

Timed-automata abstraction of switched dynamical systems using control invariants^{*}

Patricia Bouyer¹, Nicolas Markey¹,
Nicolas Perrin^{2,3}, Philipp Schlehuber-Caissier²

¹LSV – CNRS & ENS Cachan – Université Paris-Saclay – France

²Sorbonne Universités, UPMC Univ Paris 06, UMR 7222, ISIR, F-75005, Paris, France

³CNRS, UMR 7222, ISIR, F-75005, Paris, France

Abstract. The development of formal methods for control design is an important challenge with potential applications in a wide range of safety-critical cyber-physical systems. Focusing on switched dynamical systems, we propose a new abstraction, based on time-varying regions of invariance (*control funnels*), that models behaviors of systems as timed automata. The main advantage of this method is that it allows for the automated verification and reactive controller synthesis *without discretizing* the evolution of the state of the system. Efficient and analytic constructions are possible in the case of linear dynamics whereas bounding funnels with conjectured properties based on numerical simulations can be used for general nonlinear dynamics. We demonstrate the potential of our approach with three examples.

1 Introduction

Verification and synthesis are notoriously difficult for hybrid dynamical systems, i.e. systems that allow abrupt changes in continuous dynamics. For instance, reachability is already undecidable for 2-dimensional piecewise-affine maps [16], or for 3-dimensional dynamical systems with piecewise-constant derivatives [2].

To enable automated logical reasoning on switched dynamical systems, most methods tend to entirely discretize the dynamics, for example by approximating the behavior of the system with a finite-state machine. Alternatively, early work pointed out links between hybrid and timed systems [22], and several methods have been designed to create formal abstractions of dynamical systems that do not rely on a discretization of time. In [13], a finite maneuver automaton is constructed from a library of motion primitives, and motion plans correspond to timed words. In [18, 14], switched controller synthesis and stochastic optimal control are realized via metric temporal logic (MTL) or metric-interval temporal logic (MITL) specifications. In [25, 21], grid-based abstractions and timed automata are used for motion planning or to check timed properties of dynamical systems. In [27], a subdivision of the state space created from sublevel sets of Lyapunov functions

^{*} This work has been partly supported by ERC Starting grant EQualIS (FP7-308087) and by European FET project Cassting (FP7-601148)

leads to an abstraction of dynamical systems by timed automata that enables verification and falsification of safety properties. The same kind of abstraction is used in [26] for controller design via timed games, but the update map of the timed games obtained is such that synthesis cannot be realized using existing tools. In [10], the state space of each mode of a piecewise-affine hybrid system is partitioned into polytopes, and thanks to control laws that prevent the system from exiting through a given facet, or that force the system to exit through a facet in finite time, reactive control problems can be solved as timed games on timed automata.

Our contribution is a novel timed-automata abstraction of switched dynamical systems based on a particular kind of time-varying regions of invariance: control funnels. Recent results have shown that these invariants are very useful for robust motion planning and control [29, 20, 19], and that funnels or similar concepts related to the notion of Lyapunov stability can be used for formal verification of hybrid systems [15, 12], and for reactive controller synthesis [11].

The paper is organized as follows: Section 2 describes how control funnels, in particular for trajectory tracking controllers, can be used to create timed transition systems that abstract the behavior of a given switched dynamical system, and as a result can potentially allow for the use of verification tools to solve Reach-Avoid problems for this kind of systems. In Section 3, we show how these timed transition systems can be encoded as timed automata. In Section 4, we consider the case of linear dynamics and introduce the notion of fixed-size LQR funnel. In Section 5, we present two examples of application and efficient algorithms that manipulate the LQR funnels. In the first one, a timed game is solved by the tool Uppaal-Tiga [5] for the synthesis of a controller that can reactively adjust the phase of a sine wave controlled in acceleration. In the second example, we show that, using our timed-automata abstraction with LQR funnels along constant velocity trajectories, a non-trivial solution to a pick-and-place problem can be computed by the model checker Uppaal [6]. In Section 6, we introduce *bounding funnels* using conjectured properties, i.e. funnels obtained without formal proofs, for example via numerical simulations. We then present an example of application solving a Reach-Avoid problem for a nonlinear and non-holonomic system, a modified version of the Dubins' car. Section 7 concludes and presents avenues for future work.

This paper extends [9] by the introduction of *bounding funnels*, see Section 6, enlarging the class of dynamical systems to which our method can be applied in practice. An example demonstrating the usefulness of bounding funnels is provided in the same section.

2 Graphs of control funnels

2.1 Control funnels

Consider a controlled dynamical system governed by the following equation:

$$\dot{\mathbf{x}} = f(\mathbf{x}, u(\mathbf{x}, t)), \quad (1)$$




Fig. 1. An example of control funnel for a controller tracking a reference trajectory. The dashed line is a trajectory of the controlled system in the state space. On the right side, switching transitions between control funnels are depicted.

where $\mathbf{x} \in \mathbb{R}^d$ is the state of the system (which can contain velocities¹), $t \in \mathbb{R}^+$ is a real (clock) value corresponding to an internal controller time (without loss of generality we restrict ourselves to nonnegative time values), $u: \mathbb{R}^d \times \mathbb{R}^+ \rightarrow \mathbb{R}^k$ is the control input function, and f is a continuously differentiable function from $\mathbb{R}^d \times \mathbb{R}^k$ to \mathbb{R}^d (which ensures the uniqueness of the solution for given initial conditions). Assuming that the function u is fixed, we also use the following notation for Equation (1):

$$\dot{\mathbf{x}} = f_u(\mathbf{x}, t). \quad (2)$$

It is worth noting that, since t is an internal controller time, it can have a discontinuous evolution with discrete resets to any value in \mathbb{R}^+ . However, except for these resets, the controller time is assumed to continuously increase at a constant rate (with respect to the reference real time).

A *control funnel* for the above dynamical system is a function $\mathcal{F}: I \rightarrow 2^{\mathbb{R}^d}$ such that $I \subseteq \mathbb{R}^+$ and for any solution $\mathbf{x}(t)$ of (2) with no reset of the controller time t , the following property holds:

$$\forall t_1 \in I, \forall t_2 \in I, (t_2 > t_1) \Rightarrow [\mathbf{x}(t_1) \in \mathcal{F}(t_1) \Rightarrow \mathbf{x}(t_2) \in \mathcal{F}(t_2)]. \quad (3)$$

Equation (3) defines a property known as positive invariance, and the funnel \mathcal{F} corresponds to a time-varying region of invariance.

Example 1. A typical example of a control funnel based on a trajectory tracking controller (that is, a control funnel asymptotically converging towards a reference trajectory in the state space) is shown in Fig. 1.

Example 2. For a concrete example, consider the simple system whose trajectories are of the form $e^{-t} \cdot \mathbf{x}_0$. Then any set $W \subseteq \mathbb{R}^d$ defines a control funnel $\mathcal{F}_W: t \mapsto \{e^{-t} \cdot \mathbf{w} \mid \mathbf{w} \in W\}$.

¹ In this paper, we mostly consider state spaces that describe the position and velocity of systems controlled in acceleration. The continuity of trajectories in the state space ensures that the position is always a continuously differentiable function of time.

The notion of funnel was popularized by Mason [23], and it usually specifically refers to operations that eliminate uncertainty (as is the case in the example of Fig. 1) by collapsing a large set of initial conditions into a smaller set of final conditions (see for instance [29]). In our case, the control funnel may or may not reduce uncertainty, and it is important to note that the set $\mathcal{F}(t)$ does not have to decrease in size over time. This more general concept is closer to the definition of *viability tubes* [4], but we nevertheless use the term control funnel as some reduction of uncertainty is often essential to the usefulness of our abstractions. We address the computation of control funnels in Section 4, and leave them as relatively abstract objects for now.

2.2 Formalizing the Reach-Avoid problem for controlled systems

Let us suppose that we have a finite set U of control laws $u_1(\mathbf{x}, t)$, $u_2(\mathbf{x}, t)$, \dots , $u_n(\mathbf{x}, t)$ that respectively set the dynamics of a given system to $\dot{\mathbf{x}} = f_{u_1}(\mathbf{x}, t)$, $\dot{\mathbf{x}} = f_{u_2}(\mathbf{x}, t)$, \dots , $\dot{\mathbf{x}} = f_{u_n}(\mathbf{x}, t)$.

We say that the system can switch to the control law $u_i(\mathbf{x}, t)$ at some state $\tilde{\mathbf{x}}$ whenever there is $t_0 \in \mathbb{R}^+$ and a solution $\mathbf{x}(t)$ of $\dot{\mathbf{x}} = f_{u_i}(\mathbf{x}, t)$ with initial condition $\tilde{\mathbf{x}} = \mathbf{x}(t_0)$. Typically, if $u_i(\mathbf{x}, t)$ corresponds to a trajectory tracking controller, t_0 identifies the point of the trajectory where the tracking is triggered.

Informally, the Reach-Avoid problem asks, given a finite set of control laws as above, an initial point \mathbf{x}_0 , a target zone $T_f \subseteq \mathbb{R}^d$, and a zone to avoid $\Omega \subseteq \mathbb{R}^d$ (also called obstacle), whether there exists a sequence of control-law switches that generates a trajectory reaching from \mathbf{x}_0 to T_f while avoiding the obstacle Ω . Several variants of this problem can be considered, that vary on the objective (for instance some tasks can be expressed as ω -regular objectives) which could also be solved using our approach, however we focus here on a pure reachability with avoidance objective.

More formally, the Reach-Avoid problem asks for a finite sequence of time values $t_0^1 < t_1^1$, $t_0^2 < t_1^2$, \dots , $t_0^P < t_1^P$, a finite sequence of control laws indices k_1, \dots, k_P , and a finite sequence $\mathbf{x}_1, \dots, \mathbf{x}_P \in T_f$ of points in \mathbb{R}^d , such that:

- (a) for every $1 \leq p \leq P$, if \mathbf{x}^p is the unique solution to $\dot{\mathbf{x}} = f_{u_{k_p}}(\mathbf{x}, t)$ with initial condition $\mathbf{x}^p(t_0^p) = \mathbf{x}_{p-1}$, then $\mathbf{x}^p(t_1^p) = \mathbf{x}_p$.
- (b) for every $1 \leq p \leq P$, for every $t_0^p \leq t \leq t_1^p$, $\mathbf{x}^p(t) \notin \Omega$.

Intuitively, this means that we can switch conveniently between all the control laws, causing discrete changes in the system dynamics, and ensure the global (reachability with avoidance) objective. The continuous trajectory generated by the solution above is the concatenation of the trajectory portions $\{\mathbf{x}^p(t) \mid t_0^p \leq t \leq t_1^p\}$ for $1 \leq p \leq P$.

2.3 Reach-Avoid objectives on graphs of control funnels

We now explain how the Reach-Avoid problem can be abstracted using timed transition systems based on control funnels.

For each control law $u_i(\mathbf{x}, t)$, we assume that we have a finite set of control funnels $\mathcal{F}_i^1, \mathcal{F}_i^2, \dots, \mathcal{F}_i^{m_i}$, respectively defined over $I_i^1 \subseteq \mathbb{R}^+$, $I_i^2 \subseteq \mathbb{R}^+$, \dots , $I_i^{m_i} \subseteq \mathbb{R}^+$. We assume that for every $1 \leq i \leq n$, for every $1 \leq j \leq m_i$, for every $t \in I_i^j$, it holds $\mathcal{F}_i^j(t) \cap \Omega = \emptyset$, which means that trajectories contained in these funnels always avoid the obstacle Ω .

Consider a control law switch at position $\tilde{\mathbf{x}}$ to law $u_i(\mathbf{x}, t)$ with clock value t_0 . If there exists a control funnel \mathcal{F}_i^j such that $t_0 \in I_i^j$, and $\tilde{\mathbf{x}} \in \mathcal{F}_i^j(t_0)$, then we know that the state of the system will remain inside $\mathcal{F}_i^j(t)$ for any $t > t_0$ in I_i^j (as long as control law $u_i(\mathbf{x}, t)$ is used). To always keep the system inside one of the control funnels, we can impose sufficient conditions on the switches. For instance, if the state is inside $\mathcal{F}_i^j(t_0)$, and if for some future clock value t_1 , there exists a control funnel \mathcal{F}_k^l and $t_2 \in I_k^l$ such that $\mathcal{F}_i^j(t_1) \subseteq \mathcal{F}_k^l(t_2)$, then when the clock value is t_1 we can safely switch to the control law $u_k(\mathbf{x}, t)$ while setting the clock to t_2 . Indeed, we know that the state of the system at the switch instant will be inside $\mathcal{F}_k^l(t_2)$, and therefore it will remain inside $\mathcal{F}_k^l(t)$ after the switch. Such transitions from a funnel to another are illustrated on the right side of Fig. 1. It is worth noting that similar transitions could be achieved with, instead of control funnels, controller specifications as introduced in [17]. For some control funnels \mathcal{F}_i^j and \mathcal{F}_i^k associated to the same control law, it may be the case (see Section 4) that when funnel \mathcal{F}_i^j is entered at time t , then at any time $t' \geq t + h_i^{j \rightarrow k}$ (where $h_i^{j \rightarrow k}$ is a constant), the state of the system is inside $\mathcal{F}_i^k(t')$ which is itself contained in $\mathcal{F}_i^j(t')$. In that case, we say that the funnel \mathcal{F}_i^k $h_i^{j \rightarrow k}$ -absorbs the funnel \mathcal{F}_i^j .

These rules for navigating between control funnels give to the set of control funnels the structure of an infinite graph, or, more precisely, of a timed transition system with real-valued clocks. One of the clocks of this timed transition system is c_t , the *controller clock*. We add two other clocks: a global clock c_g , and a local clock c_h .

The *funnel system* $\mathcal{T}_{U,F}$ associated with the family of laws $U = (u_i(\mathbf{x}, t))_{1 \leq i \leq n}$ and the family of funnels $F = ((\mathcal{F}_i^j, I_i^j))_{1 \leq i \leq n, 1 \leq j \leq m_i}$ is defined as follows. The configurations are pairs (\mathcal{F}_i^j, v) where v assigns a non-negative real value to each of the clocks c_t, c_g and c_h , with $v(c_t) \in I_i^j$, and its transition set contains three types of elements:

- the *time-elapsing transitions*: $(\mathcal{F}_i^j, v) \rightarrow (\mathcal{F}_i^j, v + \Delta)$ whenever $[v(c_t), v(c_t) + \Delta] \subseteq I_i^j$ (where $v + \Delta$ denotes the valuation that maps each clock c to $v(c) + \Delta$);
- the *switching transitions*: $(\mathcal{F}_i^j, v) \rightarrow (\mathcal{F}_k^l, v')$ whenever $v'(c_g) = v(c_g), v'(c_h) = 0, v(c_t) \in I_i^j, v'(c_t) \in I_k^l$, and $\mathcal{F}_i^j(v(c_t)) \subseteq \mathcal{F}_k^l(v'(c_t))$;
- the *absorbing transitions*: $(\mathcal{F}_i^j, v) \rightarrow (\mathcal{F}_i^k, v')$ whenever \mathcal{F}_i^k $h_i^{j \rightarrow k}$ -absorbs \mathcal{F}_i^j , $v(c_h) \geq h_i^{j \rightarrow k}, v'(c_h) = 0, v'(c_g) = v(c_g)$ and $v'(c_t) = v(c_t)$.

A run in this timed transition system is a finite sequence of configurations $((\mathcal{F}_{i_0}^{j_0}, v_0), (\mathcal{F}_{i_1}^{j_1}, v_1), \dots, (\mathcal{F}_{i_P}^{j_P}, v_P))$ such that $v_0(c_h) = v_0(c_g) = 0, v_0(c_t) \in I_{i_0}^{j_0}$,




Fig. 2. Run of a funnel system with three control funnels: $r = ((\mathcal{F}_1^1, v_0^1), (\mathcal{F}_1^1, v_1^1), (\mathcal{F}_2^1, v_0^2), (\mathcal{F}_2^1, v_1^2), (\mathcal{F}_3^1, v_0^3), (\mathcal{F}_3^1, v_1^3), (\mathcal{F}_4^1, v_0^4), (\mathcal{F}_4^1, v_1^4))$, with: $\forall 1 \leq i \leq 4, v_0^i(c_t) = t_0^i, v_1^i(c_t) = t_1^i, v_0^i(c_h) = 0, v_1^i(c_h) = t_1^i - t_0^i, v_1^i(c_g) = v_0^i(c_g) + v_1^i(c_h)$, and $v_0^i(c_g) = 0$, and $\forall 2 \leq i \leq 4, v_0^i(c_g) = v_1^{i-1}(c_g)$.

and all the transitions $(\mathcal{F}_{i_p}^{j_p}, v_p) \rightarrow (\mathcal{F}_{i_{p+1}}^{j_{p+1}}, v_{p+1})$ for $0 \leq p < P$ are valid transitions that belong to $\mathcal{T}_{U,F}$.

Such a run is of total duration $v_P(c_g)$, and it corresponds to a schedule of control-law switches that results from the following rules: initially, the control law is set to $u_{i_0}(\mathbf{x}, t)$, and the controller clock c_t is set to $v_0(c_t)$. For every time-elapsing transition $(\mathcal{F}_i^j, v) \rightarrow (\mathcal{F}_i^j, v + \Delta)$, the same control law $u_i(\mathbf{x}, t)$ is kept for a duration of Δ time units, and for every switching transition $(\mathcal{F}_i^j, v) \rightarrow (\mathcal{F}_k^l, v')$, the control law is switched from $u_i(\mathbf{x}, t)$ to $u_k(\mathbf{x}, t)$, with an initialization of the controller clock to $v'(c_t)$. Absorbing transitions are discarded, as they just correspond to an instantaneous change of funnels for the same control law. Let us denote this sequence of switches by r . Then, it is fundamental to notice that, for every $\mathbf{x} \in \mathcal{F}_{i_0}^{j_0}(v_0(c_t))$, if we follow the schedule of control-law switches just described, then the system remains inside control funnels and reaches at the end of the run a unique point of \mathbb{R}^d , which we denote $r(\mathbf{x})$. The trajectory going from \mathbf{x} to $r(\mathbf{x})$ is also uniquely defined.

The funnel system $\mathcal{T}_{U,F}$ satisfies the following property:

Theorem 1. *Let $r = ((\mathcal{F}_{i_0}^{j_0}, v_0), (\mathcal{F}_{i_1}^{j_1}, v_1), \dots, (\mathcal{F}_{i_P}^{j_P}, v_P))$ be a run in $\mathcal{T}_{U,F}$. If $\mathbf{x} \in \mathcal{F}_{i_0}^{j_0}(v_0(c_t))$, then $r(\mathbf{x}) \in \mathcal{F}_{i_P}^{j_P}(v_P(c_t))$.*

In some sense, the funnel system $\mathcal{T}_{U,F}$ is a correct abstraction of trajectories that can be generated by the set of control laws: if $\mathbf{x}_0 \in \mathcal{F}_{i_0}^{j_0}(v_0(c_t))$ and $\mathcal{F}_{i_P}^{j_P}(v_P(c_t)) \subseteq T_f$, then such a run witnesses a solution to the Reach-Avoid problem. However, this abstraction is obviously not complete.

Example 3 (An example with obstacle). The example in Fig. 2 shows a run with three control laws $u_1(\mathbf{x}, t)$, $u_2(\mathbf{x}, t)$ and $u_3(\mathbf{x}, t)$, three control funnels \mathcal{F}_1^1 , \mathcal{F}_2^1 and \mathcal{F}_3^1 , and an obstacle in the state space. The domains of definition of the control funnels I_1^1 , I_2^1 and I_3^1 are such that for all $\alpha \in \{1, 2, 3\}$ and all $t \in I_\alpha^1$, $\mathcal{F}_\alpha^1(t)$ does not intersect the obstacle.

With the previous property, any run in the corresponding funnel system leads to a trajectory that avoids the obstacle. The example of Fig. 2, where reaching $\mathcal{F}_1^1(t_1^4)$ from $\mathcal{F}_1^1(t_0^1)$ requires a series of switches between the different control funnels, shows the potential interest of automated verification in timed transition systems, as it can result in the generation of obstacle-free dynamic trajectories via appropriate control law switches.

3 Reduction to timed automata

Timed automata [1] are a timed extension of finite-state automata, with a well-understood theory. They provide an expressive formalism for modelling and reasoning about real-time systems, and enjoy decidable reachability properties; much efforts have been invested over the last 20 years for the development of efficient algorithms and tools for their automatic verification (such as the tool Uppaal [6], which we use in this work).

Let C be a finite set of real-valued variables called *clocks*. A *clock valuation* over a finite set of clocks C is a mapping $v: C \rightarrow \mathbb{R}^+$. We write \mathbb{R}^C for the set of clock valuations over C . If $\Delta \in \mathbb{R}^+$, we write $v + \Delta$ for the clock valuation defined by $(v + \Delta)(c) = v(c) + \Delta$ for every $c \in C$. A *clock constraint* over C is a boolean combination of expressions of the form $c \sim \alpha$ where $\alpha \in \mathbb{Q}$, and $\sim \in \{\leq, <, =, >, \geq\}$. We denote by $\mathcal{C}(C)$ the set of clock constraints over C . We write $v \models g$ if v satisfies g (defined in a natural way). A reset of the clocks is an element res of $(\mathbb{Q} \cup \{\perp\})^C$ (which we may write $R(C)$), and if v is a valuation, its image by res , denoted $\text{res}(v)$, is the valuation mapping c to $v(c)$ whenever $\text{res}(c) = \perp$, and to $\text{res}(c) \in \mathbb{Q}$ otherwise.

We define a slight extension of timed automata with rational constants, general boolean combinations of clock constraints and extended clock resets; those timed automata are as expressive as standard timed automata [7], but they will be useful for expressing funnel systems. A *timed automaton* is a tuple $\mathcal{A} = (L, L_0, L_F, C, E, \text{Inv})$ where L is a finite set of locations, $L_0 \subseteq L$ is a set of initial locations, $L_F \subseteq L$ is a set of final locations, C is a finite set of clocks, $E \subseteq L \times \mathcal{C}(C) \times R(C) \times L$ is a finite set of edges, and $\text{Inv}: L \rightarrow \mathcal{C}(C)$ is an invariant labelling function.

A configuration of \mathcal{A} is a pair $(\ell, v) \in L \times \mathbb{R}^C$ such that $v \models \text{Inv}(\ell)$, and the timed transition system generated by \mathcal{A} is given by the following two rules:

- *time-elapsing transition*: $(\ell, v) \rightarrow (\ell, v + \Delta)$ whenever $v + \delta \in \text{Inv}(\ell)$ for every $0 \leq \delta \leq \Delta$;
- *switching or absorbing transition*: $(\ell, v) \rightarrow (\ell', v')$ whenever there exists $(\ell, g, \text{res}, \ell') \in E$ such that $v \models g \wedge \text{Inv}(\ell)$, $v' = \text{res}(v)$, and $v' \in \text{Inv}(\ell')$.

A run in \mathcal{A} is a sequence of consecutive transitions. The most fundamental result about timed automata is the following:

Theorem 2 ([1]). *Reachability in timed automata is PSPACE-complete.*

We consider again the family of control laws $U = (u_i(\mathbf{x}, t))_{1 \leq i \leq n}$, and the family of funnels $F = ((\mathcal{F}_i^j, I_i^j))_{1 \leq i \leq n, 1 \leq j \leq m_i}$, as in the previous section. For every pair $1 \leq i, k \leq n$, and every $1 \leq j \leq m_i$ and $1 \leq l \leq m_k$, we select finitely many tuples (switch, $[\alpha, \beta], (i, j), \gamma, (k, l)$) with $\alpha, \beta, \gamma \in \mathbb{Q}$ such that $[\alpha, \beta] \subseteq I_i^j$, $\gamma \in I_k^l$, and for every $t \in [\alpha, \beta]$, $\mathcal{F}_i^j(t) \subseteq \mathcal{F}_k^l(\gamma)$. This allows us to under-approximate the possible switches between funnels. For every $1 \leq i \leq n$, for every pair $1 \leq j, k \leq m_i$, we select at most one tuple (absorb, $\nu, (i, j, k)$) such that $\nu \in \mathbb{Q}$ and $\mathcal{F}_i^k(t)$ ν -absorbs $\mathcal{F}_i^j(t)$. This allows us to under-approximate the possible absorbing transitions. For every $1 \leq i \leq n$ and every $1 \leq j \leq m_i$, we fix a finite set of tuples (initial, $\alpha, (i, j)$) such that $\alpha \in \mathbb{Q}$ and $\mathbf{x}_0 \in \mathcal{F}_i^j(\alpha)$. This allows us to under-approximate the possible initialization to a control funnel containing the initial point \mathbf{x}_0 . For every $1 \leq i \leq n$ and $1 \leq j \leq m_i$, we fix finitely many tuples (invariant, $S_{i,j}, (i, j)$), where $S_{i,j} \subseteq I_i^j$ is a finite set of closed intervals with rational bounds. This allows us to under-approximate the definition set of the funnels. Finally, for every $1 \leq i \leq n$ and $1 \leq j \leq m_i$, we fix finitely many tuples (target, $[\alpha, \beta], (i, j)$), where $\alpha, \beta \in \mathbb{Q}$ and $[\alpha, \beta] \subseteq I_i^j \cap \{t \mid \mathcal{F}_i^j(t) \subseteq T_f\}$. This allows us to under-approximate the target zone. We denote by K the set of all tuples we just defined.

We can now define a timed automaton that conservatively computes the runs generated by the funnel system. It is defined by $\mathcal{A}_{U,F,K} = (L, L_0, L_F, C, E, \text{Inv})$ with:

- $L = \{\mathcal{F}_i^j \mid 1 \leq i \leq n, 1 \leq j \leq m_i\} \cup \{\text{init}, \text{stop}\}$; $L_0 = \{\text{init}\}$; $L_F = \{\text{stop}\}$;
- $C = \{c_t, c_g, c_h\}$;
- E is composed of the following edges:
 - for every (initial, $\alpha, (i, j)$) $\in K$, we have an edge (init, true, res, \mathcal{F}_i^j) in E , with $\text{res}(c_t) = \alpha$ and $\text{res}(c_g) = \text{res}(c_h) = 0$;
 - for every (switch, $[\alpha, \beta], (i, j), \gamma, (k, l)$) $\in K$, we have an edge ($\mathcal{F}_i^j, \alpha \leq c_t \leq \beta$, res, \mathcal{F}_k^l) with $\text{res}(c_t) = \gamma$, $\text{res}(c_h) = 0$ and $\text{res}(c_g) = \perp$;
 - for every (target, $[\alpha, \beta], (i, j)$) $\in K$, we have an edge ($\mathcal{F}_i^j, \alpha \leq c_t \leq \beta$, res, stop) in E , with $\text{res}(c_t) = \text{res}(c_g) = \text{res}(c_h) = \perp$;
 - for every (absorb, $\nu, (i, j, k)$) $\in K$, we have an edge ($\mathcal{F}_i^j, c_h \geq \nu$, res, \mathcal{F}_i^k) with $\text{res}(c_h) = 0$ and $\text{res}(c_t) = \text{res}(c_g) = \perp$;
- for every (invariant, $S_{i,j}, (i, j)$) $\in K$, we let $\text{Inv}(\mathcal{F}_i^j) \triangleq \bigvee_{[\alpha, \beta] \in S_{i,j}} (\alpha \leq c_t \leq \beta)$.

We easily get the following result:

Theorem 3. *Let (init, v_0) \rightarrow (ℓ_1, v_1) \rightarrow \dots \rightarrow (ℓ_P, v_P) \rightarrow (stop, v_P) be a run in $\mathcal{A}_{U,F,K}$ such that v_0 assigns 0 to every clock. Then $r = ((\ell_1, v_1), \dots, (\ell_P, v_P))$ is a run of the funnel system $\mathcal{T}_{U,F}$ that brings \mathbf{x}_0 to $r(\mathbf{x}_0) \in T_f$ while avoiding the obstacle Ω .*

This shows that the reachability of stop in $\mathcal{A}_{U,F,K}$ implies that there exists an appropriate schedule of control law switches that safely brings the system to the target zone. Of course, the method is not complete, not all schedules can be obtained using the timed automaton $\mathcal{A}_{U,F,K}$. But if $\mathcal{A}_{U,F,K}$ is precise enough, it will be possible to use automatic verification techniques for dynamic trajectory generation.

Remark 1. We could be more precise in the modelling as a timed automaton by using non-deterministic clock resets (while taking care of decidability issues) [7]. But non-deterministic resets are not implemented in Uppaal, which is the reason why we use deterministic resets only.

Remark 2. As we show with some examples in Section 5, our timed automata abstraction can be used for other types of objectives than just reachability with avoidance. In particular, the approach can be extended to take *external events* into account (e.g. moving obstacles), by adapting our construction above using *timed games* [3]. Timed games extend timed automata with special *uncontrollable* transitions whose occurrence is decided by an opponent, and cannot be forced nor prevented; in this context, we look for *strategies*, that have to take into account those events and react appropriately when they occur. Moving obstacles would be modelled by letting the opponent player maintain a list of *bad funnels* at each time, and a valid strategy would adapt to these choices in order to continuously avoid those funnels.

It is worth knowing that winning strategies can be computed in exponential time in timed games, and that the tool Uppaal-Tiga [5] computes winning strategies. In Section 5.1, we give an example of application where timed games and Uppaal-Tiga are used. By using the model of weighted timed automata [8], one can additionally try to minimize the number of control switches, or the time for reaching the target.

Remark 3. The timed automaton obtained from the funnel system represents an under-approximation of all the obstacle-avoiding trajectories that the system can perform. Other constraints on the system, like logical specifications as in the example of Section 5.2, can be represented by an auxiliary automaton.

4 LQR funnels

In this section we consider the particular case of linear time-invariant stabilizable systems whose dynamics are described by the following equation:

$$\dot{\mathbf{x}} = A\mathbf{x} + B\mathbf{u}, \quad (4)$$

where $A \in \mathbb{R}^{d \times d}$ and $B \in \mathbb{R}^{d \times k}$ are two constant matrices, and $\mathbf{u} \in \mathbb{R}^k$ is the control input. We also consider reference trajectories that can be realized with controlled systems described by Eq. (4), i.e. trajectories $\mathbf{x}_{\text{ref}}(t)$ for which there exists $\mathbf{u}_{\text{ref}}(\mathbf{x}, t)$ such that $\dot{\mathbf{x}}_{\text{ref}} = A\mathbf{x}_{\text{ref}} + B\mathbf{u}_{\text{ref}}$. We can combine this equation with (4) and get $\dot{\mathbf{x}} - \dot{\mathbf{x}}_{\text{ref}} = A(\mathbf{x} - \mathbf{x}_{\text{ref}}) + B(\mathbf{u} - \mathbf{u}_{\text{ref}})$, which rewrites

$$\dot{\mathbf{x}}_{\Delta} = A\mathbf{x}_{\Delta} + B\mathbf{u}_{\Delta}. \quad (5)$$

To track \mathbf{x}_{ref} , we compute \mathbf{u}_{Δ} as an infinite-time linear quadratic regulator (LQR, see [28]), i.e. a minimization of the cost: $J = \int_0^{\infty} (\mathbf{x}_{\Delta}^T Q \mathbf{x}_{\Delta} + \mathbf{u}_{\Delta}^T R \mathbf{u}_{\Delta}) dt$, where Q and R respectively are positive-semidefinite and positive-definite matrices. The solution is $\mathbf{u}_{\Delta} = -K\mathbf{x}_{\Delta}$, with $K = R^{-1}B^T P$ and P being the unique

positive-definite matrix solution of the continuous-time algebraic Riccati equation: $PA + A^\top P - PBR^{-1}B^\top P + Q = 0$.

The dynamics can be rewritten $\dot{\mathbf{x}}_\Delta = (A - BK)\mathbf{x}_\Delta = \bar{A}\mathbf{x}_\Delta$, i.e.:

$$\dot{\mathbf{x}} = \dot{\mathbf{x}}_{\text{ref}} + \bar{A}(\mathbf{x} - \mathbf{x}_{\text{ref}}), \quad (6)$$

and the matrix \bar{A} is Hurwitz, i.e. all its eigenvalues have negative real parts. Additionally, $V: \mathbf{x}_\Delta \mapsto \mathbf{x}_\Delta^\top P \mathbf{x}_\Delta$ is a Lyapunov function ($V(\mathbf{0}) = 0$ and for all $\mathbf{x}_\Delta \neq \mathbf{0}$, it holds $V(\mathbf{x}_\Delta) > 0$ and $\dot{V}(\mathbf{x}_\Delta) < 0$). The solutions of Eq. (6) can be written: $\mathbf{x}(t) = \mathbf{x}_{\text{ref}}(t) + e^{\bar{A}(t-t_0)}\mathbf{x}_\Delta(t_0)$. Since \bar{A} is Hurwitz, the term $e^{\bar{A}(t-t_0)}$ tends to 0 exponentially fast, and the tracking asymptotically converges towards the reference trajectory $\mathbf{x}_{\text{ref}}(t)$. The Lyapunov function V can be used to define control funnels as follows. For $\alpha > 0$, we let:

$$\mathcal{F}_\alpha(t) = \{\mathbf{x}_{\text{ref}}(t) + \mathbf{x}_\Delta \mid V(\mathbf{x}_\Delta) \leq \alpha\}.$$

\mathcal{F} is a control funnel defined over \mathbb{R} : if $\mathbf{x}_\Delta(t) = \mathbf{x}(t) - \mathbf{x}_{\text{ref}}(t)$ is a solution of Eq. (5) such that $\mathbf{x}(t_1) \in \mathcal{F}_\alpha(t_1)$, then for any $t_2 > t_1$, since $V(\mathbf{x}_\Delta)$ only decreases, $V(\mathbf{x}_\Delta(t_2)) \leq V(\mathbf{x}_\Delta(t_1)) \leq \alpha$, and thus $\mathbf{x}(t_2) = \mathbf{x}_{\text{ref}}(t_2) + \mathbf{x}_\Delta(t_2) \in \mathcal{F}_\alpha(t_2)$.

$\mathcal{F}_\alpha(t)$ is a fixed d -dimensional ellipsoid centered at $\mathbf{x}_{\text{ref}}(t)$. Without going into details, it is possible to get lower bounds on the rate of decay of $V(\mathbf{x}_\Delta)$, and effectively compute $\beta > 0$ such that, for any solution $\mathbf{x}_\Delta(t)$ of Eq. (5):

$$\forall t \in \mathbb{R}, \forall \delta t \in \mathbb{R}^+, V(\mathbf{x}_\Delta(t + \delta t)) \leq e^{-\beta \cdot \delta t} V(\mathbf{x}_\Delta(t)).$$

This proves that if the system is inside the control funnel $\mathcal{F}_\alpha(t)$ at a given instant, then after letting time elapse for a duration of δt , the system will be inside the control funnel $\mathcal{F}_{\alpha e^{-\beta \cdot \delta t}}(t)$. Using the terminology of Section 2.3, this can be equivalently stated as follows: for $0 < \alpha' < \alpha$, the control funnel $\mathcal{F}_{\alpha'}(t)$ $\left[\frac{1}{\beta} \log\left(\frac{\alpha}{\alpha'}\right)\right]$ -absorbs the control funnel $\mathcal{F}_\alpha(t)$. Thanks to this property, for a given LQR controller and a reference trajectory $\mathbf{x}_{\text{ref}}(t)$, we can define a finite set of fixed-size control funnels $\mathcal{F}_{\alpha_0}(t), \mathcal{F}_{\alpha_1}(t), \dots, \mathcal{F}_{\alpha_q}(t)$, with $\alpha_0 > \alpha_1 > \dots > \alpha_q > 0$, and absorbing transitions between them in the corresponding timed automaton.

In the remainder of the article, we will only use this kind of fixed-size control funnels, which we call ‘‘LQR funnels’’. They are convenient because the larger ones can be used to ‘‘catch’’ other control funnels, and the smaller ones can easily be caught by other control funnels. Figure 3 depicts a typical sequence, where first a large control funnel (in green) catches the system, then after some time, an absorbing transition can be triggered, and finally, a new transition brings the system to a larger control funnel (in blue) on another trajectory. Besides that, testing for inclusion between fixed-size ellipsoids is easy, and therefore LQR funnels allow relatively efficient algorithms for the computation of the tuples needed for the timed-automaton reduction ((**switch**, $[\alpha, \beta]$, (i, j) , γ , (k, l)), (**invariant**, $S_{i,j}$, (i, j)), \dots , see Section 3).

It should be noted that the concepts of fixed-size control funnels and absorbing transitions, introduced here for linear systems, are also suitable for general




Fig. 3. An absorbing transition (in green) between two switching transitions.

nonlinear systems. Lyapunov functions in general, and quadratic ones in particular, can be computed via optimization, for example with Sum-of-Squares techniques as shown in [19]. By imposing specific constraints on the optimization, fixed-size control funnels with exponential convergence can be obtained inducing the same kind of absorbing transitions as introduced in the last paragraph.

5 Examples of application

5.1 Synchronization of sine waves

In this example, there is a unique reference trajectory: $x_{\text{ref}}(t) = \sin(\frac{2\pi}{\tau}t)$, for $t \in [0, \tau]$ and $\tau \in \mathbb{Q}$, and a unique LQR controller tracking this trajectory. We define two fixed size LQR funnels \mathcal{F}^1 (the large one) and \mathcal{F}^2 (the small one) defined over $[0, \tau]$ such that \mathcal{F}^2 γ -absorbs \mathcal{F}^1 for some $\gamma \in \mathbb{Q}$. The size of \mathcal{F}^1 is computed so that an upper bound on the acceleration is always ensured, as long as the state of the system remains inside the control funnel.

The set $\mathcal{F}^1(\tau/2)$ contains the smaller control funnel $\mathcal{F}^2(t)$ for a range of time values $[\alpha, \beta]$ for some $\alpha < \frac{\tau}{2} \in \mathbb{Q}$ and $\beta > \frac{\tau}{2} \in \mathbb{Q}$. This allows switching transitions from \mathcal{F}^2 to \mathcal{F}^1 with abrupt modifications of the controller clock c_t . Together with the absorbing transition and “cyclic transitions” that come from the equalities $\mathcal{F}^1(0) = \mathcal{F}^1(\tau)$ and $\mathcal{F}^2(0) = \mathcal{F}^2(\tau)$, it results in an abstraction by the timed automaton shown on the left side of Fig. 4. The goal is to synchronize the controlled signal to a fixed signal $\sin(\frac{2\pi}{\tau}t + \varphi_0)$. The phase φ_0 is initially unknown, which we model using an adversary: we use a new clock c'_t , and an opponent transition as in the timed automaton on the right side of Fig. 4.

With these two timed automata, we can use the tool Uppaal-Tiga to synthesize a controller that reacts to the choice of the adversary, and performs adequate switching transitions until $c_t = c'_t$. It is even possible to generate a strategy that guarantees that the synchronization can always be performed in a bounded amount of time. We show in Fig. 5 a trajectory generated by the synthesized reactive controller. In this example, the phase chosen by the adversary is such that it is best to accelerate the controlled signal. Therefore, the controller uses twice the switching transition from \mathcal{F}^2 to \mathcal{F}^1 with a reset of the controller clock from α to $\tau/2$ (① and ② in Fig. 5). Between these switching transitions, an absorbing transition is taken to go back to the control funnel \mathcal{F}^2 (Ⓐ in Fig. 5). After the




Fig. 4. On the left: the timed automaton for the controlled signal (the system). On the right: the timed automaton used to model the target signal with an initially unknown phase φ_0 . The opponent transition (dashed) is the one used to set φ_0 .

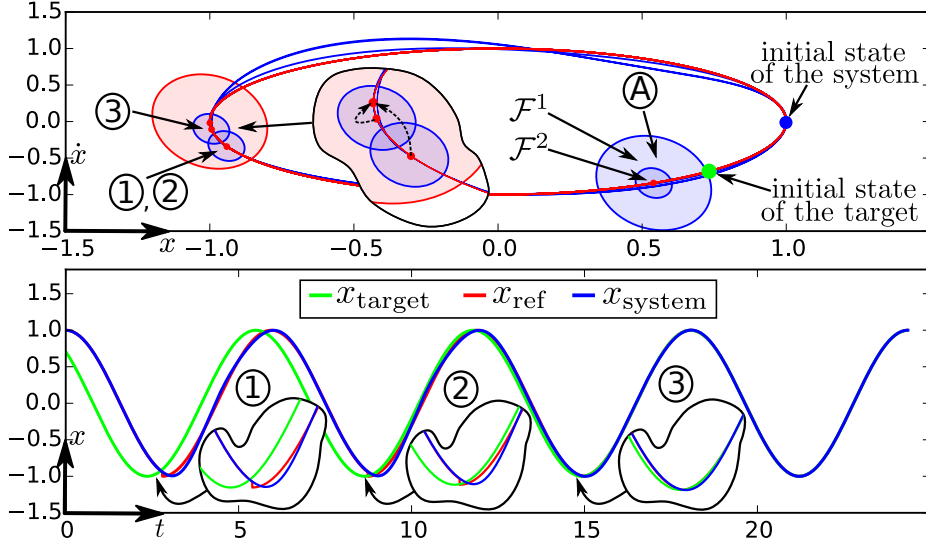


Fig. 5. The reactive controller performs three switching transitions to exactly adjust its signal to that of the target signal.

first two switching transitions, the remaining gap ε between c_t and c'_t is smaller than $\frac{\tau}{2} - \alpha$, and therefore the controller waits a bit longer (until $\frac{\tau}{2} - \varepsilon$) to perform the switching transition that exactly synchronizes the two signals ((3) in Fig. 5).

This example shows that our abstraction can be used for reactive controller synthesis via timed games. The main advantage of our approach over methods based on full discretization is that, since a continuous notion of time is kept in our abstraction, the reactive strategy is theoretically able to *exactly* synchronize the controlled signal to any real value of φ_0 . One of our hopes is that extensions of this result can lead to a general formal approach for signal processing.

5.2 A 1D pick-and-place problem

In this second example, we show that timed-automata abstractions based on control funnels can be used to perform non-trivial planning. We propose a





Fig. 6. The figure on the left shows the set-up. The black dots correspond to the position of the lanes. On the right are shown some LQR funnels along the constant velocity reference trajectories in the state space.

one-dimensional pick-and-place scenario. The set-up consists of a linear system controlled in acceleration moving along a straight line. On this line, four positions are defined as lanes (see Fig. 6). On three of these lanes (1, 2 and 3), packages arrive that have to be caught at the right time by the system and later delivered to lane 0. The system has limited acceleration and velocity, and can carry at most two packages at a time.

The LQR funnels in this example are constructed based on 12 reference trajectories. The first four have different constant positive velocities ($\mathbf{x}_{\text{ref}}^i$ with $i \in \{1, \dots, 4\}$, the fastest one being $\mathbf{x}_{\text{ref}}^4$, and the slowest one $\mathbf{x}_{\text{ref}}^1$). The next four are the same trajectories with negative velocities. On each of these reference trajectories, five different control funnels of constant size are defined (\mathcal{F}_i^j for $j \in \{0, \dots, 4\}$, the largest one being \mathcal{F}_i^0). The control funnels with negative constant velocity are the mirror image of those with positive velocity. Additionally, four stationary trajectories $\mathbf{x}_{\text{ref}}^{Lk}$ (with $k \in \{0, \dots, 3\}$) at the positions of the lanes are defined. The controllers associated to these trajectories simply stabilize the system at lane positions. For each of these trajectories a small ($j = 0$) and a large ($j = 1$) control funnel are constructed. They are denoted by \mathcal{F}_{Lk}^j . By construction, neighboring trajectories (e.g. $\mathbf{x}_{\text{ref}}^3$ and $\mathbf{x}_{\text{ref}}^2$ or $\mathbf{x}_{\text{ref}}^1$ and $\mathbf{x}_{\text{ref}}^{-1}$) are connected, meaning that for two neighboring trajectories $\mathbf{x}_{\text{ref}}^i$ and $\mathbf{x}_{\text{ref}}^k$, $\forall t \in I_i, \exists t' \in I_k$ s.t. $\mathcal{F}_i^4(t) \subset \mathcal{F}_k^0(t')$ (see Fig. 6). This allows the system to reach a higher or lower velocity without the need of an explicitly defined acceleration trajectory. While the abstraction based on these control funnels does not represent all the possible behaviors of the system (it is not complete), switching between different velocity references allows the system to perform a great variety of trajectories with continuous and bounded velocity and bounded acceleration.

To fully specify the timed-automata abstraction, the tuples defining the transition guards must be computed (see Section 3). Here, the regions of invariance defining the funnels are identically-shaped ellipses (only translated along a reference trajectory and scaled), thus the test for inclusion is computationally very cheap. Therefore, many points can be tested for inclusion on each trajectory, as depicted in Fig. 7, which leads to precise ranges for the switching transitions. Since the funnels are fixed sets translated along reference trajectories, knowing velocity or acceleration bounds on these references, and using offsets in the

Fig. 7. In order to define the tuples (switch, $[\alpha, \beta]$, (i, j) , γ , (k, l)) (see Section 3), N regularly spaced points are chosen in \mathbf{x}_{ref}^k (defining the ellipses $\mathcal{F}_k^l(t_n)$ for $n \in \{1, \dots, N\}$), and for each n , we set $\gamma = t_n$, and if a point $\mathbf{x}_{ref}^i(t)$ such that $\mathcal{F}_i^j(t) \subset \mathcal{F}_k^l(\gamma)$ is found, an incremental search is performed to define a range $[\alpha, \beta]$ such that $\forall t \in [\alpha, \beta]$, $\mathcal{F}_i^j(t) \subset \mathcal{F}_k^l(\gamma)$.



inclusion tests, we can ensure inclusion on the whole range of a switching transition with only a finite number of inclusion tests.

We consider an example where three packages respectively arrive on lanes 3, 2 and 1 at times $t_{arrive}^1 = 40$, $t_{arrive}^2 = 111$ and $t_{arrive}^3 = 122$ (corresponding equality tests on c_g can be used to refer to these moments in the timed automaton abstraction). The goal is to find a trajectory that catches all the packages and delivers them to lane 0. At the moment of the catch ($c_g = t_{arrive}^p$), the reference \mathbf{x}_{ref}^i tracked by the system must be exactly at the correct position (i.e. on the lane of the arriving package). Depending on the reference trajectory, this corresponds to a particular value of c_t . We add the following constraints on the catches: an upper bound on velocity such that the system cannot be tracking \mathbf{x}_{ref}^4 , \mathbf{x}_{ref}^3 , \mathbf{x}_{ref}^{-3} or \mathbf{x}_{ref}^{-4} when it catches a package, and a bound on uncertainty such that the system must be in a small control funnel to catch a package. Using additional constructions in our timed automaton abstraction (for example a bounded counter that keeps track of the number of packages being carried by the system), it is easy to specify these constraints and the objective as a reachability specification that can be checked by Uppaal. Uppaal outputs a timed word that corresponds to the schedule of control-law switches and the trajectory shown on Fig. 8, which successfully catches the packages and delivers them to lane 0.

The two upper graphs of Fig. 8 show the evolution of the system in its state space and some of the regions of invariance when taking a *switching transition* (colored ellipses). The green dots mark positions at which *absorbing transitions* take place ($\mathcal{F}_i^j \rightarrow \mathcal{F}_i^{j+1}$). Purple crosses represent a package. The lower graph compares the evolution of the position of the real system with the reference. One can see that even though the reference velocity can only take seven different values, a relatively smooth trajectory is realized. Before catching the first package, the system switches from \mathcal{F}_1^4 to \mathcal{F}_{L3}^0 ①. It then converges to \mathcal{F}_{L3}^1 ② just before the catch. The difference between the real system position and the reference is very small at that point in time. The system then switches to \mathcal{F}_{-1}^0 ③ in order to return to lane 0. It is interesting to notice that the system chooses to return to lane 0 after having picked only one package, therefore adopting a non-greedy strategy. This is because it wouldn't have time to perform a delivery to lane 0 between the arrival of the second and third packages.

When the second package arrives on lane 2, the system catches it while being in \mathcal{F}_{-1}^4 ④. This is again a non-trivial behavior: in order to get both the second and the third packages, the system has to first go a little bit further than




Fig. 8. Execution of a succeeding control strategy given as a timed word.

lane 2 so as to be able to catch the two packages without violating the limit on acceleration. A slight adjustment of the reference position ⑤ has to be done to catch the third package exactly on time ⑥. After that, the system performs a local acceleration ⑦ to reach lane 0 as soon as possible, and delivers the two packages.

6 Bounding funnels with conjectured properties for nonlinear systems

Many systems encountered in reality can be described as switched *nonlinear* systems. In this section, we propose a method to treat this class of systems, introducing the concept of *bounding funnels*, and using conjectured properties that are empirically verified. This approach is then used to solve a Reach-Avoid problem for a modified version of the Dubins' car, a nonlinear and non-holonomic system.

6.1 Introducing bounding funnels with conjectured properties

The main problem encountered when trying to construct control funnels for nonlinear systems, is the difficulty to design a control law that comes with a valid, monotonic Lyapunov function. There exist approaches for certain subclasses of nonlinear dynamics, like semidefinite programming for polynomial Lyapunov functions and systems with polynomial dynamics as done in [24]. In [19], it is shown how to use sum-of-squares optimization to handle nonlinear systems by using time-dependent polynomial approximations. It is an interesting approach, but its high computational complexity and the conservativeness introduced restrain its usability. We propose a different approach: bounding funnels with conjectured properties. Bounding funnels enlarge the concept of regular funnels by weakening some of the required assumptions. The properties of these funnels are as hard to guarantee as the properties of regular funnels, but due to the weakened assumptions they are more likely to be true. We propose to conjecture these properties based on numerical simulations. With these bounding funnels, the control sequence obtained is guaranteed to satisfy a given specification provided that the conjectures hold for the nonlinear dynamics under all circumstances that can occur.

Bounding funnels The concept of bounding funnels relies on a modified concept of positive invariance, which, together with the conservative approximation of convergence time, makes funnels suitable for timed automata reduction. The property of positive invariance described by Equation (3) (in Section 2) is closely linked to the concept of monotonic Lyapunov functions. For general nonlinear systems this property is very difficult to obtain. However, if a system converges asymptotically to the origin, it also eventually stays inside any neighborhood of the origin. Or, to put it differently, if $V^*(\mathbf{x}, t)$ is a Lyapunov function for the dynamical system $\dot{\mathbf{x}} = f(\mathbf{x}, t)$, then the system will also converge, possibly non-monotonically, with respect to every other Lyapunov function candidate $V'(\mathbf{x}, t)$. This property is very useful since it allows us to use functions with simple level sets, like ellipsoids, to construct our funnels, no matter what dynamical system is treated. For a bounding funnel $\mathcal{F}_i^j : I_i^j \subseteq \mathbb{R}^+ \rightarrow 2^{\mathbb{R}^d}$, the property of positive invariance is weakened in the sense that to each inner funnel \mathcal{F}_i^j we associate an outer funnel $\mathcal{F}_i^{O,j} : I \rightarrow 2^{\mathbb{R}^d}$ such that the following property holds:

$$\forall t_1 \in I_i^j, \forall t_2 \in I_i^j, (t_2 > t_1 \text{ and } \mathbf{x}(t_1) \in \mathcal{F}_i^j(t_1)) \Rightarrow \mathbf{x}(t_2) \in \mathcal{F}_i^{O,j}(t_2). \quad (7)$$

Informally, the outer funnel, for which $\forall t \in I_i^j, \mathcal{F}_i^j(t) \subseteq \mathcal{F}_i^{O,j}(t)$ holds, is chosen such that the trajectories of any initial position in \mathcal{F}_i^j will not leave $\mathcal{F}_i^{O,j}$. This modification is necessary due to the possibly non-monotonic convergence. Consequently, if the actual initial state of the system is inside $\mathcal{F}_i^j(t_0)$, the initial state of the timed automaton corresponds to the associated outer funnel $\mathcal{F}_i^{O,j}$. A *switching transition* (see Section 3) in a bounding-funnel system has the form: $(\mathcal{F}_i^j, v) \rightarrow (\mathcal{F}_k^{O,l}, v')$ whenever $v'(c_g) = v(c_g)$, $v'(c_h) = 0$, $v(c_t) \in I_i^j$, $v'(c_t) \in I_k^l$,

and $\mathcal{F}_i^j(v(c_t)) \subseteq \mathcal{F}_k^l(v'(c_t))$, where $\mathcal{F}_k^{O,l}$ denotes the bounding funnel associated with \mathcal{F}_k^l . In some cases (for example with fixed size inner and outer funnels), there exists a minimal duration that implies convergence from the outer to the inner funnel, i.e. a constant $h_i^{O,j \rightarrow j}$ such that \mathcal{F}_i^j $h_i^{O,j \rightarrow j}$ -absorbs $\mathcal{F}_i^{O,j}$.

To put the concept of bounding funnels in perspective, a regular funnel is a bounding funnel with $\mathcal{F}_i^{O,j}(t) = \mathcal{F}_i^j(t)$, $\forall t \in I_i^j$, and the absorption time $h_i^{O,j \rightarrow j}$ equal to zero.

Conjecturing the properties As stated above, proving the convergence and the weak positive invariance is a complex problem. Therefore we replace the formal guarantees by conjectures based on, for example, numerical simulations. This allows to use general optimization methods to simultaneously find a control law and suitable outer/inner funnel shapes in the sense that the outer funnel is as small as possible while achieving a good convergence time. To define the conjectures, sufficiently many initial points in \mathcal{F}_i^j can be numerically evaluated, and the convergence time $h_i^{j \rightarrow k}$ is defined as an upper bound of the maximal time needed to arrive and stay inside \mathcal{F}_i^k . The outer funnel can be taken as an ellipsoid with minimized volume under the constraint that (7) must hold.

This loss of guarantees may at first seem to be a very serious drawback, as obtaining certified behaviors is one of the main objectives of this work. Nevertheless, we argue that performing formal synthesis with such conjectured properties of the control laws can lead to interesting results. Indeed, after a controller has been synthesized with our approach, if an execution fails to verify the specification, we know that it can only be because at least one conjecture does not hold and therefore one or more properties of the bounding funnel are violated. We can even raise flags during execution to pinpoint the faulty bounding funnel or even the violated conjecture itself. This structure, where the logic of the controller is proven, but some "atomic" properties are only conjectured, is similar to formally verified cryptographic protocols, where the security depends on how reliable some cryptographic primitives are. It helps keeping safety issues localized, and therefore makes it easier to improve the global behavior with confidence by performing isolated tests of the validity of each funnel. Moreover, formally proven funnels are true funnels only in the mathematical model, and therefore, as far as runs on the real system are concerned, they are in fact conjectured as well.

6.2 Reach-Avoid problem for a modified Dubins' car

We use the above introduced bounding funnel concept to perform path planning for a modified Dubins' car. A Dubins' car is a simplified model of an automobile that evolves on a 2D plane, its state is defined by its position (denoted by \mathbf{p}) and its heading (denoted θ_p). The heading is given as the angle between the global \mathbf{x}_g -axis and the local \mathbf{x}_c -axis of the car. The current linear velocity of the car,




Fig. 9. Modified Dubins' car with controlled inputs v_p and ω_p , and reference trajectory (indexed by r).

denoted v_p , always points in the current direction of \mathbf{x}_c , so

$$\dot{\mathbf{p}} = \begin{pmatrix} \cos(\theta_p) \\ \sin(\theta_p) \end{pmatrix} v_p.$$

In this example we control directly the velocity v_p as well as the turning rate $\omega_p = \dot{\theta}_p$, but both control inputs must be continuous and bounded. The statespace of the system is the concatenation of its position with respect to global frame and the heading:

$$\begin{pmatrix} \mathbf{p} \\ \theta_p \end{pmatrix}.$$

The dynamics of the system is

$$\begin{pmatrix} \dot{\mathbf{p}} \\ \dot{\theta}_p \end{pmatrix} = \begin{pmatrix} v_p \cos(\theta_p) \\ v_p \sin(\theta_p) \\ \omega_p \end{pmatrix}.$$

We impose positive upper and lower bounds on the current velocity as well as bounds on the curvature of the resulting trajectory, so that the control law introduced afterwards always has to satisfy

$$0 < v_m \leq v_p \leq v_M \tag{8a}$$

$$-c_M \leq \omega_p/v_p \leq c_M. \tag{8b}$$

To create a (conjectured) funnel we must first define reference trajectories and a control law. For the reference trajectory we use a continuously differentiable curve defined on an interval $I \subseteq \mathbb{R}^+$ denoted by

$$\begin{pmatrix} \mathbf{r}(t) \\ \theta_r(t) \end{pmatrix}$$

satisfying the conditions² (8). In addition the curve has to be admissible, so it must hold that:

$$\forall t \in I : \dot{\mathbf{r}}(t) = \begin{pmatrix} \cos(\theta_r(t)) \\ \sin(\theta_r(t)) \end{pmatrix} v_r(t).$$

Every such curve can be used as a reference. The frame attached to the reference point is indexed by r . The angle between the global \mathbf{x}_g and the local \mathbf{x}_r axes (see Fig. 9) is denoted θ_r .

This nonlinear, non-holonomic dynamical system requires relatively complex control laws. Therefore we propose the following scheme:

$$\begin{pmatrix} v_p \\ \omega_p \end{pmatrix} = \begin{pmatrix} v_r \\ \omega_r \end{pmatrix} - \begin{pmatrix} \alpha \Delta x_c \\ \beta(\Delta\theta + \zeta \tanh(\gamma \Delta y_r)) \end{pmatrix} \quad (9)$$

where α , β and γ denote parameters in \mathbb{R}^+ , ζ is a parameter in $]0, \pi/2]$, Δx_c the projection of $\mathbf{p} - \mathbf{r}$ onto the \mathbf{x}_c -axis, Δy_r the projection of $\mathbf{p} - \mathbf{r}$ onto the \mathbf{y}_r -axis and $\Delta\theta = \theta_p - \theta_r$. The resulting values are then saturated to respect the constraints in (8) (for example if $v_p > v_M$, $v_p = v_M$; if $\omega_p/v_p > c_M$, $\omega_p = c_M v_p$).

In this control law the term $\zeta \tanh(\gamma \Delta y_r)$ is introduced to cope with the error in the orthogonal direction to the motion (Δy_c), which is not directly controllable due to the non-holonomy. We verify empirically the convergence properties of this control law: see Fig. 10.

For the bounding funnels we keep the ellipsoidal shaped funnels introduced in Sec. 4 and extend them with the introduction of outer funnels:

$$\mathcal{F}_i^j(t) = \left\{ \begin{pmatrix} \mathbf{r} \\ \theta_r \end{pmatrix} + \begin{pmatrix} \Delta \mathbf{r} \\ \Delta \theta_r \end{pmatrix} \mid V_i^j(\Delta \mathbf{r}, \Delta \theta_r) \leq \alpha_i^j \right\} \quad (10)$$

$$\mathcal{F}_{O,i}^j(t) = \left\{ \begin{pmatrix} \mathbf{r} \\ \theta_r \end{pmatrix} + \begin{pmatrix} \Delta \mathbf{r} \\ \Delta \theta_r \end{pmatrix} \mid V_i^j(\Delta \mathbf{r}, \Delta \theta_r) \leq \alpha_i^{O,j} \right\} \quad (11)$$

where $V_i^j(\Delta \mathbf{r}, \Delta \theta_r) = [\Delta \mathbf{r}^\top, \Delta \theta_r] \cdot P_i^j \cdot [\Delta \mathbf{r}^\top, \Delta \theta_r]^\top$ is a quadratic function defined by the symmetric and positive matrix P_i^j and $\alpha_i^{O,j} \geq \alpha_i^j \in \mathbb{R}^+$ are constants defining the size of the funnel. Note that in this example we have freely chosen the outer funnel to be a scaled version of the inner funnel, but it is perfectly possible to choose different shapes for P_i^j and $P_i^{O,j}$.

The objective is to find a timed sequence of transitions between reference trajectories that bring the system from an initial region $\Omega_0 = \mathcal{F}_{i_0}^{j_0}(t_0)$ to a final region $\Omega_1 = \mathcal{F}_{i_1}^{j_1}(t_1)$. Requiring the model checker to supply the fastest trace (i.e. a sequence with minimum time elapsed on the global clock c_g), we expect the kind of solutions as depicted in Fig. 11 for different sets of regions.

The reference trajectories used to construct the funnel system form a regular grid: The first layer is composed of $2N_D + 1$ trajectories with \mathbf{x}_r parallel to \mathbf{x}_g

² One should always choose the reference trajectory such that there exists a margin between the reference values and the limit imposed by (8) since otherwise the convergence is very slow.

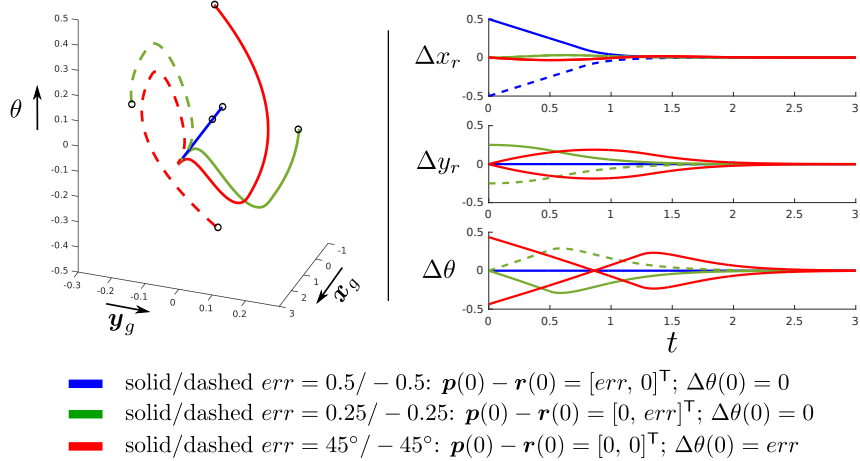


Fig. 10. Trajectories of the system for initial offsets in only one dimension and a reference trajectory of the form $\mathbf{r}(t) = [t, 0]^T$, $\theta_r(t) = 0$. An initial offset only in the \mathbf{x}_r -direction is corrected without inducing an error in the other components, since this direction is directly controllable. An initial error in the \mathbf{y}_r -direction induces an error in the heading in order to be corrected and vice versa.

of the form

$$-N_D \leq i \leq N_D : \mathbf{x}_r^i(t) = \begin{pmatrix} 0 \\ i \delta D \\ 0 \end{pmatrix} + \begin{pmatrix} tv_r - N_D \delta D \\ 0 \\ 0 \end{pmatrix}$$

defined on the interval $I^i = [0, (2N_D \delta D)/v_r]$. The other layers are formed by rotating the first layer around the θ axis, considering a 3D Cartesian representation of the statespace. We use N_A such layers, each of the trajectories having the form

$$-N_D \leq i \leq N_D, 1 \leq j \leq N_A : \mathbf{x}_r^{i,j}(t) = R_\theta(\alpha_j) \cdot \left(\begin{pmatrix} 0 \\ i \delta D \\ \alpha_j \end{pmatrix} + \begin{pmatrix} tv_r - N_D \delta D \\ 0 \\ 0 \end{pmatrix} \right)$$

with $\alpha_j = (2\pi j)/N_A$, $1 \leq j \leq N_A$ and $R_\theta(\alpha_j)$ denoting the rotation matrix corresponding to a rotation of angle α_j around the θ -axis

On each of these reference trajectories three funnels of different sizes are defined. The funnels defined on the reference $\mathbf{x}_r^{i,j}$ are denoted $\mathcal{F}_{i,j}^0$ (the 'large' funnel), $\mathcal{F}_{i,j}^1$ (the *small* funnel connecting to the layer $j+1$), $\mathcal{F}_{i,j}^2$ (the *small* funnel connecting to the layer $j-1$) and $\mathcal{F}_{i,j}^{O,x}$ (the associated outer funnels). Their respective sizes are chosen such that transitions exist from each small funnel $\mathcal{F}_{i,j}^{1/2}$ to its direct neighbors in the same layer $\mathcal{F}_{i\pm 1,j}^0$ (on the parallel reference trajectories) and to all the large funnels in the layer above $\mathcal{F}_{i,j}^1 \rightarrow \mathcal{F}_{k,j+1}^{O,0}$ and below $\mathcal{F}_{i,j}^2 \rightarrow \mathcal{F}_{k,j-1}^{O,0}$.




Fig. 11. Depiction of two instances of the problem and an optimal reference trajectory (solid black line). The dashed black line shows the qualitative evolution of Ω_0 . In Problem 2 a reference trajectory leading directly from Ω_0 to Ω_1 would violate the curvature or velocity bounds.

Fig. 12. Depiction of the first three layers of reference trajectories for $\theta_r^{i,j} = \alpha_j \in [0^\circ, 60^\circ, 120^\circ]$.

In this example, we chose $N_D = 12$ and $N_A = 6$ resulting in a total of $25 \cdot 7 \cdot 3 = 525$ funnels. Due to the symmetry of the funnel system the number of available transitions can be approximated: on any reference trajectory, only the small funnels have outgoing transitions. Each of the two small funnels $F_{i,j}^{1/2}$ is connected to the large funnels $F_{i\pm 1,j}^0$ on each of the 125 possible transition points. Furthermore there is an average of six transitions between a small funnel $F_{i,j}^{1/2}$ and any of the large funnels in the layer above ($F_{k,j+1}^0$) or below ($F_{k,j-1}^0$). So in total the automaton has $25 \cdot 7 \cdot 2(2 \cdot 125 + 25 \cdot 6) = 140.000$ transitions between 525 states.

This funnel system allows to conveniently switch the heading direction and specific funnels needed to attain a certain direction can easily be added.

As pointed out above, the convergence time is approximated using numerical simulations. In order to find a suitable ellipsoid and the corresponding control law parameters, the following optimization is performed: we fix *a priori* a diagonal matrix $D_L = \text{diag}([0.4^2 \ 0.4^2 \ (80\pi/180)^2])$ which is suited for $\alpha_{j+1} - \alpha_j = 60^\circ$. This diagonal matrix is rotated during optimization (parametrized via three Euler angles) in order to minimize the convergence time to the two small funnels

defined on the same trajectory. The small funnels $\mathcal{F}_{i,j}^{1/2}$ have the same shape as the funnels $\mathcal{F}_{i,j+1/-1}^0$, but are scaled by the factor $(80/15)^2$.

The optimization resulted in a minimal convergence time of 3.4 for the optimized funnel shape and control parameters as shown in Fig. 13.

In the two examples we consider that the initial region is centered around $\theta = 45^\circ$ and the desired final region is centered around $\theta = -45^\circ$. The decisive difference between the two problems is the distance (in \boldsymbol{x}_g -direction) between the regions as shown in Fig. 11.

The results obtained using the funnel system described above are shown in Fig. 14 and 15. The generated reference trajectories are qualitatively similar to the optimal ones shown in Fig. 11. The resulting system trajectories satisfy the specifications and are time-optimal (for the funnel system considered, not for the general case).

As shown in this example, bounding funnels (with conjectured properties) are a promising method to perform certified planning for general nonlinear systems. The advantage of this approach lies not only in the ability to treat nonlinear systems, but also in the possibility to adapt the funnel shape with respect to the needed convergence time without the additional constraint of monotonic convergence.

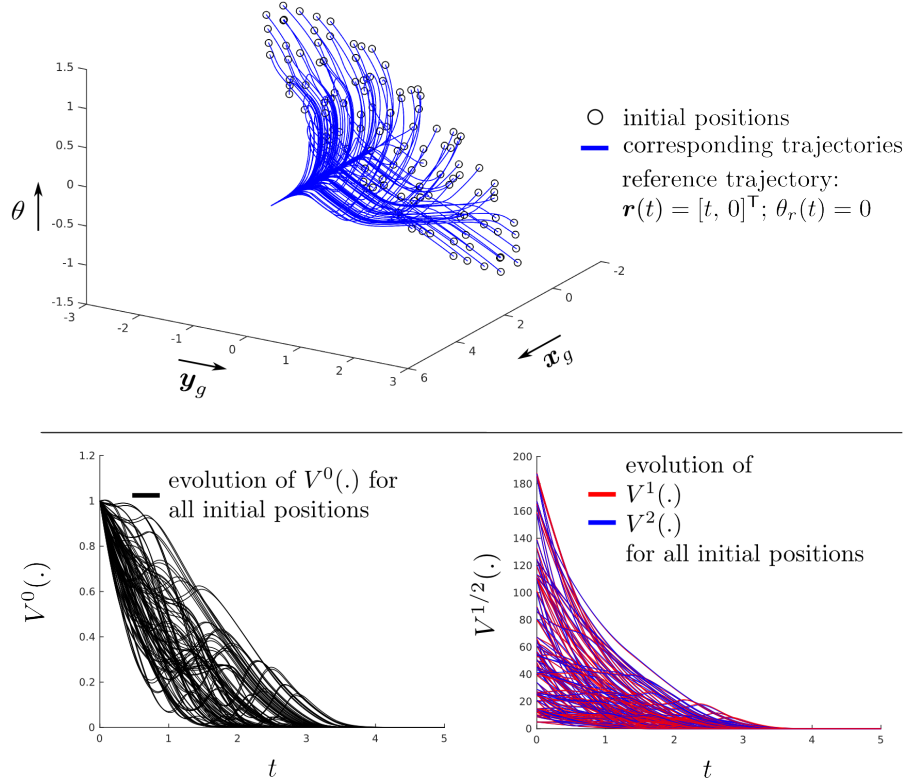


Fig. 13. On the left, the trajectories for initial states distributed on the surface of the optimized funnel shape \mathcal{F}^0 are shown. The control parameters are $\alpha = 4.43$, $\beta = 7.94$, $\gamma = 2.94$ and $\zeta = 4.57$. The dynamics induced by these parameters are denoted $f(\cdot)$. The second image depicts the evolution of $V^0(\Delta\mathbf{r}, \Delta\theta_r)$ with the large funnel \mathcal{F}^0 being defined as $V^0(\Delta\mathbf{r}, \Delta\theta_r) \leq \alpha_i^0 = 1.0$. The maximal value encountered is 1.0, so the associated outer funnel $\mathcal{F}^{0,0}$ can be chosen equal to \mathcal{F}^0 . Note that even-though $\mathcal{F}^0(t) = \mathcal{F}^{0,0}(t)$ the convergence is highly non-monotonic. The third image shows the evolution of V^1 (red) and V^2 (blue). After a time of 3.4 all states have converged to the small funnels \mathcal{F}^1 and \mathcal{F}^2 .




Fig. 14. Time-optimal solution found for problem 1. At every switching transition from $\mathcal{F}_{i,j}(\tau)$ to $\mathcal{F}_{k,l}^O(\tau')$, we distribute states over the surface of $\mathcal{F}_{i,j}(\tau)$ and show their trajectories $\mathbf{p}(t)$ until the next switching transition.




Fig. 15. Time-optimal solution found for problem 2.

7 Conclusion and future work

We have presented a timed-automata abstraction of switched dynamical systems based on control funnels, i.e. time-varying regions of invariance. Applying verification tools (such as Uppaal) on this abstraction, one can solve Reach-Avoid problems or more complicated problems with timing requirements. In the example of Section 5.2, we are able to generate a solution for a non-trivial pick-and-place problem. Using bounding funnels with conjectured properties, we extended our approach to treat nonlinear systems for which obtaining a formal certificate of invariance is beyond the state of the art. Synthesis of controllers that react to the environment can be done by solving timed games, and in the example of Section 5.1 we use Uppaal-Tiga to generate a controller that can reactively adjust the phase of a signal controlled in acceleration.

To go further, we could be more precise in our abstraction by extending timed automata with more features (we already mentioned non-deterministic clock updates in Section 3, Remark 1), and study the related decidability and algorithmic issues. We could also exploit the specific structure of the timed automata used in our abstraction and design dedicated verification and synthesis algorithms. Indeed, the timed automata of our model have three clocks, and there is non-determinism for only one of them (c_t). This makes us believe that we could potentially outperform the general algorithms of Uppaal and Uppaal-Tiga and solve more complex problems. Finally, in this quest to scale our approach up to larger models and more complicated system dynamics, bounding funnels and methods to obtain reasonable conjectures should be further investigated and exploited.

The abstraction based on conjectured properties makes it possible to increase the confidence in the global behavior via isolated testing on each of the control laws. Indeed, the switching behavior is proven correct as long as each individual conjecture holds. In fact, we believe that the use of conjectured properties could serve as an interface between existing numerical and optimization methods for dynamical systems and formal verification tools.

References

1. Alur, R., Dill, D.L.: A theory of timed automata. *Theor. Computer Science* 126(2), 183–235 (Apr 1994)
2. Asarin, E., Maler, O., Pnueli, A.: Reachability analysis of dynamical systems having piecewise-constant derivatives. *Theor. Computer Science* 138(1), 35–65 (Feb 1995)
3. Asarin, E., Maler, O., Pnueli, A., Sifakis, J.: Controller synthesis for timed automata. In: *SSSC’98*. pp. 469–474. Elsevier (1998)
4. Aubin, J.P.: Viability tubes. In: *Modelling and Adaptive Control*, pp. 27–47. Springer (1988)
5. Behrmann, G., Cougnard, A., David, A., Fleury, E., Larsen, K.G., Lime, D.: UPPAAL-Tiga: Time for playing games! In: *CAV’07*. LNCS, vol. 4590, pp. 121–125. Springer (Jul 2007)
6. Behrmann, G., David, A., Larsen, K.G., Håkansson, J., Pettersson, P., Yi, W., Hendriks, M.: Uppaal 4.0. In: *QEST’06*. pp. 125–126. IEEE (Sep 2006)

7. Bouyer, P., Dufourd, C., Fleury, E., Petit, A.: Updatable timed automata. *Theor. Computer Science* 321(2-3), 291–345 (Aug 2004)
8. Bouyer, P., Fahrenberg, U., Larsen, K.G., Markey, N.: Quantitative analysis of real-time systems using priced timed automata. *Communications of the ACM* 54(9), 78–87 (Sep 2011), <http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/BFLM-cacm11.pdf>
9. Bouyer, P., Markey, N., Perrin, N., Schlehber-Caissier, P.: Timed-automata abstraction of switched dynamical systems using control funnels. In: *FORMATS’15*. LNCS, vol. 9268, pp. 60–75. Springer (Sep 2015)
10. David, A., Grunnet, J.D., Jessen, J.J., Larsen, K.G., Rasmussen, J.I.: Application of model-checking technology to controller synthesis. In: *FMCO’10*. LNCS, vol. 6957, pp. 336–351. Springer (Nov 2012)
11. DeCastro, J., Kress-Gazit, H.: Synthesis of nonlinear continuous controllers for verifiably-correct high-level, reactive behaviors. *IJRR* 34(3), 378–394 (Mar 2014)
12. Duggirala, P.S., Mitra, S., Viswanathan, M.: Verification of annotated models from executions. In: *EMSOFT’13*. pp. 1–10. IEEE (Sep 2013)
13. Frazzoli, E., Dahleh, M.A., Feron, E.: Maneuver-based motion planning for nonlinear systems with symmetries. *IEEE Trans. Robotics* 21(6), 1077–1091 (Dec 2005)
14. Fu, J., Topcu, U.: Computational methods for stochastic control with metric interval temporal logic specifications. *Tech. Rep. 1503.07193*, ArXiv (Mar 2015)
15. Julius, A.A., Pappas, G.J.: Trajectory based verification using local finite-time invariance. In: *HSCC’09*. LNCS, vol. 5469, pp. 223–236. Springer (Apr 2009)
16. Koiran, P., Cosnard, M., Garzon, M.: Computability with low-dimensional dynamical systems. *Theor. Computer Science* 132(1), 113–128 (Sep 1994)
17. Le Ny, J.L., Pappas, G.J.: Sequential composition of robust controller specifications. In: *ICRA’12*. pp. 5190–5195. IEEE (May 2012)
18. Liu, J., Prabhakar, P.: Switching control of dynamical systems from metric temporal logic specifications. In: *ICRA’14*. pp. 5333–5338. IEEE (May 2014)
19. Majumdar, A., Ahmadi, A.A., Tedrake, R.: Control design along trajectories with sums of squares programming. In: *ICRA’13*. pp. 4054–4061. IEEE (May 2013)
20. Majumdar, A., Tedrake, R.: Robust online motion planning with regions of finite time invariance. In: *WAFR’12*. *STAR*, vol. 86, pp. 543–558. Springer (Jun 2013)
21. Maler, O., Batt, G.: Approximating continuous systems by timed automata. In: *FMSB’08*. LNBI, vol. 5054, pp. 77–89. Springer (Jun 2008)
22. Maler, O., Manna, Z., Pnueli, A.: From timed to hybrid systems. In: *Real-time: theory in practice*. LNCS, vol. 600, pp. 447–484. Springer (1992)
23. Mason, M.T.: The mechanics of manipulation. In: *ICRA’85*. vol. 2, pp. 544–548. IEEE (Mar 1985)
24. Prajna, S., Papachristodoulou, A., Wu, F.: Nonlinear control synthesis by sum of squares optimization: A Lyapunov-based approach. In: *Control Conference, 2004. 5th Asian*. vol. 1, pp. 157–165. IEEE (2004)
25. Quottrup, M.M., Bak, T., Zamanabadi, R.I.: Multi-robot planning : a timed automata approach. In: *ICRA’04*. vol. 5, pp. 4417–4422. IEEE (Apr 2004)
26. Sloth, C., Wisniewski, R.: Timed game abstraction of control systems. *Tech. Rep. 1012.5113*, ArXiv (Dec 2010)
27. Sloth, C., Wisniewski, R.: Complete abstractions of dynamical systems by timed automata. *Nonlinear Analysis: Hybrid Systems* 7(1), 80–100 (Feb 2013)
28. Sontag, E.D.: *Mathematical control theory: deterministic finite dimensional systems*. Springer (1998)
29. Tedrake, R., Manchester, I.R., Tobenkin, M., Roberts, J.W.: LQR-trees: Feedback motion planning via sums-of-squares verification. *IJRR* 29(8), 1038–1052 (Jul 2010)