

Realising Deterministic Behavior from Multiple Non-Deterministic Behaviors

Thomas Ströder

Department of Computer Science
Aachen University of Technology
Ahorn Str. 55
D-52056, Aachen, Germany
thomas.stroeder@rwth-aachen.de

Maurice Pagnucco

National ICT Australia and ARC Centre
of Excellence in Autonomous Systems
School of Computer Science and Engineering
The University of New South Wales
Sydney, NSW, 2052, Australia
morri@cse.unsw.edu.au

Abstract

This paper considers the problem of composing or scheduling several (non-deterministic) behaviors so as to conform to a specified target behavior as well as satisfying constraints imposed by the environment in which the behaviors are to be performed. This problem has already been considered by several works in the literature and applied to areas such as web service composition, the composition of robot behaviors and co-ordination of distributed devices. We develop a sound and complete algorithm for determining such a composition which has a number of significant advantages over previous proposals: a) our algorithm is different from previous proposals which resort to dynamic logic or simulation relations, b) we realized an implementation in Java as opposed to other approaches for which there are no known implementations, c) our algorithm determines all possible schedulers at once, and d) we can use our framework to define a notion of approximation when the target behavior cannot be realized.

Building and developing re-usable modules is one of the cornerstones of computer science. Furthermore, building on previously established infrastructure has allowed us to construct elaborate and sophisticated structures from skyscrapers and aeroplanes through to the world-wide web. Developing the components that go into these structures is only part of the problem however. Once we have them in place, we need to develop methods for piecing them together so as to achieve the desired outcome.

In this paper we consider the problem of composing behaviors. This problem has already attracted some attention in the recent literature [Berardi *et al.*, 2008; Calvanese *et al.*, 2008; de Giacomo and Sardina, 2007; Sardina and de Giacomo, 2007; Sardina *et al.*, 2008; Sardina and de Giacomo, 2008; Berardi *et al.*, 2006b; 2006a] with several proposals being put forward. More precisely, we consider the problem of composing or scheduling several (non-deterministic) behaviors so as to conform to a specified (deterministic) target behavior as well as satisfying constraints imposed by the environment in which the behaviors are to be performed. These behaviors are

abstractions that can represent a variety of mechanisms such as programs, robot actions, capabilities of software agents or physical devices, etc. As such, solutions to this problem have a wide field of applicability from composing web services [Berardi *et al.*, 2008] through to co-ordinating multiple robots or software agents [de Giacomo and Sardina, 2007; Sardina and de Giacomo, 2007; Sardina *et al.*, 2008; Sardina and de Giacomo, 2008]. The closest work to this paper is that of Sardina *et al.* [2008] which proposes a regression based technique to solving this problem where we present a progression based technique here and briefly consider the possibility of approximating the target behaviour.

For example, consider an urban search and rescue setting with three types of robots. *Scout robots* can search for victims and report their location. *Diagnosis robots* can assess victims and determine whether their condition requires special transportation. If not, this robot can guide victims to safety. *Rescue robots* can carry immobile victims to safety. We will elaborate this example using our framework in this paper.

This paper provides four main contributions that improve on previous approaches to this problem: 1) we provide a sound and complete algorithm for solving the behavior composition problem which works in the way of a forward search—this is in contrast to the proposal of Sardina et al. [2008] which can be seen as a backward search; 2) we have realized an implementation of our algorithm in Java which is the first known implementation of a solution to this problem; 3) our algorithm determines all possible schedulers for the target behavior (as can the proposals in [Berardi et al., 2008; Sardina et al., 2008]); and, 4) our approach allows the definition of approximate solutions to the behavior composition problem which can be used when the target behavior cannot be realized by the available behaviors in the given environment.

1 Background

The basic components of our framework are:

Environment which provides an abstract notion of the observable effects of actions and preconditions for actions.

Behaviors are skills or capabilities and are essentially programs that entities can perform.

Target behavior is the desired behavior that we would like to carry out.

Each of these components can be specified abstractly using automata as we shall see below. We now look at these components in more detail and introduce some other concepts that we require. In doing so we use the formulation of de Giacomo and Sardina [2007] and reproduce the important definitions from their paper that we require here with one modification. We allow multiple initial states where they do not. Otherwise the definitions are largely as in their paper and interested readers are pointed to this article for further details and intuitions behind this formulation of the behavior composition problem.

We begin with a definition of the *environment*. Note that incomplete knowledge about actions as well as failure of actions are dealt with by allowing the environment and the behaviours to be non-deterministic.

Definition 1.1 ([de Giacomo and Sardina, 2007])

An environment $\mathcal{E} = (\mathcal{A}, E, I_{\mathcal{E}}, \delta_{\mathcal{E}})$ is characterized by the components:

- $\mathcal{A}$ a finite set of shared actions
- E a finite set of possible environment states
- $I_{\mathcal{E}} \subseteq E$ the set of initial states of the environment
- $\delta_{\mathcal{E}} \subseteq E \times \mathcal{A} \times E$ the transition relation among states: $\delta_{\mathcal{E}}(e, a, e')$ holds when action a performed in state e may lead the environment to successor state e' .

Example 1.1 Considering the urban search and rescue example introduced earlier. In the following example environment the set of actions $\mathcal{A} = \{\text{diagnose, report, return, search, special, transport}\}$ where each of the individual actions has the obvious interpretation given the description above. The set of states $E = \{e_1, e_2, e_3, e_4\}$, the initial state $I_{\mathcal{E}} = \{e_1\}$ and the transitions $\delta_{\mathcal{E}}$ are given in the following automaton.




Figure 1: Environment

Next we turn to *behaviors* of which there are two types: *available behaviors* that specify the capabilities of entities that are under our control when it comes to scheduling and the *target behavior* that specifies the desired functionality.

Definition 1.2 ([de Giacomo and Sardina, 2007])

Behavior $\mathcal{B} = (S, I_{\mathcal{B}}, G, \delta_{\mathcal{B}}, F)$ over an environment $\mathcal{E}$ consists of:

- S a finite set of behavior states
- $I_{\mathcal{B}} \subseteq S$ the set of initial states of the behavior

- G a set of guards, which are Boolean functions $g : E \rightarrow \{\text{true}, \text{false}\}$ for environment states E of $\mathcal{E}$
- $\delta_{\mathcal{B}} \subseteq S \times G \times \mathcal{A} \times S$ transition relation, where $\mathcal{A}$ is the set of actions of $\mathcal{E}$; the $G \times \mathcal{A}$ component is the label of the transition
- $F \subseteq S$ set of final states in which the behavior may halt.

Note that available behaviors may also be non-deterministic and so may not be fully under the scheduler's control while the target behavior is deterministic.

Example 1.2 Continuing our example, the automaton in Figure 2 specifies the target behavior that we wish to obtain. $S = \{t_1, t_2, t_3, t_4, t_5, t_6\}$, $I_{\mathcal{B}} = \{t_1\}$, $F = \{t_1\}$, transitions are represented by the arcs of the automaton.




Figure 2: Target behavior

The following definitions will assist in our development of a scheduler that will compose the available behaviors according to the environment so as to achieve the target behavior. Runs describe transitions that can be realized by a behavior in the environment.

Definition 1.3 ([de Giacomo and Sardina, 2007]) For behavior $\mathcal{B} = (S, I_{\mathcal{B}}, G, \delta_{\mathcal{B}}, F)$ and environment $\mathcal{E} = (\mathcal{A}, E, I_{\mathcal{E}}, \delta_{\mathcal{E}})$, runs of $\mathcal{B}$ on $\mathcal{E}$ are (possibly infinite) alternating sequences: $(s^0, e^0)a^1(s^1, e^1)a^2 \dots$, where $s^0 \in I_{\mathcal{B}}$ and $e^0 \in I_{\mathcal{E}}$, and for every i $(s^i, e^i)a^{i+1}(s^{i+1}, e^{i+1})$ is such that:

- there is a transition $(e^i, a^{i+1}, e^{i+1}) \in \delta_{\mathcal{E}}$
- there is a transition $(s^i, g^{i+1}, a^{i+1}, s^{i+1}) \in \delta_{\mathcal{B}}$, such that $g^{i+1}(e^i) = \text{true}$.

If the run is finite, i.e., $(s^0, e^0)a^1 \dots a^l(s^l, e^l)$, then $s^l \in F$.

Traces are legal sequences of actions according to runs.

Definition 1.4 ([de Giacomo and Sardina, 2007]) A trace is a sequence of pairs (g, a) , where $g \in G$ is a guard of $\mathcal{B}$ and $a \in \mathcal{A}$ an action, of the form $t = (g^1, a^1) \cdot (g^2, a^2) \dots$ such that there is a run $(s^0, e^0)a^1(s^1, e^1)a^2 \dots$, where $g^i(e^{i-1}) = \text{true}$ for all i . If trace $t = (g^1, a^1) \dots (g^l, a^l)$ is finite, then there is a finite run $(s^0, e^0)a^1 \dots a^l(s^l, e^l)$ with $s^l \in F$. We call l the length of t .

A system, then, consists of the available behaviors that we can compose to realize the target behavior and the environment they must satisfy.

Definition 1.5 ([de Giacomo and Sardina, 2007]) System $S = (\mathcal{B}_1, \dots, \mathcal{B}_n, \mathcal{E})$ is formed by an environment $\mathcal{E}$ and n non-deterministic available behaviors $\mathcal{B}_i$. A system configuration is a tuple $(s_1, \dots, s_n, e)$ specifying a snapshot of the system where behavior $\mathcal{B}_i$ is in state s_i

and environment $\mathcal{E}$ in state e . The system has a specific component, called the scheduler able to activate, stop, and resume behaviors at each time point.

Example 1.3 The component behaviors are given by the automata in Figure 3. The environment and target behavior are given in the previous examples.




Figure 3: Component behaviors

The behavior composition problem can now be specified. We are provided with a system $\mathcal{S} = (\mathcal{B}_1, \dots, \mathcal{B}_n, \mathcal{E})$ and a deterministic target behavior $\mathcal{B}_0$ over $\mathcal{E}$. Our task is to schedule the available behaviors so that they realize the target behavior. At each point in time the scheduler can activate one of the available behaviors so that it can perform one action. In doing so, the scheduler ensures that this action is executable in the environment. We now formalize these notions.

Definition 1.6 (Ide Giacomo and Sardina, 2007) Let $\mathcal{S} = (\mathcal{B}_1, \dots, \mathcal{B}_n, \mathcal{E})$ be a system, $\mathcal{E} = (\mathcal{A}, E, I_{\mathcal{E}}, \delta_{\mathcal{E}})$ the environment, $\mathcal{B}_i = (\mathcal{S}_i, I_{\mathcal{B}_i}, G_i, \delta_i, F_i)$ available behaviors and $\mathcal{B}_0 = (\mathcal{S}_0, \{s_0^0\}, G_0, \delta_0, F_0)$ the target behavior. A system history is an alternating sequence of system configurations and actions $h = (s_1^0, \dots, s_n^0, e^0) \cdot a^1 \cdot (s_1^1, \dots, s_n^1, e^1) \cdots (s_1^{l-1}, \dots, s_n^{l-1}, e^{l-1}) \cdot a^l \cdot (s_1^l, \dots, s_n^l, e^l)$ where:

- $s_i^0 \in I_{\mathcal{B}_i}$ for $i \in \{1, \dots, n\}$, i.e., behaviors start in an initial state
- $e^0 \in I_{\mathcal{E}}$, i.e., the environment starts in an initial state
- at each step $0 \leq k \leq l$, there is an $i \in \{1, \dots, n\}$ where $(s_i^k, g_i^{k+1}, a^{k+1}, s_i^{k+1}) \in \delta_i$ and for all $j \neq i$, $s_j^{k+1} = s_j^k$, i.e., at each step in the history only one of the behaviors— $\mathcal{B}_i$ —has made a (legal) transition
- at each step $0 \leq k \leq l$, we have $(e^k, a^{k+1}, e^{k+1}) \in \delta_{\mathcal{E}}$, i.e., the environment also makes a legal transition.

Definition 1.7 (Ide Giacomo and Sardina, 2007) A scheduler is a function $P : \mathcal{H} \times \mathcal{A} \rightarrow \{1, \dots, n, u\}$ that, for history $h \in \mathcal{H}$ ($\mathcal{H}$ is the set of all system histories as defined above) and action $a \in \mathcal{A}$, returns the behavior (via its index) scheduled to perform the action.

de Giacomo and Sardina [2007] allow the scheduler to also return the value u to specify that no behavior can perform the action after the history. We now formalize what constitutes a solution to the behavior composition problem.

Definition 1.8 (Ide Giacomo and Sardina, 2007) Let $t = (g^1, a^1) \cdot (g^2, a^2) \cdots$ be a trace of the target behavior. A scheduler program P realizes the trace t iff for all l and all system histories $h \in \mathcal{H}_{t,P}^l$ (see below) such that $g^{l+1}(e_h^l) = \text{true}$ in the last environment state e_h^l of h , we have that

$P(h, a^{l+1}) \neq u$ and $\mathcal{H}_{t,P}^{l+1}$ is non-empty. The set of system histories $\mathcal{H}_{t,P}^l$ is inductively defined:

- $\mathcal{H}_{t,P}^0 = \bigcup_{e_0 \in I_{\mathcal{E}}, s_{10} \in I_{\mathcal{B}_1}, \dots, s_{n0} \in I_{\mathcal{B}_n}} \{(s_{10}, \dots, s_{n0}, e_0)\}$
- $\mathcal{H}_{t,P}^{l+1}$ is the set of $l+1$ -length system histories of the form $h \cdot a^{l+1} \cdot (s_1^{l+1}, \dots, s_n^{l+1}, e^{l+1})$ such that:
 - $h \in \mathcal{H}_{t,P}^l$, where $(s_1^l, \dots, s_n^l, e^l)$ is the last system configuration in h
 - a^{l+1} is an action where $P(h, a^{l+1}) = i$, with $i \neq u$, i.e., the scheduler specifies action a^{l+1} at system history h to be executed in behavior $\mathcal{B}_i$
 - $(s_i^l, g, a^{l+1}, s_i') \in \delta_i$ with $g(e^l) = \text{true}$, i.e., behavior $\mathcal{B}_i$ evolves from its current state s_i^l to state s_i' wrt the (current) environment state e^l
 - $(e^l, a^{l+1}, e^{l+1}) \in \delta_{\mathcal{E}}$, i.e., the environment may evolve from current state e^l to e^{l+1}
 - $s_i^{l+1} = s_i'$ and $s_j^{l+1} = s_j^l$, for $j \neq i$, i.e., only behavior $\mathcal{B}_i$ is allowed to perform a step.

A scheduler program P realizes target behavior $\mathcal{B}_0$ if it realizes all its traces.

2 The expansion algorithm

Our algorithm for computing a scheduler program continually passes through several phases. States are *expanded*, possibly *instantiated* by matching them with existing states and those states violating the constraints imposed by the target behavior are *deleted*. The idea is to start with a state representing all possible initial configurations and expand states in a way that every configuration reachable from a configuration in the current state by performing an action according to a transition in the target behavior is represented by one of the expanded children of the current states. To avoid calculating an infinite graph we use an instance relation to ‘close’ the graph back to an already existing state. de Giacomo and Sardina [2007] show that when a target behavior can be realized it can be done so with a finite number of states therefore justifying our ability to produce finite graphs (albeit with loops). We could expand states that end up in a configuration where we are not able to perform all actions possible in the target behavior. These configurations (i.e., the states representing them) have to be deleted from our graph. This is achieved by a *marking algorithm*. Since the marking algorithm can mark states where there are instances of the marked states in the graph, expansion and marking steps have to be performed alternately until either no states with instances in the graph are marked or the initial state is marked (i.e., so no scheduler exists). We now formalize the notions we require for our algorithm.

States. A state for our expansion algorithm differs from those of behaviors and the environment and must contain a number of properties. First it represents a state in the target behavior. Second the states and the relations between them represent a set of system histories and thus a state contains a system configuration that can be seen as a predecessor step in every system history belonging to that certain state. Third we need to know which behavior should execute which action in the scheduler program. Therefore a state contains an action that is

Definition 2.1 A state for our expansion algorithm wrt a given system $S = (\mathcal{B}_1, \dots, \mathcal{B}_n, \mathcal{E})$ and target behavior $\mathcal{B}_0$ over $\mathcal{E}$ is an element $s \in \mathbb{N}_0 \times \{\text{true}, \text{false}\} \times \text{States}$ where $\text{States} = S_0 \times \{0, \dots, n\} \times \mathcal{A}' \times S'_1 \times \dots \times S'_n \times S'_\mathcal{E} \times 2^{S_1 \times \dots \times S_n \times S_\mathcal{E} \times 2^{\mathcal{A}'}}$.

Successor Actions. We are interested in being able to perform all actions possible in a current state of the target behavior. Therefore, given a behavior B over an environment $\mathcal{E}$ we define the set of *successor actions* for a state $s \in S$ in the environment state $e \in E$ as $\text{SuccessorActions}(s, e) := \{a \in \mathcal{A} \mid \exists s' \in S : \exists g \in G : \delta_B(s, g, a, s') \wedge g(e)\}$.

Instance Relation. To obtain a finite graph representing the scheduler program, we need an *instance relation* that can be used to match states with existing ones and so add edges in the graph to already existing states. Instances represent the same state in the target behavior. Given two states $s^1 = (m^1, l^1, (s_0^1, b^1, a^1, s_1^1, \dots, s_n^1, e^1, O^1))$ and $s^2 = (m^2, l^2, (s_0^2, b^2, a^2, s_1^2, \dots, s_n^2, e^2, O^2))$ wrt a given system

Complete Expansion Step Algorithm. Given a system $\mathcal{S} = (\mathcal{B}_1, \dots, \mathcal{B}_n, \mathcal{E})$ and target behavior $\mathcal{B}_0$ over $\mathcal{E}$, a finite graph $SG = (V, E)$ with $V \subseteq \mathbb{N}_0 \times \{true, false\} \times States$ and $E \subseteq V \times V$ and a set of states to expand $TODO \subseteq V$, the complete expansion step algorithm is:

Line 5 of the algorithm checks whether the state s currently under consideration for expansion has an instance. If so, line 6 adds an instance edge in a way that the instance father is not an instance child of another instance father again. Otherwise we obtain the children of the state to be expanded at line 8 and add each one to the graph SG (lines 11–13).

- $l = \text{true}$,
- if s_0 is a final state of $\mathcal{B}_0$ then $s_1, \dots, s_n$ are final states of $\mathcal{B}_1, \dots, \mathcal{B}_n$,
- if s has a parent state, then the parent s is legal and either
 - s is an instance of another state: $\exists (s, s') \in E : \text{instance}(s, s')$, or
 - s can fulfill all its obligations by its children:

$$\forall o = (s'_1, \dots, s'_n, e', A') \in O : \forall a' \in A' : \exists c = (m^c, l^c, (s_0^c, b^c, a^c, s_1^c, \dots, s_n^c, e^c, O^c)) \in V : E(s, c) \wedge l^c \wedge e' = e^c \wedge s'_1 = s_1^c \wedge \dots \wedge s'_n = s_n^c \wedge a' = a^c$$

The parent state of s is the only state s_p with $E(s_p, s) \wedge \neg \text{instance}(s_p, s)$ (only the root state has no parent state). Consequently we call s *illegal* iff s is not legal. Intuitively, an illegal state is either not achievable by performing a trace or at some point while performing a trace achieving this state, there is at least one trace with the same beginning up to that point which cannot be performed by the system anymore. We will be interested in a special set of illegal states in the graph SG , namely those illegal states with a *true* label and at least one child which is not an instance father of that state (when this child is labeled with *true*). These are exactly the states in which we cannot perform every trace of the target behaviour anymore. So we define the set $\text{Illegal}(SG) := \{v = (m, l, (s_0, b, a, s_1, \dots, s_n, e, O)) \in V \mid v \text{ is illegal}, l = \text{true}, \exists v' = (m', l', (s'_0, b', a', s'_1, \dots, s'_n, e', O')) \in V : E(v, v') \wedge \neg \text{instance}(v, (m', \text{true}, (s'_0, b', a', s'_1, \dots, s'_n, e', O')))\}$.

Furthermore we will be interested in those illegal states with a *true* label that have become illegal just because all states of which it is an instance have become illegal or simply have not been expanded yet. Thus we define the set $\text{Remainder}(SG) := \{v = (m, l, (s_0, b, a, s_1, \dots, s_n, e, O)) \in V \mid v \text{ is illegal}, l = \text{true}, \forall v' = (m', l', (s'_0, b', a', s'_1, \dots, s'_n, e', O')) \in V : E(v, v') \Rightarrow \text{instance}(v, (m', \text{true}, (s'_0, b', a', s'_1, \dots, s'_n, e', O')))\}$.

Marker Algorithm The marker algorithm simply marks all illegal states from our special set after each expansion step so that they can be subsequently deleted. Thus, given a system $S = (B_1, \dots, B_n, \mathcal{E})$ and target behavior B_0 over $\mathcal{E}$ and a finite graph $SG = (V, E)$ with $V \subseteq \mathbb{N}_0 \times \{\text{true}, \text{false}\} \times \text{States}$ and $E \subseteq V \times V$ we have:

```

1: procedure mark( $SG, S, B_0$ )
2: while  $\text{Illegal}(SG) \neq \emptyset$  do
3:   for  $s = (m, l, (s_0, b, a, s_1, \dots, s_n, e, O)) \in \text{Illegal}(SG)$  do
4:      $l := \text{false}$ 
5:   end for
6: end while

```

This algorithm simply determines the illegal states and marks each corresponding vertex in SG with the label *false*.

Complete Expansion Algorithm. For initializing our expansion algorithm we need to compute all possible start configurations. Thus, given a system $S = (B_1, \dots, B_n, \mathcal{E})$ and target behavior $B_0 = (S_0, s_0^0, G_0, \delta_{B_0}, F_0)$ over $\mathcal{E}$ we define the root state $s_{\text{root}} := (0, \text{true}, (s_0^0, 0, \text{start}, \text{null}, \dots, \text{null}, \text{null}, O_{\text{root}}))$ where $O_{\text{root}} := \{(s_1, \dots, s_n, e, \text{SuccessorActions}(s_0^0, e)) \mid s_1 \in I_{B_1}, \dots, s_n \in I_{B_n}, e \in I_{\mathcal{E}}\}$. Now the complete expansion algorithm is:

```

1: procedure expansion_algorithm( $S, B_0$ )
2:  $SG := (\{s_{\text{root}}\}, \emptyset)$ 
3:  $TODO := \{s_{\text{root}}\}$ 
4: while  $s_{\text{root}}$  is labeled with true  $\wedge TODO \neq \emptyset$  do
5:   expansion_step( $SG, S, B_0, TODO$ )
6:   mark( $SG, S, B_0$ )
7:    $TODO := \text{Remainder}(SG)$ 
8: end while
9: if  $s_{\text{root}}$  is labeled with true then

```

```

10: return  $SG$ 
11: else
12:   return “computation failed”
13: end if

```

The core of the algorithm (lines 4–8) repeatedly applies the expansion step and marker algorithms. Provided the node corresponding to the initial state in the graph SG is not labeled *false*, there is at least one scheduler capable of realizing the target behavior.

Theorem 2.1 *The algorithms presented above satisfy the following formal properties:*

1. *the expansion step algorithm and instance relations preserve the possibility of realizing all traces of the target behavior in all ways possible for the system*
2. *the marker algorithm marks as false exactly those states achievable by performing a trace, where at some point while performing this trace there is at least one trace with the same beginning up to that point which cannot be performed by the system any more*
3. *each of the three algorithms terminates.*

Histories Represented by Calculated Graph. The intuitive connection between the graph calculated by our algorithm and the schedulers it represents lies in the system histories that can be simulated by the graph. These histories correspond directly to traces of the target behavior and thus we can read schedulers realizing the target behavior from our graph by looking at the system histories we can simulate with it. The simulation of histories is done by performing the respective action with the respective behavior in every state we reach by following the edges of the graph (and doing nothing on instance edges). Formally, the finite graph $SG = (V, E)$ calculated by the expansion algorithm represents all system histories h with $h = (s_1^0, \dots, s_n^0, e^0) \cdot a^1 \cdot (s_1^1, \dots, s_n^1, e^1) \cdot \dots \cdot (s_1^{l-1}, \dots, s_n^{l-1}, e^{l-1}) \cdot a^l \cdot (s_1^l, \dots, s_n^l, e^l)$ where we have $\exists (s_1^0, \dots, s_n^0, e^0, A) \in O_{\text{root}}, s^0 = s_{\text{root}}$ and there exist $s^1 = (m^{s^1}, \text{true}, (s_0^{s^1}, b^{s^1}, a^{s^1}, s_1^{s^1}, \dots, s_n^{s^1}, e^{s^1}, O^{s^1})), \dots, s^l = (m^{s^l}, \text{true}, (s_0^{s^l}, b^{s^l}, a^{s^l}, s_1^{s^l}, \dots, s_n^{s^l}, e^{s^l}, O^{s^l})) \in V$ such that $(s_1^i, \dots, s_n^i, e^i, A^i) \in O^{s^i}, (s_1^{i-1}, \dots, s_n^{i-1}, e^{i-1}) = (s_1^{s^i}, \dots, s_n^{s^i}, e^{s^i}), a^i \in A^{s^{i-1}}$ and $((s^{i-1}, s^i) \in E \wedge a^i = a^{s^i} \wedge \forall s' \in V : \neg (E(s^i, s') \wedge \text{instance}(s^i, s')) \vee \exists s_{\text{inst}} \in V : ((s^{i-1}, s_{\text{inst}}) \in E \wedge (s_{\text{inst}}, s^i) \in E \wedge \text{instance}(s_{\text{inst}}, s^i) \wedge a^i = a_{\text{inst}}))$ for all $i \in \{1, \dots, l\}$. The set of all histories represented by SG is denoted by $\mathcal{H}(SG)$.

Represented Schedulers. A finite graph $SG = (V, E)$ calculated by the expansion algorithm represents all functions $P : \mathcal{H} \times \mathcal{A} \rightarrow \{1, \dots, n, u\}$ with $P(h, a) \in \text{Next}(h, a)$ if $h \in \mathcal{H}(SG)$, where $\text{Next}(h, a) := \{b \in \{1, \dots, n\} \mid \exists s_{\text{next}} = (m, \text{true}, (s_0, b, a, s_1, \dots, s_n, e, O)) \in V : (\text{reached}(h, SG), s_{\text{next}}) \in E\}$ if this set is not empty and $\text{Next}(h, a) := \{1, \dots, n, u\}$ otherwise for all $h \in \mathcal{H}$ and $a \in \mathcal{A}$.

Example 2.1 *Completing our example, the scheduler in Figure 4 gives the realization of the target behavior.*




Figure 4: Result

In general, more than one scheduler may be possible and the expansion algorithm calculates a graph representing all schedulers realizing the target behavior.

Theorem 2.2 *Given a system $\mathcal{S} = (\mathcal{B}_1, \dots, \mathcal{B}_n, \mathcal{E})$ and target behavior $\mathcal{B}_0$ over $\mathcal{E}$, the expansion algorithm returns a finite graph SG representing all schedulers which realize the target behavior if there exists at least one such scheduler. Otherwise it returns “computation failed”.*

Approximations. The method presented here can easily be used to calculate approximations by changing the marker step in such a way that it doesn’t mark all states where the simulation of the target behavior by the system is not possible in every case (and so would become illegal). In this way we can obtain results even in situations where there is no exact solution because the available behaviors are incapable of realizing the target behavior in the environment. Instead of not marking all states necessary for an exact solution, we can also extract information about which action must be performable in a certain configuration to realize the target behaviour. This information can be used to extend the available behaviours efficiently such that they are capable of realizing a desired target behaviour. Approximations can also be used to provide faster solutions at the expense of incompleteness. To compare the closeness of some approximations to an exact solution we can calculate the expected number of system configurations reached by simulating an arbitrary trace of the target behavior with a certain maximum length in which we cannot perform the next action of the trace. To do this we have to weight the edges of the target behavior with probabilities for the respective actions. Due to space restrictions we do not expand on this here but leave it to future work.

3 Conclusions

We have provided a sound and complete algorithm for solving the behavior composition problem. While there are existing algorithms in the literature, our algorithm provides a number of fundamental advantages. Firstly, our algorithm works in the way of a forward search (progression) in contrast to the proposal of [Sardina *et al.*, 2008] which is based on regression. Also, it does not require abstract implementations like that of [de Giacomo and Sardina, 2007] in terms of propositional dynamic logic which itself is difficult to implement.¹ Secondly, as a result of its simplicity, we implemented our

¹In fact, this is where our research started. In attempting to implement the algorithm of [de Giacomo and Sardina, 2007] we realized that there are no suitable PDL provers that would make this a straightforward task.

algorithm in Java and have tested it on a number of representative problems. It is the only existing implementation of an algorithm for the behavior composition problem that we know of. While our algorithm is EXPTIME-complete like other existing approaches, the method we have used allows us to easily define the notion of an approximate solution and provide a metric for determining how close an approximation is to the target behavior. Approximations can be used when the target behavior is not able to be realized by the available behaviors in the given environment and also to determine how far the solution falls short of the target. Furthermore, if we wish to improve on the EXPTIME-complete worst-case time complexity of our algorithm, we could adopt approximations at the expense of incomplete solutions. Finally, our algorithm computes all possible schedulers capable of realizing the target behavior (as do the proposals in [Berardi *et al.*, 2008; Sardina *et al.*, 2008]).

In future work we will more fully develop the notion of approximation. One interesting avenue for research would be to develop an anytime version of our algorithm so that approximations can be developed and incrementally refined. This should not require a large modification to our existing algorithm.

References

- [Berardi *et al.*, 2006a] D. Berardi, D. Calvanese, G. De Giacomo, and M. Mecella. Automatic web service composition: Service-tailored vs. client-tailored approaches. In *Proc. of AISC-06*, 2006.
- [Berardi *et al.*, 2006b] D. Berardi, D. Calvanese, G. De Giacomo, and M. Mecella. Composing web services with nondeterministic behavior. In *Proc. of ICWS-06*, 2006.
- [Berardi *et al.*, 2008] D. Berardi, F. Cheikh, G. de Giacomo, and F. Patrizi. Automatic service composition via simulation. *International Journal of Foundations of Computer Science*, 19(2):429–451, 2008.
- [Calvanese *et al.*, 2008] D. Calvanese, G. De Giacomo, M. Lenzerini, M. Mecella, and F. Patrizi. Automatic service composition and synthesis: The Roman model. *Bull. of the IEEE Computer Society Technical Committee on Data Engineering*, 31(3):18–22, 2008.
- [de Giacomo and Sardina, 2007] G. de Giacomo and S. Sardina. Automatic synthesis of new behaviors from a library of available behaviors. In *Proc. of IJCAI-07*, pages 1866–1871, 2007.
- [Sardina and de Giacomo, 2007] S. Sardina and G. de Giacomo. Automatic synthesis of a global behavior from multiple distributed behaviors. In *Proc. of AAI-07*, pages 1063–1069, 2007.
- [Sardina and de Giacomo, 2008] S. Sardina and G. de Giacomo. Realizing multiple autonomous agents through scheduling of shared devices. In *Proc. of ICAPS-08*, pages 304–311, 2008.
- [Sardina *et al.*, 2008] S. Sardina, F. Patrizi, and G. de Giacomo. Behavior composition in the presence of failure. In *Proc. of KR-08*, pages 640–650, 2008.