On Quantization and Optimal Control of Dynamical Systems with Symmetries

Emilio Frazzoli[1]
Aeronautical and Astronautical Engineering
University of Illinois at Urbana-Champaign
Urbana, IL 61801, United States
frazzoli@uiuc.edu, http://brunello.aae.uiuc.edu

Francesco Bullo[2]
Coordinated Science Laboratory
University of Illinois at Urbana-Champaign
Urbana, IL 61801, United States
bullo@uiuc.edu, http://motion.csl.uiuc.edu

Abstract

This paper outlines some progress in the study of quantized control systems. We consider the dynamical systems arising from robotics and autonomous vehicles applications and describe a procedure to transcribe their dynamics into finite-state automata. The transcription procedure extends known previous results. We consider optimal control problems and present conditions under which linear programming algorithms are helpful in trajectory planning.

1 Introduction

A new and promising direction of research in control theory is centered on the purposeful introduction of state and/or control quantization in the design of control systems. The purpose of quantization is, in general, a reduction of the complexity of the control task, mainly in terms on computation and communications requirements, but also in terms of practical implementation of sensors, actuators, and control logics.

In this paper we are concerned with the impact of quantization of control strategies on the complexity of motion planning problems for dynamical systems with symmetries, such as robotic mechanisms and autonomous vehicles. Even though our main concern for this paper is computational complexity, we remark that an additional critical advantage of the proposed approach is the quantization of information characterizing the control action. This advantage is expected to make this set of techniques particularly attractive for applications including communications over bandwidth-limited channels. Several approaches to the solution of this kind of problems have been recently developed, based on the choice of a finite number of elementary control actions, or control laws, which are combined to generate more complex behaviors and ultimately achieve the desired objective.

[1]Partially funded by NSF grant CCR-0133869.
[2]Partially funded by ARO grant DAAD 190110716 and NSF grant IIS-0118146.

We will develop a unified framework which includes as special cases several different techniques, based on the work reported in [1, 2, 3]. The notion of "control quanta" was first introduced by Marigo *et al.* [4, 1], as a technique for steering nonlinear driftless systems. The so-called "maneuver automaton", developed by Frazzoli *et al.* [2, 5] is exploits "trim trajectories" and "relative equilibria" in mechanics to obtain a finite-state automaton description of a vehicle dynamics. The notion of decoupling vector field was introduced by Bullo and Lynch [3] for mechanical control systems to formalize equivalence notions between kinematic and dynamical systems.

In this paper, we investigate the symmetry assumptions enabling this quantization and abstraction of the dynamics. We then review and summarize conditions guaranteeing local and global controllability of the quantized system. Finally, we provide some novel observations on the optimal planning problem.

Our modeling paradigm is related to efforts to develop languages and reactive behaviors in robots, see [6, 7, 8], and to efforts to characterize hierarchical abstractions of control systems, see [9]. However, our focus is on controllability analysis, and on feasible and optimal trajectory design. Our transcription procedure is also somehow related classic (direct and indirect) transcription methods in trajectory optimization such as collocation, shooting, and differential inclusion; see [10, 11].

2 Problem Formulation

In this paper we consider smooth nonlinear control systems of the form

$$\frac{d}{dt}x = f(x, u) \tag{1}$$

where x takes value in the smooth manifold $\mathcal{M}$. In what follows, we shall tacitly assume that all relevant quantities are smooth.

The basic problem that we want to address is the following: given an initial state x_{initial}, and a goal state

x_{target}, find an input signal $t \in [0, T] \mapsto u^*(t)$ such that the solution to the dynamical system (1) with initial condition $x(0) = x_{\text{initial}}$ and with $u = u^*$ satisfies $x(T) = x_{\text{target}}$. Moreover, we also address the optimal control problem, i.e., we search for a controlled trajectory satisfying the desired boundary conditions while minimizing a cost functional of the form

$$J(x(\cdot), u(\cdot)) = \int_0^T c(x(t), u(t))\, dt. \tag{2}$$

We refer to these control problems as the motion planning or the steering problem and, when specifically requiring optimality, as the trajectory optimization problem. The basic steering problem is the motivation for a wide body of research in nonlinear controllability. It is well known that the set of points reachable by a control system is an open set in the configuration space if and only if the so-called Lie algebra rank condition (LARC) is satisfied; see [12]. The LARC is a necessary and sufficient condition for local nonlinear controllability (STLC) for driftless systems, and sufficient tests are available for STLC of systems with drift; see [13].

3 Transcriptions of control systems and the role of symmetries

Our approach to the motion planning problem is based on the selection of a finite number of elementary motions, which we for now loosely call motion primitives. Such building blocks are then combined in such a way as to produce more complex trajectories, satisfying the boundary conditions imposed by the steering problem, and approximating (by providing non-tight upper bounds) solutions to the optimal control problem.

Let X be a vector field on $\mathcal{M}$. We let $\Phi^X_{0,t}(x)$ denote the flow of X from time 0 to time $t > 0$ from initial condition $x \in \mathcal{M}$. A curve $\gamma : [0, T] \to \mathcal{M}$ is a **controlled trajectory** for the control system (1) if there exists a time-dependent control law $u_\gamma : [0, T] \times \mathcal{M} \mapsto \mathbb{R}^m$ such that γ is a solution to the closed-loop system induced by $u = u_\gamma$.

Let $\mathcal{N}$ be a submanifold of $\mathcal{M}$. A pair $(X, \mathcal{N})$ is a **trim motion** for the control system (1) if $\mathcal{N}$ is invariant under the flow of X and if, for all $x \in \mathcal{N}$, the flow map $t \mapsto \Phi^X_{0,t}(x)$ is a controlled trajectory for (1).

A pair $(X, \mathcal{N})$ is a **feasible or decoupling motion** for the control system (1) if $\mathcal{N}$ is invariant under the flow of X and if, for all $x \in \mathcal{N}$ and for all $T > 0$, there exists a controlled trajectory $\gamma : [0, T] \to \mathcal{M}$ for the control system (1) such that $\gamma(T) = \Phi^X_{0,T}(x)$. Clearly, a trim motion is a feasible motion, but the opposite is not true.

The control system (1) is **controllable from submanifold $\mathcal{X}$ to submanifold $\mathcal{Y}$** if for all $x \in \mathcal{X}$, there exists a finite-duration controlled trajectory with initial condition x to a point in $\mathcal{Y}$. This is trivially true if $\mathcal{X} \subset \mathcal{Y}$.

Finally, let us exploit these definitions to provide an alternative description of the evolution of the control system (1). Assume that we have computed a collection $\{(X_i, \mathcal{N}_i), i \in \{1, \ldots, n\}\}$ of feasible motions for the control system (1), and we have characterized when is the control system controllable from any submanifold $\mathcal{N}_i$ to any other submanifold $\mathcal{N}_j$ (including the case $j = i$). Then, we could restrict the dynamics of the control system to this set of feasible motions and to the finite-duration trajectories switching among them. Clearly, such a description of the control system (1) would not be as rich as the original differential equations. However, such a description could be advantageous when (i) the feasible motions and the switching trajectories are easily computable and representable via finite dimensional numerical objects, (ii) the resulting low-complexity representation maintains some basic properties of the original control system, in particular controllability.

3.1 Symmetries and invariant control problems

To reduce the complexity of the automaton representation we introduce symmetries in the system dynamics and we exploit optimal control problems and feasible motion primitives that are invariant. Let us define the required invariance notions in this section.

We assume that the state can be partitioned, at least locally, into the Cartesian product of two manifolds $\mathcal{M} = \mathcal{G} \times \mathcal{Z}$, where $\mathcal{G}$ is a Lie group with identity element e. Accordingly, we write the generic point $x \in \mathcal{M}$ as the pair $(g, z) \in \mathcal{G} \times \mathcal{Z}$. Following conventions from differential geometry, $\mathcal{Z}$ is the base space, $\mathcal{G}$ is the fiber, and their product $\mathcal{M}$ is a principal fiber bundle; see [14]. We let $\mathfrak{g}$ denote the Lie algebra of $\mathcal{G}$.

We also assume that the dynamics of the control system (1) are **invariant** with respect to the left action of $\mathcal{G}$ onto $\mathcal{M} \times \mathcal{U}$

$$\begin{aligned}\Phi : \mathcal{G} \times \mathcal{M} \times \mathcal{U} &\to \mathcal{M} \times \mathcal{U}\\ (h, x, u) &\mapsto \Phi_h(x, u) = ((hg, z), u).\end{aligned}$$

Invariance is equivalent to the following statement. Given any trajectory $t \mapsto (\gamma(t), u(t)) \in \mathcal{M} \times \mathcal{U}$ solution to equation (1), the trajectory $t \mapsto \Phi_h(\gamma(t), u(t))$ is also a solution to equation (1) for all $h \in \mathcal{G}$.

For invariant control systems defined over principal bundles, we consider feasible motions $(X, \mathcal{N})$ of the following form. We assume the vector field X is left invariant (hence characterized by a vector in the Lie algebra $\mathfrak{g}$ of the group $\mathcal{G}$) and leaves the base variables unchanged. We assume the submanifold $\mathcal{N}$ is of the form $\mathcal{G} \times \{z_\mathcal{N}\}$, for some $z_\mathcal{N} \in \mathcal{Z}$. We call such feasible motions **invariant**.

Finally, we consider control problems that have invariance properties. In particular, for invariant control systems, we shall consider motion planning problems with initial and goal states of the form $x_{\text{initial}} = (g_{\text{initial}}, z_0)$ and $x_{\text{target}} = (g_{\text{target}}, z_0)$ where z_0 is a fixed base point in $\mathcal{Z}$. Furthermore, we shall restrict our attention to optimal control problems with incremental cost functions c that are invariant with respect to the group action Φ, i.e., $c = c \circ \Phi_g$, for all $g \in \mathcal{G}$. We call such motion planning and trajectory optimization problems **invariant**.

4 Motion Primitive Automata

In what follows, let us assume that we have computed a collection $\{(X_i, \mathcal{N}_i), i \in \{1, \ldots, n\}\}$ of invariant feasible motions for the invariant control system (1), and we have characterized when is the control system controllable from any invariant submanifold $\mathcal{N}_i$ to any other invariant submanifold $\mathcal{N}_j$. Before proceeding, let us present some remarkable properties of invariant feasible motions. First, the flow along the vector field X_i can be computed in closed-form by means of the matrix exponential map on $\mathcal{G}$. Integral curves of X_i through the identity element e are one-parameter subgroups of $\mathcal{G}$, and the value of X_i on the manifold $\mathcal{G} \times \mathcal{Z}$ is uniquely determined by a Lie algebra element, say $\xi_i \in \mathfrak{g}$. Second, the system (1) is controllable from the submanifold $\mathcal{N}_i = \mathcal{G} \times \{z_i\}$ to the submanifold $\mathcal{N}_j = \mathcal{G} \times \{z_j\}$ if and only if a controlled trajectory exists connecting one point in $\mathcal{N}_i$ to one point in $\mathcal{N}_j$.

Following the nomenclature in [5], we call **maneuver** a finite-time controlled trajectory that connects the submanifolds $\mathcal{N}_i$ and $\mathcal{N}_j$ corresponding to two invariant feasible motions. Note that the net effect on the state of each maneuver is a displacement on the fiber and a possibly a jump on the base space. Specifically, we let $g_{ij} \in \mathcal{G}$ denote the displacement corresponding to a maneuver from feasible motion i to j.

Maneuvers and integral curves of invariant feasible motions are called **motion primitives**. The elementary control action can be constructed as the combination of a (finite- or zero-length) motion along a feasible motion, and a maneuver originating at the same feasible motion. In this setting, all possible control strategy can be represented as the outputs of a finite state, timed, automaton called the **motion primitive automaton** (MPA). It is convenient to depict the allowable primitives, and the rules for their sequential combination through a directed graph $\text{MPA}(V, E)$, such as the one shown in Fig. (1), in which the vertices represent the families of motion primitives generated by invariant feasible motions, and the edges represent the (finite or trivial) motion primitives corresponding to the transition between them. Each edge $i \in E$ is labelled by the corresponding maneuver. Each vertex $j \in V$ is labelled by a vector $\xi_j \in \mathfrak{g}$.

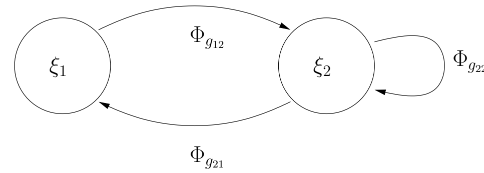

Figure 1: A simple motion primitive automaton.

Motion plans are thus generated by: (1) choosing "how long", in terms of time or distance, the system must stay in the current vertex, and follow the corresponding feasible motion, and (2) the edge to be taken to switch to the next vertex, and feasible motion. Under our assumption, it is immediate to translate sequences of controls on the MPA to the actual state evolution of the controlled physical system. For example, the permanence for "time" τ on node j followed by a maneuver from vertex j to vertex i would result in a displacement on the fiber equal to a right translation by $\exp(\xi_j \tau) g_{ji}$.

Following the procedure outlined in these paragraphs, we can define a motion description language, whose atoms, or elementary symbols, are couples $\alpha = (\tau, p) \in \mathbb{R}_0^+ \times E$; a motion plan ω is a sequence of atoms, i.e., $\omega = \{\alpha_1, \alpha_2, \ldots, \alpha_N\}$. In general, not all possible motion plans are feasible, since not all motion primitives can be combined arbitrarily. In other words, feasible motion plans belong to a subset Ω^* of the free monoid A^*, i.e., $\Omega^* \subseteq A^*$. The feasible sequences can be identified with feasible paths on the graph MPA, with arbitrary, non-negative values of τ_i. In other words, the MPA encodes the syntax of our language, or the rules for symbol concatenation. Related efforts resulted in syntax-free languages, i.e., all possible concatenations of symbols were allowed [15, 7].

In the following subsections we illustrate how some methodologies developed in recent literature fit in this framework.

4.1 Control Quanta

The concept of "control quanta" was first introduced by Marigo *et al.* [4, 1], as a technique for steering nonlinear driftless systems, with a certain structure that leads to the existence of symmetries. Their method consisted of selecting a finite number of control histories on a finite time support (the control quanta). Between the execution of control quanta, the controls are set to zero, thus resulting in an equilibrium point for the system. The execution of each control quanta results in a cycle with no net change in the base variables. In this case, the MPA is depicted by one node and several edges.

The single node corresponds to the trivial vector field $\xi_0 = 0$, whereas each edge corresponds to a control quantum.

4.2 Steering through decoupling vector fields

The notion of decoupling vector field was introduced by Bullo and Lynch [3] for second-order nonlinear control systems. A vector field Ξ is decoupling if any integral curve of Ξ is a kinematic motion, i.e., it is possible to find a time scaling $t \mapsto s(t)$ such that the path is compatible with the system's dynamics (1). In [3, 16] conditions are given for kinematic controllability of a mechanical systems, and motion planning strategies are developed which alternate kinematic motions along decoupling vector fields (and corresponding time scalings for dynamic feasibility, under, e.g., control saturation constraints) to steer the system to arbitrary configurations. The switch between different vector fields occurs at zero velocity, and clearly results in no change in the state variables. For simplicity we will restrict ourselves to the case of invariant vector fields. Moreover, for reasons which will be made clear when constructing a transcription in terms of the MPA of the cost functional, we distinguish between positive and negative motions along a vector field (we need to impose that $\dot{s}$ is non-negative). In our MPA language, the technique introduced by Bullo and Lynch consists of a graph with two vertices for each vector field, corresponding to motions along ξ and $-\xi$. All edges correspond to the identity "maneuver".

4.3 Maneuver Automata

The so-called "Maneuver Automaton", developed by Frazzoli *et al.* [2, 5] is a less general version of the framework proposed in this paper, where the feasible motions are restricted to represent a special class of trajectories for the dynamical systems. These trajectories are known as "trim trajectories" in the aerospace community, and "relative equilibria" in mechanics, and correspond to steady state behaviors for the system. In other words, along trim trajectories, the base variables and control inputs are constant, while the state evolves along the given invariant vector field: $\dot{g}(t) = g(t)\xi$, $\dot{z}(t) = 0$, $\dot{u}(t) = 0$. Each relative equilibrium $j \in V$ is hence described by a triplet (ξ_j, z_j, u_j). In this case, both the maneuvers and the vector fields are non-trivial in general. Because of the way they are defined, the rate of change of the base variables along vector fields is invariant; the only tunable parameter is then the "dwelling" or "coasting" time τ along vector fields, which is constrained to be non-negative.

5 Motion Planning on MPA

Given an initial condition such that the state evolves according to one of the vector fields included in the MPA, all motions generated by arbitrary switching on the graph (and permanence at nodes) result in a feasible trajectory for the original, continuous, dynamical system (1). Hence, the MPA can be considered a **transcription** of the dynamics of the system, under the additional constraint that system trajectories are forced to be a combination of a finite number of primitives. This transcription is to be compared and contrasted with classic (direct and indirect) transcription methods in trajectory optimization such as collocation, shooting, and differential inclusion; see [10, 11].

When the system is in a state corresponding to the inception of one of the feasible motions, the full state of the system can be written as a function of $(g, v) \in \mathcal{G} \times V$ (which we can call the **hybrid state**, due to its mixed continuous/discrete nature). If the hybrid state is (g, v), then the continuous state is equal to (g, z_v), and the control is (u_v). After the execution of a motion along the feasible motion measured by the scalar τ, and a maneuver $p \in E$, the hybrid state is updated according to:

$$(g, v) \to (g \exp(\xi_v \tau) g_p, \mathrm{Next}(p)) \qquad (3)$$

where $\mathrm{Next}(p)$ is the target vertex of the edge p. Equation (3) describes a new dynamical system, and implements system (1), in the sense that every trajectory generated by (3) is feasible for (1). Because of the introduction of additional constraints (the limitation to combination of motion primitives), the set of all possible trajectories for (1) is strictly larger than the set of trajectories generated by (3). However, the computational complexity of many motion planning problems is drastically reduced, as we will show in the following.

5.1 Controllability

As a first step in our analysis, we need to make sure that some desirable properties of the system (1) are not lost in the MPA transcription (3). In particular, we want to make sure that controllability is retained. We loosely define controllability as the condition ensuring existence of a finite-time trajectory connecting two arbitrary states on $\mathcal{G} \times V$, i.e., on the Cartesian product between the fiber and the vertices set. Given a sequence of coasting times and feasible edge transitions $\omega = \{(\tau, p)_i, i = 1, \ldots, N\}$, with $\tau_i \geq \epsilon > 0$, consider the map $M_\omega : \mathbb{R}^N \to \mathcal{G}$ which gives the displacement on the group at the end of the sequence of primitives, when the coasting times are perturbed by amounts $\delta_i < \tau_i$.

Theorem 1 (MPA controllability [5]) *An* MPA *transcription of the system is controllable if and only if: (1) the* MPA *graph is connected, and (2) there exists a fixed-point motion plan $\hat{\omega} \in \Omega^*$, such that $M_{\hat{\omega}}(0) = e$, (3) for any $\epsilon > 0$, the set $\{\hat{M}_\omega(\delta\tau) | \; \|\delta\tau\| < \epsilon\}$ has an open interior.*

We have a few remarks regarding this result.

Remark 2 *Controllability of the* MPA *implies the following notion fiber controllability of the system* (1)*: the original control system can be steered from any initial point of the form* $(g_{initial}, z_i)$ *to any* (g_{target}, z_j) *for all* i, j *in the set of feasible motions. This fiber controllability notion corresponds to the notion of kinematic controllability in [3]. Theorem (1) can be regarded as a simple test for checking controllability of a complicated dynamical system. For example, two corollaries of Theorem (1) given in [5] characterize the minimum sets of motion primitives which ensure configuration controllability for planar robots, and for aircraft-like robots (i.e., for systems with fiber diffeomorphic to* $SE(2)$ *or* $SE(2) \times \mathbb{R}$*).*

Remark 3 *The last condition in Theorem (1) can be recast in differential terms, on the tangent space at the identity* $T_e\mathcal{G}$*, or equivalently on the Lie algebra* $\mathfrak{g}$*, and eventually translates into a Lie Algebra Rank Condition on the vector fields generated by* $\delta\tau$*.*

Remark 4 *In the Control Quanta case, since the only vector field being considered is the null vector field velocity, the* MPA *transcription is not controllable according to our definition. The nature and the topology of reachable sets of such control systems have been extensively studied by Marigo et al. [4, 17, 18]; an interesting result is that in most cases of interest these reachable sets are composed entirely by either accumulation points, or by isolated points (in which case the reachable set has the structure of a lattice).*

5.2 Optimal control

If the MPA transcription is indeed controllable, one can design feasible solutions that steer the system from any initial hybrid state $(g_{\text{initial}}, v_{\text{initial}})$ to any target state $(g_{\text{target}}, v_{\text{target}})$. In the general case, the optimal control problem can be cast as a non-convex program, which must be solved numerically. However, according to the complexity of the primitives composing the MPA, this problem might be tractable in closed form, or through efficient numerical procedures.

In this paper we examine the problem of designing a trajectory that minimizes a cost functional of the form (2). Before proceeding further, we need to develop a transcription of the cost functional into the MPA language. This is easily done for invariant cost functionals. Examples of such cost functionals are those arising in minimum-time, minimum-length, minimum control effort problems, as well as cost functionals which do not depend on the fiber variables. Note that minimum-time problems for the MPA correspond to minimum-time problems for the original control system provided no time-scaling is performed in the transcription. This is not the case for transcriptions of second-order systems based on decoupling vector fields.

Given a motion plan ω that satisfies the boundary conditions on the hybrid states, the cost functional can be rewritten as:

$$J = \sum_{i=1}^{N} \left(\Gamma_{p_i} + \gamma_{v_i}\tau_i\right), \tag{4}$$

where we indicate with v_i the source of the edge p_i. Let τ and p be the vectors with components τ_i and p_i. Given the cost expression (4), the optimal control problem can be recast as the following optimization problem:

$$\begin{aligned} \omega^* = \quad & (\tau, p)^* = \arg\min_{\omega\in\Omega^*} \sum_{i=1}^{|\omega|} \left(\Gamma_{p_i} + \gamma_{v_i}\tau_i\right) \\ \text{s.t.:} \quad & \tau \geq 0 \qquad\qquad\qquad\qquad (\text{or } \tau \geq \epsilon > 0) \\ & M_\omega(0) = \prod_{i=1}^{|\omega|} \exp(\xi_{v_i}\tau_i) g_{p_i} = g_{\text{initial}}^{-1}\, g_{\text{target}}, \end{aligned} \tag{5}$$

with the additional constraint that ω encode a path of length $|\omega|$ on the MPA graph, starting at v_{initial} and ending at v_{target}. We notice that this optimization problem includes a combinatorial aspect, in the choice of the path p on the MPA graph, and in the length of the motion plan. However, once we fix a path on the graph, we get a smooth, generally non-convex optimization problem in the coasting variables τ.

Theorem 5 (Existence of optimal solutions) *If the* MPA *is controllable, and the cost of any non-trivial motion plan is bounded away from zero, i.e., if* $\inf\{J(\omega)|\ \omega \in \Omega^*, |\omega| \geq 1\} = \epsilon > 0$*, then there exists a solution of the optimal control problem (5).*

Proof: Since the MPA is controllable, there exists a finite-length motion plan $\bar{\omega} \in \Omega^*$ which satisfies the boundary conditions, with finite cost. Since the cost of any motion plan with one or more symbols is at least ϵ, there is a finite number of maneuver sequences (i.e. feasible edge transitions) on the MPA graph which could have a cost smaller than $J(\bar{\omega})$, e.g., the set of all maneuver sequences of length bounded by $J(\bar{\omega})/\epsilon$. For each of these maneuver sequences, problem (5), with the addition of the constraint $\sum \gamma_v \tau \leq J(\bar{\omega})$, is a smooth optimization problem over a compact domain: such a problem will either have an (attained) optimal solution, or be unfeasible. Hence, in addition to the solution candidate $\bar{\omega}$ obtained from the controllability theorem, there will be a finite number of additional candidates for an optimal solution. The candidate with the smallest cost is the optimal solution. ∎

Note that the optimal solution is not necessarily unique. Also, note that the assumption on the cost of non-trivial motion plans to be bounded away from zero is needed to avoid infinite sequences of primitives (resulting in Zeno automata, and chattering). Hence, to obtain solutions to optimal control problems when using technique based on decoupling vector fields, we need to impose a minimum length on each vector field. Another solution is that of assigning a positive cost to each switch between vector fields (i.e., to each edge in the graph).

5.3 Optimization on Translational Primitives

In certain cases, the optimal control problem can be solved through a sequence of linear programs. (A similar approach, leading to a Mixed-Integer Linear Programming formulation, was presented for the control quanta case in [19].) Once we fix the maneuver sequence p, the cost function to be minimized is linear in the coasting variables τ. In the case in which fiber $\mathcal{G}$ can be expressed as the Cartesian product of an arbitrary number of copies of the Euclidean group $SE(3)$ and its subgroups, the map M_ω is affine in the coasting variables that correspond to purely translational vector fields. We refer to such coasting variables as "linear."

It is therefore possible to write the optimal control problem with respect to the linear coasting variables as a linear program. Hence we can outline the following algorithm: For increasing values of L, consider all maneuver sequences $p \in E^L$ on the MPA of length L. The cost of the corresponding motion plans can be lower bounded by B_L, with $B_{L+1} - B_L \geq \epsilon > 0$ for some ϵ. For each of these paths, solve the linear program (5) for $\tau \in \mathbb{R}^L$. The solution to each of these LPs is used to update an upper bound on the cost, initialized to $+\infty$; from controllability, we know that there exists a finite path length L such that the LP is feasible, hence U_L will be finite for some finite L. The sequence of updates to the upper bound U_i is non-increasing, while the sequence of updates to the lower bound is strongly increasing: this ensures that the upper and lower bound sequences will converge in a finite number of steps. Completeness of the algorithm is ensured by the same arguments used in proving Theorem 5.

While it is possible to efficiently solve for the coasting times τ, the problem retains its combinatorial nature in the maneuver sequences p. However, in most cases a feasible solution only requires a very small number of maneuvers, typically of the order of the dimension of the fiber $\mathcal{G}$. This means that a feasible solution can be computed exactly in a very short time, while improved solutions (with bounds providing an indication on their quality) can be computed if additional computation time is available.

Within this optimization problem it is also possible to include any convex linear constraints on the coasting times and on the intermediate group variables. For example, it is possible to require for coasting times to be upper and lower bounded, and for the fiber states to evolve inside a convex polytope.

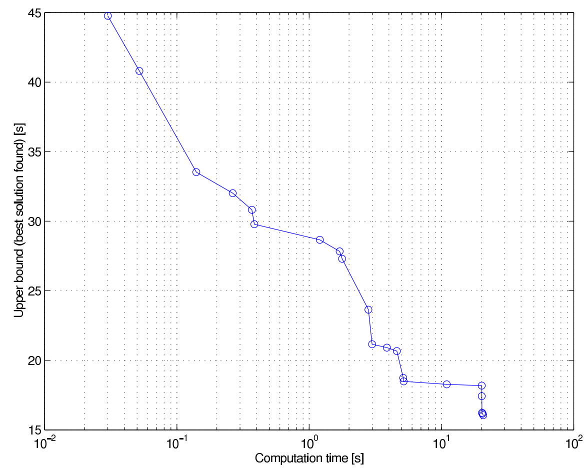

Figure 2: Computation times vs. cost of the solution for the helicopter motion planning example.

6 Example

As an example, we present a simple case of motion planning on a high-fidelity simulation of a helicopter model, as discussed in [20]. The algorithm outlined in the preceding section was run on a 700 MHz Pentium III machine, running Windows XP. The program was written in C, using the `lp_solve` library by Michel Berkelaar (the library is written in ANSI C and freely available for download, e.g. from `ftp://ftp.es.ele.tue.nl/pub/lp_solve/`). The helicopter starts flying eastbound at the maximum allowed speed (i.e., 8 m/s), at coordinates (North, East)=(25m,55m). The helicopter is requested to fly over the origin, northbound, at 8 m/s in the shortest possible time. In Figure 2 a plot of the computed upper bound on the optimal solution is presented, as a function of the computation time. As it can be seen, a feasible solution, with a cost of 44.8 seconds, is computed extremely quickly, i.e. in 3/100 of a second. As more time is available for computation, better solutions are found: the optimal solution, with a cost of 16.0 seconds, is found after about 20 seconds. A plot of the ground traces corresponding to the first computed solution and the best one is given in Figure 3. As it can be seen, the final position of the helicopter is very close to the intended target, even after the execution of several moderately aggressive maneuvers, involving high speeds and bank angles.

7 Conclusion

In this paper we have presented some ideas on quantization of dynamical control systems and their application to optimal control problems. We have transcribed a class of trajectory optimization problems into nonlinear programs over a motion primitive automaton.

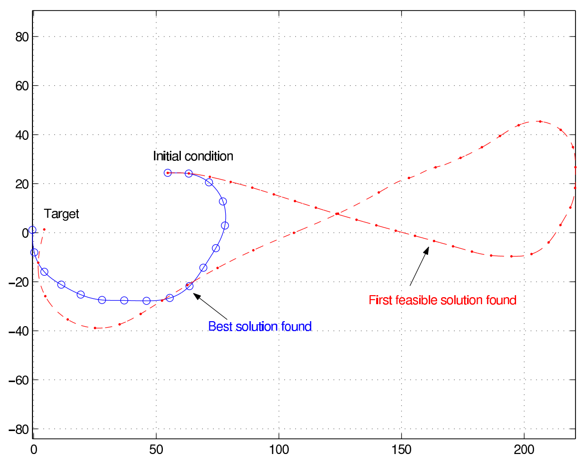

Figure 3: Ground trace of the first computed solution and of the best computed solution. The tick marks along the trajectories represent one-second time increments.

References

[1] A. Bicchi, A. Marigo, and B. Piccoli, "On the reachability of quantized control systems," *IEEE Trans. Automatic Ctrl*, vol. 47, no. 4, pp. 546–63, 2002.

[2] E. Frazzoli, M. A. Dahleh, and E. Feron, "A hybrid control architecture for aggressive maneuvering of autonomous helicopters," in *Proc CDC*, (Phoenix, AZ), pp. 2471–6, Dec. 1999.

[3] F. Bullo and K. M. Lynch, "Kinematic controllability for decoupled trajectory planning in underactuated mechanical systems," *IEEE Trans. Robotics & Automation*, vol. 17, no. 4, pp. 402–412, 2001.

[4] A. Marigo and A. Bicchi, "Steering driftless nonholonomic systems by control quanta," in *Proc CDC*, (Tampa, FL), pp. 4164–9, Dec. 1998.

[5] E. Frazzoli, *Robust Hybrid Control for Autonomous Vehicle Motion Planning*. PhD thesis, MIT, Cambridge, MA, June 2001.

[6] R. A. Brooks, "A robust layered control-system for a mobile robot," *IEEE Journal of Robotics and Automation*, vol. 2, no. 1, pp. 14–23, 1986.

[7] V. Manikonda, P. S. Krishnaprasad, and J. Hendler, "Languages, behaviors, hybrid architectures and motion control," in *Mathematical Control Theory* (J. Baillieul and J. C. Willems, eds.), New York, NY: Springer Verlag, 1998.

[8] R. C. Arkin, *Behavior-Based Robotics*. Cambridge, MA: MIT Press, 1998.

[9] G. J. Pappas, G. Lafferriere, and S. S. Sastry, "Hierarchically consistent control systems," *IEEE Trans. Automatic Ctrl*, vol. 45, no. 6, pp. 1144–60, 2000.

[10] B. A. Conway and K. M. Larson, "Collocation versus differential inclusion in direct optimization," *J. Guidance Control and Dynamics*, vol. 21, no. 5, pp. 780–85, 1998.

[11] J. T. Betts, "Survey of numerical methods for trajectory optimization," *J. Guidance Control and Dynamics*, vol. 21, no. 2, pp. 193–207, 1998.

[12] R. Hermann and A. J. Krener, "Nonlinear controllability and observability," *IEEE Trans. Automatic Ctrl*, vol. 22, pp. 728–740, 1977.

[13] H. J. Sussmann, "A general theorem on local controllability," *SIAM JCO*, vol. 25, no. 1, pp. 158–194, 1987.

[14] S. Kobayashi and K. Nomizu, *Foundations of Differential Geometry. Vol. I*, vol. 15 of *Interscience Tracts in Pure and Applied Mathematics*. New York, NY: Interscience Publishers, 1963.

[15] R. W. Brockett, "Hybrid models for motion control systems," in *Essays in Control: Perspectives in the Theory and its Applications*, pp. 29–53, Boston, MA: Birkhäuser, 1993.

[16] F. Bullo and A. D. Lewis, "Kinematic controllability and motion planning for the snakeboard," *IEEE Trans. Robotics & Automation*, vol. 19, no. 3, pp. 494–498, 2003.

[17] A. Marigo, B. Piccoli, and A. Bicchi, "Reachability analysis for a class of quantized control systems," in *Proc CDC*, (Sydney, Australia), pp. 3963–8, Dec. 2000.

[18] A. Bicchi, A. Marigo, and B. Piccoli, "Quantized control systems and discrete nonholonomy," in *IFAC Workshop on Lagrangian and Hamiltonian Methods for Nonlinear Control*, pp. 19–26, Mar. 2000.

[19] S. Pancanti, L. Leonardi, L. Pallottino, and A. Bicchi, "Optimal control of quantized input systems," in *Hybrid Systems: Computation and Control V* (C. Tomlin and M. Greenstreet, eds.), Lecture Notes in Computer Science, Springer Verlag, 2002.

[20] E. Frazzoli, "Maneuver-based motion planning and coordination for single and multiple uavs," in *AIAA Conf. on UAV Systems, Technologies, and Operations*, AIAA Paper 2002-3472, (Portsmouth, VA), 2002.