

A Randomised Kinodynamic Planner for Closed-chain Robotic Systems

Ricard Bordalba, Lluís Ros, and Josep M. Porta

Abstract—Kinodynamic RRT planners are effective tools for finding feasible trajectories in many classes of robotic systems. However, they are hard to apply to systems with closed-kinematic chains, like parallel robots, collaborative arms manipulating an object, or legged robots keeping their feet in contact with the environment. The state space of such systems is an implicitly-defined manifold that complicates the design of the sampling and steering procedures, and leads to trajectories that drift from the manifold if standard integration methods are used. To address these issues, this paper presents a kinodynamic RRT planner that constructs an atlas of the state space incrementally, and uses this atlas to generate random states, and to dynamically steer the system towards such states. The steering method exploits the atlas charts to compute locally optimal policies based on linear quadratic regulators. The atlas also allows the integration of the equations of motion as a differential equation on the manifold, which eliminates any drift from such manifold and results in accurate trajectories. To the best of our knowledge, this is the first kinodynamic planner that explicitly takes closed kinematic chains into account. We illustrate the planner performance in significantly complex tasks, including planar and spatial robots that have to lift or throw a load at a given velocity using torque-limited actuators.

Index terms—Kinodynamic motion planning, loop-closure constraint, closed kinematic chain, atlas, manifold, LQR, steering.

I. INTRODUCTION

SINCE its formalization in the early nineties [1], the kinodynamic planning problem remains as one of the most challenging open problems in robotics. The problem entails finding feasible trajectories connecting given states of a robot, each defined by a configuration and a velocity of the underlying mechanical system. To ensure feasibility, the trajectory should: 1) fulfil all kinematic constraints of the system, including holonomic ones, like loop-closure or end-effector constraints, or nonholonomic ones, like rolling contact or velocity limit constraints; 2) be compliant with the equations of motion of the robot; 3) avoid the collisions with obstacles in the environment; and 4) be executable with the limited force capacity of the actuators. In certain applications, moreover, the trajectory should also be optimal in some sense, minimising, for example, the time or control effort required for its execution.

The ability to plan such trajectories is key in a robotic system. Above all, it endows the system with a means to convert higher-level commands—like “move to a certain location”, or “throw the object at a given speed”—into appropriate reference signals that can be followed by the actuators. By accounting for the robot dynamics and force limitations at the planning stage, moreover, the motions are easier to control, and they often look more graceful, or physically natural [2], as they tend to exploit gravity, inertia, and centripetal forces to the benefit of the task.

The kinodynamic planning problem can be viewed as a full motion planning problem in the state space, as opposed to a purely kinematic problem that only requires the planning of a path in configuration space (C-space). This makes the problem harder, as the dimension of the state space is twice that of the C-space. Moreover, the obstacle region is virtually larger, involving states that correspond to an actual collision, but also those from which a future collision is inevitable due to the system momentum. The planning of steering motions is considerably more difficult as well. While direct motions suffice in the C-space, steering motions in the state space need to conform to the vector fields defined by the equations of motion and to the actuator limits of the robot.

Among all kinodynamic planning techniques, rapidly-exploring random trees (RRTs) have emerged as one of the most successful planning paradigms to date [7]. RRTs make intensive use of sampling and dynamic simulations to grow trajectory trees over the state space until the start and goal states get connected. The efficiency of the approach is remarkable, especially in view of its simplicity and relative ease of implementation. The technique is fairly general and, with proper extensions, can even converge to minimum-cost motions [8, 9]. However, existing RRT methods also suffer from a main limitation: they assume that the robot state can be described by means of independent generalised coordinates. This makes them applicable to open-chain robots, or to robots with explicit state space parametrisations, but they have problems in dealing with general mechanisms with closed-kinematic chains. Such chains arise frequently in today’s robots and manipulation systems (Fig. 1), which explains the growing interest they arouse in the recent literature [10]–[16].

Unlike in the open-chain case, the state space of a closed-chain robot is not flat anymore. Instead, it is a nonlinear manifold defined implicitly by a system of equations that, in general, cannot be solved in closed form. This manifold is a zero-measure set in a larger ambient space, which complicates the design of sampling and steering methods to explore the manifold efficiently. Moreover, if the dynamic model of the

This work has been partially funded by the Spanish Ministry of Science, Innovation, and Universities under project DPI2017-88282-P

Ricard Bordalba, Lluís Ros, and Josep M. Porta are with the Institut de Robòtica i Informàtica Industrial (CSIC-UPC). Llorens Artigas 4-6, 08028 Barcelona, Catalonia. E-mails: rbordalba@iri.upc.edu, ros@iri.upc.edu, and porta@iri.upc.edu.




Fig. 1. Example systems involving closed kinematic chains. The chains may be intrinsic to the robot structure, as in parallel robots (left picture), or they may result from manipulation constraints during a task, as in multi-limb systems transporting an object, or keeping feet attached to the environment (right pictures). From left to right: A Delta parallel robot [3], the Atlas robot from Boston Dynamics lifting a heavy load [4], the Robonaut 2 robot with two legs clamped to the International Space Station [5], and the SpiderFab Bot, a conceptual design for self-fabricating space systems [6]. Pictures courtesy of ABB, Boston Dynamics, NASA, and Tethers Unlimited, Inc (respectively).

robot is not properly handled, the planned trajectories may deviate substantially from the manifold, leading to undesired violations of the kinematic constraints, or to failure to reach the goal. Forward singularities may also complicate the planning and control of motions across certain surfaces of the state space [17].

The purpose of this paper is to extend the planner in [7] to cope with the previous complications. As we shall see, by constructing an atlas of the state space in parallel to the RRT, one can define proper sampling and steering methods that deal with closed kinematic chains effectively, while producing feasible trajectories even across forward singularities. An early version of our planner was presented in [18]. In contrast to [18], we here develop a steering method based on linear quadratic regulators (LQR), which greatly increases the planner efficiency in comparison to the randomised strategy used in [18]. New challenging test cases are also reported for demonstration, including tasks that require the throwing of objects at a given velocity, and bimanual manipulations of heavy loads, which were difficult to solve with [18]. It is worth noting that, while some path planning approaches have previously dealt with closed kinematic chains [10]–[12, 15, 19]–[23], none of them has considered the dynamics of the system into the planner. Our kinodynamic planner, in fact, can also be seen as an extension of the work in [12] to cope with dynamic constraints.

The rest of the paper is organised as follows. Section II reviews the state of the art on kinodynamic planning to better place our work into context. Section III formally states the problem we confront, enumerating our assumptions and the various constraints intervening. Section IV explains why most RRT approaches, while powerful, are limited in some way or another, and would be difficult to extend to cope with closed chains and dynamic constraints simultaneously. Sections V and VI present effective sampling, simulation, and steering methods that allow us to describe, in Section VII, our planner

implementation. Sections VIII and IX respectively examine the completeness properties of the planner and its practical performance in illustrative situations. Section X finally provides the paper conclusions and discusses several points requiring further attention.

II. RELATED WORK

A. C-space approaches

The sheer complexity of kinodynamic planning is usually managed by decomposing the problem into two simpler problems [24]. Initially, the dynamic constraints of the robot are neglected and a collision-free path in the C-space is sought that solely satisfies the kinematic constraints. Then, a time-parametric trajectory constrained to the previous path is designed while accounting for the dynamic constraints and force limits of the actuators. Although many techniques can be used to compute the path, like probabilistic roadmaps or randomised tree techniques among others [24, 25], the trajectory is usually obtained with the time-scaling method in [26] or its later improvements [27]–[30]. This method regards the path as a function $\mathbf{q} = \mathbf{q}(s)$ in which \mathbf{q} is the robot configuration and s is some path parameter, and then finds a monotonic time scaling $s = s(t)$ such that $\mathbf{q}(t) = \mathbf{q}(s(t))$ connects the start and goal configurations in minimum time. The method is fast and elegant, as it exploits the bang-bang nature of the solution in the (s, \dot{s}) plane, and robust implementations have recently been obtained [31].

The previous approach generates a trajectory that is only time-optimal for the computed path, but makes the problem more tractable, so it can be solved in systems with many degrees of freedom like humanoids, legged robots, or mobile robot formations [32]. Its lack of completeness, moreover, can be alleviated by improving the trajectory a posteriori using optimization techniques [33]–[35]. Time scaling methods, in addition, have recently been extended to compute the feasible velocities at the end of a path, given an initial range of

velocities [36], which can be combined with randomised planners to generate graceful dynamic motions [32].

It must be noted that, despite their advantages, the previous methods essentially work in the C-space, which makes them limited in some way or another. For instance, path planning approaches cannot generate swinging paths in principle, and such paths may be required in highly dynamic tasks like lifting a heavy load under strict torque limitations. In other approaches, start or goal states with nonzero velocity cannot be specified, which is necessary in, e.g., catching or throwing objects at a certain speed and direction. Time scaling methods, moreover, require the robot to be fully actuated. While this is rarely an issue in robot arms or humanoids under contact constraints [13, 14], parallel robots with passive joints are underactuated at forward singularities [17]. These configurations are problematic when managed in the C-space as they can only be traversed under particular velocities and accelerations. As it will turn out, however, the previous limitations do not apply if robot trajectories are directly planned in the state space.

B. State space approaches

Existing techniques for planning in the state space can roughly be grouped into optimization and randomised approaches. On the one hand, optimization approaches can be applied to remarkably-complex problems [37]–[41]. An advantage is that they can accommodate a wide variety of kinematic and dynamic constraints. For instance, differential constraints describing the robot dynamics can be enforced by discretising the trajectory into different knot points using an Euler method, or any higher-order method if more accuracy is needed. However, there is a trade-off between the number of knot points, the integration method adopted, and the computational cost of the optimization problem. In systems with closed kinematic chains, moreover, the discretization of the differential equations produces knot points that easily drift from the state space manifold, which results in unwanted link disassemblies or contact losses. In [38], differential constraints were approximated explicitly by means of polynomial functions while guaranteeing third-order integration accuracy in constrained systems. Even so, the problem size becomes huge for long time horizons or systems with many degrees of freedom [32]. Good discussions on the advantages and pitfalls of optimization-based techniques can be found in [42] and [13]. On the other hand, randomised approaches like the standard RRT [7] can cope with differential constraints in relatively high-dimensional problems, and guarantee to find a solution when it exists and enough computing time is available. A main issue, however, is that exact steering methods are not available for nonlinear dynamical systems. The usual RRT method tries to circumvent this problem by simulating random actions for a given time, and then selecting the action that gets the system closest to the target [7]. For particular systems, better solutions exist though. For instance, the approach in [43] assumes double integrator dynamics and exploits the fact that the minimum time problem has an efficient solution in this case. The resulting planner is fast, but the full dynamics of the system can only be coped via feedback

linearisation, which requires the inverse dynamic problem to be solvable. The method in [44] linearises the system dynamics and uses an infinite-horizon LQR controller to define a steering method, but such a controller can only be used to reach zero-velocity states. In contrast, [45], [46], and [47] use finite-horizon LQR controllers that can converge to arbitrary states. As designed, however, the previous steering methods cannot be applied to robots with closed kinematic chains, as they assume the state coordinates to be independent. Our steering approach is similar to the one in [47], but extended to cope with such chains.

III. PROBLEM FORMULATION

To formally state our problem, let us describe the robot configuration by means of a tuple \mathbf{q} of n_q generalised coordinates, which determine the positions and orientations of all links at a given instant of time. We restrict our attention to robots with closed kinematic chains, in which \mathbf{q} must satisfy a system of n_e nonlinear equations

$$\Phi(\mathbf{q}) = \mathbf{0} \quad (1)$$

enforcing the closure conditions of the chains. The C-space of the robot is then the set

$$\mathcal{C} = \{\mathbf{q} : \Phi(\mathbf{q}) = \mathbf{0}\},$$

which may be quite complex in general. In this paper, however, we assume that the Jacobian $\Phi_{\mathbf{q}}(\mathbf{q}) = \partial\Phi/\partial\mathbf{q}$ is full rank for all $\mathbf{q} \in \mathcal{C}$, so \mathcal{C} is a smooth manifold of dimension $d_{\mathcal{C}} = n_q - n_e$ without C-space singularities [17]. This assumption is not too restrictive, as these singularities are often removed by mechanical designers (e.g., by setting appropriate joint limits), and it does not rule out forward or inverse singularities [17], which can be crossed naturally by our planner.

By differentiating Eq. (1) with respect to time, we obtain the velocity equation of the robot

$$\Phi_{\mathbf{q}}(\mathbf{q}) \cdot \dot{\mathbf{q}} = \mathbf{0}, \quad (2)$$

which characterises the feasible vectors $\dot{\mathbf{q}}$ at a given $\mathbf{q} \in \mathcal{C}$.

Let $\mathbf{F}(\mathbf{x}) = \mathbf{0}$ denote the system formed by Eqs. (1) and (2), where $\mathbf{x} = (\mathbf{q}, \dot{\mathbf{q}}) \in \mathbb{R}^{n_x}$ is the state vector of the robot, with $n_x = 2n_q$. While path planning approaches operate in \mathcal{C} , kinodynamic planning problems are better represented in the state space

$$\mathcal{X} = \{\mathbf{x} : \mathbf{F}(\mathbf{x}) = \mathbf{0}\}. \quad (3)$$

It can be shown that, since $\Phi_{\mathbf{q}}(\mathbf{q})$ is full rank in our case, \mathcal{X} is also a smooth manifold, but of dimension $d_{\mathcal{X}} = 2d_{\mathcal{C}}$. This implies that the tangent space of \mathcal{X} at \mathbf{x} ,

$$\mathcal{T}_{\mathbf{x}}\mathcal{X} = \{\dot{\mathbf{x}} \in \mathbb{R}^{n_x} : \mathbf{F}_{\mathbf{x}}(\mathbf{x}) \dot{\mathbf{x}} = \mathbf{0}\}, \quad (4)$$

is well-defined and $d_{\mathcal{X}}$ -dimensional for any $\mathbf{x} \in \mathcal{X}$.

We encode the forces and torques of the actuators into an action vector $\mathbf{u} = (u_1, \dots, u_{n_u}) \in \mathbb{R}^{n_u}$. Given a starting state $\mathbf{x}_s \in \mathcal{X}$, and the vector \mathbf{u} as a function of time, $\mathbf{u} = \mathbf{u}(t)$, the time evolution of the robot is determined by a system of differential-algebraic equations of the form

$$\begin{cases} \mathbf{F}(\mathbf{x}) = \mathbf{0}, \\ \dot{\mathbf{x}} = \mathbf{g}(\mathbf{x}, \mathbf{u}). \end{cases} \quad (5)$$

$$(6)$$

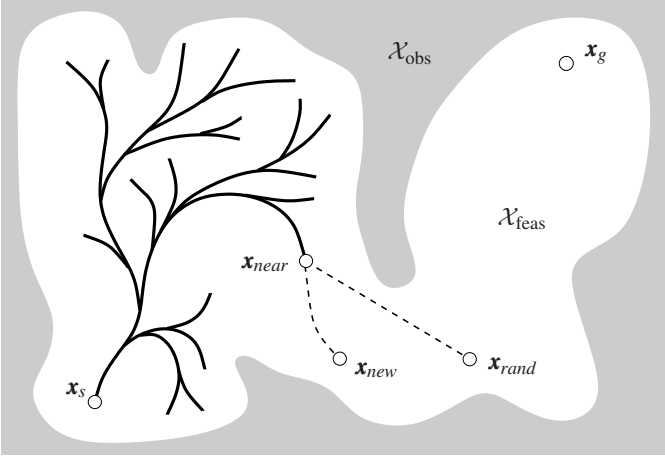


Fig. 2. Expansion of a unidirectional RRT [7].

In this system, Eq. (5) forces the states \mathbf{x} to remain in \mathcal{X} , and Eq. (6) models the dynamics of the robot, which can be described using the multiplier form of the Euler-Lagrange equations for example [48]. For each value of \mathbf{u} , Eq. (6) defines a vector field over \mathcal{X} , which can be used together with Eq. (5) to simulate the robot motion forward in time using proper integration tools [49].

To model the fact that the actuator forces are limited, we will assume that \mathbf{u} can only take values inside the box

$$\mathcal{U} = [-l_1, l_1] \times [-l_2, l_2] \times \dots \times [-l_{n_u}, l_{n_u}] \quad (7)$$

of \mathbb{R}^{n_u} , where l_i denotes the limit force or torque of the i -th motor. Along a trajectory, moreover, the robot cannot incur in collisions with itself or with the environment, and should fulfil any limits imposed on \mathbf{q} and $\dot{\mathbf{q}}$. This reduces the feasible states \mathbf{x} to those lying in a subset $\mathcal{X}_{\text{feas}} \subseteq \mathcal{X}$.

With the previous definitions, the problem we confront can be stated as follows: Given a kinematic and dynamic model of the robot, a geometric model of the environment, and two states \mathbf{x}_s and \mathbf{x}_g of $\mathcal{X}_{\text{feas}}$, find a control policy $\mathbf{u} = \mathbf{u}(t)$ lying in \mathcal{U} for all t such that the trajectory $\mathbf{x} = \mathbf{x}(t)$ determined by Eqs. (5) and (6) for $\mathbf{x}(0) = \mathbf{x}_s$ fulfils $\mathbf{x}(t_g) = \mathbf{x}_g$ at some time $t_g > 0$, with $\mathbf{x}(t) \in \mathcal{X}_{\text{feas}}$ for all $t \in [0, t_g]$.

IV. LIMITATIONS OF PRIOR RRT METHODS

Observe that the previous formulation is more general than the one assumed in earlier RRT planners. In particular, the approaches in [10, 12, 19, 22, 23] are purely kinematic, so they only consider Eq. (1), and neglect Eqs. (2), (6), and the force bounds in (7). As a result, they only compute paths in \mathcal{C} , and such paths might be unfeasible dynamically. In contrast, kinodynamic approaches like [7]–[9, 44, 46, 47] consider Eq. (6) and the bounds in (7), but not Eqs. (1) and (2), which impedes the handling of robots with closed kinematic chains. While [24] proposed a few extensions to help RRTs cope with such chains, we next see that these lead to unsatisfactory results.

Recall from [24] that a usual RRT is initialised at \mathbf{x}_s , and is extended by applying four steps repeatedly [see Fig. 2]: 1) a guiding state $\mathbf{x}_{\text{rand}} \in \mathcal{X}$ is randomly selected; 2) the RRT




Fig. 3. Generation of a guiding sample according to [24].

state \mathbf{x}_{near} that is closest to \mathbf{x}_{rand} is determined according to some metric; 3) a steering method is used to compute the action $\mathbf{u} \in \mathcal{U}$ that brings the system as close as possible to \mathbf{x}_{rand} in the absence of obstacles; and 4) the movement that results from applying \mathbf{u} during some time Δt is obtained by integrating Eq. (6). This yields a new state \mathbf{x}_{new} , which is added to the RRT if it lies in $\mathcal{X}_{\text{feas}}$, or it is discarded otherwise. In the former case, \mathbf{u} is stored in the new edge connecting \mathbf{x}_{near} to \mathbf{x}_{new} . The process terminates when a tree node is close enough to \mathbf{x}_g . It is worth noting that, in many implementations, steps 3) and 4) are repeated with \mathbf{x}_{new} playing the role of \mathbf{x}_{near} , as long as \mathbf{x}_{new} gets closer to \mathbf{x}_{rand} .

Three problems arise when applying the previous method to closed kinematic chains. First, the points \mathbf{x}_{rand} are difficult to obtain in general, as \mathcal{X} may be a manifold without explicit parametrizations. To circumvent this issue, [24, Sec. 7.4.1] proposes to randomly pick \mathbf{x}_{rand} from the larger ambient space \mathbb{R}^{n_x} (Fig. 3) and use, as a guiding state, the point $\mathbf{x}'_{\text{rand}}$ that results from projecting \mathbf{x}_{rand} onto the tangent space of \mathcal{X} at \mathbf{x}_{near} . However, while $\mathbf{x}'_{\text{rand}}$ is easy to compute, its pulling effect on the RRT may be small. The ambient space could be large in comparison to \mathcal{X} , resulting in points $\mathbf{x}'_{\text{rand}}$ that might often be close to \mathbf{x}_{near} , which diminishes the exploration bias of the RRT. This effect was analysed in [12] and [23]. A second problem concerns the dynamic simulation of robot motions. Existing RRT methods would only use Eq. (6) to generate such motions on the grounds that Eq. (5) is implicitly accounted for by Eq. (6) [24, Sec. 13.4.3.1]. However, from multibody dynamics it is known that the motion of a closed-chain robot can only be predicted reliably if Eq. (5) is actively used during the integration of Eq. (6) [50]. Otherwise, the inevitable errors introduced when discretising Eq. (6) will make the trajectory $\mathbf{x}(t)$ increasingly drift from \mathcal{X} as the simulation progresses. Such a drift may even be large enough to prevent the connection of \mathbf{x}_s with \mathbf{x}_g [18]. The use of Baumgarte stabilization to compensate this drift [51] is also

problematic, as it may lead to instabilities [52] or fictitious energy increments, and its stabilising parameters are not easy to tune. A third problem, finally, concerns the steering method. A shooting strategy based on simulating random actions from \mathcal{U} was proposed in [7], but this technique is inefficient when n_u is large, as the number of samples needed to properly represent \mathcal{U} grows exponentially with n_u . The lack of a good steering strategy is a general problem of RRT methods, but it is more difficult to address when closed kinematic chains are present.

Purely kinematic planners like [10, 12, 19, 22, 23] do not perform dynamic simulations, and employ direct steering motions between configurations. Moreover, most of them sample in ambient space. Thus, the previous three problems would also arise when trying to generalise them to cope with dynamic constraints. Among such planners, however, the one in [12] is more amenable for generalisation, as it employs atlas machinery that is applicable to mechanisms of general loop topology. Our goal in this paper is to show that, precisely, such a machinery can be extended to cope with the more general problem of Section III. As we shall see, once an atlas of \mathcal{X} is obtained, we will have the necessary tools to 1) sample the \mathcal{X} manifold directly instead of its ambient space \mathbb{R}^{n_x} ; 2) integrate Eqs. (5) and (6) as a true differential-algebraic equation so as to ensure driftless motions on \mathcal{X} ; and 3) define a proper steering method for closed kinematic chains. We develop these tools in the following two sections, and later use them as basic building blocks in our planner implementation.

V. MAPPING AND EXPLORING THE STATE SPACE

A. Atlas construction

Formally, an atlas of \mathcal{X} is a collection of charts mapping \mathcal{X} entirely, where each chart c is a local diffeomorphism ϕ_c from an open set $V_c \subset \mathcal{X}$ to an open set $P_c \subseteq \mathbb{R}^{d_x}$ [Fig. 4(a)]. The V_c sets can be thought of as partially-overlapping tiles covering \mathcal{X} , in such a way that every $\mathbf{x} \in \mathcal{X}$ lies in at least one set V_c . The point $\mathbf{y} = \phi_c(\mathbf{x})$ provides the local coordinates, or parameters, of \mathbf{x} in chart c . Since each map ϕ_c is a diffeomorphism, its inverse map $\psi_c = \phi_c^{-1}$ also exists, and gives a local parametrisation of V_c .

For particular manifolds, ϕ_c and ψ_c can be defined in closed form. However, we propose to use the tangent space parametrization [53] to define them for any manifold. Under this parametrization, the map $\mathbf{y} = \phi_c(\mathbf{x})$ around a given $\mathbf{x}_c \in \mathcal{X}$ is obtained by projecting \mathbf{x} orthogonally to $\mathcal{T}_{\mathbf{x}_c}\mathcal{X}$ [Fig. 4(b)], so this map takes the form

$$\mathbf{y} = \mathbf{U}_c^\top (\mathbf{x} - \mathbf{x}_c), \quad (8)$$

where \mathbf{U}_c is an $n_x \times d_x$ matrix whose columns provide an orthonormal basis of $\mathcal{T}_{\mathbf{x}_c}\mathcal{X}$. The inverse map $\mathbf{x} = \psi_c(\mathbf{y})$ is implicitly determined by the system of nonlinear equations

$$\left. \begin{aligned} \mathbf{F}(\mathbf{x}) &= \mathbf{0} \\ \mathbf{U}_c^\top (\mathbf{x} - \mathbf{x}_c) - \mathbf{y} &= \mathbf{0} \end{aligned} \right\} \quad (9)$$

which, for a given \mathbf{y} , can be solved for \mathbf{x} using the Newton-Raphson method when \mathbf{x} is close to \mathbf{x}_c .




Fig. 4. (a) Two neighbouring charts of \mathcal{X} , labelled c and k , together with their maps ϕ_c and ϕ_k , and inverse maps ψ_c and ψ_k . (b) Using the tangent space parametrization, ϕ_c is defined by the projection of \mathbf{x} onto $\mathcal{T}_{\mathbf{x}_c}\mathcal{X}$.

Assuming that an atlas has been created, the problem of sampling \mathcal{X} boils down to generating random points \mathbf{y}_{rand} in the P_c sets, as they can always be projected to \mathcal{X} using the map $\mathbf{x}_{rand} = \psi_c(\mathbf{y}_{rand})$. Also, the atlas allows the conversion of the vector field defined by Eq. (6) into one on the P_c sets of the charts. The time derivative of Eq. (8), $\dot{\mathbf{y}} = \mathbf{U}_c^\top \dot{\mathbf{x}}$, gives the relationship between the two vector fields, and allows writing

$$\dot{\mathbf{y}} = \mathbf{U}_c^\top \mathbf{g}(\psi_c(\mathbf{y}), \mathbf{u}) = \tilde{\mathbf{g}}(\mathbf{y}, \mathbf{u}), \quad (10)$$

which is Eq. (6), but expressed in local coordinates. This equation still takes the full dynamics into account, and forms the basis of geometric methods for the integration of differential-algebraic equations as ordinary differential equations on manifolds [54, 55]. Given a state \mathbf{x}_k and an action \mathbf{u}_k , \mathbf{x}_{k+1} is estimated by obtaining $\mathbf{y}_k = \phi_c(\mathbf{x}_k)$, then computing \mathbf{y}_{k+1} using a discrete form of Eq. (10), and finally getting $\mathbf{x}_{k+1} = \psi_c(\mathbf{y}_{k+1})$. The procedure guarantees that \mathbf{x}_{k+1} will lie on \mathcal{X} by construction, thus making the integration compliant with all kinematic constraints in Eq. (5).

B. Incremental atlas and RRT expansion

One could use the methods in [53] to construct a full atlas of the implicitly-defined state space and then use its local




Fig. 5. Thresholds determining the extension of the P_c set of the chart at \mathbf{x}_c . While \mathbf{y}_k lies in P_c , \mathbf{y}_{k+1} does not because it violates Eqs. (11)-(13).

parametrisations to implement a kinodynamic RRT planner. However, the construction of a full atlas is only feasible for low-dimensional state spaces. On the other hand, only part of the atlas is necessary to solve a given motion planning problem. For these reasons, as in [12], we combine the construction of the atlas and the expansion of the RRT. In this approach, a partial atlas is used to both generate random states and grow the RRT branches. As described next, new charts are also created as the RRT branches reach unexplored regions of the state space.

Suppose that \mathbf{x}_k and \mathbf{x}_{k+1} are two consecutive states along an RRT branch and let \mathbf{y}_k and \mathbf{y}_{k+1} be their local coordinate vectors in $\mathcal{T}_{\mathbf{x}_c}\mathcal{X}$. Then, a new chart at \mathbf{x}_k is created if Eq. (9) cannot be solved for \mathbf{x}_{k+1} using the Newton-Raphson method, or if any of the following conditions is met

$$\|\mathbf{x}_{k+1} - (\mathbf{x}_c + \mathbf{U}_c \mathbf{y}_{k+1})\| > \varepsilon, \quad (11)$$

$$\frac{\|\mathbf{y}_{k+1} - \mathbf{y}_k\|}{\|\mathbf{x}_{k+1} - \mathbf{x}_k\|} < \cos \alpha, \quad (12)$$

$$\|\mathbf{y}_{k+1}\| > \rho, \quad (13)$$

where ε , α , and ρ are user-defined thresholds (Fig. 5). These conditions are introduced to ensure that the P_c sets of the created charts capture the overall shape of \mathcal{X} with sufficient detail. The first condition limits the maximal distance between $\mathcal{T}_{\mathbf{x}_c}\mathcal{X}$ and the manifold \mathcal{X} . The second condition ensures a bounded curvature in the part of \mathcal{X} that is covered by a chart, as well as a smooth transition between neighbouring charts. The third condition finally guarantees the generation of new charts as the RRT grows, even for almost flat manifolds.

C. Chart coordination

Since the charts will be used to generate samples on \mathcal{X} , it is important to reduce the overlap between new charts and those already present in the atlas. Otherwise, the areas of \mathcal{X} covered by several charts would be oversampled. To avoid this problem, the P_c set of each chart is initialised as a ball of radius σ centred at the origin of $\mathbb{R}^{d_{\mathcal{X}}}$. This ball is progressively bounded as new neighbouring charts are created around the chart. If, while growing an RRT branch, a neighbouring chart




Fig. 6. Half planes added to trim the P_c and P_k sets of two neighboring charts. Note that $\mathbf{y}_k = \boldsymbol{\varphi}_c(\mathbf{x}_k)$ and $\mathbf{y}_c = \boldsymbol{\varphi}_k(\mathbf{x}_c)$.

is created at a point \mathbf{x}_k with parameter vector \mathbf{y}_k in P_c , the following inequality

$$\mathbf{y}^\top \mathbf{y}_k - \frac{\|\mathbf{y}_k\|^2}{2} \leq 0 \quad (14)$$

is added as a bounding half-plane of P_c (Fig. 6). An analogous inequality is added to the P_k set of the chart at \mathbf{x}_k , but using $\mathbf{y}_c = \boldsymbol{\varphi}_k(\mathbf{x}_c)$ instead of \mathbf{y}_k in Eq. (14). Note that the radius σ of the initial ball must be larger than ρ to guarantee that the RRT branches covered by chart c will eventually trigger the generation of new charts, i.e., to guarantee that Eq. (13) will eventually hold. Also, since Eq. (13) forces the norm of \mathbf{y}_k to be limited by ρ , the half-plane defined by Eq. (14) will be guaranteed to clip P_c . Consequently, the P_c sets of those charts surrounded by neighbouring charts will be significantly smaller than the P_c sets of the charts at the exploration border of the atlas. As we shall see in Section VII-A, this favours the growth of the tree towards unexplored regions of \mathcal{X} .

VI. A STEERING METHOD

Our planner can adopt different strategies to steer the system from \mathbf{x}_{near} to \mathbf{x}_{rand} . A simple one, called randomised steering, consists in simulating several random actions in \mathcal{U} , and then choosing the one that brings the robot closest to \mathbf{x}_{rand} . This strategy was proposed in [7] and is the one we adopted in our early version of the planner [18]. As explained in Section IV, however, this approach becomes inefficient as the dimension of \mathcal{U} increases. To amend this problem we next propose another strategy, called LQR steering, based on linear quadratic regulators. While LQR techniques are classical steering methods for control systems [56], they assume the state coordinates to be independent, so they are applicable to open chain robots only. However, we next show that, using the atlas charts, they can be extended to the closed chain case. The idea is to exploit system linearisations at the various chart centres so as to obtain a sequence of control laws bringing the robot from \mathbf{x}_{near} to \mathbf{x}_{rand} .

A. System linearisation at a chart centre

To apply LQR techniques to our steering problem, we must first linearise our system model at the chart centres \mathbf{x}_c and null action $\mathbf{u} = \mathbf{0}$. To do so, note that we cannot linearise Eq. (6), as this would disregard the fact that the \mathbf{x} variables are coupled by Eq. (5). We must instead linearise Eq. (10),

which expresses Eq. (6) in the independent \mathbf{y} coordinates of $\mathcal{T}_{\mathbf{x}_c}\mathcal{X}$. Since the point $\mathbf{x} = \mathbf{x}_c$ corresponds to $\mathbf{y} = \mathbf{0}$ in the local coordinates of $\mathcal{T}_{\mathbf{x}_c}\mathcal{X}$, the sought linearisation is

$$\dot{\mathbf{y}} = \underbrace{\frac{\partial \tilde{\mathbf{g}}}{\partial \mathbf{y}} \bigg|_{\mathbf{y}=\mathbf{0}}}_{\mathbf{A}} \mathbf{y} + \underbrace{\frac{\partial \tilde{\mathbf{g}}}{\partial \mathbf{u}} \bigg|_{\mathbf{u}=\mathbf{0}}}_{\mathbf{B}} \mathbf{u} + \underbrace{\tilde{\mathbf{g}}(\mathbf{0}, \mathbf{0})}_{\mathbf{c}}, \quad (15)$$

which can be written as

$$\dot{\mathbf{y}} = \mathbf{A}\mathbf{y} + \mathbf{B}\mathbf{u} + \mathbf{c}. \quad (16)$$

This system will be assumed to be controllable hereafter.

Observe that, in Eq. (16), the term

$$\mathbf{c} = \tilde{\mathbf{g}}(\mathbf{0}, \mathbf{0}) = \mathbf{U}_c^\top \mathbf{g}(\mathbf{x}_c, \mathbf{0})$$

is not null in principle, because $(\mathbf{x}, \mathbf{u}) = (\mathbf{x}_c, \mathbf{0})$ is not necessarily an equilibrium point of the system in Eq. (10). Moreover, by applying the chain rule and using the fact that $\frac{\partial \Psi}{\partial \mathbf{y}} \big|_{\mathbf{y}=\mathbf{0}} = \mathbf{U}_c$ [48], the \mathbf{A} and \mathbf{B} terms can be written as:

$$\mathbf{A} = \frac{\partial \tilde{\mathbf{g}}}{\partial \mathbf{y}} \bigg|_{\mathbf{y}=\mathbf{0}} = \mathbf{U}_c^\top \frac{\partial \mathbf{g}}{\partial \mathbf{x}} \bigg|_{\mathbf{x}=\mathbf{x}_c} \mathbf{U}_c,$$

and

$$\mathbf{B} = \frac{\partial \tilde{\mathbf{g}}}{\partial \mathbf{u}} \bigg|_{\mathbf{u}=\mathbf{0}} = \mathbf{U}_c^\top \frac{\partial \mathbf{g}}{\partial \mathbf{u}} \bigg|_{\mathbf{u}=\mathbf{0}}.$$

Notice, therefore, that \mathbf{A} , \mathbf{B} , and \mathbf{c} can exactly be obtained by evaluating the original function $\mathbf{g}(\mathbf{x}, \mathbf{u})$ and its derivatives $\partial \mathbf{g} / \partial \mathbf{x}$ and $\partial \mathbf{g} / \partial \mathbf{u}$ at $(\mathbf{x}, \mathbf{u}) = (\mathbf{x}_c, \mathbf{0})$. In those robots in which these derivatives are not easy to obtain in closed form, \mathbf{A} and \mathbf{B} can always be approximated numerically using finite differences.

B. Steering on a single chart

Suppose now that both \mathbf{x}_{near} and \mathbf{x}_{rand} lie in a same chart c centred at $\mathbf{x}_c \in \mathcal{X}$ (Fig. 7). In this case, the problem of steering the robot from \mathbf{x}_{near} to \mathbf{x}_{rand} can be reduced to that of steering the system in Eq. (16) from $\mathbf{y}_{near} = \boldsymbol{\varphi}_c(\mathbf{x}_{near})$ to $\mathbf{y}_{rand} = \boldsymbol{\varphi}_c(\mathbf{x}_{rand})$. This problem can be formulated as follows: Find the control policy $\mathbf{u}(t) = \mathbf{u}^*(t)$ and time $t_f = t_f^*$ that minimise the cost function

$$J(\mathbf{u}(t), t_f) = \int_0^{t_f} \left(1 + \mathbf{u}(t)^\top \mathbf{R} \mathbf{u}(t) \right) dt, \quad (17)$$

subject to the constraints

$$\dot{\mathbf{y}} = \mathbf{A}\mathbf{y} + \mathbf{B}\mathbf{u} + \mathbf{c}, \quad (18)$$

$$\mathbf{y}(0) = \mathbf{y}_{near}, \quad (19)$$

$$\mathbf{y}(t_f) = \mathbf{y}_{rand}. \quad (20)$$

In Eq. (17), the unit term inside the integral penalises large values of t_f , while the term $\mathbf{u}(t)^\top \mathbf{R} \mathbf{u}(t)$ penalises high control actions. In this term, \mathbf{R} is a symmetric positive-definite matrix that is fixed beforehand.

The problem just formulated is known as the fixed final state optimal control problem [56]. We shall solve this problem in two stages. Initially, we will obtain $\mathbf{u}^*(t)$ assuming that t_f is fixed, and then we will find a time t_f that leads to a minimum of $J(\mathbf{u}(t), t_f)$.




Fig. 7. When \mathbf{x}_{near} and \mathbf{x}_{rand} are covered by a same chart, the steering of the system can be reduced to a steering problem in $\mathcal{T}_{\mathbf{x}_c}\mathcal{X}$.

C. Fixed final state and fixed final time problem

If t_f is fixed, we can find the optimal action $\mathbf{u}(t) = \mathbf{u}^*(t)$ by applying Pontryagin's minimum principle. Since the function $\mathbf{u}^\top(t) \mathbf{R} \mathbf{u}(t)$ is convex, this principle provides necessary and sufficient conditions of optimality in our case [57]. To apply the principle, we first define the Hamiltonian function

$$H(\mathbf{y}, \mathbf{u}, \boldsymbol{\lambda}) = 1 + \mathbf{u}^\top \mathbf{R} \mathbf{u} + \boldsymbol{\lambda}^\top (\mathbf{A}\mathbf{y} + \mathbf{B}\mathbf{u} + \mathbf{c}), \quad (21)$$

where $\boldsymbol{\lambda} = \boldsymbol{\lambda}(t)$ is an undetermined Lagrange multiplier. Then, the corresponding state and costate equations are

$$\dot{\mathbf{y}} = \frac{\partial H}{\partial \boldsymbol{\lambda}} = \mathbf{A}\mathbf{y} + \mathbf{B}\mathbf{u} + \mathbf{c}, \quad (22)$$

$$\dot{\boldsymbol{\lambda}} = -\frac{\partial H}{\partial \mathbf{y}} = -\mathbf{A}^\top \boldsymbol{\lambda}. \quad (23)$$

For $\mathbf{u} = \mathbf{u}^*(t)$ to be an optimal control policy, H must be at a stationary point relative to \mathbf{u} , i.e., it must be

$$\frac{\partial H}{\partial \mathbf{u}} \bigg|_{\mathbf{u}=\mathbf{u}^*(t)} = \mathbf{R} \mathbf{u}^*(t) + \mathbf{B}^\top \boldsymbol{\lambda} = \mathbf{0}, \quad (24)$$

and thus,

$$\mathbf{u}^*(t) = -\mathbf{R}^{-1} \mathbf{B}^\top \boldsymbol{\lambda}(t). \quad (25)$$

Since Eq. (23) is decoupled from Eq. (22), its solution can be found independently. It is

$$\boldsymbol{\lambda}(t) = e^{\mathbf{A}^\top(t_f-t)} \boldsymbol{\lambda}(t_f), \quad (26)$$

where $\boldsymbol{\lambda}(t_f)$ is still unknown.

To find $\boldsymbol{\lambda}(t_f)$, let us consider the closed-form solution of Eq. (22) for $\mathbf{u} = \mathbf{u}^*(t)$:

$$\mathbf{y}(t) = e^{\mathbf{A}t} \mathbf{y}(0) + \int_0^t e^{\mathbf{A}(t-\tau)} (\mathbf{B} \mathbf{u}^*(\tau) + \mathbf{c}) d\tau. \quad (27)$$

If we evaluate this solution for $t = t_f$ and take into account Eqs. (25) and (26), we arrive at the expression

$$\mathbf{y}(t_f) = \mathbf{r}(t_f) - \mathbf{G}(t_f) \boldsymbol{\lambda}(t_f), \quad (28)$$

where

$$\mathbf{r}(t_f) = e^{\mathbf{A}t_f} \mathbf{y}(0) + \int_0^{t_f} e^{\mathbf{A}(t_f-\tau)} \mathbf{c} d\tau, \quad (29)$$

and

$$\begin{aligned} \mathbf{G}(t_f) &= \int_0^{t_f} \mathbf{e}^{\mathbf{A}(t_f-\tau)} \mathbf{B} \mathbf{R}^{-1} \mathbf{B}^\top \mathbf{e}^{\mathbf{A}^\top(t_f-\tau)} d\tau \\ &= \int_0^{t_f} \mathbf{e}^{\mathbf{A}\tau} \mathbf{B} \mathbf{R}^{-1} \mathbf{B}^\top \mathbf{e}^{\mathbf{A}^\top\tau} d\tau. \end{aligned} \quad (30)$$

Given that $\mathbf{y}(t_f)$ is known from Eq. (20), we can solve Eq. (28) for $\boldsymbol{\lambda}(t_f)$ to obtain

$$\boldsymbol{\lambda}(t_f) = \mathbf{G}(t_f)^{-1} (\mathbf{r}(t_f) - \mathbf{y}(t_f)). \quad (31)$$

Now, substituting Eq. (31) into (26), and the result into Eq. (25), we finally obtain the optimal control policy for the fixed final state and fixed final time problem:

$$\mathbf{u}^*(t) = -\mathbf{R}^{-1} \mathbf{B}^\top \mathbf{e}^{\mathbf{A}^\top(t_f-t)} \mathbf{G}(t_f)^{-1} (\mathbf{r}(t_f) - \mathbf{y}(t_f)). \quad (32)$$

Note that this is an open-loop policy, as \mathbf{u}^* depends on t only. The values $\mathbf{r}(t_f)$ and $\mathbf{G}(t_f)$ in this policy can be obtained by computing the integrals in Eqs. (29) and (30) numerically. The matrix $\mathbf{G}(t_f)$ is known as the weighted continuous reachability Gramian, and since the system is controllable, it is symmetric and positive-definite for $t > 0$ [47], which ensures that $\mathbf{G}(t_f)^{-1}$ always exists.

D. Finding the optimal time t_f

To find a time t_f^* for which the cost J in Eq. (17) attains a minimum value, we substitute the optimal policy in Eq. (32) into Eq. (17), and take into account Eq. (30), obtaining

$$J(t_f) = t_f + [\mathbf{y}(t_f) - \mathbf{r}(t_f)]^\top \mathbf{G}(t_f)^{-1} [\mathbf{y}(t_f) - \mathbf{r}(t_f)]. \quad (33)$$

The time t_f^* is thus the one that minimises $J(t_f)$ in Eq. (33). Assuming that t_f^* lies inside a specified time window $[0, t_{\max}]$, this time can be computed approximately by evaluating $\mathbf{r}(t_f)$, $\mathbf{G}(t_f)$ and $J(t_f)$ using Eqs. (29), (30), and (33) for $t_f = 0$ to $t_f = t_{\max}$, and selecting the t_f value for which $J(t_f)$ is minimum.

Finally, the values t_f^* , $\mathbf{r}(t_f^*)$, and $\mathbf{G}(t_f^*)$ can be used to evaluate the optimal control policy in Eq. (32). By applying this policy to the full nonlinear system of Eq. (6) during t_f^* seconds, we will follow a trajectory ending in some state \mathbf{y}'_{rand} close to \mathbf{y}_{rand} . This trajectory can be recovered on the \mathcal{X} space by means of the Ψ_c map and, if it lies in \mathcal{X}_{feas} , the corresponding branch can be added to the RRT.

E. Steering over multiple charts

If \mathbf{x}_{rand} is not covered by the chart c of \mathbf{x}_{near} , we can iteratively apply the steering process as shown in Fig. 8(a). To this end, we compute $\mathbf{y}_{rand} = \Phi_c(\mathbf{x}_{rand})$ and drive the system from $\mathbf{y}_{near} = \Phi_c(\mathbf{x}_{near})$ towards \mathbf{y}_{rand} on $\mathcal{T}_{\mathbf{x}_c} \mathcal{X}$, projecting the intermediate states \mathbf{y} to \mathcal{X} via Ψ_c . Eventually, we will reach some state $\mathbf{x}_k \in \mathcal{X}$ that is in the limit of the V_c set of the current chart (see the conditions in Sec. V-B). At this point, we generate a chart at \mathbf{x}_k and linearise the system again. We then use this linearisation to recompute the optimal control policy to go from \mathbf{x}_k to \mathbf{x}_{rand} . Such a “linearise and steer” process can be repeated as needed, until the system gets closely enough to \mathbf{x}_{rand} .

Although the previous procedure is often effective, it can also fail in some situations. As shown in Fig. 8(b), the initial steering on chart c might bring the system from \mathbf{x}_{near} to \mathbf{x}_k but, due to the position of \mathbf{x}_{rand} , a new control policy computed at \mathbf{x}_k would steer the system back to \mathbf{x}_{near} , leading to a back-and-forth cycle not converging to \mathbf{x}_{rand} . Such limit cycles can be detected however, because the time t_f^* will no longer decrease eventually. As shown in Fig. 8(c), moreover, the steering procedure can sometimes reach \mathbf{y}_{rand} , but we might find that $\Psi_c(\mathbf{y}_{rand}) \neq \mathbf{x}_{rand}$ because, due to the curvature of \mathcal{X} , several states can project to the same point on a given tangent space. Even so, such situations do not prevent the connection of \mathbf{x}_s with \mathbf{x}_g , as the steering algorithm is to be used inside a higher-level RRT planner. The implementation of such a planner is next addressed.

VII. PLANNER IMPLEMENTATION

Algorithm 1 gives the top-level pseudocode of the planner. At this level, the algorithm is almost identical to the RRT planner in [7]. The only difference is that, in our case, we construct an atlas A of \mathcal{X} to support the lower-level sampling, simulation, and steering tasks. The atlas is initialised with one chart centred at \mathbf{x}_s and another chart centred at \mathbf{x}_g (line 1). As in [7], the algorithm implements a bidirectional RRT where a tree T_s is rooted at \mathbf{x}_s (line 2) and another tree T_g is rooted at \mathbf{x}_g (line 3). Initially, a random state is sampled (\mathbf{x}_{rand} in line 5), the nearest state in T_s is determined (\mathbf{x}_{near} in line 6), and then T_s is extended with the aim of connecting \mathbf{x}_{near} with \mathbf{x}_{rand} using the CONNECT method (line 7). This method reaches a state \mathbf{x}_{new} that, due to the presence of obstacles or to a failure of the steering procedure, may be different from \mathbf{x}_{rand} . Then, the state in T_g that is nearest to \mathbf{x}_{new} is determined (\mathbf{x}'_{near} in line 8) and T_g is extended from \mathbf{x}'_{near} with the aim of reaching \mathbf{x}_{new} (line 9). This extension generates a new state \mathbf{x}'_{new} . After this step, the trees are swapped (line 10) and, if the last connection was unsuccessful, i.e., if \mathbf{x}_{new} and \mathbf{x}'_{new} are not closer than a user-provided threshold (line 11), lines 5 to 10 are repeated again. If the connection was successful, a solution trajectory is reconstructed using the paths from \mathbf{x}_{new} and \mathbf{x}'_{new} to the roots of T_s and T_g (line 12). Different metrics

Algorithm 1: The top-level pseudocode of the planner

PLAN TRAJECTORY($\mathbf{x}_s, \mathbf{x}_g$)
input : The query states, \mathbf{x}_s and \mathbf{x}_g .
output: A trajectory connecting \mathbf{x}_s and \mathbf{x}_g .

```

1  $A \leftarrow \text{INITATLAS}(\mathbf{x}_s, \mathbf{x}_g)$ 
2  $T_s \leftarrow \text{INITRRT}(\mathbf{x}_s)$ 
3  $T_g \leftarrow \text{INITRRT}(\mathbf{x}_g)$ 
4 repeat
5    $\mathbf{x}_{rand} \leftarrow \text{SAMPLE}(A, T_s)$ 
6    $\mathbf{x}_{near} \leftarrow \text{NEARESTSTATE}(T_s, \mathbf{x}_{rand})$ 
7    $\mathbf{x}_{new} \leftarrow \text{CONNECT}(A, T_s, \mathbf{x}_{near}, \mathbf{x}_{rand})$ 
8    $\mathbf{x}'_{near} \leftarrow \text{NEARESTSTATE}(T_g, \mathbf{x}_{new})$ 
9    $\mathbf{x}'_{new} \leftarrow \text{CONNECT}(A, T_g, \mathbf{x}'_{near}, \mathbf{x}_{new})$ 
10   $\text{SWAP}(T_s, T_g)$ 
11 until  $\|\mathbf{x}_{new} - \mathbf{x}'_{new}\| < \beta$ 
12 RETURN( $\text{TRAJECTORY}(T_s, \mathbf{x}_{new}, T_g, \mathbf{x}'_{new})$ )
```




Fig. 8. (a) Steering towards states not covered by the chart of x_{near} . (b) Cyclic behavior of the steering method. (c) Convergence to y_{rand} but not to x_{rand} .

Algorithm 2: Generate a random state x_{rand} .

SAMPLE(A, T)

input : The atlas A and the tree T to be extended.

output: A guiding sample x_{rand} .

```

1 repeat
2    $c \leftarrow \text{RANDOMCHARTINDEX}(A, T)$ 
3    $y_{rand} \leftarrow \text{RANDOMONBALL}(\sigma)$ 
4 until  $y_{rand} \in P_c$ 
5  $x_{rand} \leftarrow \Psi_c(y_{rand})$ 
6 if  $x_{rand} = \text{NULL}$  then
7    $x_{rand} \leftarrow x_c + U_c y_{rand}$ 
8 RETURN( $x_{rand}$ )

```

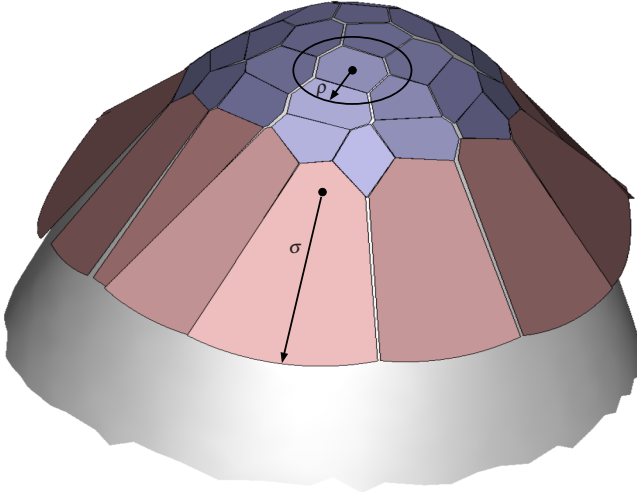


Fig. 9. A partial atlas of a paraboloid, with its inner and border charts colored in blue and red, respectively. Black dots indicate chart centers.

can be used to determine the distance between two states without affecting the overall structure of the planner. As in [7], we use Euclidean distance for simplicity.

A. Sampling

The **SAMPLE** method is described in Algorithm 2. Initially, one of the charts covering the tree T is selected at random with uniform distribution (line 2). A vector y_{rand} of parameters is then randomly sampled also with uniform distribution inside a ball of radius σ centred at the origin of $\mathbb{R}^{d_{\mathcal{X}}}$ (line 3). Chart

Algorithm 3: Try to connect x_{near} with x_{rand} .

CONNECT(A, T, x_{near}, x_{rand})

input : An atlas A , a tree T , the state x_{near} from which T is to be extended, and the guiding sample x_{rand} .

output: The new state x_{new} .

```

1  $x_{new} \leftarrow x_{near}$ 
2  $t_{fp}^* \leftarrow \infty$ 
3 repeat
4    $c \leftarrow \text{CHARTINDEX}(x_{near})$ 
5    $(u^*, t_f^*) \leftarrow \text{LQRPOLICY}(A, x_{near}, x_{rand})$ 
6   if  $t_f^* \leq t_{fp}^*$  then
7      $t_{fp}^* \leftarrow t_f^*$ 
8      $(x_{new}, u_{new}) \leftarrow \text{SIMULATE}(A, c, x_{near}, x_{rand}, u^*, t_f^*)$ 
9     if  $x_{new} \in \mathcal{X}_{feas}$  and  $x_{new} \neq x_{near}$  then
10       $T \leftarrow \text{ADDEDGE}(T, x_{near}, u_{new}, x_{new})$ 
11       $x_{near} \leftarrow x_{new}$ 
12 until  $t_f^* > t_{fp}^*$  or  $\|\varphi_c(x_{new}) - \varphi_c(x_{rand})\| \leq \delta$  or  $x_{new} \notin \mathcal{X}_{feas}$ 
13  $x_{new} \leftarrow x_{near}$ 
14 RETURN( $x_{new}$ )

```

selection and parameter sampling are repeated until y_{rand} falls inside the P_c set for the selected chart. This process generates a sample y_{rand} with uniform distribution over the union of the P_c sets covering T . Note here that the P_c set of a chart in the interior of the atlas is included in a ball of radius ρ , while the P_c set of a chart at the border of the atlas is included inside a ball of radius $\sigma > \rho$ (Fig. 9). If we fix ρ but increase σ the overall volume of the border charts increases, whereas that of the inner charts stays constant. Therefore, by increasing σ we can increase the exploration bias of the algorithm. This bias is analogous to the Voronoi bias in standard RRTs [58]. After generating a valid sample, the method then attempts to compute the point $x_{rand} = \Psi_c(y_{rand})$ (line 5) and returns this point if the Newton method implementing Ψ_c is successful (line 8). Otherwise, it returns the ambient space point corresponding to y_{rand} (line 7). This point lies on $\mathcal{T}_{x_c}\mathcal{X}$, instead of on \mathcal{X} , but it still provides a guiding direction to steer the tree towards unexplored regions of \mathcal{X} .

B. Tree extension

Algorithm 3 attempts to connect a state x_{near} to a state x_{rand} using LQR steering. The analogous procedure using ran-

Algorithm 4: Simulate an action.

SIMULATE($A, c, \mathbf{x}_k, \mathbf{x}_{rand}, \mathbf{u}^*, t_f^*$)**input** : An atlas, A , the chart index c , the state from where to start the simulation, \mathbf{x}_k , the state to approach, \mathbf{x}_{rand} , the policy to simulate, \mathbf{u}^* , and the optimal time t_f^* to simulate.**output**: The last state in the simulation and the executed control sequence.

```
1  $t \leftarrow 0$ 
2  $\mathbf{u}_k \leftarrow \emptyset$ 
3  $\text{VALIDSTATE} \leftarrow \text{TRUE}$ 
4 while  $\text{VALIDSTATE}$  and  $\|\boldsymbol{\varphi}_c(\mathbf{x}_k) - \boldsymbol{\varphi}_c(\mathbf{x}_{rand})\| > \delta$  and  $|t| < t_f^*$ 
  do
5    $(\mathbf{x}_{k+1}, \mathbf{y}_{k+1}, h) \leftarrow \text{NEXTSTATE}(\mathbf{x}_k, \mathbf{y}_k, \mathbf{u}^*(t), \mathbf{F}, \mathbf{x}_c, \mathbf{U}_c, \delta)$ 
6   if  $\mathbf{x}_{k+1} \notin \mathcal{X}_{feas}$  then
7      $\mathbf{x}_k \leftarrow \mathbf{x}_{k+1}$ 
8      $\text{VALIDSTATE} \leftarrow \text{FALSE}$ 
9   else
10    if  $\|\mathbf{x}_{k+1} - (\mathbf{x}_c + \mathbf{U}_c \mathbf{y}_{k+1})\| > \varepsilon$  or
11       $\|\mathbf{y}_{k+1} - \mathbf{y}_k\| / \|\mathbf{x}_{k+1} - \mathbf{x}_k\| < \cos(\alpha)$  or
12       $\|\mathbf{y}_{k+1}\| > \rho$  then
13       $\text{ADDCARTOATLAS}(A, \mathbf{x}_k)$ 
14       $\text{VALIDSTATE} \leftarrow \text{FALSE}$ 
15    else
16       $\mathbf{x}_k \leftarrow \mathbf{x}_{k+1}$ 
17       $\mathbf{u}_k \leftarrow \mathbf{u}_k \cup \{(\mathbf{u}(t), h)\}$ 
18       $t \leftarrow t + h$ 
19      if  $\mathbf{y}_{k+1} \notin P_c$  then
20       $\text{VALIDSTATE} \leftarrow \text{FALSE}$ 
21 RETURN( $\mathbf{x}_k, \mathbf{u}_k$ )
```

domised steering is available in [18]. The algorithm implements a loop where, initially, the optimal policy \mathbf{u}^* and time t_f^* to connect these two states are computed (line 5). The policy is a function of time given by Eq. (32). If t_f^* is lower than the optimal time t_{fp}^* obtained in the previous iteration, the policy is used to simulate the evolution of the system from \mathbf{x}_{near} (line 8). The simulation produces a state \mathbf{x}_{new} which, if it is feasible and different from \mathbf{x}_{near} , is added to the tree. This involves the creation of an edge between \mathbf{x}_{near} and \mathbf{x}_{new} (line 10), which stores the control sequence \mathbf{u}_{new} executed in the simulation. The loop is repeated until t_f^* is larger than t_{fp}^* (line 12), or \mathbf{x}_{rand} is reached with accuracy δ in parameter space, or the next state is unfeasible.

Algorithm 4 summarises the procedure used to simulate a given policy $\mathbf{u}^*(t)$ from a particular state \mathbf{x}_k . The simulation progresses while the new state is valid, the target state is not reached with accuracy δ in parameter space, and the integration time t is lower than t_f^* (line 4). A state is not valid if it is not in \mathcal{X}_{feas} (line 8), if it is not in the validity area of the chart (line 14), or if it is not included in the current P_c set (line 20), i.e., it is parametrised by a neighbouring chart. In the first case, both the simulation and the connection between states are stopped. In the last two cases the simulation is stopped, but the connection continues after recomputing the optimal policy, on a newly created chart (line 13) or on the neighbouring chart, respectively.

The key procedure in the simulation is the **NEXTSTATE**

method (line 5), which provides the next state \mathbf{x}_{k+1} , given the current state \mathbf{x}_k and the action $\mathbf{u}^*(t)$ at time t . The elements of $\mathbf{u}^*(t)$ are saturated to their bounds in Eq. (7) if such bounds are surpassed. Then, the simulation is implemented by integrating Eq. (6) using local coordinates as explained in Section V-A. Any numerical integration method, either explicit or implicit, could be used to discretise Eq. (10). We here apply the trapezoidal rule as it yields an implicit integrator whose computational cost (integration and projection to the manifold) is similar to the cost of using an explicit method of the same order [49]. Using this rule, Eq. (10) is discretised as

$$\mathbf{y}_{k+1} = \mathbf{y}_k + \frac{h}{2} \mathbf{U}_c^\top (\mathbf{g}(\mathbf{x}_k, \mathbf{u}) + \mathbf{g}(\mathbf{x}_{k+1}, \mathbf{u})), \quad (34)$$

where h is the integration time step. The value \mathbf{x}_{k+1} in Eq. (34) is unknown but, since it must satisfy Eq. (9), it must fulfil

$$\begin{aligned} \mathbf{F}(\mathbf{x}_{k+1}) &= \mathbf{0}, \\ \mathbf{U}_c^\top (\mathbf{x}_{k+1} - \mathbf{x}_c) - \mathbf{y}_{k+1} &= \mathbf{0}. \end{aligned} \quad (35)$$

Now, substituting Eqs. (34) into Eq. (35) we obtain

$$\begin{aligned} \mathbf{F}(\mathbf{x}_{k+1}) &= \mathbf{0}, \\ \mathbf{U}_c^\top (\mathbf{x}_{k+1} - \frac{h}{2} (\mathbf{g}(\mathbf{x}_k, \mathbf{u}) + \mathbf{g}(\mathbf{x}_{k+1}, \mathbf{u})) - \mathbf{x}_c) - \mathbf{y}_k &= \mathbf{0}, \end{aligned} \quad (36)$$

where \mathbf{x}_k , \mathbf{y}_k , and \mathbf{x}_c are known and \mathbf{x}_{k+1} is the unknown to be determined. We could use a Newton method to solve this system, but the Broyden method is preferable as it avoids the computation of the Jacobian of the system at each step. Potra and Yen [49] gave an approximation of this Jacobian that allows finding \mathbf{x}_{k+1} in only a few iterations. For backward integration, i.e., when extending the RRT with root at \mathbf{x}_g , the time step h in Eq. (36) must simply be negative.

C. Setting the planner parameters

The planner depends on eight parameters: the three parameters ε , α , and ρ controlling chart creation, the radius σ used for sampling, the tolerances δ and β measuring closeness between states and trees, respectively, and the LQR steering parameters \mathbf{R} and t_{max} . All of them are positive reals except \mathbf{R} , which must be a $n_u \times n_u$ symmetric positive-definite matrix.

Parameters ε , α , and ρ appear, respectively, in Eqs. (11), (12), and (13). Parameter α bounds the angle between neighboring charts. This angle should be small, otherwise the V_c sets for neighboring charts might not overlap, impeding a smooth transition between the charts [53]. Such problematic areas can be detected and patched [12], but this process introduces inefficiencies. Thus, we suggest to keep this parameter below $\pi/6$. Parameter ε is only relevant if the distance between the manifold and the tangent space becomes large without a significant change in curvature, which rarely occurs. Since this distance is computed in ambient space, if on average we wish to tolerate an error of e in each dimension, we should set $\varepsilon \simeq e\sqrt{n_\chi}$. In our test cases we have used $\varepsilon = 0.05\sqrt{n_\chi}$. Finally, ρ only plays a relevant role on almost flat manifolds. The only restriction to consider is that ρ must be smaller than σ to ensure the eventual creation of new charts. Following [12], we suggest to set $\rho = d_\chi/2$. With this value, charts are generally created before the numerical process implementing Eq. (9)




Fig. 10. Example tasks used to illustrate the performance of the planner. From left to right, and columnwise: weight lifting, weight throwing, conveyor switching, and truck loading. The robots involved are, respectively, a four-bar robot, a five-bar robot, a Delta robot, and a double-arm manipulation system. The top and bottom rows show the start and goal states for each task. In the goal state of the second task, and in the start state of the third task, the load is moving at a certain velocity indicated by the red arrow. The velocity of the remaining start and goal states is null. In all robots, the motor torques are limited to prevent the generation of direct trajectories to the goal.

fails and before Eqs. (11) and (12) hold. In this way, the paving of the manifold tends to be more regular.

As explained in Section VII-A, the sampling radius σ used in line 3 of Algorithm 2 controls the exploration bias of the algorithm. The role of σ is equivalent to that of the parameter used in standard RRTs to limit the sampling space (e.g., the boundaries of a 2D maze where a mobile robot is set to move). A too large σ complicates the solution of problems with narrow corridors. Thus, we propose to set $\sigma = 2\rho$ since this a moderate value that still creates a strong push towards unexplored regions, specially in large dimensional state spaces. If necessary, existing techniques to automatically tune this parameter [59] could be adapted to kinodynamic planning.

Parameter δ appears in line 12 of Algorithm 3 and in line 4 of Algorithm 4. An equivalent parameter is present in the standard RRT algorithm [7]. If two states are closer than δ , they are considered to be close enough so that the transition between them is not problematic. Thus, this parameter is used as an upper bound of the distance between consecutive states along an RRT branch. Therefore, the value of h in Eq. (36) is adjusted so that $\|\varphi_c(\mathbf{x}_k) - \varphi_c(\mathbf{x}_{k+1})\| < \delta$. Moreover, to correctly detect the transition between charts, δ must be significantly smaller than ρ . With these considerations in mind, we propose to set $\delta \simeq 0.02\rho$.

Parameter β appears in line 11 of Algorithm 1 and is the

tolerated error in the connection between trees. This parameter is also used in the standard RRT algorithm. A small value may unnecessarily complicate some problems, specially if the steering algorithm is not very precise (like in randomised steering), and a large value may produce unfeasible solutions. We suggest to use $\beta = 0.1\sqrt{n\chi}$, but this value has to be tuned according to the particularities of the obstacles in the environment.

Matrix \mathbf{R} in Eq. (17) is used in the standard LQR to penalise the control effort employed and is typically initialised using the Bryson rule [60]. Finally, t_{max} fixes the time window over which $J(t_f)$ in Eq. (33) is to be minimised. Ideally, it should be slightly larger than t_f^* . A much larger value would result in a waste of computational resources and a too low value would produce sub-optimal policies. We propose to set this parameter to a fraction of the expected trajectory time t_g .

VIII. PROBABILISTIC COMPLETENESS

In its fully randomised version, i.e., when using randomised steering instead of LQR steering, the planner is probabilistically complete. A formal proof of this point would replicate the same arguments used in [61] with minor adaptations, so we only sketch the main points supporting the claim.

Assume that the action to execute is selected at random from \mathcal{U} , with a random time horizon. Then, in the part




Fig. 11. Solution trajectories for the four test cases. The shown trails depict earlier positions of the load during a same time span. A longer trail, therefore, corresponds to a higher velocity of the load. See https://youtu.be/_DMzK5SGzQ for an animated version of this figure.

of \mathcal{X} already covered by a partial atlas, we are in the same situation as the one considered in [61, Section IV]: \mathcal{X} is a smooth manifold, we have a procedure to sample \mathcal{X} , Euclidean distance is used to determine nearest neighbours, and the motion of the system is governed by a differential equation depending on the state and the control inputs. The main relevant difference is that our sample distribution is uniform in tangent space, but not on \mathcal{X} . However, the difference between the uniform distribution in parameter space and the actual distribution on the manifold is bounded by parameter α [62]. Thus, the probability bounds given in [61] may need to be modified, but their proof would still hold. Thus, under the same mild conditions assumed in [61] (i.e., Lipschitz-continuity conditions), our planner with randomised steering is probabilistically complete in, at least, the part of the manifold already covered by the atlas. This implies that the planner will be probabilistically complete provided it is able to extend the

atlas to cover \mathcal{X} completely. But this will certainly be achieved in the limit, as the procedure described in Section V-A ensures that new charts are generated each time the RRT branches approach the border of the subset of \mathcal{X} covered up to a given moment. The reasoning in [61] can also be used to provide a formal proof that the tree will eventually reach such border regions, just by defining goal areas in them. Moreover, as shown in [53], the expansion of the atlas will stop when the atlas has no border, i.e., when it fully covers \mathcal{X} . The chart coordination procedure described in Section V-C may leave uncovered areas on the manifold, of size $O(\alpha)$ [53]. However, such areas can be detected during tree extension, and can be eliminated by slightly enlarging the P_c sets of the charts around them, as described in [12].

In principle, the use of LQR steering instead of randomised steering can only result in better performance, as it should facilitate the connection between the balls used in [61, The-




Fig. 12. Geometry of the four-bar mechanism in Fig. 10, left column. For each coordinate system, only the x axis is depicted.

orem 2] to cover the solution trajectory: connecting them using LQR steering should be easier than doing so with randomised steering. However, a formal proof of this point would require to provide error bounds for the LQR steering policies analogous to those in [61, Lemma 3] for randomly-selected constant actions. As in [61], the obtention of such bounds remains as an open problem however, so we only conjecture the planner to be probabilistically complete if LQR steering is used. Even so, note we could always retain the probabilistic completeness by using randomised steering once in a while, instead of using LQR steering exclusively.

IX. PLANNING EXAMPLES

The planner has been implemented in C and it has been integrated into the CUIK suite [63]. We next analyse its performance in planning four tasks of increasing complexity (Fig. 10). The first two tasks involve planar single-loop mechanisms, which are simple enough to illustrate key aspects of the planner, like the performance of the steering method, the traversal of singularities, or the ability to plan trajectories towards states of nonzero velocity. The third and fourth tasks, on the other hand, show the planner performance in spatial robots of considerable complexity. In all cases the robots are subject to gravity and viscous friction in all joints, and their action bounds l_i in Eq. (7) are small enough so as to impede direct trajectories between \mathbf{x}_s and \mathbf{x}_g . This complicates the problems and forces the generation of swinging motions to reach the goal. Following Section VII-C, we have fixed $\cos(\alpha) = 0.9$, $\varepsilon = 0.05 \sqrt{n_{\mathcal{X}}}$, $\rho = d_{\mathcal{X}}/2$, $\sigma = 2\rho$, $\delta = 0.02\rho$, $\beta = 0.1 \sqrt{n_{\mathcal{X}}}$. Matrix \mathbf{R} in Eq. (17) has been set to be diagonal, with $\mathbf{R}_{i,i} = 1/l_i^2$, and we use $t_{max} = 1.5$. The planner performance, however, does not depend on these parameters exclusively, but also on the peculiarities of each problem, like the torque limits of the actuators, the system masses, or




Fig. 13. A partial atlas of \mathcal{X} used to plan the lifting of a weight with the four-bar robot. The red and green trees are rooted at \mathbf{x}_s and \mathbf{x}_g respectively, and they are grown towards each other in parallel with the atlas. Each polygon in dark blue corresponds to the P_c set of a given chart. To allow a clearer visualization of the atlas, we have used $\rho = 0.5$ to obtain the plot, so the shown charts are actually smaller than those used by the planner. See https://youtu.be/_DMzK5SGzQ for an animated version of this figure.

the presence of obstacles. An example of the trajectories we obtain can be seen in Fig. 11 and in the companion video of this paper (also available in https://youtu.be/_DMzK5SGzQ). The complete set of geometric and dynamic parameters of all examples is provided in <http://www.iri.upc.edu/cuik>.

Table I summarises the problem dimensions and performance statistics for the four mentioned tasks. For each task we provide the number of generalised coordinates in \mathbf{q} (n_q), the number of loop-closure constraints (n_e), the dimension of the state space ($d_{\mathcal{X}}$), and the dimension of the action space (n_u). The specific formulations used for Eqs. (1) and (2) are given in [48]. The table also provides the average over twenty runs of the number of samples and charts required to solve the problem, and the planning time in seconds using a MacBook Pro with an Intel i9 octa-core processor running at 2.93 GHz. The column “Success rate” gives the percentage of planner runs that were able to solve each problem in at most one hour. Statistics for both the randomised steering strategy in [18] and the LQR steering strategy of this paper are given for comparison. The randomised strategy employs $2n_u$ random actions from \mathcal{U} , which are applied during 0.1 seconds in accordance with [18]. As seen in the table, the LQR strategy is more efficient than the randomised strategy, as it requires a smaller number of samples and charts, and less time, to find a solution. In fact, the success rate of the randomised strategy is only 55% in the truck loading task. Further details on the four tasks are next provided.

A. Weight lifting

The first task to be planned consists in lifting a heavy load with a four-bar robot (Fig. 10, left column). The robot involves four links cyclically connected with revolute joints from which only joint J_1 is actuated (Fig. 12). The relative angle with the

TABLE I
PROBLEM DIMENSIONS AND PERFORMANCE STATISTICS FOR THE EXAMPLE TASKS.

Example task	n_q	n_e	$d_{\mathcal{X}}$	n_u	Randomised steering				LQR steering		
					No. samples	No. charts	Plan. Time (s)	Success Rate	No. samples	No. charts	Plan. Time (s)
Weight lifting	4	3	2	1	525	94	0.89	100%	66	46	0.17
Weight throwing	5	3	4	2	17085	5358	359.95	100%	1250	607	9.83
Conveyor switching	15	12	6	3	24247	256	420.13	100%	1809	165	40.89
Truck loading	10	6	8	10	11600	1282	1649.86	55%	9779	1269	105.47

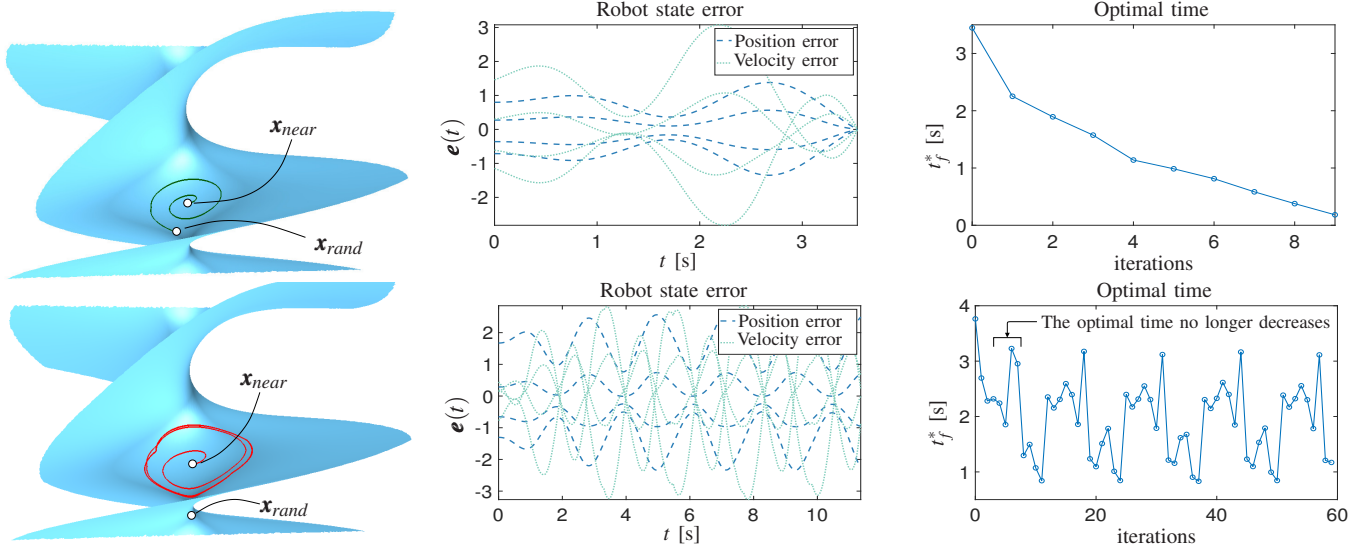


Fig. 14. Steering the four-bar robot from \mathbf{x}_{near} to \mathbf{x}_{rand} . Top: The LQR strategy allows the planner to connect \mathbf{x}_{near} and \mathbf{x}_{rand} . Bottom: The strategy enters a limit cycle and is never able to reach \mathbf{x}_{rand} . The right plot shows that t_f^* no longer decreases after six iterations, so it would be aborted at this point.

following link is denoted by q_i , and the robot configuration is then given by $\mathbf{q} = (q_1, q_2, q_3, q_4)$.

Figure 13 shows the shape of \mathcal{X} when projected to the space defined by q_1 , \dot{q}_1 , and \dot{q}_2 , with the start and goal states indicated. To design a trajectory connecting \mathbf{x}_s with \mathbf{x}_g , the planner constructs the partial atlas that is shown in the figure. Since the motor torque at J_1 is limited, quasi static trajectories near the straight line from \mathbf{x}_s to \mathbf{x}_g are impossible, and the robot is deemed to perform pendulum-like motions to be able to reach the goal. This translates into the spirally tree trajectories that we observe in the figure. A typical trajectory returned by the planner can be seen in Fig. 11, top row.

This example can be used to illustrate the performance of the LQR steering strategy. Fig. 14-top, shows an example in which this strategy successfully finds a trajectory connecting \mathbf{x}_{near} with \mathbf{x}_{rand} , i.e., the error $\mathbf{e}(t) = \mathbf{x}(t) - \mathbf{x}_{rand}$ converges to $\mathbf{0}$. Whenever a new chart is created during the simulation or when a chart is revisited, the policy is recomputed and, in this example, t_f^* monotonically decreases. In contrast, Fig. 14-bottom shows another example in which the process tends to a limit cycle like the one in Fig. 8(b), and is never able to reach the goal, i.e., the error $\mathbf{e}(t)$ is never $\mathbf{0}$. The steering method in Algorithm 3 would stop as soon as t_f^* no longer decreases.

In Fig. 15 we also show the performance of the LQR strategy for states \mathbf{x}_{rand} that are progressively further away from \mathbf{x}_{near} . We have generated 5 batches of 100 random samples, where the samples in each batch are at tangent space

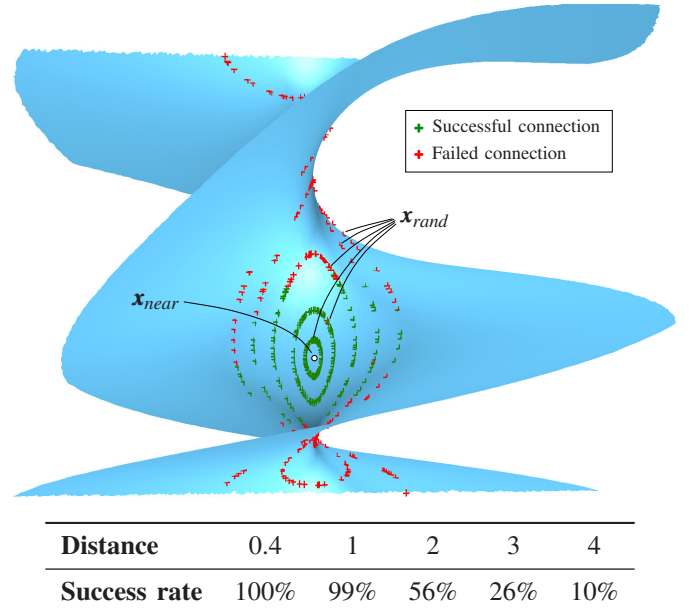


Fig. 15. Success rate of the LQR steering strategy for states \mathbf{x}_{rand} that are increasingly far from \mathbf{x}_{near} .

distances of 0.4, 1, 2, 3, and 4 from \mathbf{x}_{near} . As a reference, the distance from \mathbf{x}_s to \mathbf{x}_g is 3.7 in this example. The states \mathbf{x}_{rand} that could be connected to \mathbf{x}_{near} are shown in green, while

those that could not be shown in red. As expected for a local planner, the closer \mathbf{x}_{rand} from \mathbf{x}_{near} , the higher the probability of success of the steering process.

B. Weight throwing

The second task involves a five-bar robot. It consists in throwing a given object from a certain position at a prescribed velocity (indicated with the red arrow in Fig. 10, second column). This shows the planner ability to reach goal states \mathbf{x}_g with nonzero velocity, which would be difficult to achieve with conventional C-space approaches.

The computed trajectory can be seen in the second row of Fig. 11. The robot first lifts the object to the right until it achieves a zero-velocity position (second snapshot), to later move it back to the left along a nearly-circular path (remaining snapshots). Almost two turns of this path are completed in order to reach the launch point with the required momentum (last snapshot).

The task also illustrates the planner capacity to traverse forward singularities, which are configurations in which the robot is locally underactuated. These configurations are difficult to manage, as they can only be crossed under specific velocities and accelerations fulfilling certain geometric conditions [17, 64]. However, since our planner trajectories result from simulating control policies $\mathbf{u}(t)$ using forward dynamics, they naturally satisfy the mentioned conditions at the singularities, and are thus kinematically and dynamically feasible even in such configurations. In particular, a five-bar robot is known to exhibit a forward singularity when its two distal links happen to be aligned [65]. In the trajectory shown in Fig. 11 this occurs in the third and sixth snapshots. From the companion video we see that the robot passes through these configurations in a smooth and predictable manner with no difficulty. Note that, while such a trajectory would be difficult to execute using classical computed-torque controllers [66], recent LQR controllers for closed kinematic chains have no trouble in accomplishing this task [67].

C. Conveyor switching

In the previous tasks the robot was a single-loop mechanism in an obstacle-free environment. To exemplify the planner in a multi-loop mechanism surrounded by obstacles, we next apply it to a conveyor switching task on a Delta robot (Fig. 10, third column).

The task to be planned consists in picking a loudspeaker from a conveyor belt moving at a certain speed, to later place it inside a static box on a second belt. Obstacles play a major role in this example, as the planner has to avoid the collisions of the robot with the conveyor belts, the boxes, and the supporting structure, while respecting the joint limits. In fact, around 70% of branch extensions are stopped due to collisions in this example. Moreover, the boxes have thin walls that require us to set $\beta = 0.05\sqrt{n_{\chi}}$ to avoid obtaining unfeasible solutions. An example of a resulting trajectory can be seen in Fig. 11, third row, and in its companion video. Given the velocity of the moving belt, the planner is forced to reduce the initial momentum of the load before it can place it

inside the destination box. The trajectory follows an ascending path that converts the initial momentum into potential energy, to later move the load back to the box on the goal location.

D. Truck loading

The fourth task involves two 7-DOF Franka Emika arms moving a gas bottle cooperatively. The task consists in lifting the bottle onto a truck while avoiding the collisions with the surrounding obstacles (a conveyor belt, the ground, and the truck). The first and last joints in each arm are held fixed during the task, and the goal is to compute control policies for the remaining joints, which are all actuated. The weight of the bottle is twice the added payload of the two arms, so in this example the planner allows the system to move much beyond its static capabilities.

The example also illustrates that the randomised steering strategy performs poorly when n_u is large. In this case, $n_u = 10$, which is notably higher than in the previous examples. Note that the number of random actions needed to properly represent \mathcal{U} should be proportional to its volume, so it should grow exponentially with n_u in principle. To alleviate the curse of dimensionality, however, [7] proposes to simulate only $2n_u$ random actions for each branch extension. Our implementation adopts this criterion but, like [7], it then shows a poor exploration capacity when n_u is large, resulting in the excessive planning times reported for the truck loading task (Table I). We have also tried to simulate 2^{n_u} random actions, instead of just $2n_u$, but then the gain in exploration capacity does not outweigh the large computational cost of simulating the actions. In contrast, the LQR strategy only computes one control policy per branch extension, so an increase in n_u does not affect the planning time dramatically (Table I, last column). Using this strategy, the planner obtained trajectories like the one shown in Fig. 11, bottom row, in which we see that, in order to gain momentum, the robot is moved backwards before it lifts the bottle onto the truck.

X. CONCLUSIONS

This paper has proposed a randomised planner to compute dynamically-feasible trajectories for robots with closed kinematic chains. The state space of such robots is an intricate manifold that poses three major hurdles to the planner design: 1) the generation of random samples on the manifold; 2) the accurate simulation of robot trajectories along the manifold; and 3) the steering of the system towards random states. The three issues have been addressed by constructing an atlas of the manifold in parallel to the RRT. The result is a planner that can explore the state space in an effective manner, while conforming to the vector fields defined by the equations of motion and the force bounds of the actuators. In its fully randomised version (i.e., using randomised steering), the planner is probabilistically complete. We also conjecture it is probabilistically complete if LQR steering is used, but proving this point remains open so far. The examples in the paper show that the planner can solve significantly complex problems that require the computation of swinging motions between start

and goal states, under restrictive torque limitations imposed on the motors.

Several points should be considered in further improvements of this work. Note that, as usual in a randomised planner, our control policies are piecewise continuous, so the planned trajectories are smooth in position, but not in velocity or acceleration. Therefore, to reduce control or vibration issues in practice, a post-processing should be applied to obtain twice differentiable trajectories. The trajectories should also be optimised in some sense, minimising the time or control effort required for its execution. Trajectory optimization tools like those in [38], [41], or [42] might be very helpful to both ends. Another sensitive point is the metric employed to measure the distance between states. This is a general concern in any motion planner, but it is more difficult to address in our context as such metric should not only consider the vector flows defined by the equations of motion, but also the curvature of the state space manifold defined by the loop-closure constraints. Using a metric derived from geometric insights provided by such constraints might result in substantial performance improvements. Another point deserving attention would be the monitoring of constraint forces during the extension of the RRT. While such forces result in no motion, they do stress the robot parts unnecessarily and should be kept under admissible bounds. Note that these forces can be computed as the simulations proceed, since they can be inferred, e.g., from the values of the Lagrange multipliers involved in the equations of motion [68]. The ability to impose bounds on constraint forces would also allow computing trajectories in closed kinematic chains induced by unilateral contacts, like those that arise when a hand moves an object in contact with a surface. Such contacts could be maintained along a trajectory by setting pertinent signed bounds on the constraint forces arising.

REFERENCES

- [1] B. Donald, P. Xavier, J. Canny, and J. Reif, "Kinodynamic motion planning," *Journal of the ACM*, vol. 40, no. 5, pp. 1048–1066, 1993.
- [2] S.-H. Lee, J. Kim, F. C. Park, M. Kim, and J. E. Bobrow, "Newton-type algorithms for dynamics-based robot movement optimization," *IEEE Transactions on robotics*, vol. 21, no. 4, pp. 657–667, 2005.
- [3] I. Bonev, "Delta parallel robot - the story of success," *Newsletter*, available at <http://www.parallelimc.org>, 2001.
- [4] S. Feng, E. Whitman, X. Xinjilefu, and C. G. Atkeson, "Optimization based full body control for the atlas robot," in *IEEE-RAS International Conference on Humanoid Robots*, 2014, pp. 120–127.
- [5] M. A. Diftler, J. S. Mehling, M. E. Abdallah, N. A. Radford, L. B. Bridgwater, A. M. Sanders, R. S. Askew, D. M. Linn, J. D. Yamokoski, F. A. Permenter, B. K. Hargrave, R. Platt, R. T. Savely, and R. O. Ambrose, "Robonaut 2 - the first humanoid robot in space," in *2011 IEEE International Conference on Robotics and Automation*, 2011, pp. 2178–2183.
- [6] R. P. Hoyt, "Spiderfab: An architecture for self-fabricating space systems," in *AIAA Space 2013 Conference and Exposition*, 2013, p. 5509.
- [7] S. M. LaValle and J. J. Kuffner, "Randomized kinodynamic planning," *International Journal of Robotics Research*, vol. 20, no. 5, pp. 378–400, 2001.
- [8] Y. Li, Z. Littlefield, and K. E. Bekris, "Asymptotically optimal sampling-based kinodynamic planning," *The International Journal of Robotics Research*, vol. 35, no. 5, pp. 528–564, 2016.
- [9] K. Hauser and Y. Zhou, "Asymptotically optimal planning by feasible kinodynamic planning in a state-cost space," *IEEE Transactions on Robotics*, vol. 32, no. 6, pp. 1431–1443, 2016.
- [10] D. Berenson, S. Srinivasa, and J. J. Kuffner, "Task space regions: A framework for pose-constrained manipulation planning," *International Journal of Robotics Research*, vol. 30, no. 12, pp. 1435–1460, 2011.
- [11] J. M. Porta, L. Jaillet, and O. Bohigas, "Randomized path planning on manifolds based on higher-dimensional continuation," *The International Journal of Robotics Research*, vol. 31, no. 2, pp. 201–215, 2012.
- [12] L. Jaillet and J. M. Porta, "Path planning under kinematic constraints by rapidly exploring manifolds," *IEEE Transactions on Robotics*, vol. 29, no. 1, pp. 105–117, 2013.
- [13] K. Hauser, "Fast interpolation and time-optimization with contact," *The International Journal of Robotics Research*, vol. 33, no. 9, pp. 1231–1250, 2014.
- [14] Q.-C. Pham and O. Stasse, "Time-optimal path parameterization for redundantly actuated robots: A numerical integration approach," *IEEE/ASME Transactions on Mechatronics*, vol. 20, no. 6, pp. 3257–3263, 2015.
- [15] Z. Kingston, M. Moll, and L. E. Kavraki, "Decoupling constraints from sampling-based planners," in *International Symposium of Robotics Research*, 2017.
- [16] —, "Sampling-based methods for motion planning with constraints," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 1, pp. 159–185, 2018.
- [17] O. Bohigas, M. Manubens, and L. Ros, *Singularities of robot mechanisms: Numerical computation and avoidance path planning*, ser. Mechanisms and Machine Science. Springer, 2017, vol. 41.
- [18] R. Bordenal, L. Ros, and J. M. Porta, "Randomized kinodynamic planning for constrained systems," in *IEEE International Conference on Robotics and Automation*, 2018, pp. 7079–7086.
- [19] J. H. Yakey, S. M. LaValle, and L. E. Kavraki, "Randomized path planning for linkages with closed kinematic chains," *IEEE Transactions on Robotics and Automation*, vol. 17, no. 6, pp. 951–959, 2001.
- [20] L. Han and N. M. Amato, "A kinematics-based probabilistic roadmap method for closed chain systems," in *Algorithmic and Computational Robotics - New Directions*, 2000, pp. 233–246.
- [21] J. Cortés, T. Siméon, and J.-P. Laumond, "A random loop generator for planning the motions of closed kinematic chains using PRM methods," in *IEEE International Conference on Robotics and Automation*, 2002, pp. 2141–2146.
- [22] M. Stilman, "Task constrained motion planning in robot joint space," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2007, pp. 3074–3081.
- [23] B. Kim, T. T. Um, C. Suh, and F. C. Park, "Tangent bundle RRT: A randomized algorithm for constrained motion planning," *Robotica*, vol. 34, no. 1, pp. 202–225, 2016.
- [24] S. M. LaValle, *Planning algorithms*. New York: Cambridge University Press, 2006.
- [25] H. Choset, K. M. Lynch, S. Hutchinson, G. A. Kantor, W. Burgard, L. E. Kavraki, and S. Thrun, *Principles of robot motion: theory, algorithms, and implementations*, ser. Intelligent Robotics and Autonomous Agents. MIT Press, 2005.
- [26] J. E. Bobrow, S. Dubowsky, and J. Gibson, "Time-optimal control of robotic manipulators along specified paths," *The International Journal of Robotics Research*, vol. 4, no. 3, pp. 3–17, 1985.
- [27] K. Shin and N. McKay, "Minimum-time control of robotic manipulators with geometric path constraints," *IEEE Transactions on Automatic Control*, vol. 30, no. 6, pp. 531–541, 1985.
- [28] F. Pfeiffer and R. Johanni, "A concept for manipulator trajectory planning," *IEEE Journal on Robotics and Automation*, vol. 3, no. 2, pp. 115–123, 1987.
- [29] J.-J. Slotine and H. S. Yang, "Improving the efficiency of time-optimal path-following algorithms," *IEEE Transactions on Robotics and Automation*, vol. 5, no. 1, pp. 118–124, 1989.
- [30] Z. Shiller and H.-H. Lu, "Computation of path constrained time optimal motions with dynamic singularities," *Journal of Dynamic Systems, Measurement, and Control*, vol. 114, no. 1, pp. 34–40, 1992.
- [31] Q.-C. Pham, "A general, fast, and robust implementation of the time-optimal path parameterization algorithm," *IEEE Transactions on Robotics*, vol. 30, no. 6, pp. 1533–1540, 2014.
- [32] Q.-C. Pham, S. Caron, P. Lertkultanon, and Y. Nakamura, "Admissible velocity propagation: Beyond quasi-static path planning for high-dimensional robots," *The International Journal of Robotics Research*, vol. 36, no. 1, pp. 44–67, 2017.
- [33] J. E. Bobrow, "Optimal robot path planning using the minimum-time criterion," *IEEE Journal on Robotics and Automation*, vol. 4, no. 4, pp. 443–450, 1988.

- [34] Z. Shiller and S. Dubowsky, "Robot path planning with obstacles, actuator, gripper, and payload constraints," *The International Journal of Robotics Research*, vol. 8, no. 6, pp. 3–18, 1989.
- [35] Z. Shiller and S. Dubowsky, "On computing the global time-optimal motions of robotic manipulators in the presence of obstacles," *IEEE Transactions on Robotics and Automation*, vol. 7, no. 6, pp. 785–797, Dec 1991.
- [36] Q.-C. Pham, S. Caron, and Y. Nakamura, "Kinodynamic planning in the configuration space via admissible velocity propagation," in *Robotics: Science and Systems*, 2013.
- [37] M. Kalakrishnan, S. Chitta, E. Theodorou, P. Pastor, and S. Schaal, "STOMP: Stochastic trajectory optimization for motion planning," in *IEEE International Conference on Robotics and Automation*, 2011, pp. 4569–4574.
- [38] M. Posa, S. Kuindersma, and R. Tedrake, "Optimization and stabilization of trajectories for constrained dynamical systems," in *IEEE International Conference on Robotics and Automation*, 2016, pp. 1366–1373.
- [39] J. Schulman, D. Y. J. Ho, A. Lee, I. Awwal, H. Bradlow, J. Pan, S. Patil, K. Goldberg, and P. Abbeel, "Motion planning with sequential convex optimization and convex collision checking," *The International Journal of Robotics Research*, vol. 33, no. 9, pp. 1251–1270, 2014.
- [40] M. Zucker, N. Ratliff, A. D. Dragan, M. Pivtoraiko, M. Klingensmith, C. M. Dellin, J. A. Bagnell, and S. S. Srinivasa, "CHOMP: Covariant Hamiltonian optimization for motion planning," *The International Journal of Robotics Research*, vol. 32, no. 9-10, pp. 1164–1193, 2013.
- [41] J. T. Betts, *Practical methods for optimal control and estimation using nonlinear programming*. SIAM Publications, 2010.
- [42] —, "Survey of numerical methods for trajectory optimization," *Journal of Guidance, Control, and Dynamics*, vol. 21, no. 2, pp. 193–207, 1998.
- [43] T. Kunz and M. Stilman, "Probabilistically complete kinodynamic planning for robot manipulators with acceleration limits," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2014, pp. 3713–3719.
- [44] A. Perez, R. Platt, G. Konidaris, L. Kaelbling, and T. Lozano-Perez, "LQR-RRT*: Optimal sampling-based motion planning with automatically derived extension heuristics," in *IEEE International Conference on Robotics and Automation*, 2012, pp. 2537–2542.
- [45] R. Tedrake, "LQR-Trees: Feedback motion planning on sparse randomized trees," in *Robotics: Science and Systems*, 2009.
- [46] G. Goretkin, A. Perez, R. Platt, and G. Konidaris, "Optimal sampling-based planning for linear-quadratic kinodynamic systems," in *IEEE International Conference on Robotics and Automation*, 2013, pp. 2429–2436.
- [47] D. J. Webb and J. van den Berg, "Kinodynamic RRT*: Asymptotically optimal motion planning for robots with linear dynamics," in *IEEE International Conference on Robotics and Automation*, 2013, pp. 5054–5061.
- [48] R. Bordenal, L. Ros, and J. M. Porta, "A randomized kinodynamic planner for closed-chain robotic systems," Institut de Robòtica i Informàtica. CSIC-UPC. Technical Report, 2019, <https://bit.ly/2vk4MSO>.
- [49] F. A. Potra and J. Yen, "Implicit numerical integration for Euler-Lagrange equations via tangent space parametrization," *Journal of Structural Mechanics*, vol. 19, no. 1, pp. 77–98, 1991.
- [50] L. R. Petzold, "Numerical solution of differential-algebraic equations in mechanical systems simulation," *Physica D: Nonlinear Phenomena*, vol. 60, no. 1–4, pp. 269–279, 1992.
- [51] R. Featherstone and D. E. Orin, "Dynamics," in *Springer Handbook of Robotics*. Springer, 2016, pp. 37–66.
- [52] W. Blajer, "Methods for constraint violation suppression in the numerical simulation of constrained multibody systems - A comparative study," *Computer Methods in Applied Mechanics and Engineering*, vol. 200, no. 13-16, pp. 1568–1576, 2011.
- [53] M. E. Henderson, "Multiple parameter continuation: computing implicitly defined k-manifolds," *International Journal of Bifurcation and Chaos*, vol. 12, no. 3, pp. 451–476, 2002.
- [54] E. Hairer, "Geometric integration of ordinary differential equations on manifolds," *BIT Numerical Mathematics*, vol. 41, no. 5, pp. 996–1007, 2001.
- [55] E. Hairer, C. Lubich, and G. Wanner, "Geometric numerical integration: structure-preserving algorithms for ordinary differential equations," 2006.
- [56] F. L. Lewis, D. Vrabie, and V. L. Syrmos, *Optimal control*. John Wiley & Sons, 2012.
- [57] D. P. Bertsekas, *Dynamic programming and optimal control*. Athena Scientific, 2005.
- [58] J. Kuffner and S. M. LaValle, "RRT-connect: An efficient approach to single-query path planning," in *IEEE International Conference on Robotics and Automation*, vol. 2, 2000, pp. 995–1001.
- [59] J. M. Porta and L. Jaillet, "Sampling strategies for path planning under kinematic constraints," *CoRR*, vol. abs/1407.2544, 2014.
- [60] A. E. Bryson and Y.-C. Ho, *Applied Optimal Control: Optimization, Estimation and Control*. Taylor Francis, 1975.
- [61] M. Kleinbort, K. Solovey, Z. Littlefield, K. E. Bekris, and D. Halperin, "Probabilistic completeness of RRT for geometric and kinodynamic planning with forward propagation," *IEEE Robotics and Automation Letters*, vol. 4, no. 3, pp. 277–283, 2019.
- [62] L. Jaillet and J. M. Porta, "Efficient asymptotically-optimal path planning on manifolds," *Robotics and Autonomous Systems*, vol. 61, no. 8, pp. 797–807, 2013.
- [63] J. M. Porta, L. Ros, O. Bohigas, M. Manubens, C. Rosales, and L. Jaillet, "The Cuik Suite: Analyzing the motion of closed-chain multibody systems," *IEEE Robotics and Automation Magazine*, vol. 21, no. 3, pp. 105–114, 2014.
- [64] S. Briot and V. Arakelian, "Optimal force generation in parallel manipulators for passing through the singular positions," *The International Journal of Robotics Research*, vol. 27, no. 8, pp. 967–983, 2008.
- [65] F. Bourbonnais, P. Bigras, and I. A. Bonev, "Minimum-time trajectory planning and control of a pick-and-place five-bar parallel robot," *IEEE/ASME Transactions on Mechatronics*, vol. 20, no. 2, pp. 740–749, 2015.
- [66] F. Aghili, "A unified approach for inverse and direct dynamics of constrained multibody systems based on linear projection operator: applications to control and simulation," *IEEE Transactions on Robotics*, vol. 21, no. 5, pp. 834–849, 2005.
- [67] R. Bordenal, J. M. Porta, and L. Ros, "A singularity-robust LQR controller for parallel robots," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2018, pp. 270–276.
- [68] R. Featherstone, *Robot dynamics algorithms*. Kluwer, Norwell, MA, 1987.