

Supremal Realizability of Behaviors with Uncontrollable Exogenous Events*

Nitin Yadav

RMIT University
Melbourne, Australia
nitin.yadav@rmit.edu.au

Paolo Felli

Sapienza Univ. di Roma
Roma, Italy
felli@dis.uniroma1.it

Giuseppe De Giacomo

Sapienza Univ. di Roma
Roma, Italy
degiamoco@dis.uniroma1.it

Sebastian Sardina

RMIT University
Melbourne, Australia
sebastian.sardina@rmit.edu.au

Abstract

The behavior composition problem involves the automatic synthesis of a controller able to “realize” (i.e., implement) a desired target behavior specification by suitably coordinating a set of already available behaviors. While the problem has been thoroughly studied, one open issue has resisted a principled solution: *if the target specification is not fully realizable, is there a way to realize it “at best”?* In this paper we answer positively, by showing that there exists a unique supremal realizable target behavior satisfying the specification. More importantly we give an effective procedure to compute such a target. Then, we introduce exogenous events, and show that the supremal can again be computed, though this time, into two variants, depending on the ability to observe such events.

1 Introduction

The *behavior composition problem* amounts to whether a set of available, though partially controllable, behavior modules (e.g., smart robotic devices such as automatic blinds and lights, audio and screen devices, video cameras) can be suitably coordinated (i.e., composed) in a way that it appears as if a desired but non-existent target behavior (e.g., a house entertainment or surveillance system) is being executed. The problem is appealing in that with computers now present in everyday devices like phones, cars and planes or places like homes, offices and factories, the trend is to build embedded complex systems from a collection of simple components. Indeed, the composition problem has been studied in various areas of computer science, including (web) services [Balbiani *et al.*, 2008], AI reasoning about action [Sardina *et al.*, 2008; Stroeder and Pagnucco, 2009; De Giacomo *et al.*, 2013], verification [Lustig and Vardi, 2009], and robotics [Bordignon *et al.*, 2007]. From an AI perspective, a behavior refers to the abstract operational model of a device or program, and is generally represented as a nondeterministic transition system.

The classical behavior composition task has been extensively investigated in the recent literature (see [De Giacomo *et*

al., 2013] for an extensive review). However, one open issue has resisted principled solutions: *if the target specification is not fully realizable, is there a way to realize it “at best”?* Stroeder and Pagnucco [2009] were the first to highlight this issue and proposed a search-based method that could eventually be adapted to compute approximate solutions “close” to the perfect one. However, they did not detail what those “approximations” look like. Then, Yadav and Sardina [2012] developed an account of “approximate” composition where the task is to return an alternative target closest to the original one, but fully solvable. While their proposal, based on the formal notion of simulation, comes as a principled generalization of the classical framework, it did not provide ways to actually compute such solutions for the general case, but only for the special case of deterministic behaviors.

In this paper, we present a novel technique to *effectively build* the largest realizable fragment—the “*supremal*”—of a given target specification for the general composition case in which available behaviors may be nondeterministic. The technique relies on two simple and well-known operations over transition systems (or state models), namely, cross product and belief-level state construction. In doing so, we provide an elegant result on the uniqueness of such fragments.

Then, we investigate—inspired by work on AI reasoning about action [Reiter, 2001] and on discrete event systems [Cassandras and Lafortune, 2006]—the composition task in the presence of *exogenous events*. These are special events that behaviors may *spontaneously generate*, such as the light bulb of a projector fusing when turned on. Importantly, such events are *uncontrollable* and their occurrence cannot be disabled. As a result, we obtain a strictly more general composition framework. We demonstrate that the supremal realizable target can again be defined and computed. However, this time, solutions come into two variants, depending on the ability of the target’s user to observe such events. If exogenous events can be observed by the user, then the supremal fragment may be *conditional* on such events (e.g., if the projector’s light bulb fuses, the user may only request changing the bulb). Otherwise, the supremal ought to be *comformant* to all possible exogenous events that may ensue.

2 Preliminaries

We briefly review the classical composition framework [Sardina *et al.*, 2008; Stroeder and Pagnucco, 2009; De Giacomo

*The last author acknowledges the support of the Australian Research Council under grant DP120100332.

et al., 2013]. We omit wlog the environment for simplicity.

Behavior A behavior represents the operational logic of a device or a program and it is modeled using a finite transition system. Formally, a *behavior* is a tuple $\mathcal{B} = \langle B, \mathcal{A}, b_0, \varrho \rangle$, where:

- B is the finite set of behavior’s states;
- \mathcal{A} is the set of actions;
- $b_0 \in B$ is the initial state;
- $\varrho \subseteq B \times \mathcal{A} \times B$ is the behavior’s transition relation, where $\langle b, a, b' \rangle \in \varrho$, or $b \xrightarrow{a} b'$ in \mathcal{B} , denotes that action a executed in behavior state b may lead the behavior to successor state b' .

Observe that since behaviors may be nondeterministic, one cannot know beforehand what actions will be available to execute after an action is performed in a state, as the next set of applicable actions would depend on the successor state in which the behavior happens to be in. Hence, we say that nondeterministic behaviors are only *partially controllable*. A *deterministic* behavior is one where the successor state is always uniquely determined—a *fully controllable* behavior.

Available System The system stands for a collection of behaviors that are at disposal. Technically, an *available system* is a tuple $\mathcal{S} = \langle \mathcal{B}_1, \dots, \mathcal{B}_n \rangle$, where $\mathcal{B}_i = \langle B_i, \mathcal{A}_i, b_{i0}, \varrho_i \rangle$, for $i \in \{1, \dots, n\}$, is a, possibly nondeterministic, behavior, called an *available behavior* in the system. To refer to the behavior that emerges from the joint execution of available behaviors, the notion *enacted system behavior* is used in the literature [De Giacomo et al., 2013]. The *enacted system behavior* of a system \mathcal{S} is a tuple $\mathcal{E}_\mathcal{S} = \langle S, \mathcal{A}_\mathcal{S}, s_0, \delta \rangle$, where:

- $S = B_1 \times \dots \times B_n$ is the finite set of $\mathcal{E}_\mathcal{S}$ ’s states; when $s = \langle b_1, \dots, b_n \rangle$, we denote b_i by $\text{beh}_i(s)$, for $i \leq n$;
- $\mathcal{A}_\mathcal{S} = \bigcup_{i=1}^n \mathcal{A}_i$ is the set of actions of $\mathcal{E}_\mathcal{S}$;
- $s_0 = \langle b_{10}, \dots, b_{n0} \rangle \in S$ is $\mathcal{E}_\mathcal{S}$ ’s initial state;
- $\delta \subseteq S \times \mathcal{A}_\mathcal{S} \times \{1, \dots, n\} \times S$ is $\mathcal{E}_\mathcal{S}$ ’s transition relation, where $\langle s, a, k, s' \rangle \in \delta$, or $s \xrightarrow{a,k} s'$ in $\mathcal{E}_\mathcal{S}$, iff:
 - $\text{beh}_k(s) \xrightarrow{a} \text{beh}_k(s')$ in \mathcal{B}_k ; and
 - $\text{beh}_i(s) = \text{beh}_i(s')$, for $i \in \{1, \dots, n\} \setminus \{k\}$.

Target Specification A *target behavior specification* $\mathcal{T} = \langle T, \mathcal{A}_T, t_0, \varrho_T \rangle$ is a behavior specification that stands for the desired, though not directly available, functionality to be obtained. Following in particular [Yadav and Sardina, 2012], the idea is that the user agent is meant to request transitions in \mathcal{T} in a step-by-step fashion, and the action in the chosen transition is suitably delegated by a controller to one of the available behavior of the system.

So, informally, the behavior composition task can be stated as follows: Given a system \mathcal{S} and a target behavior \mathcal{T} , is it possible to (partially) control the available behaviors in \mathcal{S} in a step-by-step manner—by instructing them on which action to execute next and observing, afterwards, the outcome in the

behavior used—so as to “realize” the desired target behavior. In other words, by adequately controlling the system, it appears as if one was actually executing the target behavior.

Consider the presentation room scenario depicted in Figure 1, ignoring all dashed transitions. There are two available behaviors, a projector $\mathcal{B}_\mathcal{P}$ and a speaker system $\mathcal{B}_\mathcal{A}$. The projector allows setting of the SOURCE and WARMUP of the device in any order, followed by turning it OFF. The speaker on the other hand can simply be toggled on/off. The question then is whether these two devices are enough to be able to run the desired target behavior \mathcal{T} . The answer, in this case, is *yes*.

Exact Compositions via Simulation Though technically involved, one can formally define when a so-called *controller*, a function taking a run of the system and the next action request and outputting the index of the available behavior where the action is being delegated, realizes the target behavior; see [De Giacomo and Sardina, 2007; De Giacomo et al., 2013]. Such controllers are called *exact compositions*, solutions to the composition problem guaranteeing the complete realization of the target in the system. In our presentation room example, there exists an exact composition for target behavior \mathcal{T} in available system $\langle \mathcal{B}_\mathcal{P}, \mathcal{B}_\mathcal{A} \rangle$.

An interesting and much used result links exact compositions to the formal notion of simulation [Milner, 1971]. A simulation relation captures the behavioral equivalence of two transition systems. Intuitively, a (transition) system S_1 “simulates” another system S_2 , denoted $S_2 \preceq S_1$, if S_1 is able to *match* all of S_2 ’s moves. Thus, Sardina et al. [2008] defined a so-called *ND-simulation* (nondeterministic simulation) relation between (the states of) the target behavior \mathcal{T} and (the states of) the enacted system $\mathcal{E}_\mathcal{S}$, denoted \preceq_{ND} , and prove that there exists an exact composition for a target behavior \mathcal{T} on an available system \mathcal{S} iff $\mathcal{T} \preceq_{\text{ND}} \mathcal{E}_\mathcal{S}$, that is, the enacted system can ND-simulate the target behavior. While in this paper we do not really need the details of ND-simulation, the plain notion of simulation plays a key role, so let us introduce it formally. Let $\mathcal{T}_i = \langle S_i, \mathcal{A}, s_{i0}, \varrho_i \rangle$, where $i \in \{1, 2\}$, be two transition systems. A *simulation relation* of \mathcal{T}_2 by \mathcal{T}_1 is a relation $\text{Sim} \subseteq S_2 \times S_1$ such that $\langle s_2, s_1 \rangle \in \text{Sim}$ iff: $\forall a, s'_2. \langle s_2, a, s'_2 \rangle \in \varrho_2 \Rightarrow \exists s'_1 \langle s_1, a, s'_1 \rangle \in \varrho_1 \wedge \langle s'_2, s'_1 \rangle \in \text{Sim}$. We say that a state $s_2 \in S_2$ is *simulated* by a state $s_1 \in S_1$ (or s_1 simulates s_2), denoted $s_2 \preceq s_1$, iff there exists a simulation relation Sim of \mathcal{T}_2 by \mathcal{T}_1 such that $\langle s_2, s_1 \rangle \in \text{Sim}$. Observe that relation \preceq is itself a simulation relation (of \mathcal{T}_2 by \mathcal{T}_1), and in fact, it is the largest simulation relation, in that all simulation relations are contained in it. We say that a transition system \mathcal{T}_1 *simulates* another transition system \mathcal{T}_2 , denoted $\mathcal{T}_2 \preceq \mathcal{T}_1$, if it is the case that $s_{20} \preceq s_{10}$. Two transition systems are said to be *simulation equivalent*, denoted $\mathcal{T}_1 \sim \mathcal{T}_2$, whenever they simulate each other.

Approximated Compositions The classical composition task described above has been extensively studied in the literature and various extensions—e.g., distributed and multi-target composition, composition under uncertainty—have been developed (see [De Giacomo et al., 2013] for a comprehensive review). However, such frameworks may prove

insufficient for composition instances admitting no exact solutions (i.e., unsolvable instances)—a mere “no solution” answer may be highly unsatisfactory in many settings.

The first one to concretely deal with this issue was Stroeder and Pagnucco [2009]. In their work, they claimed that their search-based method “can easily be used to calculate approximations,” that is, controllers that may not qualify as exact solutions but come “close” (enough) to them. They argue approximations are useful when no exact solution exists and when one is willing to trade faster solutions at the expense of incompleteness (of target realizability). Nonetheless, the authors did not provide a semantics of what these “approximations” are and what “closeness” means, both were left as important future work.

Later, Yadav and Sardina [2012] looked closer at a composition framework that can better accommodate unsolvable instances. In doing so, however, they proposed to focus on approximating the target, rather than the controller. To that end, they defined, based on the notion of simulation, what they called *target approximations*, namely, alternative target behaviors that are “contained” in the original target while enjoying exact composition solutions. In turn, they defined the *optimal target approximation* as that one which is “closest” possible to the original target (and that is fully realized by some controller). In fact, they showed that such an optimal target is unique. They also provided a technique to compute such a solution, but only for the special case of deterministic behaviors. Here, we will provide a general technique as well as the complexity characterization of the problem.

3 Supremal Realizable Target Behavior

We adopt the approach of Yadav and Sardina [2012] to define the notion of realizable target from a target specification. We do not call it “approximation” in light of the extension with exogenous events that we study later. Indeed, once exogenous events are mentioned in the target specification, such a specification is not directly a target behavior anymore.

Formally, we say that behavior \tilde{T} is a *realizable target behavior* (RTB) of target specification \mathcal{T} in system \mathcal{S} iff

1. $\tilde{T} \preceq \mathcal{T}$ (that is, \tilde{T} is “contained” in \mathcal{T});
2. $\tilde{T} \preceq_{\text{ND}} \mathcal{E}_{\mathcal{S}}$, i.e., there is an exact composition for \tilde{T} on \mathcal{S} (that is, it is fully realizable).

Notice that we elected “simulation” as the measure for comparing target behaviors. In particular if $\mathcal{T}_1 \preceq \mathcal{T}_2$ this means that an agent can mimic the behavior \mathcal{T}_1 by suitably choosing the transitions to traverse in \mathcal{T}_2 . If \mathcal{T}_1 and \mathcal{T}_2 are simulation equivalent (i.e., $\mathcal{T}_1 \preceq \mathcal{T}_2$ and $\mathcal{T}_2 \preceq \mathcal{T}_1$) then the agent can mimic exactly one behavior using the other one, hence from the point of view of the agent the two behaviors are identical.

A behavior \tilde{T} is “the” *supremal realizable target behavior* (SRTB) of target \mathcal{T} on system \mathcal{S} iff \tilde{T} is a RTB of \mathcal{T} in \mathcal{S} and there is no RTB \tilde{T}' such that $\tilde{T} \prec \tilde{T}'$, that is, \tilde{T} is the largest realizable fragment of \mathcal{T} . It can be shown that the SRTB is unique up to simulation equivalence.

We provide a simple and elegant characterization of SRTB as follows. Let $\mathcal{T}_1 \cup \mathcal{T}_2 = \langle T, \mathcal{A}, t_{10}, \varrho \rangle$, where $\mathcal{T}_i = \langle T_i, \mathcal{A}_i, t_{i0}, \varrho_i \rangle$ have disjoint states, be the resulting unified

(target) behavior where \mathcal{T}_2 's initial state is merged with \mathcal{T}_1 's: (i) $\tilde{T} = T_1 \cup (T_2 \setminus \{t_{20}\})$; (ii) $\mathcal{A} = \mathcal{A}_1 \cup \mathcal{A}_2$; and $\varrho = \varrho_1 \cup \varrho_2|_{t_{20}}^{t_{10}}$ ($\varrho|_t^{t'}$ is relation ϱ with all states t replaced with t').

Theorem 1 *Let \tilde{T}_1 and \tilde{T}_2 be two RTB for target specification \mathcal{T} in system \mathcal{S} . Then $\tilde{T}_1 \cup \tilde{T}_2$ is an RTB for \mathcal{T} in \mathcal{S} too.*

In words, RTBs are closed under union. With this in mind, it is not hard to see that one can build the largest RTB—the supremal—by taking the union of all realizable targets.

Theorem 2 *Let \mathcal{S} be a system and \mathcal{T} be a target. Then the SRTB \mathcal{T}^* of \mathcal{T} in \mathcal{S} is: $\mathcal{T}^* = \bigcup_{\tilde{T} \text{ is a RTB of } \mathcal{T} \text{ in } \mathcal{S}} \tilde{T}$.*

Notice that any \tilde{T} which is simulation equivalent to \mathcal{T}^* is also “the” SRTB (we focus on semantics not syntax here).

Obviously, it remains to be seen if the SRTB can actually be computed and represented finitely. This is what we do next. Our technique to synthesize the SRTB relies on two simple operations on transition systems, namely, a specific synchronous product and a *conformance* enforcing procedure. Roughly speaking, the technique is as follows:

1. We take the *synchronous product* of the enacted system $\mathcal{E}_{\mathcal{S}}$ and the target \mathcal{T} , yielding the structure $\mathcal{F}_{\langle \mathcal{S}, \mathcal{T} \rangle}$.
2. We modify $\mathcal{F}_{\langle \mathcal{S}, \mathcal{T} \rangle}$ to *enforce conformance* on its states which cannot be distinguished by the user of the target.

In fact the second step is needed only when the system includes *nondeterministic available behaviors*.

Full enacted system The full enacted system models the behavior that emerges from joint parallel execution of the enacted system and the target. Formally, given the enacted system $\mathcal{E}_{\mathcal{S}} = \langle \mathcal{S}, \mathcal{A}_{\mathcal{S}}, s_0, \delta \rangle$ for a system $\mathcal{S} = \langle \mathcal{B}_1, \dots, \mathcal{B}_n \rangle$ and a target specification $\mathcal{T} = \langle T, \mathcal{A}_T, t_0, \varrho_T \rangle$, the *full enacted system* of \mathcal{T} and \mathcal{S} , denoted by $\mathcal{T} \times \mathcal{E}_{\mathcal{S}}$, is a tuple $\mathcal{F}_{\langle \mathcal{S}, \mathcal{T} \rangle} = \langle F, \mathcal{A}_{\mathcal{F}}, f_0, \gamma \rangle$, where:

- $F = \mathcal{S} \times T$ is the finite set of $\mathcal{F}_{\langle \mathcal{S}, \mathcal{T} \rangle}$'s states; when $f = \langle s, t \rangle$, we denote s by $\text{sys}(f)$ and t by $\text{tgt}(f)$;
- $f_0 = \langle s_0, t_0 \rangle \in F$, is $\mathcal{F}_{\langle \mathcal{S}, \mathcal{T} \rangle}$'s initial state;
- $\mathcal{A}_{\mathcal{F}} = \mathcal{A}_{\mathcal{S}} \cup \mathcal{A}_T$ (note that we allow for $\mathcal{A}_{\mathcal{S}} \neq \mathcal{A}_T$);
- $\gamma \subseteq F \times \mathcal{A}_{\mathcal{F}} \times \{1, \dots, n\} \times F$ is $\mathcal{F}_{\langle \mathcal{S}, \mathcal{T} \rangle}$'s transition relation, where $\langle f, a, k, f' \rangle \in \gamma$, or $f \xrightarrow{a,k} f'$ in $\mathcal{F}_{\langle \mathcal{S}, \mathcal{T} \rangle}$ iff
 - $\text{tgt}(f) \xrightarrow{a} \text{tgt}(f')$ in \mathcal{T} ; and
 - $\text{sys}(f) \xrightarrow{a,k} \text{sys}(f')$ in $\mathcal{E}_{\mathcal{S}}$.

Observe that the transition relation of the full enacted system requires both the enacted system and the target to evolve *jointly*: the full enacted system is the *synchronous product* of the target specification and the enacted system.

As expected, the synchronous product (once we project out the indexes $\{1, \dots, n\}$) is simulated by both the enacted system and the target (i.e., $\mathcal{T} \times \mathcal{E}_{\mathcal{S}} \preceq \mathcal{E}_{\mathcal{S}}$ and $\mathcal{T} \times \mathcal{E}_{\mathcal{S}} \preceq \mathcal{T}$). If the system includes only deterministic available behavior the regular simulation $\mathcal{T} \times \mathcal{E}_{\mathcal{S}} \preceq \mathcal{E}_{\mathcal{S}}$ suffices to conclude that the composition exists (ND-simulation is not needed in this case) [Sardina *et al.*, 2008]. Hence, by Theorem 2, if available behaviours are deterministic $\mathcal{T} \times \mathcal{E}_{\mathcal{S}}$ is included in, and

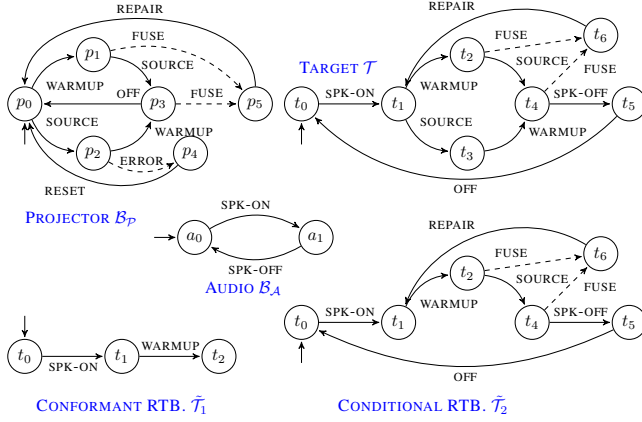


Figure 1: Media room scenario consisting of a projector, speaker and a target specification (see text for details). Dashed transitions denote uncontrollable exogenous events.

simulation by, \mathcal{T}^* . The converse can be shown along the line suggested in [Yadav and Sardina, 2012]. Hence:

Theorem 3 Let $\mathcal{S} = \langle \mathcal{B}_1, \dots, \mathcal{B}_n \rangle$ be a deterministic available system and $\mathcal{T} = \langle T, \mathcal{A}_T, t_0, \varrho_T \rangle$ a target specification behavior. Then, $\mathcal{F}_{\langle \mathcal{S}, \mathcal{T} \rangle}$ is the SRTB of \mathcal{T} in \mathcal{S} .

Also from the construction of $\mathcal{F}_{\langle \mathcal{S}, \mathcal{T} \rangle}$ we can conclude that building the SRTB of \mathcal{T} in \mathcal{S} can be done in exponential time in the number of behaviors and polynomial in the number of states in each behavior.

When we consider nondeterministic available modules, and hence resort to ND-simulation, this is not true anymore. Indeed, there are examples where $\mathcal{T} \times \mathcal{E}_S \preceq_{\text{ND}} \mathcal{E}_S$ does not hold due to the nondeterminism present in \mathcal{E}_S . In those cases, the full enacted system is a sort of target behavior in which agent transition requests are *conditional* on the nondeterministic execution of available behaviors. However, the agent using the target is not meant to have observability on such contingencies. Figure 2 shows one such case. Take product $\mathcal{T} \times \mathcal{E}_S$ as a candidate for SRTB. After fulfilling transition request $q_0 \xrightarrow{a} q_2$ using module \mathcal{B}_1 , the next request $q_2 \xrightarrow{c} q_0$ can only be honored if \mathcal{B}_1 happens to evolve to state b_2 , but this is not guaranteed. Therefore, $\mathcal{T} \times \mathcal{E}_S$ cannot be realized by \mathcal{B}_1 and hence it is not an RTB of \mathcal{T} in \mathcal{S} .

What we need, is the target to be *conformant*, i.e., independent of conditions on the available behaviors states. Hence inspired by the literature on planning under uncertainty we construct a sort of belief states, and in turn, the *belief level full enacted system*. The idea behind generating the belief states is to track the states where the enacted system could evolve. Given a full enacted system $\mathcal{F}_{\langle \mathcal{S}, \mathcal{T} \rangle} = \langle F, \mathcal{A}_F, f_0, \gamma \rangle$ for a target $\mathcal{T} = \langle T, \mathcal{A}_T, t_0, \varrho_T \rangle$ and a system $\mathcal{S} = \langle \mathcal{B}_1, \dots, \mathcal{B}_n \rangle$ where $\mathcal{B}_i = \langle B_i, \mathcal{A}_i, b_{i0}, \varrho_i \rangle$ for $i \leq n$, the *belief-level full enacted system* is a tuple $\mathcal{K}_{\langle \mathcal{S}, \mathcal{T} \rangle} = \langle Q, \mathcal{A}_K, q_0, \delta_K \rangle$, where:

- $Q = 2^{(B_1 \times \dots \times B_n)} \times T$ is $\mathcal{K}_{\langle \mathcal{S}, \mathcal{T} \rangle}$'s set of states; when $q = \{\{s_1, \dots, s_\ell\}, t\} \in Q$ we denote $\{s_1, \dots, s_\ell\}$ by $\text{sys}(q)$ and t by $\text{tgt}(q)$;




Figure 2: Instance where full enacted system is not a RTB

- $q_0 = \langle \{s_0\}, t_0 \rangle$ such that $f_0 = \langle s_0, t_0 \rangle$, is the initial state;
- $\delta_K \subseteq Q \times \mathcal{A} \times Q$ is $\mathcal{K}_{\langle \mathcal{S}, \mathcal{T} \rangle}$'s transition relation, where $\langle \langle S, t \rangle, a, \langle S', t' \rangle \rangle \in \delta_K$ iff:
 - there exists a set $\text{Indx} = \{\langle s_1 : k_1 \rangle, \dots, \langle s_\ell : k_\ell \rangle\}$ such that $\{s_1, \dots, s_\ell\} = S$; and $\langle s_i, t \rangle \xrightarrow{a, k_i} \langle s', t' \rangle$ in $\mathcal{F}_{\langle \mathcal{S}, \mathcal{T} \rangle}$ for all $i \leq \ell$; that is, the action a must be executable from all enacted system states in S ; and
 - $S' = \bigcup_{\langle s_i, t \rangle \in \text{Indx}} \{s' \mid \langle \langle s, t \rangle, a, i, \langle s', t' \rangle \rangle \in \gamma\}$; that is, S' should contain all successors of enacted system states in S resulting from action a .

Observe, $\mathcal{K}_{\langle \mathcal{S}, \mathcal{T} \rangle}$ is nondeterministic with respect to target evolutions and different behavior delegations. Note also that $\mathcal{K}_{\langle \mathcal{S}, \mathcal{T} \rangle}$ can be built in time $2^{O(|\mathcal{B}|^n)}$ where $|\mathcal{B}|$ is the number of states of the largest behavior in \mathcal{S} , and n is the number of available behaviors in \mathcal{S} . Observe, however, that $\mathcal{K}_{\langle \mathcal{S}, \mathcal{T} \rangle}$ can be computed *on-the-fly* in a step-wise fashion: given the current belief state q we can generate the next possible states without looking at any other state in Q .

Next, we show that $\mathcal{K}_{\langle \mathcal{S}, \mathcal{T} \rangle}$ is the finite representation of SRTB \mathcal{T}^* of target \mathcal{T} in system \mathcal{S} (see Theorem 2).

Theorem 4 Let \mathcal{S} be an available system and \mathcal{T} a target specification behavior. Then, $\mathcal{K}_{\langle \mathcal{S}, \mathcal{T} \rangle}$ is the SRTB of \mathcal{T} in \mathcal{S} .

PROOF (SKETCH). Let \mathcal{T}^* be the SRTB of \mathcal{T} in \mathcal{S} . Suppose \mathcal{T}^* is not simulated by $\mathcal{K}_{\langle \mathcal{S}, \mathcal{T} \rangle}$. Therefore, there exists a finite trace τ' of a RTB \mathcal{T}' contained in \mathcal{T}^* whose last transition cannot be simulated by any trace of $\mathcal{K}_{\langle \mathcal{S}, \mathcal{T} \rangle}$. Observe, there exists a trace τ of \mathcal{T} such that τ simulates τ' . Moreover, there exists a set of traces Γ in \mathcal{S} , induced as a result of realizing τ' . Using τ and Γ we construct a trace τ_K of $\mathcal{K}_{\langle \mathcal{S}, \mathcal{T} \rangle}$ such that, each state $\tau_K[i]$ comprises of $\tau[i]$ and set of states $\Gamma[i]$, where $x[i]$ denotes the i -th state of trace x , and $\Gamma[i]$ is the union of i^{th} state of each trace in Γ . Such a trace τ_K is a legal trace of $\mathcal{K}_{\langle \mathcal{S}, \mathcal{T} \rangle}$ and simulates τ' , thereby, contradicting our assumption. Since, by construction, $\mathcal{K}_{\langle \mathcal{S}, \mathcal{T} \rangle}$ is a RTB of \mathcal{T} in \mathcal{S} , by definition it is contained in \mathcal{T}^* . Therefore, \mathcal{T}^* and $\mathcal{K}_{\langle \mathcal{S}, \mathcal{T} \rangle}$ are simulation equivalent. \square

We note some similarities in the use of belief-level behaviors with the work in [De Giacomo *et al.*, 2009] for composition under partial observability of the available behaviors. There the *controller* required to be conformant, here instead the *target behavior* must be so.

4 Composition with Exogenous Events

With an effective technique to synthesize the supremal realizable target at hand, we now turn to the second contribution

of this paper. Inspired by discrete event systems [Cassandras and Lafortune, 2006] and reasoning about action work for dynamic systems [Reiter, 2001], we show here how to accommodate *exogenous uncontrollable events* into the composition framework in a parsimonious manner. In doing so, it will come clear how robust and elaboration tolerant the definition of SRTBs and the technique to compute them are.

Let us return to our presentation room example in Figure 1. Suppose that when the projector’s light bulb is on—after WARMUP has been executed—it may FUSE anytime and requires the device to be repaired. Similarly, if a source is set before warming up the projector, an ERROR may occur and the projector will need to be RESET. The occurrence of both events—FUSE and ERROR—is outside the control of the client or the controller, they occur *spontaneously*. Hence, they are akin to *exogenous events* in reasoning about action literature [Reiter, 2001] and *uncontrollable events* in discrete event systems [Cassandras and Lafortune, 2006].

Next, we extend the classical composition framework from Section 2 with *exogenous events*. To that end, we assume that the set of actions \mathcal{A} in a behavior is partitioned into *domain* (\mathcal{A}^C) and *exogenous* (\mathcal{A}^U) events, that is, $\mathcal{A} = \mathcal{A}^C \cup \mathcal{A}^U$ and $\mathcal{A}^C \cap \mathcal{A}^U = \emptyset$. Furthermore, as standard in discrete event systems, we assume exogenous events to be deterministic.¹

We note that exogenous events play an inherently different role in available behaviors than nondeterminism. Exogenous (uncontrollable) events may happen *anytime* from a relevant state (e.g., p_1 in \mathcal{B}_P), which allows modeling of concepts such as delayed uncertainty. Moreover, whereas nondeterminism is *not* observable to the target’s user (in fact, the user agent is not even aware of the internal logic of available behaviors), exogenous events may be. Hence, the user of the projector room may be able to observe the light bulb fusing.

When it comes to the target specification, exogenous event transitions represent those transitions that are accounted—accepted—by the target but outside the controller of the user of the target. Thus, when the target is in state t_2 , it only allows one exogenous event, namely, event FUSE, whose occurrence will cause the target to evolve to state t_6 where its user is only allowed to request repairing the projector.

Since the user may be able to observe exogenous events, we can now consider—unlike standard composition—two types of composition solutions. Following planning terminology, a *conditional* SRTB is one that assumes the user is able to observe exogenous events, whereas a *conformant* SRTB is one where such events are non-observable to the user.

In this section, we formally define conditional and conformant solution concepts and explain how to generalize the technique developed in Section 3 to compute such solutions.

Enacted and full enacted system The formal definitions of the enacted system and the full enacted system remain same, except we assume the action set to be partitioned into controllable actions and uncontrollable exogenous events. In addition, given a full enacted system $\mathcal{F}_{\langle S, \mathcal{T} \rangle} = \langle F, \mathcal{A}_{\mathcal{F}}^C \cup \mathcal{A}_{\mathcal{F}}^U, f_0, \gamma \rangle$ for an enacted system $\mathcal{E}_S = \langle S, \mathcal{A}_S^C \cup \mathcal{A}_S^U, s_0, \delta \rangle$

¹Should this not be the case, we can model the various outcomes with different uncontrollable exogenous events.

and a target specification $\mathcal{T} = \langle T, \mathcal{A}_T^C \cup \mathcal{A}_T^U, t_0, \varrho_T \rangle$, we define set $\Delta_{\langle S, \mathcal{T} \rangle}$ as those states in $\mathcal{F}_{\langle S, \mathcal{T} \rangle}$ from where prohibited exogenous events may fire. Formally,

$$\Delta_{\langle S, \mathcal{T} \rangle} = \{ \langle s, t \rangle \mid \langle s, \alpha, k, s' \rangle \in \delta, \forall t' \langle t, \alpha, t' \rangle \notin \varrho_T : \alpha \in \mathcal{A}_S^U \}.$$

4.1 Conditional SRTBs

When it comes to formally defining conditional SRTBs, interestingly, the definition of SRTBs from the classical framework (see Section 3) fits as is. However, we need to define exact solutions in the context of exogenous events. We do this by extending the ND-simulation relation in the light of exogenous events.

A transition system $\tilde{\mathcal{T}} = \langle \tilde{T}, \tilde{\mathcal{A}}_{\tilde{\mathcal{T}}}^C \cup \tilde{\mathcal{A}}_{\tilde{\mathcal{T}}}^U, \tilde{t}_0, \tilde{\varrho}_{\tilde{\mathcal{T}}} \rangle$ is a *conditional-RTB* for a target $\mathcal{T} = \langle T, \mathcal{A}_T^C \cup \mathcal{A}_T^U, t_0, \varrho_T \rangle$ in system \mathcal{S} with enacted system $\mathcal{E}_S = \langle S, \mathcal{A}_S^C \cup \mathcal{A}_S^U, s_0, \delta \rangle$ iff $\tilde{\mathcal{T}} \preceq \mathcal{T}$ and $\langle \tilde{t}_0, s_0 \rangle \in \mathcal{C}$ where $\mathcal{C} \subseteq \tilde{T} \times S$ is the *conditional simulation relation* of $\tilde{\mathcal{T}}$ by \mathcal{E}_S such that $\langle \tilde{t}, s \rangle \in \mathcal{C}$ iff:

1. $\forall \tilde{t}' \forall a \in \tilde{\mathcal{A}}_{\tilde{\mathcal{T}}}^C \exists k \forall s' (\langle \tilde{t}', a, \tilde{t}' \rangle \in \tilde{\varrho}_{\tilde{\mathcal{T}}} \Rightarrow \langle s, a, k, s' \rangle \in \delta)$ such that $\langle \tilde{t}', s' \rangle \in \mathcal{C}$; and
2. $\forall \alpha \in \mathcal{A}_S^U, \forall k (\langle s, \alpha, k, s' \rangle \in \delta \Rightarrow \langle \tilde{t}, \alpha, \tilde{t}' \rangle \in \tilde{\varrho}_{\tilde{\mathcal{T}}})$ such that $\langle \tilde{t}', s' \rangle \in \mathcal{C}$.

The first condition (analogous to ND-simulation) requires all controllable actions of the RTB to be *feasible*. The second defines how uncontrollable exogenous events should be treated: since they are uncontrollable, their occurrences must be allowed in the target. If we want to prevent the occurrence of some exogenous event this can only be done by cutting some controllable action ahead of exogenous event’s possible occurrence. This is related to the notion of *controllability* in discrete event systems [Wonham and Ramadge, 1987].

As usual, a conditional RTB is *supremal* iff it is not strictly simulated by any other conditional RTB. Consider our media room example (Figure 1), $\tilde{\mathcal{T}}_2$ is conditioned on FUSE and prohibits ERROR. Indeed, while realizing $\tilde{\mathcal{T}}_2$, it is guaranteed that ERROR will never occur.

Computing conditional SRTBs When it comes to computing conditional-SRTBs, we modify the belief level construction to allow for exogenous events. Notice that exogenous events are considered to be observable in this case, so we can use their occurrence to refine the belief states in the belief-level full enacted system. This leads to the following definition: given a belief level full enacted system $\mathcal{K}_{\langle S, \mathcal{T} \rangle} = \langle Q, \mathcal{A}_{\mathcal{K}}^C \cup \mathcal{A}_{\mathcal{K}}^U, q_0, \delta_{\mathcal{K}} \rangle$ for full enacted system $\mathcal{F}_{\langle S, \mathcal{T} \rangle} = \langle F, \mathcal{A}_{\mathcal{F}}^C \cup \mathcal{A}_{\mathcal{F}}^U, f_0, \gamma \rangle$, the *conditional belief-level full enacted system* is a tuple $\mathcal{K}_{\langle S, \mathcal{T} \rangle}^C = \langle Q^C, \mathcal{A}_{\mathcal{K}}^C \cup \mathcal{A}_{\mathcal{K}}^U, q_0, \delta_{\mathcal{K}}^C \rangle$, where:

- $Q^C = Q \setminus \{ \langle S, t \rangle \mid \langle s, t \rangle \in \Delta_{\langle S, \mathcal{T} \rangle}, s \in S \}$; that is, prohibited exogenous events should never occur;
- $\delta_{\mathcal{K}}^C \subseteq Q \times \mathcal{A} \times Q$ is $\mathcal{K}_{\langle S, \mathcal{T} \rangle}^C$ ’s transition relation, where $\langle \langle S, t \rangle, a, \langle S', t' \rangle \rangle \in \delta_{\mathcal{K}}^C$ iff:
 - $a \in \mathcal{A}_{\mathcal{F}}^C$ and $\langle \langle S, t \rangle, a, \langle S', t' \rangle \rangle \in \delta_{\mathcal{K}}$; that is, action a should be executable from all enacted states; and
 - $a \in \mathcal{A}_{\mathcal{F}}^U$ and $S' = \{ s' \mid \langle \langle s, t \rangle, a, k, \langle s', t' \rangle \rangle \in \gamma, s \in S \}$; we revise belief state if an exogenous event occurs.

Next result shows that the conditional belief-level full enacted system is the SRTB in this context.

Theorem 5 *Let \mathcal{S} be an available system and \mathcal{T} a target spec. Then, $\mathcal{K}_{\langle \mathcal{S}, \mathcal{T} \rangle}^C$ is the conditional-SRTB of \mathcal{T} in \mathcal{S} .*

4.2 Conformant SRTBs

Conformant solutions guarantee realizability in absence of any observation over exogenous events. For example, the conformant solution $\tilde{\mathcal{T}}_1$ in Figure 1 contains a very restricted subset of the target as, if the bulb is fused then the projector cannot be operated again without a REPAIR. Solutions of such type are stricter, promising execution irrespective of which uncontrollable events occur. This provides robustness in modelling as one can still prevent unacceptable conditions under non-observability at runtime. We say a RTB to be conformant if it does not include any exogenous event, that is, $\mathcal{A}_{\mathcal{T}}^U = \emptyset$. Note, the target specification (problem input) is allowed to have exogenous events, however, a conformant RTB must have compiled them away. More precisely, a transition system $\tilde{\mathcal{T}} = \langle \tilde{T}, \tilde{\mathcal{A}}_{\mathcal{T}}^C, \tilde{t}_0, \tilde{\varrho}_{\mathcal{T}} \rangle$ is a *conformant-RTB* for a target $\mathcal{T} = \langle T, \mathcal{A}_{\mathcal{T}}^C \cup \mathcal{A}_{\mathcal{T}}^U, t_0, \varrho_{\mathcal{T}} \rangle$ in system \mathcal{S} with enacted system $\mathcal{E}_{\mathcal{S}} = \langle S, \mathcal{A}_{\mathcal{S}}^C \cup \mathcal{A}_{\mathcal{S}}^U, s_0, \delta \rangle$ iff $\tilde{\mathcal{T}} \preceq \mathcal{T}$ and $\langle \tilde{t}_0, s_0 \rangle \in \mathcal{Z}$ where $\mathcal{Z} \subseteq \tilde{T} \times S$ is the *conformant simulation relation* of $\tilde{\mathcal{T}}$ by $\mathcal{E}_{\mathcal{S}}$ such that $\langle \tilde{t}, s \rangle \in \mathcal{Z}$ iff:

1. $\forall \tilde{t} \forall a \exists k \forall s' (\langle \tilde{t}, a, k, s' \rangle \in \tilde{\varrho}_{\mathcal{T}} \Rightarrow \langle s, a, k, s' \rangle \in \delta)$ such that $\langle \tilde{t}', s' \rangle \in \mathcal{Z}$;
2. $\forall \alpha \in \mathcal{A}_{\mathcal{S}}^U, \forall k (\langle s, \alpha, k, s' \rangle \in \delta \Rightarrow \langle \tilde{t}, s' \rangle \in \mathcal{Z})$; and
3. $\forall \alpha \in \mathcal{A}_{\mathcal{S}}^U, \forall k (\langle s, \alpha, k, s' \rangle \in \delta \wedge \langle \tilde{t}, t \rangle \in \preceq \Rightarrow \langle t, \alpha, t' \rangle \in \varrho_{\mathcal{T}})$ such that $\langle \tilde{t}, t' \rangle \in \preceq$.

The first condition is analogous to the usual ND-simulation one. The second condition requires occurring of exogenous events should retain the simulation relation. The third condition enforces only permitted exogenous events to ever occur in the system. As usual, a conformant RTB is *supremal* iff it is not strictly simulated by any other conformant RTB.

Computing conformant SRTBs Conformant solutions require realizability guarantee irrespective of any nondeterministic or exogenous evolution. In order to include them in the belief-level system we first define what we call as the ε -closure of a state. That is, where all could the system be as a result of an exogenous event from that state. Formally, given a full enacted system $\mathcal{F}_{\langle \mathcal{S}, \mathcal{T} \rangle} = \langle F, \mathcal{A}_{\mathcal{F}}^C \cup \mathcal{A}_{\mathcal{F}}^U, f_0, \gamma \rangle$ and a state $f \in F$, the ε -closure of f , denoted by $\varepsilon(f)$, is defined recursively as follows:

1. $f \in \varepsilon(f)$, that is, the state itself is in the closure;
2. $\forall \alpha \in \mathcal{A}_{\mathcal{F}}^U, \forall f' \in \varepsilon(f) (f \xrightarrow{\alpha, k} f' \in \gamma \Rightarrow f' \in \varepsilon(f))$, that is, all exogenous event reachable states are included; and
3. Nothing else except for 1 and 2 should be in $\varepsilon(f)$.

Next, we re-define the belief level full enacted system to accommodate exogenous events. Here, we consider the ε -closure in both the initial state and the transition relation. Given a full enacted system $\mathcal{F}_{\langle \mathcal{S}, \mathcal{T} \rangle} = \langle F, \mathcal{A}_{\mathcal{F}}^C \cup$

$\mathcal{A}_{\mathcal{F}}^U, f_0, \gamma \rangle$ for a target $\mathcal{T} = \langle T, \mathcal{A}_{\mathcal{T}}, t_0, \varrho_{\mathcal{T}} \rangle$ and a system $\mathcal{S} = \langle \mathcal{B}_1, \dots, \mathcal{B}_n \rangle$, the *conformant belief-level full enacted system* is a tuple $\mathcal{K}_{\langle \mathcal{S}, \mathcal{T} \rangle}^Z = \langle Q, \mathcal{A}_{\mathcal{F}}^C, q_0, \delta_K \rangle$, where:

- $Q = 2^{(B_1 \times \dots \times B_n \times T)} \setminus \{S \mid S \in \Delta_{\langle \mathcal{S}, \mathcal{T} \rangle}, s \in S\}$;
- $q_0 = \varepsilon(f_0)$ is $\mathcal{K}_{\langle \mathcal{S}, \mathcal{T} \rangle}^Z$'s initial state;
- $\delta_K \subseteq Q \times \mathcal{A} \times Q$, where $\langle S, a, S' \rangle \in \delta_K$ iff :
 - there exists a set $\text{Idx} = \{\langle s_1 : k_1 \rangle, \dots, \langle s_\ell : k_\ell \rangle\}$ such that $\{s_1, \dots, s_\ell\} = S$; $s_i \xrightarrow{a, k_i} s'_i$ in $\mathcal{F}_{\langle \mathcal{S}, \mathcal{T} \rangle}$ for all $i \leq \ell$; and for all $i, j \leq \ell$ if $\text{tgt}(s_i) = \text{tgt}(s_j)$, then $\text{tgt}(s'_i) = \text{tgt}(s'_j)$; and
 - $S' = \bigcup_{\langle s : i \rangle \in \text{Idx}} \{\varepsilon(s') \mid \langle s, a, i, s' \rangle \in \gamma\}$, that is, S' should contain the ε -closure of all successors of enacted system states in S resulting from action a .

Note, the belief level full enacted system is now exponential also on the target states. Observe, if the target specification allows all exogenous events at any point then the complexity in regards to the target will no longer be exponential.

Theorem 6 *Let \mathcal{S} be an available system and \mathcal{T} a target spec. Then, $\mathcal{K}_{\langle \mathcal{S}, \mathcal{T} \rangle}^Z$ is the conformant-SRTB of \mathcal{T} in \mathcal{S} .*

5 Conclusion and Future work

In this paper, we proved that *every* classical behavior composition problem instance has an optimal supremal solution (Theorem 2) and that such supremal can be effectively built using cross-product between transition systems and belief-level state construction operations combined (Theorem 4). What is more, borrowing notions from discrete-event systems and reasoning about action, we showed how to accommodate *exogenous uncontrollable events* to obtain a more expressive composition framework (Section 4). We demonstrated that the definitions and techniques for supremal fragments can be adapted to this new framework (Theorems 5 and 6).

Many issues remain to be investigated. First, we aim to confirm the conjecture that our technique builds SRTBs that are optimal wrt worst-case complexity, thus implying that synthesis of supremals is, in general, more difficult than synthesis of exact composition controllers.

Second, our approach to realizing a target specification to the “best” possible is developed within a *strict* uncertainty context. This contrasts with the decision-theoretic approach of Yadav and Sardina [2011], where they proposed to optimize the expected reward of controllers. It would be interesting to adopt such a quantitative framework focusing on targets and devising a suitable decision-theoretic notion of SRTBs.

Finally, one may devise approaches that trade optimality for faster computation, such as restricting realizable target fragments to merely removing transitions from the original target specification, bounding its number of states, or computing it in anytime fashion. This paper provides principled means to assess the quality of such approximate approaches.

We close the paper by noting that there are some interesting links between behavior composition in AI and work on controllability in discrete event systems [Cassandras and Laforune, 2006; Sun *et al.*, 2010]. Exploring those links is indeed worth investigating as it can facilitate synergies between these two different fields.

References

- [Balbiani *et al.*, 2008] Philippe Balbiani, Fahima Cheikh, and Guillaume Feuillade. Composition of interactive web services based on controller synthesis. In *Proceedings of the IEEE Congress on Services (SERVICES)*, pages 521–528, 2008.
- [Bordignon *et al.*, 2007] M. Bordignon, J. Rashid, M. Broxvall, and Alessandro Saffiotti. Seamless integration of robots and tiny embedded devices in a PEIS-ecology. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3101–3106, 2007.
- [Cassandras and Lafortune, 2006] Christos G. Cassandras and Stephane Lafortune. *Introduction to Discrete Event Systems*. Springer, Secaucus, NJ, USA, 2006.
- [De Giacomo and Sardina, 2007] Giuseppe De Giacomo and Sebastian Sardina. Automatic synthesis of new behaviors from a library of available behaviors. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1866–1871, 2007.
- [De Giacomo *et al.*, 2009] Giuseppe De Giacomo, Riccardo De Masellis, and Fabio Patrizi. Composition of partially observable services exporting their behaviour. In *Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS)*, pages 90–97, 2009.
- [De Giacomo *et al.*, 2013] Giuseppe De Giacomo, Fabio Patrizi, and Sebastian Sardina. Automatic behavior composition synthesis. *Artificial Intelligence Journal*, 196:106–142, 2013.
- [Lustig and Vardi, 2009] Yoad Lustig and Moshe Y. Vardi. Synthesis from component libraries. In *Proceedings of the International Conference on Foundations of Software Science and Computational Structures (FOSSACS)*, volume 5504 of *Lecture Notes in Computer Science (LNCS)*, pages 395–409. Springer, 2009.
- [Milner, 1971] Robin Milner. An algebraic definition of simulation between programs. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 481–489, 1971.
- [Reiter, 2001] Ray Reiter. *Knowledge in Action. Logical Foundations for Specifying and Implementing Dynamical Systems*. The MIT Press, 2001.
- [Sardina *et al.*, 2008] Sebastian Sardina, Fabio Patrizi, and Giuseppe De Giacomo. Behavior composition in the presence of failure. In *Proceedings of Principles of Knowledge Representation and Reasoning (KR)*, pages 640–650, 2008.
- [Stroeder and Pagnucco, 2009] Thomas Stroeder and Maurice Pagnucco. Realising deterministic behaviour from multiple non-deterministic behaviours. In *Proceedings of IJCAI*, pages 936–941, 2009.
- [Sun *et al.*, 2010] Yajuan Sun, Hai Lin, Fuchun Liu, and Ben M Chen. Computation for supremal simulation-based controllable subautomata. In *Control and Automation (ICCA), 2010 8th IEEE International Conference on*, pages 1450–1455. IEEE, 2010.
- [Wonham and Ramadge, 1987] W. M. Wonham and P. J. Ramadge. On the supremal controllable sub-language of a given language. *SIAM Journal on Control and Optimization*, 25(3):637–659, 1987.
- [Yadav and Sardina, 2011] Nitin Yadav and Sebastian Sardina. Decision theoretic behavior composition. In *Proceedings of Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 575–582, 2011.
- [Yadav and Sardina, 2012] Nitin Yadav and Sebastian Sardina. Qualitative approximate behavior composition. In *Proceedings of the European Conference on Logics in Artificial Intelligence (JELIA)*, volume 7519 of *Lecture Notes in Computer Science (LNCS)*, pages 450–462. Springer, 2012.