the next 3 hours. Also it should be completed within 4 to 6 hours and the duration of execution should last at least 2 hours.



$$\begin{array}{l} R.2 : [\text{in}(P.\text{Aborted})] \\ R.3 : [\text{in}(P.\text{Completed})] \end{array}$$

$$\begin{array}{l} A.2 : [\text{in}(P.\text{Aborted})] \\ A.3 : [\text{in}(P.\text{Completed})] \\ A.4 : [\bigvee_{i=1}^k \text{in}(C_i^m.\text{Rejected})] \\ A.5 : [\bigvee_{i=1}^k \text{in}(C_i^m.\text{Aborted})] \end{array}$$

$$\begin{array}{l} C : [\text{complete_tp} = \dots \\ \wedge \text{efs} + \text{ref} \leq \dots \\ \wedge \text{mindu} \leq \dots \\ \wedge \bigwedge_{i=1}^k \text{in}(C_i^m.\text{Completed}) \end{array}$$

Fig. 6. Semantics of propagation

A. Syntax

time-annotation

= time-range

[starting-shift [earliest expression] [latest expression]]
[finishing-shift [earliest expression] [latest expression]]
[duration
[minimum expression] [maximum expression]]
reference-point (expression|now)

Within a time-annotation, expressions are used to define a variety of time points. Informally, a plan must be activated within the starting shift. It must complete within the finishing shift and its duration of execution must comply with the minimum and maximum duration. Time values are relative to the given reference point. In this paper, time annotations are abbreviated as follows


$$[[\text{ess}, \text{lss}], [\text{efs}, \text{lfs}], [\text{mindu}, \text{maxdu}], \text{ref}]$$

and we will use the underscore '_' to represent unspecified values.

[5] describes a number of static checks that a time annotation must satisfy to be considered well-formed. Here, we assume that every time annotation is well-formed.

B. Semantics overview

Changes to the plan state model which capture the additional timing constraints are displayed in Figure 7. If, dur-



$R.1 : [lss + ref < time]$
 $E.1 : \begin{bmatrix} \dots \equiv \text{automatic} \\ \wedge ess + ref \leq time \leq lss + ref \\ / \text{start_time} := time \end{bmatrix}$
 $E.2 : \begin{bmatrix} \dots \equiv \text{manual} \\ \wedge ess + ref \leq time \leq lss + ref \\ / \text{start_time} := time \end{bmatrix}$
 $A.1 : \begin{bmatrix} lfs + ref < time \\ \vee maxdu < time - start_time \end{bmatrix}$
 $C : \begin{bmatrix} complete_tp = true \\ \wedge efs + ref \leq time \leq lfs + ref \\ \wedge mindu \leq time - start_time \leq maxdu \end{bmatrix}$

Fig. 7. Semantics of time annotations for plan activation

ing the selection phase, the latest starting shift has elapsed, the plan is rejected (transition $R.1$). If, during execution, the latest finishing shift has expired, the plan is aborted (transition $A.1$). For the positive case of activating and completing in time, transitions $E.1$, $E.2$, and C have been adapted. In order to distinguish between the different reasons of failure, the states *Rejected* and *Aborted* have been refined with sub states. Variable *time* is an external variable which is incremented in every step of the environment.

Additional and adapted SOS rules are straightforward.

X. DATA ABSTRACTION

Conditions are given as temporal patterns to allow monitoring of parameters over a longer period of time. Temporal patterns are evaluated in the data abstraction unit.

A. Syntax

$temporal_pattern$
 $= \text{parameter-proposition } formula \text{ time-annotation}$
 $| \text{simple-condition } formula$
 $| temporal_pattern (\vee | \wedge) temporal_pattern$
 $| \neg temporal_pattern$

A temporal-pattern is either a parameter proposition (including a time annotation), a simple condition or several patterns combined with $\wedge$, $\vee$, or $\neg$. In contrast to parameter propositions, simple conditions are not evaluated over time.

B. Semantics overview

The underlying data abstraction unit is not described in detail here and only its purpose is summarized. The semantics of the abstraction unit is not operational, but functional in nature. As input, measurements of patient parameters are taken. The type of parameters can be very different reaching from quantitative values, like bilirubin levels in the blood, to boolean values, e.g. whether the patient is male or female. Data can be provided as a continuous stream of patient readings (high frequency domain as in artificial ventilation of premature babies) or as sporadic measurements once every month (low frequency, e.g. diabetes mellitus).

The incoming data is memorized in the patient record. Quantitative values can be abstracted to qualitative values. An example has been provided in Sect. II. More important, the abstraction unit evaluates data over a longer time period, if the data is time annotated in an ASBRU plan. The example

$bilirubin_decrease < 1 \ [[4\text{ h}, -], [-, 6\text{ h}], [-, -], now]$

requires monitoring the decrease of bilirubin level over a period of at least 4 and up to 6 hours.

As output of the abstraction unit, the truth values of conditions are provided. As evaluation of a condition may take time, the result is either *true*, *false*, or yet *unknown*. A condition is *false* only, if it cannot be satisfied in the future. Otherwise, it would be considered *unknown*.

XI. INTENTIONS

Intentions describe temporal properties of plans and can be verified as described next.

A. Syntax and semantics overview

$intentions$
 $= intentions$
 $((intermediate_state | overall_state)$
 $(avoid | maintain | achieve)$
 $temporal_pattern)^*$

An intention is to either avoid, maintain, or achieve an overall or intermediate state which is described by a temporal pattern.

Intentions can be translated into temporal logic. Details are omitted here.

B. Verification

One major goal of our project is to formally verify the operational behaviour of Asbru plans against properties which are expressed as intentions. For the example in Sect. II the task would be to verify that *bilirubin* is never equal to transfusion throughout execution – which should be easy – and if the plan completes, *bilirubin* equals to Observation – which is not so obvious.

For verification we are using the interactive theorem prover KIV. We regard automatic verification not powerful enough to deal with the data involved and therefore an interactive verifier is necessary. However it would be worthwhile to define sub tasks which could be treated with model checkers. KIV already supports the verification of parallel programs against properties expressed in temporal logic. The verification strategy is to symbolically execute programs and to use induction, if necessary [2]. A similar approach shall be applied to Asbru plans.

Because of its functional nature, the tasks of the data abstraction unit can be translated directly into algebraic specifications. The difficulty is to capture the operational, state based, parallel behaviour of the Asbru plans themselves. Encoding the SOS rules representing the formal semantics directly into proof rules has been tried but turned out to be too inefficient. Hundreds of proof steps were necessary to execute one Asbru step. In part, this is because of the explicit encoding of the plan hierarchy. A more direct representation of Asbru plans with higher level proof rules is necessary. Currently we are translating Asbru plans into parallel programs preserving the hierarchy of parent and sub plans. With this representation we are able to verify the example intentions, which are translated into temporal logic. However, this translation is only possible for a subset of features and its correctness needs to be examined still. In a future step we would like to directly support Asbru syntax and design proof rules for executing Asbru. In order to be correct, these proof rules still need to adhere to the formal semantics presented here.

XII. CONCLUSION

The formal semantics of major concepts of Asbru has been explained in this paper. We are confident that the formal semantics alone will help to better understand Asbru plans and thereby improve quality of medical protocols. Furthermore the semantics is an important link between the modelling language and the representation in KIV. An overview on how we will use KIV to verify properties formally has been given. Further research on this topic will be our next major step.

We have chosen to give the formal semantics of Asbru in the form of SOS rules. However, these rules turned out to be difficult to understand. Representing rules as transitions in statecharts resulted in a more compact and intuitive picture of plan behaviour. Even if some of the technical details of the semantics are not correctly captured within these graphics, the statecharts are very suitable for discussions and language documentation.

ACKNOWLEDGEMENTS

This work was possible only thanks to long and fruitful discussions with Silvia Miksch (TU Vienna) and her group Andreas Seyfang and Robert Kosara. Also many thanks to Frank van Harmelen, Mar Marcos (VU Amsterdam), and

Annette ten Teije (Univ. Utrecht) for their assistance and good collaboration. The work has been partially funded by the European Community as project 'Protocure' (nr. IST-2001-33049).

REFERENCES

- [1] L. Aceto, W. Fokkink, and C. Verhoef. Structural operational semantics. In J. A. Bergstra, A. Ponse, and S. A. Smolka, editors, *Handbook of Process Algebra*. Elsevier, 2001.
- [2] M. Balser, C. Duelli, W. Reif, and G. Schellhorn. Verifying concurrent systems with symbolic execution. *Journal of Logic and Computation (Special Issue)*, 2002. (to appear).
- [3] M. Balser, W. Reif, G. Schellhorn, K. Stenzel, and A. Thums. Formal system development with KIV. In T. Maibaum, editor, *Fundamental Approaches to Software Engineering*, number 1783 in LNCS. Springer, 2000.
- [4] J. Bury, J. Fox, and D. Sutton. The PROforma guideline specification language: progress and prospects. In *Proceedings of the First European Workshop, Computer-based Support for Clinical Guidelines and Protocols (EWGLP 2000)*, 2000.
- [5] G. Duftschmid and S. Miksch. Knowledge-based verification of clinical guidelines by detection of anomalies. *Artificial Intelligence in Medicine*, Special Issue: Workshop Management and Clinical Guidelines in Medicine(22(1)), 2001.
- [6] S. Herbert, C. Gordon, A. Jackson-Smale, and S. Renaud. Protocols for clinical care. *Computer Methods and Programs in Biomedicine*, 48, 1995.
- [7] S. Miksch, Y. Shahar, and P. Johnson. Asbru: A task-specific, intention-based, and time-oriented language for representing skeletal plans. In E. Motta, F. v. Harmelen, C. Pierret-Golbreich, I. Filby, and N. Wijngaards, editors, *7th Workshop on Knowledge Engineering: Methods & Languages (KEML-97)*. Milton Keynes, UK, 1997.
- [8] L. Ohno-Machado, J. Gennari, S. Murphy, N. Jain, S. Tu, D. Oliver, E. Pattison-Gordon, R. Greenes, E. Shortliffe, and G. Barnett. The GuideLine Interchange Format: A model for representing guidelines. *American Medical Association*, 5, 1998.
- [9] G. D. Plotkin. A structural approach to operational semantics. Technical Report DAIMI FN-19, Aarhus University, 1981.
- [10] A. Pnueli and M. Shalev. What is in a step: On the semantics of statecharts. In *Symposium on Theoretical Aspects of Computer Software*, pages 244–264, 1991.
- [11] Protocure – Improving medical protocols by formal methods. <http://www.protocure.org>.
- [12] A. Seyfang, R. Kosara, and S. Miksch. Asbru's reference manual, asbru version 7.2, document revision 1. Technical report, Vienna University of Technology, Institute of Software Technology, 2000.