

A Modular Framework for Motion Planning using Safe-by-Design Motion Primitives

Marijan Vukosavljev, Zachary Kroeze, Angela P. Schoellig, and Mireille E. Broucke

Abstract—We present a modular framework for solving a motion planning problem among a group of robots. The proposed framework utilizes a finite set of low level *motion primitives* to generate motions in a gridded workspace. The constraints on allowable sequences of motion primitives are formalized through a *maneuver automaton*. At the high level, a control policy determines which motion primitive is executed in each box of the gridded workspace. We state general conditions on motion primitives to obtain provably correct behavior so that a library of safe-by-design motion primitives can be designed. The overall framework yields a highly robust design by utilizing feedback strategies at both the low and high levels. We provide specific designs for motion primitives and control policies suitable for multi-robot motion planning; the modularity of our approach enables one to independently customize the designs of each of these components. Our approach is experimentally validated on a group of quadcopters.

Index Terms—Hybrid systems, motion primitives, planning and control for multiple mobile robots, aerial robotics

I. INTRODUCTION

This paper presents a modular, hierarchical framework for motion planning and control of robotic systems. While motion planning has received a great deal of attention by many researchers, because the problem is highly complex especially when there are several robotic agents working together in a cluttered environment, significant challenges remain. Hierarchy, in which the control design has several layers, is an architectural strategy to overcome this complexity. Almost all hierarchical frameworks for motion planning aim to balance flexibility in the control specification at the high level, guarantees on correctness and safety at the low level, and computational feasibility overall.

Historically motion planning was focused on high level planning algorithms, while suppressing details on the dynamic capabilities of the robots at the low level [19]. Taking full account of low level dynamics in combination with solving the high level planning problem can lead to a computationally intractable problem. Despite the wealth of available research [3], [9], [17], [28], computationally efficient solutions to the motion planning problem with tight integration of high and low levels are highly sought after.

We propose a modular hierarchical framework so that one can independently plug and play both low level controllers and

Marijan Vukosavljev, Zachary Kroeze, and Mireille E. Broucke are with the Dept. of Electrical and Computer Engineering, University of Toronto, Canada (e-mails: mario.vukosavljev@mail.utoronto.ca, zach.kroeze@mail.utoronto.ca, broucke@control.utoronto.ca). Angela P. Schoellig is with the University of Toronto Institute for Aerospace Studies (UTIAS), Canada (email: schoellig@utias.utoronto.ca). Supported by the Natural Sciences and Engineering Research Council of Canada (NSERC).

Manuscript received March 18, 2018; revisions received October 19, 2018 and March 21, 2019; accepted May 19, 2019.




Fig. 1. Crazyflie quadcopters navigate in a cluttered environment. Video results are available at <http://tiny.cc/modular-3alg>.

high level planning algorithms in order to realize a balance between flexibility at the high level, safety at the low level, and computational feasibility. To make a customizable approach feasible, we introduce three assumptions. First, the output space of the underlying dynamical system has translational symmetry, namely position invariance, a property satisfied by many robotic models [12]. Second, the output space is gridded uniformly into rectangular boxes. Finally, the control capabilities are discretized into a finite set of *motion primitives*, where the low level describes the implementation of the motion primitives while the high level selects the motion primitives. Together, these assumptions imply that motion primitives can be designed over a single box, so that they can then be reapplied to any other box.

Now we give an overview of the features and techniques we employ, and we highlight other frameworks that share those features. We provide general formulation of motion primitives for nonlinear systems so that they can be applied to multi-robot systems. We focus on reach-avoid specifications in a priori known environments, in which the system must reach a desired configuration in a safe manner [3], [9], [15], [20]. Reach-avoid offers a fairly rich behavior set so that, for instance, a fragment of linear temporal logic (LTL) can be encoded as a sequence of reach-avoid problems [33], as we also show in our applications.

As we have mentioned, we abstract the output space into rectangular regions [28] rather than more general polytopic regions [9], [11], [17], [20] in order to exploit symmetry. Motion primitives have been employed in various ways [15], [28] and we encode feasible sequences of motion primitives by a *maneuver automaton* [12]. In contrast to the motion primitive methods above, our implementation of the low-level control

design of motion primitives is based on *reach control theory* [4], [27], which provides a highly flexible and intuitive set of design tools that have two notable advantages over tracking: first, it is not necessary to find feasible open-loop trajectories to track; second, safety constraints on the system states during the execution and concatenation of motion primitives can be guaranteed by design. Finally, planning at the high level is based on standard *shortest path algorithms* [5], [19] applied to the graph arising from the synchronous product of the discrete part of the maneuver automaton and the graph arising from the output space partition. The high-level plan generates a *control policy*, which selects the motion primitives over the gridded output space. The modularity of our approach enables one to employ other closed-loop methods such as potential methods [9] or vector-field shaping [20] for low-level control design, and standard or customized graph search algorithms to generate a high-level plan.

There are three main contributions of this work. First, we provide the complete theoretical details on the requirements for the low-level control design and high-level plan, and show that these two levels operate consistently to solve the reach-avoid problem. Second, we formulate the parallel composition of maneuver automata in order to obtain correct-by-design motion primitives for a system composed of individual subsystems, such as in the case of multiple vehicles. Finally, the modularity and effectiveness of our framework is experimentally validated on a group of quadcopters in several illustrative scenarios. In particular, we feature a novel and versatile design of motion primitives based on double integrators and we show how the customizability of the high level plan generation can be used to easily trade-off solution quality with computational efficiency. This paper is an extension of our previous work [31], which now supplies all the theoretical details along with proofs on correctness, the parallel composition construction, additional approaches to generate control policies, and more elaborate experimental results.

The paper is organized as follows. In the next section we highlight our contributions relative to the literature. In Section III we present a formal problem statement. The modular framework is introduced in Section IV. We define the output transition system, the maneuver automaton, the product automaton, and the high level plan, each of which contribute to realizing a solution of the motion planning problem. In Section V, we prove that our overall methodology solves the motion planning problem. In Section VI we give the procedure for composing motion primitives. In Section VII we present specific motion primitives for a double integrator system. In Section VIII we consider several methods to generate high level plans, which are experimentally demonstrated on quadcopters. We conclude the paper in Section IX.

Notation. Let \mathbb{Z} denote the integers and \mathbb{R} denote the real numbers. Let $|\cdot|$ denote the cardinality of a set. If A is a set, we denote its power set as 2^A . If A and B are sets, let $A \setminus B$ denote the usual set difference. If there are n sets A_i , let $\prod_{i=1}^n A_i$ denote the usual cartesian product. Given a

function $f : A \rightarrow B$, the image of $A_1 \subset A$ under f and the preimage of $B_1 \subset B$ under f are defined the usual way, and are denoted as $f(A_1) \subset B$ and $f^{-1}(B_1) \subset A$, respectively. Let $\text{co}\{v_1, \dots, v_m\}$ denote the convex hull of the vectors $v_1, \dots, v_m \in \mathbb{R}^n$. Given two vectors $v, w \in \mathbb{R}^n$, we denote the component-wise multiplication (or Hadamard product) as $v \circ w$. Let $\mathcal{X}(\mathbb{R}^n)$ denote the set of globally Lipschitz vector fields on \mathbb{R}^n .

II. RELATED LITERATURE

The literature on motion planning is vast and encompasses many research communities. As such, we have categorized some common approaches and discussed how they relate to our method.

A. Graph Search and Trajectory Planning

Motion planning has often been addressed as a discrete planning problem, for which many standard graph search algorithms exist [19]. Recent work on the multi-agent reach-avoid problem has developed novel algorithms in the context of applications such as manufacturing and warehouse automation, aiming to balance computational efficiency with solution quality. For example, a centralized approach is given in [35], discretizing the workspace into a lattice and using integer linear programming to minimize the total time for robots to traverse in high densities. In [10], a sampling-based roadmap is constructed in the joint robot space using individual robot roadmaps, which is shown to be asymptotically optimal. Prioritized planning enables to safely coordinate many vehicles and is considered in a centralized and decentralized fashion in [6]. Subdimensional expansion computes mainly decentrally, but coordinates in the joint search space when agents are neighboring [32]. While such approaches typically provide various theoretical guarantees on the proposed algorithms, dynamical models and application on real robotic systems is often not considered.

The modularity of our framework is complementary, as it potentially enables existing multi-agent literature on gridded workspaces to be used directly or adapted for the generation of a high-level plan when used in conjunction with our proposed formulation of motion primitives. However, the consideration of continuous time dynamics may complicate the application of discrete planning methods in two ways. First, we must contend with constraints on successive motion primitives so that the continuous time behavior is acceptable - for example, avoiding abrupt changes in velocity. Second, we must contend with non-deterministic transitions to neighboring boxes, because motion primitives may allow more than one next box to be reached [18] - for example, modeling the joint asynchronous motion capabilities of a multi-robot system.

Trajectory tracking methods have also been applied to the formation change problem on real vehicles with complex dynamics. A sequential convex programming approach is given in [2], which computes discretized, non-colliding positional trajectories for a modest number of quadcopters.

More recently, an impressive number of quadcopters were coordinated in [25], by first computing a sequence of grid-based waypoints and then refining it into smoother piecewise polynomials. However, since these open-loop trajectories are computed offline, deviations from the computed trajectories could result in crashes. On the other hand, our approach is more robust as it is completely untimed, carefully monitoring the progress of vehicles over the grid in a reactive way based on the measured box transitions.

B. Formal Methods

A growing body of research has explored the use of formal methods in motion planning. This paper has been particularly inspired by [17], which provides a general framework for solving control problems for affine systems with LTL specifications. Their approach involves constructing a transition system over a polyhedral partition of the state space that arises from linear inequality constraints that constitute the atomic propositions of the LTL specification. Transitions between states of the transition system can occur if there exists an affine or piecewise affine feedback steering all continuous time trajectories from one polyhedral region to a contiguous one. Similar works to [17] include [9], [14], [20], which consider the simpler reach-avoid problem. Single and multi-robot applications followed shortly after in [11] and [3] respectively.

The appeal of these approaches is derived from their generality and faithful account of the low level system capabilities. On the downside, these methods generally do not scale well to larger state space dimensions, and so they would have limited applicability to large multi-robot systems. Our approach specializes these ideas by exploiting symmetry in the system dynamics and grid partition in order to strike a better balance between generality and computational efficiency. In particular, our feedback controllers are given as motion primitives, which can be designed independently of the obstacle and goal locations.

More recent works have also built on these formal method approaches, investigating more complex and realistic multi-robot problems. For example, service requests by multiple car robots in a city-like environment with communication constraints was considered in [7]. A cooperative task for ground vehicles was addressed in a distributed manner, enabling knowledge sharing amongst neighbors and reconfiguration of the motion plan in real time [13]. Tasks such as picking up objects are considered in conjunction with motion requirements in [26]. Since these works consider only fairly simple vehicle dynamics, they place greater emphasis on the synthesis of discrete plans satisfying the task specification. On the other hand, this paper considers the simpler reach-avoid problem in order to develop a formulation of motion primitives for nonlinear systems with symmetries.

C. Motion Primitives

The usage of motion primitives has become popular recently in robotics, as they serve to simplify the motion

planning problem by using predefined executable motion segments. Many variations exist, which have designed motion primitives using timed reference trajectories to control a formation of quadcopters [28], paths on a state space lattice for a mobile robots [8], [24], and funnels in the state space centered about a reference trajectory for a car [15] and a small airplane [22].

We have been inspired by ideas in [12], from which we borrowed the term “maneuver automaton”. They define a motion primitive either as an equivalence class of trajectories or a timed maneuver between two classes, whereas we define a motion primitive as a feedback controller over a polyhedral region in the state space. In our formulation, concatenations between motion primitives are possible only across contiguous boxes in the output space, which provides a strict safety guarantee during concatenation. Moreover, this enables our approach to simplify obstacle avoidance to a discrete planning problem over safe boxes as in [8], bypassing the need to concatenate motion primitive trajectories using numerical optimization techniques as in [12].

Our presentation of the maneuver automaton gives explicit constraints on the design of motion primitives so that they can be used reliably for high level planning. We have also introduced the notion of parallel composition of maneuver automata to build motion primitives for multi-robot systems. While our construction resembles existing methods of parallel composition [30], [34], we additionally prove that our construction preserves desired properties that enable consistency between the low and high levels. To the authors’ best knowledge, this paper is the first rigorous treatment of feedback-based motion primitives defined on a uniformly gridded output space.

III. PROBLEM STATEMENT

Consider the general nonlinear control system

$$\dot{x} = f(x, u), \quad y = h(x), \quad (1)$$

where $x \in \mathbb{R}^n$ is the state, $u \in \mathbb{R}^m$ is the input, and $y \in \mathbb{R}^p$ is the output. Let $\phi(\cdot, x_0)$ and $y(\cdot, x_0)$ denote the state and output trajectories of (1) starting at initial condition $x_0 \in \mathbb{R}^n$ and under some open-loop or feedback control.

Let $\mathcal{P} \subset \mathbb{R}^p$ be a feasible set in the output space and let $\mathcal{G} \subset \mathcal{P}$ be a goal set. In multi-vehicle motion planning contexts, \mathcal{P} represents the feasible joint output configurations of the system, which can arise from specifications involving obstacle avoidance, collision avoidance, communication constraints, and others. We consider the following problem.

Problem III.1 (Reach-Avoid). *We are given the system (1), a non-empty feasible set $\mathcal{P} \subset \mathbb{R}^p$ and a non-empty goal set $\mathcal{G} \subset \mathcal{P}$. Find a feedback control $u(x)$ and a set of initial conditions $\mathcal{X}_0 \subset \mathbb{R}^n$ such that for each $x_0 \in \mathcal{X}_0$ we have*

- (i) **Avoid:** $y(t, x_0) \notin \mathbb{R}^p \setminus \mathcal{P}$ for all $t \geq 0$,
- (ii) **Reach:** there exists $T \geq 0$ such that for all $t \geq T$, $y(t, x_0) \in \mathcal{G}$.




Fig. 2. Our modular framework consists of five modules.

We make an assumption regarding the outputs of the system (1) in order to exploit symmetry; see [12] for an exposition on nonlinear control systems with symmetries.

Assumption III.1. *First, we assume that there is an injective map $o : \{1, \dots, p\} \rightarrow \{1, \dots, n\}$ associating each output to a distinct state, so that $h(x) = (x_{o(1)}, \dots, x_{o(p)})$. We define the (injective) insertion map $h_o^{-1} : \mathbb{R}^p \rightarrow \mathbb{R}^n$ as $h_o^{-1}(y) = x$, which satisfies $h(x) = y$ and $x_i = 0$ for all $i \in \{1, \dots, n\} \setminus \{o(1), \dots, o(p)\}$. Second, we assume that the system has a translational invariance with respect to its outputs. That is, for all $x \in \mathbb{R}^n$, $u \in \mathbb{R}^m$ and $y \in \mathbb{R}^p$, we have $f(x, u) = f(x + h_o^{-1}(y), u)$. \triangleleft*

The assumption that the outputs of the system are a subset of the states is used in our framework to be able to design feedback controllers in the full state space that achieve desirable behavior in the output space. The second statement says that the vector field is invariant to the value of the output. In the literature this condition is called a *symmetry of the system* or *translational invariance*. This assumption is satisfied for many robotic systems, for example, when the outputs are positions. Also, we will see in Section VII that it significantly simplifies our control design.

IV. MODULAR FRAMEWORK

In this section we present our methodology to solve the motion planning problem in the form of an architecture that breaks down Problem III.1. This architecture consists of five main modules, as depicted in Figure 2.

- The *Problem Data* include the system (1) with p outputs satisfying Assumption III.1 and a reach-avoid task to be executed in the output space.
- The *Output Transition System (OTS)* is a directed graph whose nodes (called locations) represent p -dimensional boxes on a gridded output space and whose edges describe which boxes in the output space are contiguous.
- The *Maneuver Automaton (MA)* is a hybrid system whose modes correspond to so-called motion primitives. Each motion primitive is associated with a closed-loop vector field by applying a feedback law to (1). The edges of the MA represent feasible successive motion




Fig. 3. A two output ($p = 2$) example of a reach-avoid task. Shown on the left is the feasible space \mathcal{P} consisting of 15 non-obstacle boxes (not red) and the goal region \mathcal{G} (green). The output transition system (OTS), which abstracts the box regions and their neighbour connectivity, is shown on the right. Shown below, the possible offsets towards a neighbouring box are labelled using $\Sigma = \{-1, 0, 1\}^2$.

primitives. Each motion primitive generates some desired behavior of the output trajectories of the closed-loop system over a box in the output space. Because of the uniform gridding of the output space into boxes and because of the symmetry in the outputs described in Assumption III.1, motion primitives can be designed over only one canonical box Y^* .

- The *Product Automaton (PA)* is a graph which is the synchronous product of the OTS and the discrete part of the MA. It represents the combined constraints on feasible motions in the output space and feasible successive motion primitives.
- The *Hybrid Control Strategy* is a combination of low level controllers obtained from the design of motion primitives, and a high level plan on the product automaton.

Next we describe in greater detail the OTS, MA, and PA.

A. Output Transition System

The OTS provides an abstract description of the *workspace* or *output space* associated with the system (1). It serves to capture the feasible motions of output trajectories of the system (1) in a gridded output space, as in Figure 3. Specifically, we partition the output space into p -dimensional boxes with lengths $d = (d_1, \dots, d_p)$, where $d_i > 0$ is the length of i -th edge. We use a finite number of boxes to under-approximate the feasible set \mathcal{P} . Enumerating the boxes as $\{Y_1, \dots, Y_{n_L}\}$, the j -th box can be expressed in the form $Y_j := \prod_{i=1}^p [\eta_{ji} d_i, (\eta_{ji} + 1) d_i]$, where $\eta_{ji} \in \mathbb{Z}$, $i = 1, \dots, p$ are constants. We require that $\bigcup_{j=1}^{n_L} Y_j \subset \mathcal{P}$. Among these boxes, we assume there is a non-empty set of

indices $I_g \subset \{1, \dots, n_L\}$, so that we may similarly under-approximate the goal region as $\bigcup_{j \in I_g} Y_j \subset \mathcal{G} \subset \mathcal{P}$. We define a canonical p -dimensional box with edge lengths $d_i > 0$ given by $Y^* = \prod_{i=1}^p [0, d_i]$. Each box Y_j , $j = 1, \dots, n_L$ is a translation of Y^* by an amount $\eta_{ji}d_i$ along the i -th axis.

Definition IV.1. Given the lengths d and a non-empty goal index set I_g , an output transition system (OTS) is a tuple $\mathcal{A}_{\text{OTS}} = (L_{\text{OTS}}, \Sigma, E_{\text{OTS}}, L_{\text{OTS}}^g)$ with the following components:

State Space $L_{\text{OTS}} := \{l_1, \dots, l_{n_L}\} \subset \mathbb{Z}^p$ is a finite set of nodes called locations. Each location $l_j \in L_{\text{OTS}}$ is associated with a safe box $Y_j \subset \mathcal{P}$ in the output space and hence we write $l_j = (\eta_{j1}, \dots, \eta_{jp})$.

Labels $\Sigma := \{-1, 0, 1\}^p \subset \mathbb{Z}^p$ is a finite set of labels. A label $\sigma \in \Sigma$ is used to identify the offset between neighbouring boxes.

Edges $E_{\text{OTS}} \subset L_{\text{OTS}} \times \Sigma \times L_{\text{OTS}}$ is a set of directed edges where $(l_j, \sigma, l_{j'}) \in E_{\text{OTS}}$ if $j \neq j'$, $Y_j \cap Y_{j'} \neq \emptyset$, and $\sigma = l_{j'} - l_j \in \Sigma$. Thus, for each $i = 1, \dots, p$, the neighbouring box $l_{j'}$ is either one box to the left ($\sigma_i = -1$), the same box ($\sigma_i = 0$), or one box to the right ($\sigma_i = 1$). In this manner σ records the offset between contiguous boxes.

Final Condition $L_{\text{OTS}}^g = \{l_j \in L_{\text{OTS}} \mid j \in I_g\}$ denotes the set of locations associated with goal boxes. \triangleleft

Remark IV.1. We observe that the OTS is deterministic. That is, for a given $l \in L_{\text{OTS}}$ and $\sigma \in \Sigma$, there is at most one $l' \in L_{\text{OTS}}$ such that $(l, \sigma, l') \in E_{\text{OTS}}$. This follows immediately from the fact that $\sigma = l' - l$ records the offset between the neighbouring boxes.

Figure 3 shows a sample OTS for a simple scenario. The OTS locations are associated with 15 feasible boxes, including a goal box for the reach-avoid task. The OTS edges are shown as bidirectional arrows; for example, interpreting $l_1 = (0, 0)$ and $l_6 = (1, 1)$ on the grid, then $e = (l_6, (-1, -1), l_1) \in E_{\text{OTS}}$.

B. Maneuver Automaton

The maneuver automaton (MA) is a hybrid system consisting of a finite automaton and continuous time dynamics in each discrete state. The discrete states of the finite automaton correspond to *motion primitives*, while transitions between discrete states correspond to the allowable transitions between motion primitives. The continuous time dynamics are given by closed-loop vector fields (1) with a control law designed based on reach control theory (any other feedback control design method can be used).

Before presenting the MA, we first explain how this module is used in the overall framework. To solve Problem III.1, we assign motion primitives to the boxes Y_j of the partitioned output space such that obstacle regions are avoided and the goal region is eventually reached. The discrete part of the MA encodes the constraints on successive motion primitives. Such constraints might arise from a non-chattering requirement, continuity requirement, or requirement on correct switching between regions of the state space. A dynamic programming

algorithm for assignment of motion primitives on boxes is addressed in Section IV-D; the salient point about this algorithm at this stage is that it only uses the discrete part of the MA.

In contrast, the continuous time part of the MA is used both for simulation of the closed-loop dynamics to verify that the motion primitives are well designed, as well as for the implementation of the low level feedback in real-time. The motion primitives are defined only on the canonical box Y^* to simplify the design. This simplification is possible because of the translational symmetry provided by Assumption III.1 and the fact that each box Y_j is a translation of Y^* . In simulation, a given motion primitive can cause output trajectories to reach certain faces of Y^* . If a face is reached, the output trajectory is interpreted as being reset to the opposite face and another motion primitive is selected to be implemented over Y^* (of course, the real experimental output trajectories do not undergo resets but move continuously from box to box in the output space). The selection of the next motion primitive is constrained by a combination of the previous motion primitive and the face of Y^* that is reached. The discrete transitions in the MA encode these constraints.

Definition IV.2. Consider the system (1) satisfying Assumption III.1 and the box Y^* with lengths d . The maneuver automaton (MA) is a tuple $\mathcal{H}_{\text{MA}} = (Q_{\text{MA}}, \Sigma, E_{\text{MA}}, X_{\text{MA}}, I_{\text{MA}}, G_{\text{MA}}, R_{\text{MA}}, Q_{\text{MA}}^0)$, where

State Space $Q_{\text{MA}} = M \times \mathbb{R}^n$ is the hybrid state space, where $M = \{m_1, \dots, m_{n_M}\}$ is a finite set of nodes, each corresponding to a motion primitive.

Labels Σ , the same labels used in the OTS.

Edges $E_{\text{MA}} \subset M \times \Sigma \times M$ is a finite set of edges.

Vector Fields $X_{\text{MA}} : M \rightarrow \mathcal{X}(\mathbb{R}^n)$ is a function assigning a globally Lipschitz closed-loop vector field to each motion primitive $m \in M$. That is, for each $m \in M$, we have $X_{\text{MA}}(m) = f(\cdot, u_m(\cdot))$ where $u_m(\cdot)$ is a feedback controller associated with $m \in M$.

Invariants $I_{\text{MA}} : M \rightarrow 2^{\mathbb{R}^n}$ assigns a bounded invariant set $I_{\text{MA}}(m)$ to each $m \in M$. We impose that $I_{\text{MA}}(m) \subset h^{-1}(Y^*)$. The set $I_{\text{MA}}(m)$ defines the region on which the vector field $X_{\text{MA}}(m)$ is defined. Note that there is no requirement that the invariant is a closed set.

Enabling Conditions $G_{\text{MA}} : E_{\text{MA}} \rightarrow \{g_e\}_{e \in E_{\text{MA}}}$ assigns to each edge $e = (m, \sigma, m') \in E_{\text{MA}}$, a non-empty enabling or guard condition $g_e \subset \mathbb{R}^n$. We require that $g_e \subset I_{\text{MA}}(m)$. We make an additional requirement that g_e lies on a certain face of Y^* determined by the label $\sigma = (\sigma_1, \dots, \sigma_p) \in \Sigma$. Defining the face associated with σ as

$$\mathcal{F}_\sigma = \left\{ y \in Y^* \mid \begin{cases} y_i = 0, & \text{if } \sigma_i = -1 \\ y_i = d_i, & \text{if } \sigma_i = 1 \\ y_i \in [0, d_i], & \text{if } \sigma_i = 0 \end{cases} \right\},$$

we require that also $g_e \subset h^{-1}(\mathcal{F}_\sigma)$.

Reset Conditions $R_{\text{MA}} : E_{\text{MA}} \rightarrow \{r_e\}_{e \in E_{\text{MA}}}$ assigns to each edge $e = (m, \sigma, m') \in E_{\text{MA}}$ a reset map $r_e : \mathbb{R}^n \rightarrow \mathbb{R}^n$. We require that $r_e(x) = x - h_\sigma^{-1}(d \circ \sigma)$, where \circ is the Hadamard product. This definition says that the i -th output




Fig. 4. The maneuver automaton edges E_{MA} for the double integrator dynamics with $p = 1$. There are three motion primitives: Hold (\mathcal{H}), Forward (\mathcal{F}), and Backward (\mathcal{B}).




Fig. 5. The closed-loop vector fields in the (x_1, x_2) position-velocity state space for the Hold, Forward, and Backward motion primitives.

component is reset to the right face of Y^* , $x_{o(i)} = d_i$, if $\sigma_i = -1$, reset to the left face $x_{o(i)} = 0$ if $\sigma_i = 1$, and unchanged otherwise. Overall, resets of states are determined by the event $\sigma \in \Sigma$ and only affect the output coordinates in order to maintain output trajectories inside the canonical box Y^* .

Initial Conditions $Q_{\text{MA}}^0 \subset Q_{\text{MA}}$ is the set of initial conditions given by $Q_{\text{MA}}^0 = \{(m, x) \in Q_{\text{MA}} \mid x \in I_{\text{MA}}(m)\}$.

◁

Example IV.1. Suppose the system is a double integrator and the first state is the translationally invariant output y . The box Y^* is simply a segment. Let $M = \{\mathcal{H}, \mathcal{F}, \mathcal{B}\}$, where Hold (\mathcal{H}) stabilizes y , Forward (\mathcal{F}) increases y , and Backward (\mathcal{B}) decreases y . Referring to Figure 4, if \mathcal{F} is the current motion primitive and y reaches the right face of Y^* , then the event $1 \in \Sigma$ occurs and we may select \mathcal{H} or \mathcal{B} as the next motion primitive. To correctly implement the discrete evolution of the MA in the continuous state space, an invariant and feedback control must be associated with each motion primitive, while an enabling and reset condition must be associated with each edge; see Figure 5. Formal details are given in Section VII.

We now formulate assumptions on the motion primitives so that correct continuous time behavior is ensured at the low level for consistency with the high level. For each $m \in M$, define the set of possible events as

$$\Sigma_{\text{MA}}(m) := \{\sigma \in \Sigma \mid (\exists m' \in M)(m, \sigma, m') \in E_{\text{MA}}\}. \quad (2)$$

Assumption IV.1.

- (i) For all $m \in M$, $\varepsilon := (0, \dots, 0) \notin \Sigma_{\text{MA}}(m)$.
- (ii) For all $e_1, e_2 \in E_{\text{MA}}$ such that $e_1 = (m_1, \sigma, m_2)$ and $e_2 = (m_1, \sigma, m_3)$, $g_{e_1} = g_{e_2}$.
- (iii) For all $e_1, e_2 \in E_{\text{MA}}$ such that $e_1 = (m_1, \sigma_1, m_2)$ and $e_2 = (m_1, \sigma_2, m_3)$, if $\sigma_1 \neq \sigma_2$, then $g_{e_1} \cap g_{e_2} = \emptyset$.
- (iv) For all $e_1, e_2 \in E_{\text{MA}}$ such that $e_1 = (m_1, \sigma_1, m_2)$ and $e_2 = (m_2, \sigma_2, m_3)$, $r_{e_1}(g_{e_1}) \cap g_{e_2} = \emptyset$.
- (v) For all $e = (m_1, \sigma, m_2) \in E_{\text{MA}}$, $r_e(g_e) \subset I_{\text{MA}}(m_2)$.
- (vi) For all $m \in M$, if $\Sigma_{\text{MA}}(m) = \emptyset$ then for all $x_0 \in I_{\text{MA}}(m)$ and $t \geq 0$, $\phi_{\text{MA}}(t, x_0) \in I_{\text{MA}}(m)$.
- (vii) For all $m \in M$, if $\Sigma_{\text{MA}}(m) \neq \emptyset$, then for all $x_0 \in I_{\text{MA}}(m)$ there exist (a unique) $\sigma \in \Sigma_{\text{MA}}(m)$ and (a unique) $T \geq 0$ such that for all $e = (m, \sigma, m') \in E_{\text{MA}}$ and for all $t \in [0, T]$, $\phi_{\text{MA}}(t, x_0) \in I_{\text{MA}}(m)$ and $\phi_{\text{MA}}(T, x_0) \in g_e$.

◁

Condition (i) disallows tautological chattering behavior that arises by erroneously interpreting continuous evolution of trajectories in the interior of Y^* as “discrete transitions” of the MA (see Section V for definitions). Condition (ii) imposes that guard sets are independent of the next motion primitive. Since guard sets arise as the set of exit points of closed-loop trajectories from Y^* under a given motion primitive, it is reasonable that these exit points should depend only on the current motion primitive $m \in M$, and not on the choice of next motion primitive. Condition (iii) imposes that all guard sets corresponding to different labels are non-overlapping. This ensures that when the continuous trajectory reaches a guard g_e , then it is unambiguous which edge of the MA is taken next; namely $e \in E_{\text{MA}}$. Conditions (v), (vi), and (vii) are placed to guarantee that the MA is non-blocking. These conditions are based on known results in the literature [21]; see Lemma V.1. In order for condition (vii) to make sense, there must exist a unique label $\sigma \in \Sigma$ and a unique time $T \geq 0$ for an MA trajectory to reach a guard set. First, we have uniqueness of solutions since the vector fields are globally Lipschitz. Second, the unique MA trajectory can only reach one guard set by condition (iii); this in turn means there is a unique σ . Obviously there exists a unique time to reach the guard set. Conditions (vi) and (vii) work together to state that either all trajectories do not leave, or all trajectories do eventually leave. Referring to Figure 5, all closed-loop state trajectories within the invariant of \mathcal{F} reach the guard set shown in green on the right. For either choice of next feasible motion primitive, \mathcal{H} or \mathcal{B} , trajectories enter the next invariant on the left due to the reset. Finally, condition (iv) eliminates potential chattering Zeno behavior, see Remark V.1.

Remark IV.2. We make several further observations about the MA.

- (i) The MA is non-deterministic in the sense that given $m \in M$ and $\sigma \in \Sigma$, there may be multiple $m' \in M$ such that $(m, \sigma, m') \in E_{\text{MA}}$. The discrete part of the MA is non-deterministic in a second sense: for each $m \in M$, the cardinality of the set $\Sigma_{\text{MA}}(m)$ may be greater than one. The latter situation corresponds to the fact that for different initial conditions $x_1, x_2 \in I_{\text{MA}}(m)$ of the continuous part, the

associated output trajectories can reach different guard sets. In essence, which guard is enabled is interpreted, at the high level, as an uncontrollable event [34]. Remark IV.3 further illustrates these two types of non-determinism in the case of the PA.

(ii) The set of events Σ in the MA correspond to the same events Σ in the OTS. This correspondence is used in the product automaton PA, described in the next section, to synchronize transitions in the MA with transitions in the OTS. The interpretation is that when a continuous trajectory of the MA (over the box Y^*) undergoes a reset with the label $\sigma \in \Sigma$, the associated continuous trajectory of (1) in the box Y_j enters a neighboring box $Y_{j'}$ with the offset $\sigma = l_{j'} - l_j$. Obviously, this interpretation assumes that the vector of box lengths d is the same in both OTS and MA. \triangleleft

C. Product Automaton

In this section we introduce the *product automaton* (PA). It is constructed as the synchronous product of the OTS and the discrete part of the MA, namely (M, Σ, E_{MA}) . The purpose of the PA is to merge the constraints on successive motion primitives with the constraints on transitions in the OTS in order to enforce feasible and safe motions. As such, it captures the overall feasible motions of the system – any high level plan must adhere to these feasible motions.

Definition IV.3. We are given an OTS \mathcal{A}_{OTS} and an MA \mathcal{H}_{MA} satisfying Assumption IV.1. We define the product automaton (PA) to be the tuple $\mathcal{A}_{PA} = (Q_{PA}, \Sigma, E_{PA}, Q_{PA}^f)$, where

State Space $Q_{PA} \subset L_{OTS} \times M$ is a finite set of PA states. A PA state $q = (l, m) \in Q_{PA}$ satisfies the following: if there exists $\sigma \in \Sigma$ and $m' \in M$ such that $(m, \sigma, m') \in E_{MA}$, then there exists $l' \in L_{OTS}$ such that $(l, \sigma, l') \in E_{OTS}$. That is, $(l, m) \in Q_{PA}$ if all faces that can be reached by motion primitive $m \in M$ lead to a neighboring box of the box associated with location $l \in L_{OTS}$ of the OTS.

Labels Σ is the same set of labels used by the OTS and the MA.

Edges $E_{PA} \subset Q_{PA} \times \Sigma \times Q_{PA}$ is a set of directed edges defined according to the following rule. Let $q = (l, m) \in Q_{PA}$, $q' = (l', m') \in Q_{PA}$, and $\sigma \in \Sigma$. If $(l, \sigma, l') \in E_{OTS}$ and $(m, \sigma, m') \in E_{MA}$, then $(q, \sigma, q') \in E_{PA}$.

Final Condition $Q_{PA}^f \subset L_{OTS}^g \times M$ is the set of final PA states. \triangleleft

Remark IV.3. Formally an automaton is said to be non-deterministic if there exists a state with more than one outgoing edge with the same label. The PA is non-deterministic. First, consider a PA state $q = (l, m) \in Q_{PA}$. Because the MA allows for more than one feasible next motion primitive m' such that $(m, \sigma, m') \in E_{MA}$, the PA will also have multiple next PA states $q' = (l', m')$ such that $(q, \sigma, q') \in E_{PA}$. Second, there can be multiple possible labels $\sigma \in \Sigma$ such that $e = (q, \sigma, q') \in E_{PA}$ for some $q' \in Q_{PA}$. Thus, the PA inherits the two types of non-determinism of the MA that we discussed in Remark IV.2. For example, consider the PA




Fig. 6. A fragment of a generic PA, showing a state and its neighbours.

fragment in Figure 6. For the first type of non-determinism, observe that there are two PA edges $(q_1, \sigma_1, q_2) \in E_{PA}$ and $(q_1, \sigma_1, q_3) \in E_{PA}$ with the same label. For the second type, observe that there are two possible events $\sigma_1, \sigma_2 \in \Sigma$ from q_1 , each with its own set of PA edges. Note also some additional structure: since the OTS is deterministic, the box state is l_2 in both q_2 and q_3 , corresponding to the OTS edge $(l_1, \sigma_1, l_2) \in E_{OTS}$. \triangleleft

D. High-Level Plan

In this section we formulate the notion of a control policy on the PA, which gives a rule for selecting subsequent PA states by choosing the next motion primitive. Informally, the objective of the high level plan is to produce a control policy and find a set of initial PA states such that a goal PA state is eventually reached. To this end, in this section we also develop a Dynamic Programming Principle (DPP) suitable for use on the PA. Because of the two types of non-determinism of the PA, existing algorithms cannot be applied directly [5], [33]. By adapting the algorithm in [5], we obtain two formulations of the DPP, one of which is more computationally efficient as it exploits certain structure in the PA; further details are provided in Remark IV.5.

First some notation will be useful. Recall from (2), given $m \in M$, $\Sigma_{MA}(m)$ is the set of all labels $\sigma \in \Sigma$ on outgoing edges $e \in E_{MA}$ starting at m . Similarly, $\Sigma_{PA}(q)$ is the set of all labels $\sigma \in \Sigma$ on outgoing edges $e \in E_{PA}$ starting at q . That is,

$$\Sigma_{PA}(q) := \{\sigma \in \Sigma \mid (\exists q' \in Q_{PA})(q, \sigma, q') \in E_{PA}\}.$$

Now we formalize the semantics of the PA. A *state* of the PA is a pair $q = (l, m) \in Q_{PA}$ where $l \in L_{OTS}$ is a location in the OTS and $m \in M$ is a motion primitive. A *run* π of \mathcal{A}_{PA} is a finite or infinite sequence of states $\pi = q^0 q^1 q^2 \dots$, with $q^i = (l^i, m^i) \in Q_{PA}$ and for each i , there exists $\sigma^i \in \Sigma_{PA}(q^i)$ such that $(q^i, \sigma^i, q^{i+1}) \in E_{PA}$. We define the length of a run to be n_π ; for infinite runs n_π is defined to be ∞ . We consider a subset of runs $\Pi_{PA}(q)$ starting at $q \in Q_{PA}$ that satisfy one further property. If the run π is infinite, then $\pi \in \Pi_{PA}(q)$ if $q^0 = q$. Instead if the run π is finite, then $\pi \in \Pi_{PA}(q)$ if $q^0 = q$ and additionally, $\Sigma_{PA}(q^{n_\pi}) = \emptyset$. It is the latter requirement – that the last PA state of a finite run may not have outgoing edges in the PA – which is of interest. The interpretation is that we regard the event labels between PA states as uncontrollable, so if any event is possible, then it

must occur eventually. Thus without loss of generality, each run $\pi = q^0 q^1 \dots$ is the prefix of a run $\pi' \in \Pi_{\text{PA}}(q^0)$. Further elaboration is given in Remark IV.4 (ii).

Given $q \in Q_{\text{PA}}$ and $\sigma \in \Sigma_{\text{PA}}(q)$, the set of *admissible motion primitives* is

$$\mathcal{M}(q, \sigma) := \{m' \in M \mid (\exists q' = (l', m')) (q, \sigma, q') \in E_{\text{PA}}\}.$$

More generally, given $q \in Q_{\text{PA}}$ and $\Sigma_{\text{PA}}(q) = \{\sigma_1, \dots, \sigma_k\}$, the set of admissible motion primitives at q is

$$\mathcal{M}(q) := \{(m_1, \dots, m_k) \mid m_i \in \mathcal{M}(q, \sigma_i), i = 1, \dots, k\},$$

Next we introduce the notion of a control policy. Given $q \in Q_{\text{PA}}$ and $\Sigma_{\text{PA}}(q) = \{\sigma_1, \dots, \sigma_k\}$, an *admissible control assignment* at q is a vector

$$c(q) = (c(q, \sigma_1), \dots, c(q, \sigma_k)),$$

where $c(q, \sigma_i) \in \mathcal{M}(q, \sigma_i)$, or equivalently $c(q) \in \mathcal{M}(q)$. Notice that $c(q)$ is a vector whose dimension varies as a function of the cardinality of the set $\Sigma_{\text{PA}}(q)$. An *admissible control policy* $c : Q_{\text{PA}} \times \Sigma \rightarrow M$ is a map that assigns an admissible control assignment at each $q \in Q_{\text{PA}}$. Thus, for each $q \in Q_{\text{PA}}$ and $\sigma \in \Sigma_{\text{PA}}(q)$, $c(q, \sigma) \in \mathcal{M}(q, \sigma)$. The set of all admissible control policies is denoted by \mathcal{C} .

Consider an admissible control policy $c \in \mathcal{C}$ and a state $q \in Q_{\text{PA}}$. We denote the set of runs in $\Pi_{\text{PA}}(q)$ induced by c as $\Pi_c(q)$. Formally, $\pi = q^0 q^1 \dots \in \Pi_c(q)$ if $q^0 = q$, and for all $i \geq 0$ and $i < n_\pi$, $m^{i+1} = c(q^i, \sigma^i)$. Similarly, we denote the subset of runs in $\Pi_c(q)$ that eventually reach a state in Q_{PA}^f as $\Pi_c^f(q)$. Formally, $\pi \in \Pi_c^f(q)$ if there exists an integer $i \in \{0, \dots, n_\pi\}$ such that $q^i \in Q_{\text{PA}}^f$. For $\pi \in \Pi_c^f(q)$, we define

$$r_\pi := \min\{i \in \{0, \dots, n_\pi\} \mid q^i \in Q_{\text{PA}}^f\}.$$

Next we define an instantaneous cost $D_{\text{PA}} : E_{\text{PA}} \rightarrow \mathbb{R}$, which satisfies $D_{\text{PA}}(e) > 0$ for all $e \in E_{\text{PA}}$, and a terminal cost $H_{\text{PA}} : Q_{\text{PA}} \rightarrow \mathbb{R}$. Now consider the run $\pi = q^0 q^1 \dots q^{n_\pi} \in \Pi_c^f(q)$ with $q^0 = q$, $c(q^i, \sigma^i) = m^{i+1}$, and $e^i := (q^i, \sigma^i, q^{i+1}) \in E_{\text{PA}}$. We define a *cost-to-go* $J : Q_{\text{PA}} \times \mathcal{C} \rightarrow \mathbb{R}$ by

$$J(q, c) = \begin{cases} \max_{\pi \in \Pi_c(q)} \left\{ \sum_{j=0}^{r_\pi-1} D_{\text{PA}}(e^j) + H_{\text{PA}}(q^{r_\pi}) \right\}, & \Pi_c(q) = \Pi_c^f(q) \\ \infty, & \text{otherwise.} \end{cases}$$

Remark IV.4. There are several notable features of our formula for the cost-to-go.

(i) For a given $q \in Q_{\text{PA}}$, there may be multiple runs $\pi \in \Pi_c(q)$ due to the (second, non-standard type of) non-determinism of the PA. As such, we assume the worst case and take the maximum over $\Pi_c(q)$ in the cost-to-go. Moreover, we require $\Pi_c(q) = \Pi_c^f(q)$ for a finite cost-to-go so that r_π is well-defined and all runs starting at q eventually reach Q_{PA}^f .

(ii) We have assumed that finite runs must terminate on PA states that have no outgoing edges. Suppose we included in $\Pi_c(q)$ finite prefixes of (finite or infinite) runs. These necessarily would be finite runs with final PA states that have outgoing edges. Then if we take a finite or infinite run that eventually reaches a goal PA state, certain finite prefixes of




Fig. 7. At the top, a PA is depicted for a single output system having three motion primitives $M = \{\mathcal{H}, \mathcal{F}, \mathcal{B}\}$ over three boxes $L_{\text{OTS}} = \{l_j \mid j = 1, 2, 3\}$. The numbers $1, -1 \in \Sigma$ on the edges (shown as arrows) are the corresponding labels. The bottom pictures show the reduced set of transitions induced by control policies $c_1, c_2 \in \mathcal{C}$.

that run may not yet have reached a goal PA state, and we would get $\Pi_c(q) \neq \Pi_c^f(q)$ and an infinite cost-to-go. This anomaly arises from creating an artificial situation in which not all runs starting at an initial PA state reach a goal PA state because we included (unsuccessful) finite prefixes of successful runs.

(iii) The cost-to-go function also accounts for infinite runs by using the variable r_π to record the first time a goal PA state is reached and by taking the cost only over the associated prefix of the infinite run. Although our primary focus is on reach-avoid specifications, in which finite runs terminate on goal PA states with no outgoing edges, infinite runs allow us to extend our framework to a fragment of LTL where, for example, a goal PA state is reached always eventually; see Remark V.3 for further details.

Example IV.2. Consider the PA shown at the top of Figure 7 corresponding to a single output system with the three motion primitives $M = \{\mathcal{H}, \mathcal{F}, \mathcal{B}\}$ from Example IV.1 over three boxes $L_{\text{OTS}} = \{l_j \mid j = 1, 2, 3\}$. Suppose that $D_{\text{PA}}(e) = 1$ for all $e \in E_{\text{PA}}$ and that $H_{\text{PA}} = 0$ for all $q \in Q_{\text{PA}}$.

First consider the feasible control policy $c_1 \in \mathcal{C}$ with the control assignments: $c_1(q_1, 1) = \mathcal{F}$, $c_1(q_3, 1) = \mathcal{B}$, $c_1(q_5, -1) = \mathcal{F}$, and $c_1(q_7, -1) = \mathcal{B}$. The bottom left of Figure 7 shows how the control policy trims away possible edges in the PA. Now suppose that $Q_{\text{PA}}^f = \{q_7\}$. Choosing the initial condition $q_1 \in Q_{\text{PA}}$ and under the assumption that we do not include finite runs that terminate at PA states with outgoing edges, we can see that $\Pi_{c_1}(q_1)$ consists of only the single infinite run $\pi = q_1 q_3 q_7 q_5 q_1 \dots$. Even though this run is infinite, $\pi \in \Pi_{c_1}^f(q_1)$, $r_\pi = 2$, and $J(q_1, c_1) = 2$. Similarly, we compute $J(q_5, c_1) = 3$, $J(q_3, c_1) = 1$, $J(q_7, c_1) = 0$, and $J(q_2, c_1) = J(q_4, c_1) = J(q_6, c_1) = \infty$. In contrast, the feasible control policy $c_2 \in \mathcal{C}$ shown on the bottom right of Figure 7 only contains finite runs. \triangleleft

Next we define the *value function* $V : Q_{\text{PA}} \rightarrow \mathbb{R}$ to be

$$V(q) := \min_{c \in \mathcal{C}} J(q, c).$$

The value function satisfies a dynamic programming principle (DPP) that takes into account the non-determinacy of \mathcal{A}_{PA} ; see [5] where a slightly different result is proved. The proof is found in the appendix.

Theorem IV.1. *Consider $q \in Q_{\text{PA}} \setminus Q_{\text{PA}}^f$ and suppose $|\Sigma_{\text{PA}}(q)| > 0$. Then V satisfies*

$$V(q) = \min_{c(q) \in \mathcal{M}(q)} \left\{ \max_{\sigma \in \Sigma_{\text{PA}}(q)} \{D_{\text{PA}}(e) + V(q')\} \right\} \quad (3)$$

$$= \max_{\sigma \in \Sigma_{\text{PA}}(q)} \left\{ \min_{\bar{m} \in \mathcal{M}(q, \sigma)} \{D_{\text{PA}}(\bar{e}) + V(\bar{q})\} \right\}, \quad (4)$$

where $q' = (l', c(q, \sigma)) \in Q_{\text{PA}}$, $e = (q, \sigma, q') \in E_{\text{PA}}$, $\bar{q} = (\bar{l}, \bar{m}) \in Q_{\text{PA}}$, and $\bar{e} = (q, \sigma, \bar{q}) \in E_{\text{PA}}$.

Notice that for all $q \in Q_{\text{PA}} \setminus Q_{\text{PA}}^f$ such that $|\Sigma_{\text{PA}}(q)| = 0$, $V(q) = \infty$ (since there can be no paths to the goal). Also, for all $q \in Q_{\text{PA}}^f$, $V(q) = H_{\text{PA}}(q)$.

Remark IV.5. *In (3) of Theorem IV.1, it is shown that $V(q)$ can be computed using the local information of $\mathcal{M}(q)$ instead of using all of \mathcal{C} . In (4), the result is taken one step further by showing that $V(q)$ can be calculated using only $\mathcal{M}(q, \sigma)$ for each $\sigma \in \Sigma_{\text{PA}}(q)$. The benefit of (4) becomes clear when we compare the cardinality of the sets over which the minimizations occur. Given $q \in Q_{\text{PA}}$, let $\Sigma_{\text{PA}}(q) = \{\sigma_1, \dots, \sigma_k\}$. In (3) the minimization is over $\mathcal{M}(q)$, and therefore the cardinality of the minimization set is $\prod_{i=1}^k |\mathcal{M}(q, \sigma_k)|$. In (4) the minimization is over $\mathcal{M}(q, \sigma)$ for each $\sigma \in \Sigma_{\text{PA}}(q)$, and therefore the cardinality of the set is $|\mathcal{M}(q, \sigma)|$. While both (3) and (4) can be used to compute $V(q)$, in general (4) will be more computationally efficient.*

Corollary IV.1. *Consider the control policy c^* such that for all $q \in Q_{\text{PA}}$, and $\sigma \in \Sigma_{\text{PA}}(q)$*

$$c^*(q, \sigma) \in \operatorname{argmin}_{m' \in \mathcal{M}(q, \sigma)} \{D_{\text{PA}}(e) + V(q')\},$$

where $q' = (l', m')$, and $e = (q, \sigma, q')$. Then c^* is an optimal control policy such that for all $q \in Q_{\text{PA}}$, $V(q) = J(q, c^*)$.

Figure 8 shows a possible control policy for the scenario in Figure 3. Since there are two outputs, we use the motion primitives from Example IV.1 in each output; formal details are given in Section VI. The control policy was hand-computed. Notice that different routes may be taken from the same product state depending on the face reached, but ultimately the control policy ensures that all paths lead to the goal.

V. MAIN RESULTS

In this section we present our main results on a solution to Problem III.1. Our final result combines the notion of a control policy at the high level with feedback controllers executing correct continuous time behavior at the low level.




Fig. 8. This figure shows a discrete control strategy for the scenario shown in Figure 3.

First, in accordance with the reach-avoid objective (see Remark IV.4 (iii)), we assume the existence of motion primitives that can stabilize trajectories within a given box, that is, there exists $m \in M$ such that $\Sigma_{\text{MA}}(m) = \emptyset$. We restrict the final PA states to be goal OTS states equipped with such motion primitives

$$Q_{\text{PA}}^f = \{(l, m) \in L_{\text{OTS}}^g \times M \mid \Sigma_{\text{MA}}(m) = \emptyset\}. \quad (5)$$

Now suppose we have an admissible control policy $c \in \mathcal{C}$ derived using Theorem IV.1 or otherwise with Q_{PA}^f as above. We present a complete solution to Problem III.1 including an initial condition set $\mathcal{X}_0 \subset \mathbb{R}^n$, a feedback control $u(x)$, and conditions on the motion primitives so that the reach-avoid specifications of Problem III.1 are met.

First we specify the initial condition set \mathcal{X}_0 . The set of feasible initial PA states is

$$Q_{\text{PA}}^0 := \{q \in Q_{\text{PA}} \mid \Pi_c(q) = \Pi_c^f(q)\}. \quad (6)$$

That is, a feasible initial PA state satisfies that every run (induced by the control policy) starting at the PA state eventually reaches a goal PA state.

Now consider a state $x_0 \in \mathbb{R}^n$. It can be used as an initial state of the system if there is some $(l_j, m) \in Q_{\text{PA}}^0$ for which the state is both in the box Y_j and in the invariant of m . Recall that for all $y \in \mathbb{R}^p$ and $j \in \{1, \dots, n_L\}$, $y \in Y^*$ if and only if $y + d \circ l_j \in Y_j$. With this in mind, we define the set of initial states to be:

$$\mathcal{X}_0 = \bigcup_{(l_j, m) \in Q_{\text{PA}}^0} \{x + h_o^{-1}(d \circ l_j) \mid x \in I_{\text{MA}}(m)\}. \quad (7)$$

Next we specify the feedback controllers to solve Problem III.1. Consider any $q = (l_j, m) \in Q_{\text{PA}}^0$. Then for all $x \in \mathbb{R}^n$ such that $x - h_o^{-1}(d \circ l_j) \in I_{\text{MA}}(m)$, we define the feedback

$$u(x, q) := u_m(x - h_o^{-1}(d \circ l_j)). \quad (8)$$

This defines a family of feedback controllers parametrized by x , the state of (1) and by the PA state $q = (l_j, m)$. These feedbacks work in tandem with the control policy $c \in \mathcal{C}$, which effectively determines the next feasible PA state $q' \in Q_{\text{PA}}^0$. For example, suppose $q = (l_j, m) \in Q_{\text{PA}}^0$ and suppose

the label $\sigma \in \Sigma$ is measured. This event corresponds to $x \in g_e$ for $e = (m, \sigma, c(q, \sigma)) \in E_{\text{MA}}$. Let $m' := c(q, \sigma)$ and let $l' \in L_{\text{OTS}}$ be the unique location of the OTS such that $(l, \sigma, l') \in E_{\text{OTS}}$. Then the next PA state is $q' = (l', m') \in Q_{\text{PA}}^0$ and the controller that is applied in the next location $l' \in L_{\text{OTS}}$ is $u_{m'}(\cdot)$.

The main result of the paper is the following.

Theorem V.1. *Consider the system (1) satisfying Assumption III.1, the non-empty feasible set $\mathcal{P} \subset \mathbb{R}^p$, and the goal set $\mathcal{G} \subset \mathcal{P}$. Let d be the vector of box lengths such that the goal indices I^g is non-empty. Consider an associated OTS \mathcal{A}_{OTS} , an MA \mathcal{H}_{MA} satisfying Assumption IV.1, a PA \mathcal{A}_{PA} with Q_{PA}^f as in (5), and an admissible control policy $c \in \mathcal{C}$. Then the initial condition set \mathcal{X}_0 given in (7) and the feedback controllers (8) solve Problem III.1.*

In the remainder of this section we prove Theorem V.1. We now give a roadmap for these results. The verification of correctness at the low level is broken down into two steps that we now describe. First, we show that the MA is non-blocking in Lemma V.1. The key requirements are summarized in Assumption IV.1. The non-blocking condition ensures that MA trajectories continually evolve in time and stay within the invariant regions. We also put conditions to avoid chattering in which two discrete transitions can occur in immediate succession. While physical systems never undergo infinite switching in finite time, if our model predictions diverge from reality, then we have no grounds to claim that Problem III.1 is indeed solved. Second, in Lemma V.2 we show that to each closed-loop trajectory of (1) under the feedback controllers (8) and a control policy $c \in \mathcal{C}$, we can associate a unique execution of the MA (defined below) and run of the PA.

We begin by describing the semantics of the MA. These definitions are standard; see [21]. A *state* of the MA is a pair (m, x) , where $m \in M$ and $x \in \mathbb{R}^n$. Trajectories of the MA are called *executions* and are defined over hybrid time domains that identify the time intervals when the trajectory of a hybrid system is in a fixed motion primitive $m \in M$. Precisely, a *hybrid time domain* of the MA is a finite or infinite sequence of intervals $\tau = \{\mathcal{I}_0, \dots, \mathcal{I}_{n_\tau}\}$, such that

- (i) $\mathcal{I}_i = [\tau_i, \tau'_i]$, for all $0 \leq i < n_\tau$,
- (ii) if $n_\tau < \infty$, then either $\mathcal{I}_{n_\tau} = [\tau_{n_\tau}, \tau'_{n_\tau}]$ or $\mathcal{I}_{n_\tau} = [\tau_{n_\tau}, \tau'_{n_\tau})$,
- (iii) $\tau_i \leq \tau'_i = \tau_{i+1}$, for all $0 \leq i < n_\tau$.

Definition V.1. *An execution of the MA is a collection $\chi = (\tau, m(\cdot), \phi_{\text{MA}}(\cdot, x_0))$ such that*

- (i) *the initial condition of the execution satisfies: $(m(0), x_0) \in Q_{\text{MA}}^0$.*
- (ii) *the continuous evolution of the execution satisfies: for all $i \in \{0, \dots, n_\tau\}$ with $\tau_i < \tau'_i$, then for all $t \in [\tau_i, \tau'_i]$, $m(\cdot)$ is constant and $\frac{d}{dt}\phi_{\text{MA}}(t, x_0) = f(\phi_{\text{MA}}(t, x_0), u_{m(t)}(\phi_{\text{MA}}(t, x_0)))$, while for all $t \in [\tau_i, \tau'_i)$, $\phi_{\text{MA}}(t, x_0) \in I_{\text{MA}}(m(t))$.*
- (iii) *a discrete transition of the execution satisfies: for all $i \in \{0, \dots, n_\tau - 1\}$, there exists $\sigma_i \in \Sigma_{\text{MA}}(m(\tau'_i))$ such that*

$$(m(\tau'_i), \sigma_i, m(\tau_{i+1})) =: e_i \in E_{\text{MA}}, \phi_{\text{MA}}(\tau'_i, x_0) \in g_{e_i}, \text{ and } \phi_{\text{MA}}(\tau_{i+1}, x_0) = r_{e_i}(\phi_{\text{MA}}(\tau'_i, x_0)).$$

Given an execution $\chi = (\tau, m(\cdot), \phi_{\text{MA}}(\cdot, x_0))$, we associate to it the *output trajectory of the MA* given by $y_{\text{MA}}(\cdot, x_0) := h(\phi_{\text{MA}}(\cdot, x_0))$ (the subscript MA is included to avoid confusion with output trajectories $y(\cdot, x_0)$ of the physical system (1) which do not undergo resets). The *execution time* of an execution χ is defined as $\mathcal{T}(\chi) := \sum_{i=0}^{n_\tau} (\tau'_i - \tau_i) = \lim_{i \rightarrow n_\tau} \tau'_i$. An execution is called *finite* if τ is a finite sequence ending with a compact time interval. An execution is called *infinite* if either τ is an infinite sequence or if $\mathcal{T}(\chi) = \infty$. Finally, an execution is called *Zeno* if it is infinite but $\mathcal{T}(\chi) < \infty$.

Remark V.1. *There are two types of Zeno behavior. In one type that we call chattering, transitions are instantaneous. The second more subtle type is when the times between discrete transitions of the MA converge to zero, but the transitions are not instantaneous. Assumptions IV.1 (i) and (iv) ensure that we cannot have chattering. True Zeno behavior with convergent transition times is more difficult to identify in the setting when the MA is formed as a parallel composition. Fortunately, for our reach-avoid objective, the induced MA executions cannot be Zeno since there are a finite number of transitions by construction, see Lemma V.2.*

Definition V.2. *The MA is non-blocking if for all $(m(0), x_0) \in Q_{\text{MA}}^0$, the set of all infinite executions of the MA with initial condition $(m(0), x_0)$ is non-empty.*

Lemma V.1. *Under Assumption IV.1, the MA is non-blocking.*

Proof. Let $(m, x) \in Q_{\text{MA}}^0$. If $\Sigma_{\text{MA}}(m) = \emptyset$, then by Assumption IV.1 (vi), $I_{\text{MA}}(m)$ is invariant, so the trajectory $\phi_{\text{MA}}(t, x)$ starting at (m, x) remains in $I_{\text{MA}}(m)$ for all future time. Therefore, trivially, the MA is non-blocking for this initial condition. If $\Sigma_{\text{MA}}(m) \neq \emptyset$, then by Assumption IV.1 (vii), $\phi_{\text{MA}}(t, x)$ remains in $I_{\text{MA}}(m)$ until it reaches a guard set. Additionally, by Assumption IV.1 (v), the trajectory is mapped under the reset into the next invariant. By Lemma 1 of [21], the MA is again non-blocking for this initial condition. Overall, the MA is non-blocking. \square

The purpose of the Assumptions IV.1 is to guarantee consistency between low level continuous time behavior and the high level discrete plan. This consistency is formalized by way of a one-to-one correspondence between infinite MA executions and finite PA runs, both starting from the same initial condition. The proof is found in the appendix.

Lemma V.2. *Suppose we have an admissible control policy $c \in \mathcal{C}$, and we have an MA satisfying Assumption IV.1. For each $(l^0, m^0) \in Q_{\text{PA}}^0$ and $x_0 \in I_{\text{MA}}(m^0)$ there exist a unique infinite MA execution $\chi = (\tau, m(\cdot), \phi_{\text{MA}}(\cdot, x_0))$ and a unique finite PA run $\pi = q^0 q^1 \dots q^N$.*

Before we can prove Theorem V.1 we need one further preliminary result stating that because of the translational invariance of Assumption III.1, the continuous part of an MA execution has a unique correspondence to a closed-loop

trajectory of the system (1). The proof is straightforward and is omitted.

Lemma V.3. *Let $m \in M$, $x_0 \in I_{\text{MA}}(m)$, $y \in \mathbb{R}^p$, and $\tilde{x}_0 = x_0 + h_o^{-1}(y)$. Consider the trajectory $\phi(t, \tilde{x}_0)$ of (1) with the feedback control $u(x) = u_m(x - h_o^{-1}(y))$. Also consider the MA trajectory $\phi_{\text{MA}}(t, x_0)$ with feedback control $u_m(x)$. For all $t \geq 0$ such that $\phi_{\text{MA}}(t, x_0) \in I_{\text{MA}}(m)$,*

$$\phi(t, \tilde{x}_0) = \phi_{\text{MA}}(t, x_0) + h_o^{-1}(y).$$

Finally we are ready to prove Theorem V.1.

Proof of Theorem V.1. We must show that (i) output trajectories of system (1) remain within \mathcal{P} , and (ii) output trajectories eventually reach and remain within the goal set \mathcal{G} . Let $\tilde{x}_0 \in \mathcal{X}_0$. Choose any $(l_{j^0}, m^0) \in Q_{\text{PA}}^0$ such that $x_0 := \tilde{x}_0 - h_o^{-1}(d \circ l_{j^0}) \in I_{\text{MA}}(m^0)$. By Lemma V.2, we may associate a unique MA execution χ and a unique PA run π to $(l_{j^0}, m^0) \in Q_{\text{PA}}^0$ and $x_0 \in I_{\text{MA}}(m^0)$. Denote the hybrid time domain as $\tau = \{\mathcal{I}_0, \dots, \mathcal{I}_N\}$ with $\mathcal{I}_k = [\tau_k, \tau'_k]$ for $k = 0, \dots, N-1$ (with $\tau_0 = 0$) and $\mathcal{I}_N = [\tau_N, \infty)$. The last interval follows from the definition of $(l_{j^N}, m^N) \in Q_{\text{PA}}^f$ (5), since $\Sigma_{\text{MA}}(m^N) = \emptyset$ and thus Assumption IV.1 (vi) implies that we must have that $\mathcal{I}_N = [\tau_N, \infty)$. As in the proof of Lemma V.2, denote the corresponding sequence of events as $\sigma^0 \dots \sigma^{N-1}$.

Using Lemma V.3 with $y = d \circ l_{j^0}$, we have that $\phi(t, \tilde{x}_0) = \phi_{\text{MA}}(t, x_0) + h_o^{-1}(d \circ l_{j^0})$. We claim that for all $k = 0, \dots, N$ and $t \in \mathcal{I}_k$,

$$\phi(t, \tilde{x}_0) = \phi_{\text{MA}}(t, x_0) + h_o^{-1}(d \circ l_{j^k}). \quad (9)$$

Clearly the result is true for $k = 0$.

We derive two facts to assist in proving this claim. Recall that by definition of the OTS edges, we have that for all $k = 0, \dots, N-1$, $\sigma^k = l_{j^{k+1}} - l_{j^k}$. Furthermore, by rearranging, multiplying component-wise by d , and taking the preimage h_o^{-1} , we have the first fact: for all $k = 0, \dots, N-1$ that $h_o^{-1}(d \circ l_{j^{k+1}}) = h_o^{-1}(d \circ l_{j^k}) + h_o^{-1}(d \circ \sigma^k)$. Also by definition of the reset map and MA execution, we get the second fact: for all $k = 0, \dots, N-1$, $r_{e^k}(\phi_{\text{MA}}(\tau'_k, x_0)) = \phi_{\text{MA}}(\tau'_k, x_0) - h_o^{-1}(d \circ \sigma^k) = \phi_{\text{MA}}(\tau_{k+1}, x_0)$.

Returning to (9), by induction we assume that it is true for $0 \leq k < N$ and show that it is true for $k+1$. Using the above facts and (9) for k at $t = \tau'_k = \tau_{k+1}$ yields

$$\begin{aligned} \phi(\tau_{k+1}, \tilde{x}_0) &= \phi(\tau'_k, \tilde{x}_0) = \phi_{\text{MA}}(\tau'_k, x_0) + h_o^{-1}(d \circ l_{j^k}) \\ &= (\phi_{\text{MA}}(\tau_{k+1}, x_0) + h_o^{-1}(d \circ \sigma^k)) + h_o^{-1}(d \circ l_{j^k}) \\ &= \phi_{\text{MA}}(\tau_{k+1}, x_0) + h_o^{-1}(d \circ l_{j^{k+1}}). \end{aligned}$$

Applying Lemma V.3 with $y = h_o^{-1}(d \circ l_{j^{k+1}})$ at the new initial condition $\phi_{\text{MA}}(\tau_{k+1}, x_0) \in I_{\text{MA}}(m^{k+1})$, we have that for $k+1$ and for all $t \in \mathcal{I}_{k+1}$ that (9) holds. When $k+1 = N$, the induction terminates and the claim is proven.

Using (9) and projecting to the output space we conclude that for all $k = 0, \dots, N$ and $t \in \mathcal{I}_k$, $y(t, \tilde{x}_0) \in Y_{j^k}$. Since all the boxes are contained in \mathcal{P} by construction, then for all $t \geq 0$ we have (i). Moreover, since $l_{j^N} \in L_{\text{OTS}}^g$ implies the

goal box Y_{j^N} is contained in \mathcal{G} and $\mathcal{I}_N = [\tau_N, \infty)$, we have (ii). \square

Remark V.2. *The above result does not depend on the method of construction of the admissible control policy $c \in \mathcal{C}$, nor does it require the control policy to be optimal. This allows for different path planning techniques on the PA, as we show in Section VIII-B.*

Remark V.3. *The extension to a sequence of reach-avoid problems is straightforward, following the idea in [33]. First, the reach property (ii) of Problem III.1 is relaxed to $y(T, x_0) \in \mathcal{G}$. Next, suppose there is a finite sequence of goals $L_{\text{OTS}}^{g,i}$, $i = 1, \dots, n_g > 1$. In contrast to (5), we set the final PA states to be $Q_{\text{PA}}^{f,i} = \{(l, m) \in L_{\text{OTS}}^{g,i} \times M \mid \Sigma_{\text{MA}}(m) \neq \emptyset\}$ for $i = 1, \dots, n_g - 1$. Finally, one must design control policies c_i with associated initial conditions $Q_{\text{PA}}^{0,i}$ (6) such that $Q_{\text{PA}}^{f,i} \subset Q_{\text{PA}}^{0,i+1}$ for $i = 1, \dots, n_g - 1$. For $i = n_g$, one may impose solutions to remain invariant or connect back to the first goal.*

VI. PARALLEL COMPOSITION OF MOTION PRIMITIVES

In this section we describe the operation of parallel composition of two maneuver automata. By repeated application of this operation, more complex higher-dimensional MA's can be constructed by starting from simple low dimensional atomic motion primitives, such as those described in Section VII. The key challenge is to ensure that the resulting parallel composed MA satisfies Assumptions IV.1, if the two constituent MA's do. This is proved in Theorem VI.1. First we give some preliminary definitions and we fix some notation, followed by the formal definition of parallel composition of MA's.

We consider two independent systems

$$\dot{x}^j = f^j(x^j, u^j), \quad y^j = h^j(x^j), \quad (10)$$

where $x^j \in \mathbb{R}^{n^j}$, $u^j \in \mathbb{R}^{m^j}$, and $y^j \in \mathbb{R}^{p^j}$ for $j = 1, 2$. We use superscripts to identify the distinct subsystems. Assume that each system satisfies Assumption III.1. That is, for $j = 1, 2$, $y_i^j = x_i^j$, $i = 1, \dots, p^j$. Associated with each system $j = 1, 2$ is the MA

$$\mathcal{H}_{\text{MA}}^j = (Q_{\text{MA}}^j, \Sigma^j, E_{\text{MA}}^j, X_{\text{MA}}^j, I_{\text{MA}}^j, G_{\text{MA}}^j, R_{\text{MA}}^j, Q_{\text{MA}}^{0,j}). \quad (11)$$

We additionally assume that $\mathcal{H}_{\text{MA}}^1$ and $\mathcal{H}_{\text{MA}}^2$ satisfy Assumption IV.1. Denote the canonical boxes in the respective output spaces as $Y^{*,j} = \prod_{i=1}^{p^j} [0, d_i^j]$. The event sets labelling the faces of $Y^{*,j}$ are $\Sigma^j = \{-1, 0, 1\}^{p^j}$. The empty strings are denoted as $\varepsilon^j := (0, \dots, 0) \in \Sigma^j$, $j = 1, 2$, and the empty string is $\varepsilon := (\varepsilon^1, \varepsilon^2)$. Other sets are similarly denoted with a superscript to identify the system, such as the set of possible events $\Sigma_{\text{MA}}^j(m^j)$ for $m^j \in M^j$ and the output indices σ^j . For the parallel composition we also require some extra notation. First, for $j = 1, 2$ and for each $m^j \in M^j$, define the invariant set minus all the guard sets

$$I^j(m^j) := I_{\text{MA}}^j(m^j) \setminus \left(\bigcup_{\varepsilon^j = (m^j, \sigma^j, m_2^j) \in E_{\text{MA}}^j} g_{\varepsilon^j} \right). \quad (12)$$

Next, we need three sets: an augmented set of edges that includes a transition with the empty string, an augmented set of possible events for a motion primitive $m \in M^j$, and an augmented set of next feasible motion primitives. That is, for $j = 1, 2$, we define

$$\begin{aligned} \overline{E}_{\text{MA}}^j &:= E_{\text{MA}}^j \cup \{(m^j, \varepsilon^j, m_2^j) \mid m^j, m_2^j \in M^j, \\ &\quad I^j(m^j) \subset I_{\text{MA}}^j(m_2^j), \\ &\quad (\forall e_2^j = (m_2^j, \sigma_2^j, m_3^j) \in E_{\text{MA}}^j) \\ &\quad I^j(m^j) \cap g_{e_2^j} = \emptyset\} \\ \overline{\Sigma}_{\text{MA}}^j(m^j) &:= \Sigma_{\text{MA}}^j(m) \cup \{\varepsilon^j\}, \quad m^j \in M^j \\ \overline{M}^j(m^j, \sigma^j) &:= \{m_2^j \in M^j \mid (m^j, \sigma^j, m_2^j) \in \overline{E}_{\text{MA}}^j\}, \\ &\quad m^j \in M^j, \sigma^j \in \overline{\Sigma}_{\text{MA}}^j(m^j). \end{aligned}$$

We also define the products of these sets:

$$\begin{aligned} \overline{\Sigma}_{\text{MA}}(m) &:= \overline{\Sigma}_{\text{MA}}^1(m^1) \times \overline{\Sigma}_{\text{MA}}^2(m^2), \\ &\quad m = (m^1, m^2) \in M, \\ \overline{M}(m, \sigma) &:= \overline{M}^1(m^1, \sigma^1) \times \overline{M}^2(m^2, \sigma^2), \\ &\quad m = (m^1, m^2) \in M, \\ &\quad \sigma = (\sigma^1, \sigma^2) \in \overline{\Sigma}_{\text{MA}}(m). \end{aligned}$$

Finally, the canonical box in the output space of the parallel composition is $Y^* = Y^{*,1} \times Y^{*,2}$. We can now define the parallel composition of two MA's.

Definition VI.1. Consider two MA's $\mathcal{H}_{\text{MA}}^1$ and $\mathcal{H}_{\text{MA}}^2$ each satisfying Assumption IV.1. The parallel composition $\mathcal{H}_{\text{MA}}^1 \parallel \mathcal{H}_{\text{MA}}^2$ is $\mathcal{H}_{\text{MA}} = (Q_{\text{MA}}, \Sigma, E_{\text{MA}}, X_{\text{MA}}, I_{\text{MA}}, G_{\text{MA}}, R_{\text{MA}}, Q_{\text{MA}}^0)$ where

State Space $Q_{\text{MA}} = M \times \mathbb{R}^n$ with $M = M^1 \times M^2$ and $n = n^1 + n^2$.

Labels $\Sigma = \Sigma^1 \times \Sigma^2 = \{-1, 0, 1\}^p$ with $p = p^1 + p^2$.

Edges $E_{\text{MA}} \subset M \times \Sigma \times M$, where $e = (m, \sigma, m') \in E_{\text{MA}}$ if $\sigma \neq \varepsilon$, $\sigma \in \overline{\Sigma}_{\text{MA}}(m)$, and $m' \in \overline{M}(m, \sigma)$. Observe that for all $m \in M$, $\Sigma_{\text{MA}}(m) = \overline{\Sigma}_{\text{MA}}(m) \setminus \{\varepsilon\}$.

Vector Fields For all $m = (m^1, m^2) \in M$, $X_{\text{MA}}(m) = \begin{bmatrix} f^1(x^1, u_{m^1}(x^1)) \\ f^2(x^2, u_{m^2}(x^2)) \end{bmatrix}$. The state is $x := (x^1, x^2) \in \mathbb{R}^n$, the control input is $u := (u^1, u^2) \in \mathbb{R}^\mu$ where $\mu = \mu^1 + \mu^2$, and the output is $y := (y^1, y^2) \in \mathbb{R}^p$. The output map is $h(x) = \begin{bmatrix} h^1(x^1) \\ h^2(x^2) \end{bmatrix}$, with $o(i) = o^1(i)$ for $i = 1, \dots, p^1$ and $o(i) = n^1 + o^2(i - p^1)$ for $i = p^1 + 1, \dots, p$.

Invariants For all $m = (m^1, m^2) \in M$, $I_{\text{MA}}(m) = I_{\text{MA}}^1(m^1) \times I_{\text{MA}}^2(m^2)$.

Enabling and Reset Conditions Consider an edge $e = (m_1, \sigma, m_2) \in E_{\text{MA}}$, where $m_1 = (m_1^1, m_1^2) \in M$, $\sigma = (\sigma^1, \sigma^2) \in \overline{\Sigma}_{\text{MA}}(m_1)$, $m_2 = (m_2^1, m_2^2) \in \overline{M}(m_1, \sigma)$, and $e^j = (m_1^j, \sigma^j, m_2^j) \in \overline{E}_{\text{MA}}^j$ for $j = 1, 2$. If $\sigma^j \in \overline{\Sigma}_{\text{MA}}^j(m_1^j)$ and $\sigma^j = \varepsilon^j$, then we define

$$g_{e^j} := I^j(m_1^j), \quad r_{e^j}(x^j) := x^j.$$

Otherwise if $\sigma^j \in \Sigma_{\text{MA}}^j(m_1^j)$, we have $g_{e^j} = G_{\text{MA}}^j(e^j)$ and $r_{e^j} = R_{\text{MA}}^j(e^j)$, corresponding to their definitions in $\mathcal{H}_{\text{MA}}^j$.

Finally, we define $g_e = g_{e^1} \times g_{e^2}$ and $r_e(x) = \begin{bmatrix} r_{e^1}(x^1) \\ r_{e^2}(x^2) \end{bmatrix}$.

Initial Conditions $Q_{\text{MA}}^0 \subset Q_{\text{MA}}$ is the set of initial conditions given by $Q_{\text{MA}}^0 = \{(m, x) \mid (m^j, x^j) \in Q_{\text{MA}}^{0,j}, i = 1, 2\}$.

◁

First, notice that for each $\mathcal{H}_{\text{MA}}^j$ and for each $m^j \in M^j$, the definition of $\overline{E}_{\text{MA}}^j$ automatically includes self-loop edges $(m, \varepsilon^j, m) \in \overline{E}_{\text{MA}}^j$. We include such transitions with ε^j so that the parallel composition is properly constructed. For example, suppose a proper face of $Y^{*,1}$ is crossed by the first system, but no proper face of $Y^{*,2}$ is crossed by the second system. To correctly account for such possibilities, the overall transition for the composed MA must record the lack of crossing in $Y^{*,2}$ by the empty string ε^2 . Second, notice that we have allowed for additional edges with ε^j to allow for the possibility of switching to a different motion primitive over the same box $Y^{*,j}$ if the invariants overlap and are not mapped immediately to a guard set, as can be observed by the definition of $\overline{E}_{\text{MA}}^j$. Referring to Figure 8, an edge such as $((\mathcal{F}, \mathcal{H}), (1, 0), (\mathcal{H}, \mathcal{F})) \in E_{\text{MA}}$ consists of $(\mathcal{F}, 1, \mathcal{H}) \in E_{\text{MA}}^1$ and $(\mathcal{H}, 0, \mathcal{F}) \in \overline{E}_{\text{MA}}^2$, which encodes a turn from Right to Up.

The main result is now stated; the proof is in the appendix.

Theorem VI.1. We are given $\mathcal{H}_{\text{MA}}^1$ and $\mathcal{H}_{\text{MA}}^2$, two MA's that satisfy Assumption IV.1. The parallel composition $\mathcal{H}_{\text{MA}} = \mathcal{H}_{\text{MA}}^1 \parallel \mathcal{H}_{\text{MA}}^2$ defined above is an MA that also satisfies Assumption IV.1.

Remark VI.1. We have defined the event set as $\Sigma = \Sigma^1 \times \Sigma^2$, but the usual parallel composition of automata would have $\Sigma = \Sigma^1 \cup \Sigma^2$ [34]. Given the interpretation of the event set as crossing faces of Y^* , the cartesian product is the more natural choice.

VII. MOTION PRIMITIVES FOR INTEGRATOR SYSTEMS

In this section we give the formal details for the MA consisting of the three motion primitives Hold (\mathcal{H}), Forward (\mathcal{F}), and Backward (\mathcal{B}) introduced in Example IV.1. This design is able to be succinctly expressed within the MA formalism since the underlying double integrator system satisfies Assumption III.1. By exploiting the parallel composition construction from Section VI, the usefulness of this MA is demonstrated in the context of multi-robot systems in Section VIII.

Suppose the nonlinear control system is the double integrator system:

$$\dot{x}_1 = x_2, \quad \dot{x}_2 = u_2, \quad y = x_1, \quad (13)$$

where $x := (x_1, x_2) \in \mathbb{R}^2$, $u_2 \in \mathbb{R}$, and the output y is the position. Each motion primitive's invariant region is a polytopic set in the state space defined as the convex hull of vertices v_2^k , $k \in \{1, \dots, 6\}$; see Figure 5. The vertices are determined by the segment length $d > 0$, and a pre-specified maximum control value $u_2^* > 0$. Let $\bar{u}_1 := \sqrt{d u_2^*}$. The vertices are $v_2^1 = (0, -\bar{u}_1)$, $v_2^2 = (0, 0)$, $v_2^3 = (0, \bar{u}_1)$,

$v_2^4 = (d, -\bar{u}_1)$, $v_2^5 = (d, 0)$, and $v_2^6 = (d, \bar{u}_1)$. For each motion primitive $m \in M := \{\mathcal{H}, \mathcal{F}, \mathcal{B}\}$, we define an affine feedback

$$u_m(x) = K_m x + g_m. \quad (14)$$

Our specific choices are $K_{\mathcal{H}} = [-2u_2^*/d \quad -2u_2^*/\bar{u}_1]$, $K_{\mathcal{F}} = K_{\mathcal{B}} = [0 \quad -2u_2^*/\bar{u}_1]$, $g_{\mathcal{H}} = g_{\mathcal{F}} = u_2^*$, and $g_{\mathcal{B}} = -u_2^*$. These controllers are derived using reach control theory [4], [27]. One first selects control values at the vertices of the polytopes so that trajectories remain in the invariant region (for the Hold primitive) or they exit the polytope through a certain facet and not through others. In particular, we have chosen all the control values at the vertices to have magnitude u_2^* . Then the velocity vectors at the vertices are affinely extended to obtain affine feedbacks over the entire polytope, yielding the vector fields shown in Figure 5.

Now we construct the MA. The state space is $Q_{\text{MA}} = M \times \mathbb{R}^2$. The labels are $\Sigma = \{-1, 0, 1\}$. The set of edges E_{MA} are shown in Figure 4. In the context of parallel composition, one may compute that the augmented edges are

$$\bar{E}_{\text{MA}} = E_{\text{MA}} \cup \{(m, 0, m)\}_{m \in M} \cup \{(\mathcal{H}, 0, \mathcal{F}), (\mathcal{H}, 0, \mathcal{B})\}.$$

For each $m \in M$, the closed-loop vector fields are given by $[X_{\text{MA}}(m)](x) = (x_2, u_m(x))$, which are clearly globally Lipschitz. The invariants are given by the convex hull of vertices, as seen in Figure 5, and excluding the two points $(0, 0)$ and $(d, 0)$, so the invariants are clearly bounded. For example, $I_{\text{MA}}(\mathcal{H}) = \text{co}\{v_2^k\}_{k=2}^5 \setminus \{(0, 0), (d, 0)\}$. The enabling conditions are constructed by taking the convex hull of vertices of the exit facet and excluding again $(0, 0)$ or $(d, 0)$. Specifically, the edges $(\mathcal{F}, 1, \mathcal{H}), (\mathcal{F}, 1, \mathcal{B}) \in E_{\text{MA}}$ both have guard sets $g_e = \text{co}\{v_2^5, v_2^6\} \setminus \{(d, 0)\} = \{d\} \times (0, \bar{u}_1]$, as shown highlighted in green on the invariant region of \mathcal{F} in Figure 5, whereas $(\mathcal{B}, -1, \mathcal{H}), (\mathcal{B}, -1, \mathcal{B}) \in E_{\text{MA}}$ both have guard sets $g_e = \text{co}\{v_2^1, v_2^2\} \setminus \{(0, 0)\} = \{0\} \times [-\bar{u}_1, 0)$. The reset conditions are constructed according to their definition. The proof of the following result is found in the appendix.

Lemma VII.1. *The double integrator MA satisfies Assumption IV.1.*

Remark VII.1. *We noted in Remark V.1 that Zeno executions do not arise for reach-avoid specifications that, by construction, involve only finite MA executions. However, one may be interested in analyzing whether an MA is non-Zeno in its own right, independently of the high level plan or control specification for which it is used. It can be verified rather easily that the $p = 1$ double integrator MA design we have presented above is non-Zeno. The situation is considerably more complicated when considering an MA that is a parallel composition of these MA's or when considering an arbitrary MA. Generic conditions when hybrid systems have a Zeno execution have been studied in [16], [36]. However, further study of this problem is needed in our context since existing results do not apply to all the situations that can arise in our MA.*

VIII. QUADROPTER APPLICATIONS

In this section we apply our methodology to a group of quadcopters. We first explain how motion primitives can be applied to the system, how to specify the reach-avoid objective, and the overall solution pipeline. Next, we compare and contrast three algorithms for computing a control policy. Then we present experimental results on three different scenarios. Lastly, we provide a discussion.

A. Interfacing Multiple Quadcopters

The standard quadrotor dynamical model has six degrees of freedom, which can be described by the inertial linear positions (x_w, y_w, z_w) and the roll-pitch-yaw Euler angles (ϕ, θ, ψ) [23], [25]. It is well known that this system is differentially flat, relating the full state and motor inputs of the quadrotor to the flat outputs (x_w, y_w, z_w, ψ) and their derivatives [23]. Rather than specifying positional reference trajectories, we use the motion primitives from Section VII independently in the (x_w, y_w, z_w) directions to compute the linear accelerations as a feedback on the linear position and velocity states. Specifying an arbitrary yaw reference, differential flatness maps these linear accelerations to the (ϕ, θ) angles and the total vehicle thrust, which through the use of an attitude tracking controller can be converted to motor inputs [23]. Although we have avoided computing motion primitives on the high dimensional nonlinear model, our experiments show that the quadrotor is fairly well approximated as double integrators in the (x_w, y_w, z_w) directions using our proposed motion primitives.

We consider a centralized reach-avoid objective among N quadcopters. A copy of the gridded 3D workspace must be associated with each vehicle, resulting in a total of $p = 3N$ outputs. The p -dimensional MA representing the asynchronous motion capabilities of the multi-vehicle system is obtained by parallel composing p times the single-output MA from Section VII.

To specify the reach-avoid objective, we must identify the obstacle and goal boxes in $p = 3N$ dimensions. First we assume that the physical obstacles and goals for each vehicle are labelled on the physical 3D grid. Obstacle boxes in the output space correspond to any vehicle occupying a physical obstacle box or any two or more vehicles occupying the same physical box simultaneously. To avoid the effects of downwash, we do not allow vehicles to simultaneously occupy boxes that are displaced only in the z_w direction. Goal boxes in the output space correspond to all the combinations of individual vehicle 3D goal boxes. For simplicity, we assume that each vehicle has a single 3D goal box.

The multi-vehicle reach-avoid problem is solved offline using our proposed methodology. The runtime workflow is depicted in Figure 9. Each runtime component requires negligible computation, even for a large number of vehicles and outputs.




Fig. 9. Interface between multiple vehicles and the framework with $p = 3N$ outputs. The hybrid controller internal state consists of the joint state measurement of all the vehicles and includes the current (joint) box, l , and the current (joint) motion primitive, m . The internal state is updated via external state measurements (assumed to be given) and is used to compute the feedback controls.

B. Control Policy Generation

We highlight three options for generating a control policy in the context of the multi-vehicle reach-avoid problem. For each, we give some implementation details and discuss its computational complexity. These are then compared in the experiments.

1) *Exhaustive Non-Deterministic Dijkstra (NDD)*: The first strategy follows the proposed methodology of Section IV. We highlight our main implementation steps. First, we compute the OTS states and edges for the associated output space obstacle boxes described earlier. Second, the p times parallel composed MA states and edges are computed. Third, the PA states and edges are computed. Fourth, the value function V is computed using (4). This is done by initializing the value function to be zero at goal states and infinite elsewhere, and then propagating backwards along PA edges using a non-deterministic Dijkstra (NDD) algorithm [5], [33]. Once the value function is computed at all states, we compute the optimal control policy c^* using Corollary IV.1. The initial PA states (6) correspond precisely to those states $q \in Q_{PA}$ with $V(q) < \infty$.

The computational complexity grows exponentially as the number of inputs $p = 3N$ increases. Suppose that the physical grid has (n_x, n_y, n_z) boxes in the (x_w, y_w, z_w) directions. Since there are 3^p motion primitives, the number of PA states is bounded by $|Q_{PA}| < (n_x n_y n_z)^N 3^p =: k_1$. The number of edges from an OTS state is bounded by $3^p - 1$ (the neighboring directions), whereas the number of edges from a MA state is bounded by $(2^p - 1)3^p =: k_2$ (the neighboring directions times the possible next motion primitives). Since the MA neighboring directions are more restrictive, we have the number of PA edges is bounded by $|E_{PA}| < k_1 k_2$. The presence of obstacles can dramatically reduce the number of PA states and edges. The NDD algorithm generally must inspect all the PA states and edges to compute the value function. As a result, it is optimal and complete (with respect to the selected grid resolution and motion primitive capabilities), which results in the largest possible set of initial conditions \mathcal{X}_0 .

2) *Deterministic A**: In this strategy, we make two simplifying assumptions to compromise the quality of the control policy in exchange for better computational efficiency. First, we take the p times composed MA and prune out motion

primitives enabling simultaneous motion. Second, we forego computing the largest possible set of initial conditions and instead assume that a single physical initial box is specified for each vehicle. As such, it is sufficient to compute a single path of boxes in the OTS connecting the initial and goal boxes in the $p = 3N$ dimensional output space. From this path the control policy is immediately extracted, by assigning to each box the unique motion primitive leading to the next neighboring box along the path. The path is computed using a standard A* algorithm [19], which starts from the initial box and propagates outwards until the goal box is reached. The (admissible) heuristic function is chosen to be the Manhattan distance, which is the sum of distances along each output direction from the current box to the goal box.

The number of nodes that A* must investigate is bounded by the maximum number of OTS boxes, $(n_x n_y n_z)^N$, which still has exponential complexity in the number of robots. The pruned MA has $2p + 1$ motion primitives, corresponding to \mathcal{F} or \mathcal{B} in a single output component with \mathcal{H} elsewhere, plus the motion primitive $(\mathcal{H}, \dots, \mathcal{H})$. Thus from the current box, we must check the $2p$ neighboring directions to select a feasible direction, taking into account out-of-bounds and obstacle configurations. In this implementation, the OTS, MA, and PA serve more as conceptual constructs, and do not need to be precomputed explicitly as it is expensive. In the worst case, the A* algorithm may investigate all boxes; as a result, it also produces a control policy that is complete with respect to the chosen grid and pruned MA motion capabilities. The policy produced by A* is of minimal length, but may have a long runtime execution.

3) *Deterministic Greedy Search*: This strategy also makes use of the two simplifying assumptions as with A* above, but differs in how the path is constructed. In greedy (best first) search [19], the path is constructed by starting from the initial box in the output space and then extending it from the current box into any feasible neighboring direction that decreases the Manhattan distance to the goal box. Greedy search can often find a path very quickly, although not necessarily an optimal one. Moreover, since greedy search may fail to find a path, it is not complete.

C. Experimental Results

Our experimental platform is the Crazyflie 2.0; see Figure 1. We used a VICON motion capture system to obtain the state estimates of the vehicles. Our implementation was done in Python 2.7.10 and ROS Kinetic, and computations were performed on a 64-bit Lenovo ThinkPad with an 8 core 3.0 GHz Intel Xeon processor and 15.4 GiB RAM. We illustrate three different scenarios and consider the three policy generation strategies on each of them. The corresponding video results are available at <http://tiny.cc/modular-3alg>.

1) *Open Space*: The first representative scenario involves an open 3D space partitioned into a $7 \times 7 \times 2$ grid and a sparse collection of pillar-shaped obstacles. The left plot of Figure 10 compares the resulting 3D trajectories in the




Fig. 10. Experimental results for the open scenario, projected onto the (x_w, y_w) plane. In all plots, all the vehicles must swap corners of the room. The left plot compares trajectories for a single vehicle using the three different control policy generation strategies. The middle plot shows the resulting trajectories for two vehicles using non-deterministic Dijkstra. The right plot shows the resulting trajectories for four vehicles using greedy search, see also Figure 1. Although difficult to depict, the maneuvers are safe, as the trajectories do not occupy the same physical boxes at the same time.

(x_w, y_w) plane for the three strategies in the case of a single vehicle. The computation times were 40.63 milliseconds, 1.59 milliseconds, and 0.27 milliseconds for NDD, A*, and greedy search, respectively. The NDD algorithm offers the best quality control policy in that there is simultaneous motion in the different degrees of freedom whenever possible and the same policy can be used from any starting box. The A* and greedy search algorithms offer similar results to each other, with both producing an optimal path of length 14. Both yield less efficient grid-like motion that is defined only along a single path from the initial box, although a new policy can quickly be recomputed from different starting boxes. Based on simulation tests for a single vehicle, each of these algorithms scale well to larger spaces or finer grids; even NDD is able to compute a solution on a $100 \times 100 \times 10$ grid in about two minutes in the worst cases. Next we compare each strategy on more vehicles.

The middle plot of Figure 10 shows the resulting trajectories for two vehicles using NDD. The control policy was computed in about 18 minutes and is defined on about PA 180000 states. While the resulting control policy yields highly efficient motion defined over a large set of initial conditions, adding more vehicles or more boxes generally explodes the computation time and memory requirements. Thus NDD is best suited for small scenarios involving a modest number of vehicles, when one can afford to spend time precomputing the control policy.

The right plot of Figure 10 shows the resulting trajectories for four vehicles swapping corners of the room using greedy search. Since the vehicles and physical obstacles occupy a single box, greedy search performs well, as each action typically results in one vehicle making progress towards the goal. The computation time was about four milliseconds. Simulation results on a $100 \times 100 \times 10$ grid with eight vehicles placed randomly demonstrate that greedy search is usually able to find a solution on the order of one second. As one would expect, greedy search typically fails to find a solution

if long wall-like or non-convex obstacles are introduced, or if the goals are not spaced out sufficiently. Furthermore, the time to execute the entire maneuver scales with the number of vehicles.

Finally we consider the deterministic A* algorithm. Although the resulting trajectories follow a path of optimal length, they look quite similar to those found by greedy search and thus are not shown. Moreover, the method quickly becomes more computationally expensive beyond three vehicles.

2) *Channel Swapping*: The second representative scenario involves two rooms connected by a channel, defined over a $5 \times 2 \times 1$ grid, see Figure 11. Two of the vehicles must continually swap places, while the third is required to act as a gatekeeper. We specify this objective as an infinitely looping sequence of two distinct reach-avoid problems. This illustrates that reach-avoid is a useful building block for addressing more complex specifications.

The NDD algorithm produced both control policies in about 10 seconds, while the A* algorithm took about 0.03 seconds. Greedy search fails to find a solution because it is unable to coordinate the third vehicle away from its goal to make space for the other two. Since the resulting trajectories overlap in physical space, Figure 12 shows the trajectories as a function of time using the policy computed with NDD. The trajectories are highly non-trivial, but show that the objective is satisfied for at least one cycle of both reach-avoids. Although not shown, the trajectories computed using A* are similar but take a few seconds longer to execute the objective since the motion primitives are deterministic.

3) *8-Puzzle*: We conclude our experimental results with the well-known 8-puzzle. On a $3 \times 3 \times 1$ grid, eight vehicles are placed randomly and must return to an ordered configuration. For this application, the A* algorithm is the most suitable, computing the control policy in 0.32 seconds. The NDD approach would spend too much time precomputing edges in the high dimensional output space, while greedy search would never make progress. Results are available to view in




Fig. 11. This figure shows the channel swapping experiment involving three vehicles. In particular, it shows the specification for the first reach-avoid, where goals are shown as the colored boxes. For the second reach-avoid, the initial and goal boxes of vehicle 1 and 2 are swapped.




Fig. 12. Trajectories using non-deterministic Dijkstra for the channel experiment. The alternating grey and white areas reflect the size of the grid boxes. The duration for each of the two reach-avoid specifications is highlighted.

the video.

D. Discussion

Throughout the various experimental scenarios presented, we have demonstrated the modularity offered by our approach. The designer can customize their own algorithms for generating a control policy in order to trade-off solution quality with computational efficiency. Depending on the specific application scenario, a different control policy generation strategy may be more suitable.

In our analysis, all of the complexity was associated with the generation of the control policy for a given MA. The MA formalism enables us to generate control policies with no further regard to the continuous time trajectories that may result, due to the guarantees on discrete behavior encoded in the MA edges. On the other hand, the generation of a MA for an arbitrary system is a difficult challenge in its own right and is left to the discretion of the designer, although the design we have presented in Section VII can potentially be applied to control systems that are feedback-linearizable into a collection of double integrators. Taking care that the outputs are translationally invariant and that obstacle boxes can be computed, this includes end effector control of fully actuated robotic manipulators [29] and some wheeled vehicles through the use of look-ahead points [1].

Our approach offers robustness through the use of feedback-based motion primitives, as the construction of invariant regions ensures a wide range of initial conditions for which output trajectories exit through appropriate guard sets into subsequent boxes. Since the motion primitives are updated during execution based on the measured box transitions and control policy, we do not require timing estimates for completing box transitions, which can be difficult to compute. These features are advantageous under model uncertainty, which we

must contend with since we base our motion primitive design on the double integrator model rather than the more complex quadcopter model, and since aerodynamic effects arise when multiple quadcopters fly in close proximity. Our previous work also demonstrated similar robustness of operation under wind disturbances generated by a fan on a larger quadcopter [31]. Finally, we note that our framework can easily be applied to a heterogeneous team of robots; if each vehicle has its own MA, the parallel composition automatically constructs the overall MA for the multi-vehicle system.

Of course, our solution to Problem III.1 is conservative because we have restricted ourselves to a particular discretization, namely the choice of a partition into boxes and the use of motion primitives. As we have demonstrated, this is a reasonable trade-off, especially since the resolution of the output space discretization, the richness of motion primitives, and the complexity of the control policy are all design parameters.

IX. CONCLUSION

We have developed a modular, hierarchical framework for motion planning of multiple robots in known environments. It consists of several modules. An output transition system (OTS) models the allowable motions of the robots by partitioning their workspace into boxes. A set of motion primitives is designed based on reach control on polytopes. A maneuver automaton (MA) captures constraints on successive motion primitives. Finally, a control policy is generated based on the synchronous product of the OTS and the discrete part of the MA. Overall we obtain a two-level control design which is highly robust, modular, and conceptually elegant. We presented a specific maneuver automaton for the double integrator system, and we showed how this design can be composed to obtain maneuver automata for multi-robot systems. The methodology was experimentally validated on a group of quadcopters. Future work includes application of our methodology to different vehicle classes such as robotic manipulators or wheeled vehicles, and integration with more advanced multi-robot planning algorithms in dynamic environments.

APPENDIX

Proof of Theorem IV.1. First we prove (3). Consider $q = (l, m) \in Q_{\text{PA}} \setminus Q_{\text{PA}}^f$ and suppose $|\Sigma_{\text{PA}}(q)| > 0$. By the definition of J , for any $c \in \mathcal{C}$,

$$J(q, c) = \max_{e=(q, \sigma, q') \in E_{\text{PA}}} \{D_{\text{PA}}(e) + J(q', c)\}, \quad (15)$$

where $q' = (l', c(q, \sigma)) \in Q_{\text{PA}}$. Observe that given $q = (l, m) \in Q_{\text{PA}}$ and $\sigma \in \Sigma_{\text{PA}}(q)$, there exists a unique $l' \in L_{\text{OTS}}$ such that $(l, \sigma, l') \in E_{\text{OTS}}$ (since the OTS is a deterministic automaton). Therefore, when we take the maximum over $e = (q, \sigma, q') \in E_{\text{PA}}$ in (15) with $q' = (l', c(q, \sigma))$, the only free variable to maximize over is $\sigma \in \Sigma_{\text{PA}}(q)$. Therefore, (15) is equivalent to

$$J(q, c) = \max_{\sigma \in \Sigma_{\text{PA}}(q)} \{D_{\text{PA}}(e) + J(q', c)\}, \quad (16)$$

where, as before, $q' = (l', c(q, \sigma)) \in Q_{\text{PA}}$ and $e = (q, \sigma, q') \in E_{\text{PA}}$. By definition of V

$$J(q, c) \geq \max_{\sigma \in \Sigma_{\text{PA}}(q)} \{D_{\text{PA}}(e) + V(q')\}.$$

Again by definition of V

$$V(q) = \min_{c \in \mathcal{C}} J(q, c) \geq \min_{c \in \mathcal{C}} \left\{ \max_{\sigma \in \Sigma_{\text{PA}}(q)} \{D_{\text{PA}}(e) + V(q')\} \right\}.$$

Thus,

$$V(q) \geq \min_{c(q) \in \mathcal{M}(q)} \left\{ \max_{\sigma \in \Sigma_{\text{PA}}(q)} \{D_{\text{PA}}(e) + V(q')\} \right\}.$$

To prove the reverse inequality, suppose by the way of contradiction that there exists an admissible control assignment at $q \in Q_{\text{PA}}$, $\hat{c}(q) \in \mathcal{M}(q)$, such that

$$\begin{aligned} V(q) &> \max_{\sigma \in \Sigma_{\text{PA}}(q)} \{D_{\text{PA}}(\hat{e}) + V(\hat{q})\} \\ &\geq \min_{c(q) \in \mathcal{M}(q)} \left\{ \max_{\sigma \in \Sigma_{\text{PA}}(q)} \{D_{\text{PA}}(e) + V(q')\} \right\}, \end{aligned}$$

where $\hat{q} = (\hat{l}, \hat{c}(q, \sigma)) \in Q_{\text{PA}}$, $\hat{e} = (q, \sigma, \hat{q}) \in E_{\text{PA}}$, $q' = (l', c(q, \sigma)) \in Q_{\text{PA}}$, and $e = (q, \sigma, q') \in E_{\text{PA}}$. Suppose the maximum for $\hat{c}(q)$ is achieved with $\sigma^* \in \Sigma_{\text{PA}}(q)$. We define $q^* = (l^*, \hat{c}(q, \sigma^*))$, and $e^* = (q, \sigma^*, q^*) \in E_{\text{PA}}$. Then

$$V(q) > D_{\text{PA}}(e^*) + V(q^*).$$

Suppose an admissible optimal control policy for q^* to achieve $V(q^*)$ is $c^* \in \mathcal{C}$. Define a new policy $c = c^*$ on $Q_{\text{PA}} \setminus \{q\}$ and $c(q) = \hat{c}(q)$. Then

$$\begin{aligned} J(q, c) &= \max_{\sigma \in \Sigma_{\text{PA}}(q)} \{D_{\text{PA}}(e) + V(q')\} \\ &= D_{\text{PA}}(e^*) + V(q^*) < V(q), \end{aligned}$$

a contradiction. Hence, it must be that

$$V(q) \leq \min_{c(q) \in \mathcal{M}(q)} \left\{ \max_{\sigma \in \Sigma_{\text{PA}}(q)} \{D_{\text{PA}}(e) + V(q')\} \right\},$$

as desired. This proves (3).

Second we prove (4). Consider $q = (l, m) \in Q_{\text{PA}} \setminus Q_{\text{PA}}^f$ and suppose $|\Sigma_{\text{PA}}(q)| > 0$. Let $\bar{c}(q) \in \mathcal{M}(q)$ be an admissible control assignment at q such that for all $\sigma \in \Sigma_{\text{PA}}(q)$,

$$\bar{c}(q, \sigma) \in \operatorname{argmin}_{\bar{m} \in \mathcal{M}(q, \sigma)} \{D_{\text{PA}}(\bar{e}) + V(\bar{q})\}, \quad (17)$$

where $\bar{q} = (\bar{l}, \bar{m}) \in Q_{\text{PA}}$, and $\bar{e} = (q, \sigma, \bar{q}) \in E_{\text{PA}}$. We will show that

$$\begin{aligned} \max_{\sigma \in \Sigma_{\text{PA}}(q)} \{D_{\text{PA}}(\bar{e}) + V(\bar{q})\} &\leq \\ \min_{c(q) \in \mathcal{M}(q)} \left\{ \max_{\sigma \in \Sigma_{\text{PA}}(q)} \{D_{\text{PA}}(e) + V(q')\} \right\}, &\quad (18) \end{aligned}$$

where $\bar{q} = (\bar{l}, \bar{c}(q, \sigma)) \in Q_{\text{PA}}$, $\bar{e} = (q, \sigma, \bar{q}) \in E_{\text{PA}}$, $q' = (l', c(q, \sigma)) \in Q_{\text{PA}}$, and $e = (q, \sigma, q') \in E_{\text{PA}}$. Suppose the minimum and maximum on the r.h.s. are achieved with $c^*(q) \in \mathcal{M}(q)$ and $\sigma^* \in \Sigma_{\text{PA}}(q)$. Also, suppose the maximum

on the l.h.s. is achieved with $\bar{\sigma}^* \in \Sigma_{\text{PA}}(q)$. Then (18) becomes

$$D_{\text{PA}}(\bar{e}^*) + V(\bar{q}^*) \leq D_{\text{PA}}(e^*) + V(q^*), \quad (19)$$

where $\bar{q}^* = (\bar{l}^*, \bar{c}(q, \bar{\sigma}^*)) \in Q_{\text{PA}}$, $\bar{e}^* = (q, \bar{\sigma}^*, \bar{q}^*) \in E_{\text{PA}}$, $q^* = (l^*, c^*(q, \sigma^*)) \in Q_{\text{PA}}$, and $e^* = (q, \sigma^*, q^*) \in E_{\text{PA}}$. Suppose by way of contradiction that (18) does not hold. Then (19) does not hold. That is,

$$D_{\text{PA}}(\bar{e}^*) + V(\bar{q}^*) > D_{\text{PA}}(e^*) + V(q^*).$$

By the maximality of $\sigma^* \in \Sigma_{\text{PA}}(q)$ we have

$$\begin{aligned} &D_{\text{PA}}(e^*) + V(q^*) \\ &= D_{\text{PA}}((q, \sigma^*, (l^*, c^*(q, \sigma^*)))) + V((l^*, c^*(q, \sigma^*))) \\ &\geq D_{\text{PA}}((q, \bar{\sigma}^*, (\bar{l}^*, c^*(q, \bar{\sigma}^*)))) + V((\bar{l}^*, c^*(q, \bar{\sigma}^*))). \end{aligned}$$

Therefore

$$\begin{aligned} &D_{\text{PA}}(\bar{e}^*) + V(\bar{q}^*) \\ &= D_{\text{PA}}((q, \bar{\sigma}^*, (\bar{l}^*, \bar{c}(q, \bar{\sigma}^*)))) + V((\bar{l}^*, \bar{c}(q, \bar{\sigma}^*))) \\ &> D_{\text{PA}}((q, \bar{\sigma}^*, (\bar{l}^*, c^*(q, \bar{\sigma}^*)))) + V((\bar{l}^*, c^*(q, \bar{\sigma}^*))). \end{aligned}$$

This contradicts the definition of $\bar{c}(q, \bar{\sigma}^*)$ in (17). We conclude that (18) must hold.

Now consider $\bar{c}^* \in \mathcal{C}$ such that $\bar{c}^*(q) = \bar{c}(q)$ and \bar{c}^* is any admissible optimal control policy for $q' \neq q$. Using (16) we have

$$\begin{aligned} J(q, \bar{c}^*) &= \max_{\sigma \in \Sigma_{\text{PA}}(q)} \{D_{\text{PA}}(\bar{e}) + J(\bar{q}, \bar{c}^*)\} \\ &= \max_{\sigma \in \Sigma_{\text{PA}}(q)} \{D_{\text{PA}}(\bar{e}) + V(\bar{q})\}, \end{aligned}$$

where $\bar{q} = (\bar{l}, \bar{c}^*(q, \sigma)) = (\bar{l}, \bar{c}(q, \sigma)) \in Q_{\text{PA}}$ and $\bar{e} = (q, \sigma, \bar{q}) \in E_{\text{PA}}$. Now by (3) and (18) we have

$$\begin{aligned} J(q, \bar{c}^*) &= \max_{\sigma \in \Sigma_{\text{PA}}(q)} \{D_{\text{PA}}(\bar{e}) + V(\bar{q})\} \\ &\leq \min_{c(q) \in \mathcal{M}(q)} \left\{ \max_{\sigma \in \Sigma_{\text{PA}}(q)} \{D_{\text{PA}}(e) + V(q')\} \right\} = V(q). \end{aligned}$$

By definition of V , we obtain that $J(q, \bar{c}^*) = V(q)$. That is,

$$J(q, \bar{c}^*) = V(q) = \max_{\sigma \in \Sigma_{\text{PA}}(q)} \{D_{\text{PA}}(\bar{e}) + V(\bar{q})\} \quad (20)$$

with $\bar{q} = (\bar{l}, \bar{c}(q, \sigma)) \in Q_{\text{PA}}$ and $\bar{e} = (q, \sigma, \bar{q}) \in E_{\text{PA}}$. However, we know $\bar{c}^*(q) = \bar{c}(q)$, and $\bar{c}(q)$ satisfies (17). Therefore,

$$V(q) = \max_{\sigma \in \Sigma_{\text{PA}}(q)} \left\{ \min_{\bar{m} \in \mathcal{M}(q, \sigma)} \{D_{\text{PA}}(\bar{e}) + V(\bar{q})\} \right\},$$

where now $\bar{q} = (\bar{l}, \bar{m}) \in Q_{\text{PA}}$, and this proves (4). \square

Proof of Lemma V.2. Let $(l^0, m^0) \in Q_{\text{PA}}^0$ and $x_0 \in I_{\text{MA}}(m^0)$. The initial MA state of the MA execution is $(m(0), x_0) = (m^0, x_0) \in Q_{\text{MA}}^0$, and the initial PA state of the PA run π is $q^0 = (l^0, m^0)$. The hybrid time domain of χ will be denoted as $\tau = \{\mathcal{I}_i\}_{i=0}^{n_\tau}$, and is initialized as $\tau = \{\mathcal{I}_0\}$, where $\mathcal{I}_0 = \{\tau_0\}$ and $\tau_0 = 0$. With the base case $k = 0$ established, we construct the remainder of the MA execution and PA run by induction.

The run so far is $\pi = q^0 \dots q^k$, where $q^i = (l^i, m^i)$ for $i = 0, \dots, k$, and $\mathcal{I}_k = \{\tau_k\}$. Suppose $\Sigma_{\text{MA}}(m^k) = \emptyset$. Then by Assumption IV.1 (vi), $\phi_{\text{MA}}(t, x_0) \in I_{\text{MA}}(m^k)$ for all t in the extended interval $\mathcal{I}_k = [\tau_k, \infty)$. The complete PA run is $\pi = q^0 \dots q^k$ and the induction terminates. Suppose instead $\Sigma_{\text{MA}}(m^k) \neq \emptyset$. Then by Assumption IV.1 (vii), there exist unique $\sigma^k \in \Sigma_{\text{MA}}(m^k)$ and $T^k \geq 0$, such that $\phi_{\text{MA}}(t, x_0) \in I_{\text{MA}}(m^k)$ for all t in the extended interval $\mathcal{I}^k = [\tau_k, \tau'_k]$, $\tau'_k := \tau_k + T^k$. Also, for each $e = (m^k, \sigma^k, m')$ in E_{MA} , there exists a guard set g_e such that $\phi_{\text{MA}}(\tau'_k, x_0) \in g_e$. Assumption IV.1 (ii) tells us that for all such m' , the guard set is the same. Also, Assumption IV.1 (iii) ensures that σ^k is unique. Now we invoke the control policy to select a specific m' . Let $m^{k+1} := c(q^k, \sigma^k)$ so that $e^k := (m^k, \sigma^k, m^{k+1}) \in E_{\text{MA}}$ and $\phi_{\text{MA}}(\tau'_k, x_0) \in g_{e^k}$. Define $x_0^{k+1} := r_{e^k}(\phi_{\text{MA}}(\tau'_k, x_0))$. By Assumption IV.1 (iv), $x_0^{k+1} \notin g_e$ for any $e = (m^{k+1}, \sigma, m')$ in E_{MA} . The next PA state is $q^{k+1} = (l^{k+1}, m^{k+1})$, where $l^{k+1} \in L_{\text{OTS}}$ is uniquely determined through $(l^k, \sigma^k, l^{k+1}) \in E_{\text{OTS}}$, by the determinism of the OTS. The PA run so far is $\pi = q^0 \dots q^{k+1}$ and the new interval $\mathcal{I}_{k+1} = \{\tau'_k\}$ is added to τ .

The above inductive process is guaranteed to terminate with a finite PA run by definition of Q_{PA}^0 . That is, since $(l^0, m^0) \in Q_{\text{PA}}^0$ there will be a smallest N such that $(l^N, m^N) \in Q_{\text{PA}}^f$. Moreover, by definition of Q_{PA}^f (5), we have that $\Sigma_{\text{MA}}(m^N) = \emptyset$ and so the run cannot be extended further. The resulting MA execution is infinite with a finite number of intervals in the hybrid time domain τ , and it is non-blocking by Lemma V.1. \square

Proof of Theorem VI.1. We employ the following two standard facts regarding products, intersections, and subsets of sets. Formally, if A, B, C, D are sets, then

$$(A \cap C) \times (B \cap D) = (A \times B) \cap (C \times D), \quad (21)$$

$$A \subset C \text{ and } B \subset D \Rightarrow (A \times B) \subset (C \times D). \quad (22)$$

First we show that the resulting \mathcal{H}_{MA} is in fact an MA according to the definition. Clearly the composed vector fields are also globally Lipschitz and the composed invariants are bounded. The non-trivial points to show are that (a) the stacked system satisfies Assumption III.1, (b) the invariants project within the canonical box, (c) the enabling conditions lie both within the invariant and on an appropriate face determined by $\sigma \in \Sigma$, (d) the reset conditions are determined only by the event $\sigma \in \Sigma$, and (e) the initial conditions are the entire invariants. We prove each of these in turn.

(a) We show that Assumption III.1 for the stacked system holds. For the first condition, it can be verified by direct expansion that the definition of h necessarily produces the injective output map $o : \{1, \dots, p\} \rightarrow \{1, \dots, n\}$ defined earlier. For the second condition, letting $x = (x^1, x^2)$, $u = (u^1, u^2)$ and $y = (y^1, y^2)$, we must show that $f(x, u) = f(x + h^{-1}(y), u)$. First, by Assumption III.1 on each system, $f^j(x^j, u^j) = f^j(x^j + (h_{\sigma^j}^j)^{-1}(y^j), u^j)$. Second, it is easy (but tedious) to show that $h_o^{-1}(y) = ((h_{\sigma^1}^1)^{-1}(y^1), (h_{\sigma^2}^2)^{-1}(y^2))$. Putting these two facts together gives the desired result.

(b) We show that for all $m \in M$, $I_{\text{MA}}(m) \subset h^{-1}(Y^*)$. Letting $m = (m^1, m^2) \in M$, we have by the fact that each system is an MA that $I_{\text{MA}}^j(m^j) \subset (h^j)^{-1}(Y^{*,j})$ for $j = 1, 2$. It is easy (but tedious) to show that $h^{-1}(Y^*) = (h^1)^{-1}(Y^{*,1}) \times (h^2)^{-1}(Y^{*,2})$. The result then follows by applying (22).

(c) We show that for all $e = (m_1, \sigma, m_2) \in E_{\text{MA}}$, $g_e \subset h^{-1}(\mathcal{F}_\sigma) \cap I_{\text{MA}}(m_1)$. Let $e = (m_1, \sigma, m_2) \in E_{\text{MA}}$ and decompose it as $e^j = (m_1^j, \sigma^j, m_2^j) \in \overline{E}_{\text{MA}}^j$ for $j = 1, 2$. For $j = 1, 2$, if $\sigma^j \neq \varepsilon^j$, then $g_{e^j} \subset (h^j)^{-1}(\mathcal{F}_{\sigma^j}) \cap I_{\text{MA}}^j(m_1^j)$ since each system is an MA. Otherwise, if $\sigma^j = \varepsilon^j$, observe that $\mathcal{F}_{\varepsilon^j} = Y^{*,j}$ and $g_{e^j} = I^j(m_1^j) \subset I_{\text{MA}}^j(m_1^j) \subset (h^j)^{-1}(Y^{*,j})$ by construction. Consequently $g_{e^j} \subset (h^j)^{-1}(\mathcal{F}_{\sigma^j}) \cap I_{\text{MA}}^j(m_1^j)$ again. Next, by definition $g_e = g_{e^1} \times g_{e^2}$ and $I_{\text{MA}}(m) = I_{\text{MA}}^1(m_1^1) \times I_{\text{MA}}^2(m_2^2)$. It is also easy (but tedious) to show that $h^{-1}(\mathcal{F}_\sigma) = (h^1)^{-1}(\mathcal{F}_{\sigma^1}) \times (h^2)^{-1}(\mathcal{F}_{\sigma^2})$. The result then follows by applying (21) and (22).

(d) We show that for all $e = (m_1, \sigma, m_2) \in E_{\text{MA}}$, $r_e(x) = x - h_o^{-1}(d \circ \sigma)$. Let $e = (m_1, \sigma, m_2) \in E_{\text{MA}}$ and decompose it as $e^j = (m_1^j, \sigma^j, m_2^j) \in \overline{E}_{\text{MA}}^j$ for $j = 1, 2$. First, by definition $r_e(x) = (r_{e^1}(x^1), r_{e^2}(x^2))$. Then for $j = 1, 2$, if $\sigma^j \neq \varepsilon^j$, then $r_{e^j}(x^j) = x^j - (h_{\sigma^j}^j)^{-1}(d^j \circ \sigma^j)$ since each system is an MA. This is also the case when $\sigma^j = \varepsilon^j$ because $r_{e^j}(x^j) = x^j$ and $(h_{\sigma^j}^j)^{-1}(d^j \circ \varepsilon^j) = 0$. Next, since $d = (d^1, d^2)$ and $\sigma = (\sigma^1, \sigma^2)$, component-wise multiplication gives $d \circ \sigma = (d^1 \circ \sigma^1, d^2 \circ \sigma^2)$. Using $x = (x^1, x^2)$ and the decomposition $h_o^{-1}(y) = ((h_{\sigma^1}^1)^{-1}(y^1), (h_{\sigma^2}^2)^{-1}(y^2))$ established in (a) with $y = d \circ \sigma$ proves the result.

(e) We must show that $Q_{\text{MA}}^0 = \{(m, x) \in Q_{\text{MA}} \mid x \in I_{\text{MA}}(m)\}$. This follows immediately from the definitions of Q_{MA} , I_{MA} , and Q_{MA}^0 .

Next we prove that (i)-(vii) of Assumption IV.1 hold.

(i) We must show that for all $m \in M$, $\varepsilon \notin \Sigma_{\text{MA}}(m)$. This follows immediately from the definition of the edges since for all $m \in M$, $\Sigma_{\text{MA}}(m) = \overline{\Sigma}_{\text{MA}}(m) \setminus \{\varepsilon\}$.

(ii) We must show that for all $e_1, e_2 \in E_{\text{MA}}$ such that $e_1 = (m_1, \sigma, m_2)$ and $e_2 = (m_1, \sigma, m_3)$, $g_{e_1} = g_{e_2}$. To that end, we write $e_1^j = (m_1^j, \sigma^j, m_2^j) \in \overline{E}_{\text{MA}}^j$ and $e_2^j = (m_1^j, \sigma^j, m_3^j) \in \overline{E}_{\text{MA}}^j$ for $j = 1, 2$. To show that $g_{e_1} = g_{e_2}$, we must show that $g_{e_1^j} = g_{e_2^j}$ for $j = 1, 2$. Let $j \in \{1, 2\}$. If $\sigma^j = \varepsilon^j$, then by construction $g_{e_1^j} = I^j(m_1^j) = g_{e_2^j}$. Otherwise, if $\sigma^j \neq \varepsilon^j$, then $g_{e_1^j} = g_{e_2^j}$ follows from Assumption IV.1 (ii) on the j -th system.

(iii) We must show that for all $e_1, e_2 \in E_{\text{MA}}$ such that $e = (m_1, \sigma_1, m_2)$ and $e_2 = (m_1, \sigma_2, m_3)$, if $\sigma_1 \neq \sigma_2$, then $g_{e_1} \cap g_{e_2} = \emptyset$. To that end, we write $e_1^j = (m_1^j, \sigma_1^j, m_2^j) \in \overline{E}_{\text{MA}}^j$ and $e_2^j = (m_1^j, \sigma_2^j, m_3^j) \in \overline{E}_{\text{MA}}^j$ for $j = 1, 2$. If $\sigma_1 \neq \sigma_2$, then suppose w.l.o.g. that $\sigma_1^1 \neq \sigma_2^1$. To show $g_{e_1} \cap g_{e_2} = \emptyset$, by (21) it suffices to show that $g_{e_1^1} \cap g_{e_2^1} = \emptyset$. If both σ_1^1 and σ_2^1 are not equal to ε^1 , then by Assumption IV.1 (iii) $g_{e_1^1} \cap g_{e_2^1} = \emptyset$. If one of σ_1^1 or σ_2^1 is ε^1 , say σ_1^1 , then we cannot invoke Assumption IV.1 (iii). However, by definition $g_{e_1^1} = I^1(m_1^1)$ is not intersecting with any other guards, so that $g_{e_1^1} \cap g_{e_2^1} = \emptyset$, as desired.

(iv) We must show that for all $e_1, e_2 \in E_{\text{MA}}$ such that $e_1 = (m_1, \sigma_1, m_2)$ and $e_2 = (m_2, \sigma_2, m_3)$, $r_{e_1}(g_{e_1}) \cap g_{e_2} = \emptyset$.

To that end, we write $e_1^j = (m_1^j, \sigma_1^j, m_2^j) \in \bar{E}_{\text{MA}}^j$ and $e_2^j = (m_2^j, \sigma_2^j, m_3^j) \in \bar{E}_{\text{MA}}^j$ for $j = 1, 2$. By (21), it suffices to show that $r_{e_1^j}(g_{e_1^j}) \cap g_{e_2^j} = \emptyset$ for at least one of $j = 1, 2$. Since ε cannot be an event by Assumption IV.1 (i), at least one of $j_1 = 1, 2$ must have $\sigma_1^{j_1} \neq \varepsilon^{j_1}$, and at least one of $j_2 = 1, 2$ must have $\sigma_2^{j_2} \neq \varepsilon^{j_2}$. Formally, there are several cases, but they can be summarized as follows. If there is at least one matching $j = j_1 = j_2$, where both $\sigma_1^j, \sigma_2^j \neq \varepsilon^j$, then both $e_1^j, e_2^j \in E_{\text{MA}}^j$, and we invoke Assumption IV.1 (iv) to get $r_{e_1^j}(g_{e_1^j}) \cap g_{e_2^j} = \emptyset$. There are two remaining cases which are similar so we only look at one of them. Suppose $\sigma_1^1 = \varepsilon^1$ and $\sigma_2^2 = \varepsilon^2$. Then $e_1^1 \notin E_{\text{MA}}^1$ and $e_2^2 \notin E_{\text{MA}}^2$ so we cannot invoke Assumption IV.1 (iv). However, by construction $e_1^1 \in \bar{E}_{\text{MA}}^1$ implies that $r_{e_1^1}(g_{e_1^1}) = I^1(m_1^1)$ is not intersecting with any guards in E_{MA}^1 . In particular, $e_2^1 \in E_{\text{MA}}^1$ since it must be that $\sigma_2^1 \neq \varepsilon^1$, and so $r_{e_1^1}(g_{e_1^1}) \cap g_{e_2^1} = \emptyset$.

(v) We must show that for all $e = (m_1, \sigma, m_2) \in E_{\text{MA}}$, $r_e(g_e) \subset I_{\text{MA}}(m_2)$. To that end, we write $e^j = (m_1^j, \sigma^j, m_2^j) \in \bar{E}_{\text{MA}}^j$ for $j = 1, 2$. For $j = 1, 2$, if $\sigma^j \neq \varepsilon^j$, then $r_{e^j}(g_{e^j}) \subset I_{\text{MA}}^j(m_2^j)$ follows from Assumption IV.1 (v) on the individual system. This is also the case when $\sigma^j = \varepsilon^j$, since by definition of \bar{E}_{MA}^j we have $r_{e^j}(g_{e^j}) = I^j(m_1^j) \subset I_{\text{MA}}^j(m_2^j)$. Next by definition we have that $I_{\text{MA}}(m_2) = I_{\text{MA}}^1(m_2^1) \times I_{\text{MA}}^2(m_2^2)$. Also, it is easy to verify that $r_e(g_e) = r_{e^1}(g_{e^1}) \times r_{e^2}(g_{e^2})$. The result follows by applying (22).

(vi) We must show that for all $m \in M$, if $\Sigma_{\text{MA}}(m) = \emptyset$ then $I_{\text{MA}}(m)$ is invariant. Let $m = (m^1, m^2) \in M$, $x_0 = (x_0^1, x_0^2) \in I_{\text{MA}}(m)$, and suppose that $\Sigma_{\text{MA}}(m) = \emptyset$. Then for $j = 1, 2$, $\Sigma_{\text{MA}}^j(m^j) = \emptyset$. To see this, suppose one was not empty, say $\sigma^1 \in \Sigma_{\text{MA}}^1(m^1) \subset \bar{\Sigma}_{\text{MA}}^1(m^1)$, where by Assumption IV.1 (i) $\sigma^1 \neq \varepsilon^1$. By construction, $\varepsilon^2 \in \bar{\Sigma}_{\text{MA}}^2(m^2)$ and by Assumption IV.1 (i) proven above $(\sigma^1, \varepsilon^2) \in \bar{\Sigma}_{\text{MA}}(m) \setminus \{\varepsilon\} = \Sigma_{\text{MA}}(m)$, so we get a contradiction. Appealing to the Assumption IV.1 (vi) for $j = 1, 2$, for all $t \geq 0$ the individual trajectories satisfy $\phi_{\text{MA}}^j(t, x_0^j) \in I_{\text{MA}}^j(m^j)$. Thus for all $t \geq 0$, $\phi_{\text{MA}}(t, x_0) \in I_{\text{MA}}(m)$.

(vii) We must show that for all $m \in M$, if $\Sigma_{\text{MA}}(m) \neq \emptyset$ then $I_{\text{MA}}(m)$ forces all trajectories to exit in finite time through some guard. Let $m = (m^1, m^2) \in M$, $x_0 = (x_0^1, x_0^2) \in I_{\text{MA}}(m)$, and suppose that $\Sigma_{\text{MA}}(m) \neq \emptyset$. Reversing the argument used in Assumption IV.1 (vi), now we can conclude for at least one of $j = 1, 2$ that $\Sigma_{\text{MA}}^j(m^j) \neq \emptyset$. Suppose first that only $\Sigma_{\text{MA}}^1(m^1) \neq \emptyset$, then applying Assumption IV.1 (vii) for $j = 1$ and Assumption IV.1 (vi) for $j = 2$ furnishes the event $\sigma = (\sigma^1, \varepsilon^2) \in \Sigma_{\text{MA}}(m)$, with exit time T^1 . A similar argument holds if only $\Sigma_{\text{MA}}^2(m^2) \neq \emptyset$. If for both $j = 1, 2$, $\Sigma_{\text{MA}}^j(m^j) \neq \emptyset$, then Assumption IV.1 (vii) for both $j = 1, 2$ furnishes the event $\sigma = (\sigma^1, \sigma^2)$ with $\sigma^j \in \Sigma_{\text{MA}}^j(m^j)$, with exit times T^j . If $T^1 < T^2$, then the overall exit time is $T = T^1$ with event $\sigma = (\sigma^1, \varepsilon^2) \in \Sigma_{\text{MA}}(m)$. If $T^1 > T^2$, then the exit time is $T = T^2$ with event $\sigma = (\varepsilon^1, \sigma^2) \in \Sigma_{\text{MA}}(m)$. Otherwise, $T = T^1 = T^2$, and the event is $\sigma = (\sigma^1, \sigma^2) \in \Sigma_{\text{MA}}(m)$. In all the cases above, it is then easy (but tedious) to verify that for all $e = (m, \sigma, m') \in E_{\text{MA}}$ and for all $t \in [0, T]$, $\phi_{\text{MA}}(t, x_0) \in I_{\text{MA}}(m)$ and $\phi_{\text{MA}}(T, x_0) \in g_e$.

Proof of Lemma VII.1. (i) We must show that for all $m \in M$, $0 \notin \Sigma_{\text{MA}}(m)$. This is clearly true since there is no edge in E_{MA} containing the label 0.

(ii) We must show that for all $e_1, e_2 \in E_{\text{MA}}$ such that $e_1 = (m_1, \sigma, m_2)$ and $e_2 = (m_1, \sigma, m_3)$, $g_{e_1} = g_{e_2}$. This is clearly true since we have designed $g_{(\mathcal{F}, 1, \mathcal{H})} = g_{(\mathcal{F}, 1, \mathcal{F})}$ and $g_{(\mathcal{B}, -1, \mathcal{H})} = g_{(\mathcal{B}, -1, \mathcal{B})}$.

(iii) We must show that for all $e_1, e_2 \in E_{\text{MA}}$ such that $e_1 = (m_1, \sigma_1, m_2)$ and $e_2 = (m_1, \sigma_2, m_3)$, if $\sigma_1 \neq \sigma_2$, then $g_{e_1} \cap g_{e_2} = \emptyset$. This is trivially true since for all $m \in M$, $|\Sigma_{\text{MA}}(m)| < 2$.

(iv) We must show that for all $e_1, e_2 \in E_{\text{MA}}$ such that $e_1 = (m_1, \sigma_1, m_2)$ and $e_2 = (m_2, \sigma_2, m_3)$, $r_{e_1}(g_{e_1}) \cap g_{e_2} = \emptyset$. Using Assumption IV.1 (ii) above, we only need to check two cases, that is, $e_1 = e_2 = (\mathcal{F}, 1, \mathcal{F})$ and $e_1 = e_2 = (\mathcal{B}, -1, \mathcal{B})$. Both cases satisfy the condition because of the reset action on the first coordinate; for example, the first case gives $r_{e_1}(\{d\} \times (0, \bar{u}_1]) \cap \{d\} \times (0, \bar{u}_1] = \emptyset$.

(v) We must show that for all $e = (m_1, \sigma, m_2) \in E_{\text{MA}}$, $r_e(g_e) \subset I_{\text{MA}}(m_2)$. This is easily verified for all four edges in E_{MA} , for example, if $e = (\mathcal{F}, 1, \mathcal{H})$, clearly $r_{e_1}(\{d\} \times (0, \bar{u}_1]) \subset I_{\text{MA}}(\mathcal{H})$.

(vi) We must show that for all $m \in M$, if $\Sigma_{\text{MA}}(m) = \emptyset$ then $I_{\text{MA}}(m)$ is invariant. We have that only $\Sigma_{\text{MA}}(\mathcal{H}) = \emptyset$. As can be seen in Figure 5, the closed-loop vector field does not allow trajectories to cross outside of $I_{\text{MA}}(\mathcal{H})$, and therefore for all $x_0 \in I_{\text{MA}}(\mathcal{H})$, and for all $t \geq 0$, $\phi_{\text{MA}}(t, x_0) \in I_{\text{MA}}(\mathcal{H})$.

(vii) We must show that for all $m \in M$, if $\Sigma_{\text{MA}}(m) \neq \emptyset$ then $I_{\text{MA}}(m)$ forces all trajectories to exit in finite time through some guard. Consider \mathcal{F} with $\Sigma_{\text{MA}}(\mathcal{F}) = \{1\}$. As can be seen in Figure 5, for all $x_0 \in I_{\text{MA}}(\mathcal{F})$, there exists $T \geq 0$ such that for all $t \in [0, T]$, $\phi_{\text{MA}}(t, x_0) \in I_{\text{MA}}(\mathcal{F})$, and $\phi_{\text{MA}}(T, x_0) \in \{d\} \times (0, \bar{u}_1]$. Since both $g_{(\mathcal{F}, 1, \mathcal{H})} = g_{(\mathcal{F}, 1, \mathcal{F})} = \{d\} \times (0, \bar{u}_1]$, the assumption holds. A similar argument can be made for \mathcal{B} . \square

REFERENCES

- [1] B. d'Andrea-Novet, G. Bastin, and G. Campion. "Modelling and control of non-holonomic wheeled mobile robots". *International Conference on Robotics and Automation*, 1991.
- [2] F. Augugliaro, A. P. Schoellig, and R. D'Andrea. "Generation of collision-free trajectories for a quadcopter fleet: A sequential convex programming approach". *International Conference on Intelligent Robots and Systems*, 2012.
- [3] N. Ayanian and V. Kumar. "Decentralized feedback controllers for multiagent teams in environments with obstacles". *IEEE Transactions on Robotics*. Vol. 26, no. 5, pp. 878-887, 2010.
- [4] M.E. Broucke and M. Ganness. "Reach control on simplices by piecewise affine feedback". *SIAM J. Control and Optimization*. Vol. 52, no. 5, pp. 3261-3286, 2014.
- [5] M. E. Broucke, S. Di Gennaro, M. Di Benedetto, and A. Sangiovanni-Vincentelli. "Efficient solution of optimal control problems using hybrid systems". *SIAM Journal on Control and Optimization*, Vol. 43, no. 6, pp. 1923-1952, 2005.
- [6] M. Čáp, P. Novák, A. Kleiner, and M. Selecký. "Prioritized planning algorithms for trajectory coordination of robots". *IEEE Transactions on Automation Science and Engineering*, Vol. 12, no. 3, pp. 835-849, 2015.
- [7] Y. Chen, X. C. Ding, A. Stefanescu, and C. Belta. "Formal approach to the deployment of distributed robotic teams". *IEEE Transactions on Robotics*, Vol. 28, no. 1, pp. 158-171, 2012.

- [8] M. Cirillo, T. Uras, and S. Koenig. "A lattice-based approach to multi-robot motion planning for non-holonomic vehicles". *International Conference on Intelligent Robots and Systems*, 2014.
- [9] D. C. Conner, A. A. Rizzi, and H. Choset. "Composition of local potential functions for global robot control and navigation". *International Conference on Intelligent Robots and Systems*, 2003.
- [10] A. Dobson, K. Solovey, R. Shome, D. Halperin, and K. E. Bekris. "Scalable asymptotically-optimal multi-robot motion planning". *International Symposium on Multi-Robot and Multi-Agent Systems*, 2017.
- [11] G. E. Fainekos, A. Girard, H. Kress-Gazit, and G. J. Pappas. "Temporal logic motion planning for dynamic robots". *Automatica*. Vol. 45, no. 2, pp. 343-352, 2009.
- [12] E. Frazzoli, M. A. Dahleh, and E. Feron. "Maneuver-based motion planning for nonlinear systems with symmetries". *IEEE Transactions on Robotics*. Vol. 21, no. 6, pp. 1077-1091, 2005.
- [13] M. Guo and D. V. Dimarogonas. "Multi-agent plan reconfiguration under local LTL specifications". *The International Journal of Robotics Research*, Vol. 34, no. 2, pp. 218-235, 2015.
- [14] L.C.G.J.M. Habets, P.J. Collins, and J.H. van Schuppen. "Reachability and control synthesis for piecewise-affine hybrid systems on simplices". *IEEE Transactions on Automatic Control*. Vol. 51, no. 6, pp. 938-948, 2006.
- [15] D. Heß, M. Althoff, and T. Sattel. "Formal verification of maneuver automata for parameterized motion primitives". *International Conference on Intelligent Robots and Systems*. 2014.
- [16] K. H. Johansson, M. Egerstedt, J. Lygeros, and S. Sastry. "On the regularization of Zeno hybrid automata". *Systems and Control Letters*. Vol. 38, no. 3, pp. 141-150, 1999.
- [17] M. Kloetzer and C. Belta. "A fully automated framework for control of linear systems from temporal logic specifications". *IEEE Transactions on Automatic Control*. Vol. 53, no. 1, pp. 287-297, 2008.
- [18] M. Kloetzer and C. Belta, "Dealing with non-determinism in symbolic control". *International Workshop on Hybrid Systems: Computation and Control* 2008.
- [19] S. M. LaValle. *Planning Algorithms*. Cambridge University Press. 2006.
- [20] S. R. Lindemann and S. M. LaValle. "Simple and efficient algorithms for computing smooth, collision-free feedback laws over given cell decompositions". *International Journal of Robotics Research*, Vol. 28, no. 5, pp. 600-621, 2009.
- [21] J. Lygeros, K.H. Johansson, S. Sastry, and M. Egerstedt. "On the existence of executions of hybrid automata." *IEEE Conference on Decision and Control*, 1999.
- [22] A. Majumdar and R. Tedrake. "Funnel libraries for real-time robust feedback motion planning". *International Journal of Robotics Research*, Vol. 36, no. 8, pp. 947-982, 2017.
- [23] D. Mellinger and V. Kumar. "Minimum snap trajectory generation and control for quadrotors". *International Conference on Robotics and Automation*, 2011
- [24] P. Mihail, R. A. Knepper, and A. Kelly. "Differentially constrained mobile robot motion planning in state lattices". *Journal of Field Robotics*, Vol. 26, no. 3, pp. 308-333, 2009.
- [25] J. A. Preiss, W. Hönig, N. Ayanian, and G. S. Sukhatme. "Downwash-aware trajectory planning for large quadrotor teams". *International Conference on Intelligent Robots and Systems*, 2017.
- [26] V. Raman and H. Kress-Gazit. "Synthesis for multi-robot controllers with interleaved motion". *International Conference on Robotics and Automation*, 2014.
- [27] B. Roszak and M. E. Broucke. "Necessary and sufficient conditions for reachability on a simplex". *Automatica*. Vol. 42, no. 11, pp. 1913-1918, 2006.
- [28] I. Saha, R. Ramaithitima, V. Kumar, G. J. Pappas, and S. A. Seshia. "Automated composition of motion primitives for multi-robot systems from safe LTL specifications". *International Conference on Intelligent Robots and Systems*, 2014.
- [29] M.W. Spong, S. Hutchinson, and M. Vidyasagar "Robot Modeling and Control". New York: Wiley, 2005.
- [30] P. Tabuada, G. J. Pappas, and P. Lima, "Compositional abstractions of hybrid control systems". *IEEE Conference on Decision and Control*, 2001.
- [31] M. Vukosavljev, Z. Kroeze, M. E. Broucke, and A. P. Schoellig. "A framework for multi-vehicle navigation using feedback-based motion primitives". *International Conference on Intelligent Robots and Systems*, 2017.
- [32] G. Wagner and H. Choset, "Subdimensional expansion for multirobot path planning". *Artificial Intelligence* Vol. 219, no. 2, pp. 1-24, 2015.
- [33] E. M. Wolff, U. Topcu, and R. Murray. "Optimal control of non-deterministic systems for a computationally efficient fragment of temporal logic". *IEEE Conference on Decision and Control*. December 2013.
- [34] W.M. Wonham and K. Cai. *Supervisory Control of Discrete-Event Systems*. <http://www.control.utoronto.ca/~wonham/Research.html>.
- [35] J. Yu and D. Rus. "An Effective Algorithmic Framework for Near Optimal Multi-Robot Path Planning". Robotics Research, Springer Berlin/Heidelberg, Vol. 1, pp. 495-511, 2018.
- [36] J. Zhang, K. H. Johansson, J. Lygeros, and S. Sastry, "Dynamical systems revisited: hybrid systems with Zeno executions". *Hybrid Systems: Computation and Control*, 2000.



Marijan Vukosavljev received the BAsC in Aerospace Engineering in 2013 and is currently completing his Ph.D. studies in Systems Control, both from the University of Toronto. His research interests include hybrid systems, optimal control, and applications to drone swarms.



Zachary Kroeze received the BAsC in Mathematics and Engineering from Queens University in 2012, and his Ph.D. in Systems Control from the University of Toronto in 2019. His research interests include hybrid systems, reach control, and applications to self-driving cars.



Angela P. Schoellig received an M.Sc. in Engineering Science and Mechanics from the Georgia Institute of Technology, a Masters degree in Engineering Cybernetics from the University of Stuttgart, and a Ph.D. from ETH Zurich, in 2007, 2008, and 2012 respectively. She is currently an Assistant Professor at the University of Toronto Institute for Aerospace Studies where she leads the Dynamic Systems Lab. Her general areas of interest lie at the interface of robotics, controls and machine learning.



Mireille E. Broucke received the BAsC from University of Texas at Austin in 1984, and the MASc and Ph.D. from University of California, Berkeley, in 1987 and 2001 respectively. She is a Professor at the University of Toronto. Her current research interests are in the application of control theory to neuroscience.