A Reachability-based planner for sequences of acyclic contacts in cluttered environments

Steve Tonneau, Nicolas Mansard, Chonhyon Park, Dinesh Manocha, Franck Multon, Julien Pettré

To cite this version:

Steve Tonneau, Nicolas Mansard, Chonhyon Park, Dinesh Manocha, Franck Multon, et al.. A Reachability-based planner for sequences of acyclic contacts in cluttered environments. International Symposium on Robotics Research (ISSR 2015), Sep 2015, Sestri Levante, Italy. lirmm-01149666

HAL Id: lirmm-01149666
https://hal-lirmm.ccsd.cnrs.fr/lirmm-01149666

Submitted on 7 May 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Reachability-based planner for sequences of acyclic contacts in cluttered environments

S. Tonneau, N. Mansard, C. Park, D. Manocha, F. Multon, J. Pettré

Abstract Multiped locomotion in cluttered environments is addressed as the problem of planning acyclic sequences of contacts, that characterize the motion. In order to overcome the inherent combinatorial difficulty of the problem, we separate it in two subproblems: first, planning a guide trajectory for the root of the robot and then, generating relevant contacts along this trajectory. This paper proposes theoretical contributions to these two subproblems. We propose a theoretical characterization of the guide trajectory, named "true feasibility", which guarantee that a guide can be mapped into the contact manifold of the robot. As opposed to previous approaches, this property makes it possible to assert the relevance of a guide trajectory without explicitly computing contact configurations. This property can be efficiently checked by a sample-based planner (e.g. we implemented a visibility PRM). Since the guide trajectories that we characterized are easily mapped to a valid sequence of contacts, we then focused on how to select a particular sequence with desirable properties, such as robustness, efficiency and naturalness, only considered for cyclic locomotion so far. Based on these novel theoretical developments, we implemented a complete acyclic contact planner and demonstrate its efficiency by producing a large variety of movements with three very different robots (humanoid, insectoid, dexterous hand) in five challenging scenarios. The planner is very efficient in quality of the produced movements and in computation time: given a computed RB-PRM, a legged figure or a dexterous hand can generate its motion in real time. This result outperforms any previous acyclic contact planner.

Steve Tonneau, Nicolas Mansard
LAAS-CNRS, Toulouse, e-mail: `surname.lastname@laas.fr`

Chonhyon Park, Dinesh Manocha
UNC, Chapel Hill e-mail: `dm/chpark@cs.unc.edu`

Franck Multon
Unversité Rennes II e-mail: `fmulton@irisa.fr`

Julien Pettré
Inria e-mail: `julien.pettre@inria.fr`

1 Introduction

1.1 State of the art

We consider the problem of planning the acyclic sequence of contacts describing the motion of a multiped robot in a cluttered environment. Acyclic contact planning is a particular class of motion planning where every configuration of the resulting trajectory must be in contact with the environment in order to support the balance of the system. The difficulty of the problem comes both in practice from the proximity to the obstacles (that tends to make the sampling of valid configuration tedious) and in theory from the foliation of the configuration space, where zero-measure manifolds intersect in a combinatorial manner [17]. Acyclic motion planning is a problem of interest in robotics, neurosciences and biomechanics. It is also interesting for virtual character animation, from where come early contributions were proposed. The first successful implementations come from local adaptation of motion graphs [11] and ad-hoc construction of locomotion controllers [15]. These approaches can intrinsically not adapt to new situations or discover complex behaviors in unforeseen contexts. In robotics, the attention of the community first focused on the generation of cyclic locomotion patterns, in particular for bipedal walking on flat terrains [10]. While planning cyclic bipedal contacts is now mature, with existing real-time solutions [3], the problem is mostly open for more generic acyclic contacts.

The issue of planning acyclic contacts was first completely described in [6], where it is proven to require the handling of two simultaneous problems: $\mathscr{P}_1$: a relevant guide trajectory for the root of the robot in $SE(3)$; and $\mathscr{P}_2$: the planning of a discrete sequence of acyclic, balanced contact configurations along the trajectory. A key issue is to avoid combinatorial explosion when considering at the same time the possible contact configurations and the potential trajectories. This seminal paper proposes a first effective algorithm, able to handle simple situations (such as climbing scenarios), but not scalable to arbitrary environments. Following it, several papers have proposed practical contributions to apply this approach in particular situations, typically limiting the combinatorial by imposing a fixed set of possible contacts ([9], [18]).

After [6], most of the papers have explored alternative formulations to handle the combinatorial issue. Two main directions have been explored. **On one hand, local optimization of both the root trajectory $\mathscr{P}_1$ and the contact positions $\mathscr{P}_2$** has been used, to trade the combinatorial of the complete problem for a differential complexity, at the cost of local convergence. A complete example of the potential offered by such approaches is given in [13], with a successful adaptation to a real robot in [12]. To keep reasonable computation times, the method uses a simplified dynamic model for the avatar. Still, the computation time is far from interactive (about 1 minute of computation for a sequence of 20 contacts). In [7] contact planning is formulated as a mixed integer problem, with computation costs similar to [13]. Aside from the computation cost, a major drawback of these optimization based approaches is that they only offer local convergence.

On the other hand, the two problems $\mathscr{P}_1$and $\mathscr{P}_2$might be decoupled to reduce the complexity. The feasibility and interest of the decoupling is shown in [8], by manually setting up a rough guide trajectory (i.e. an ad-hoc solution to $\mathscr{P}_1$). $\mathscr{P}_2$is then addressed as the combinatorial computation of a feasible contact sequence in the neighborhood of the guide. A solution can then be found, at the cost of prohibitive computation times (several hours). Furthermore, this approach is suboptimal because the quality of the motion is conditioned by the relevance of the guide trajectory, which is not evaluated *a priori*. In [4], the authors precisely focus on automatically computing a guide trajectory with guarantees of contact feasibility, by extending key frames of the trajectory into whole-body, balanced contact configurations. Randomly sampled configurations are projected into the contact submanifold using a generalized inverse kinematics solver, a computationally expensive process (about 15 minutes are required to compute a guide trajectory in the examples presented). Moreover this explicit projection is yet an insufficient condition and does not provide strong guarantees on the feasibility of the path between two key positions in the trajectory.

1.2 Paper contribution and organization

If the optimization approach is very promising, we are interested in this paper by the sample-based methodology, more able to find complex trajectories in cluttered environment. The theoretical structure of the problem is now well understood, however there is currently no scalable method to solve it. The combinatorial of the original problem (as described in [6]) is too high to have any hope of tackling it directly. Alternative formulations are necessary to obtain practical solutions. We believe that the separation, proposed in [8], between the guide trajectory and the contact sequence is the most promising direction. However, this direction opens two theoretical questions that remain to be solved, or even to be properly formulated:

- The guide trajectory must respect a property that guarantees the existence of a contact sequence to actuate it. [1] This property has not been studied yet; the only way on validating a trajectory was proposed in [8], by explicitly computing the contacts, which is computationally not reasonable.
- There is in general an infinite combination of possible contact sequences for a given root trajectory. The selection of one particular contact sequence with interesting properties (minimum number of contact change, robustness, efficiency or naturalness) has been studied for cyclic cases, as in [9], but has not been efficiently applied to cluttered environments (e.g. [5] and [8] mostly randomly pick one contact sequence corresponding to the guide, obtaining possibly very tedious contact sequences).

---

[1] This property is related to the controllability of the root actuated by the contact forces, but for discrete bounded properties.

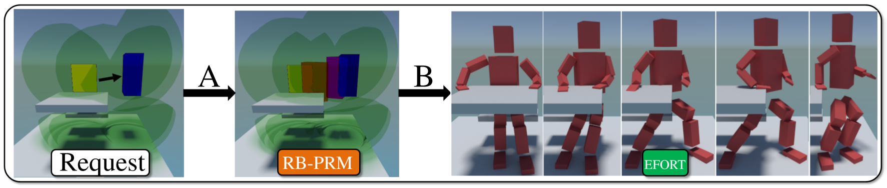

Fig. 1 Overview of our 2-stage framework. (A) Given a path request between the yellow and blue positions, a guide trajectory is computed in C_{reach} using RB-PRM. (B) The trajectory is extended into a discrete sequence of contact configurations using EFORT.

We claim that the desirable contact properties of a guide trajectory, proposed in [4], can be formulated in a space of lower dimension, which we call C_{reach}. This formulation can make the planning of a guide trajectory more efficient computationally, while providing equivalent guarantees than planning directly in the configuration space. Among the particular properties obtained when planning in C_{reach}, we would like to guarantee that any reduced trajectory can actually lead to a feasible sequence of contacts: in which case, we say that the reduced trajectory is truly feasible. It is possible in theory to guarantee that any reduced trajectory is truly feasible, even if it is more efficient in practice to approximate this property. A second positive side effect of the true-feasibility of the guide trajectory is that we can now focus on the selection of one particular sequence of contacts, for example one that minimizes the number of contacts in the sequence or maximizes the robot efficiency or style.

Based on these fundamental observations, we have implemented a very efficient acyclic contact planner. Our method is based on a probabilistic roadmap (PRM), that computes offline guide trajectories that are approximately truly feasible. The planner then resolves online the contact sequence by refining a guide trajectory computed from the PRM. Our planner was able to compute physically-consistent contact sequences for very complex systems (a humanoid, 28 joints; and an insectoid, 48 joints) in a few seconds for classical scenarios like climbing, and less than a minute for very complex problems like egress of a damaged truck. The contact planner also generalizes to planning dexterous manipulation movements, as demonstrated by preliminary results.

The contributions of the paper are twofold. We propose the first theoretical characterization of today's most efficient practical approach to sampled-based planning of acyclic contacts. And based on this characterization, we propose a very efficient and general implementation of a acyclic contact planner, that outperforms all the previously demonstrated planners.

We propose a framework to address the motion planning problem for multiped robots in cluttered environments: given a start and a goal configuration, the objective is to compute a sequence of contact configurations allowing to achieve the motion. For instance, we can consider the task of standing up, illustrated in Figure 1–right. The problem is decoupled into two sequential phases: 1) the computation of a guide trajectory for the root of the robot; 2) the computation of a discrete sequence of contact configurations allowing to achieve the motion along the trajectory. The remain-

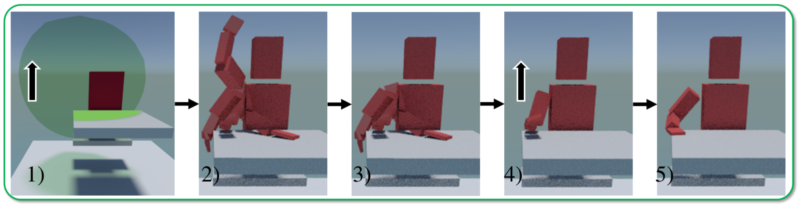

Fig. 2 Generation of a contact configuration for the right arm of a humanoid robot. 1) Selection of reachable obstacles. 2) A request is performed on a database of configurations. 3) Configurations too far from contact are eliminated. 4) The best candidate according to EFORT is chosen. 5) The final contact is achieved using inverse kinematics.

ing of this section presents the general organization of our method in Section 2. The two following sections 2 and 3 present respectively our answer to problems $\mathscr{P}_1$and $\mathscr{P}_2$. Finally, we propose a very complete experimental validation of the planner with three very different kinematic chains (humanoid, insectoid and three-finger manipulator) in various scenarios.

1.3 Computation of a guide trajectory

We first consider the problem of planning a relevant guide trajectory. The objective is to compute a trajectory of root placements which will allow contact creation. The stake of this first part is to preserve the completeness of the planner: the planner should be able to explore any possible guide trajectory, but at the same time, any computed guide trajectory must be truly feasible, i.e. must lead to a valid sequence of contacts.

An intuitive description of such placements is to see them as ”close, but not too close”: close, because a contact surface must be partially included in the range of motion of the robot (represented for the right arm in Fig. 2–1); not too close, because the robot must avoid collision (which is represented by the hull including the torso in Fig. 2–1). We define formally C_{reach}, the set of interesting root placements, in which we compute a guide trajectory with a sampling based planner, the reachability PRM (RB-PRM), is used (Figure 1–A). Planning in C_{reach} boils down to planning in SE(3), which has an acceptable practical complexity. Details are presented in Section 2.

1.4 Generating a discrete sequence of contact configurations

The second stage is to extend the guide trajectory into a sequence of contact configurations (Fig. 1–B). Thanks to the nice property that we manage to obtain for the guide, obtaining a random sequence of contact is an easy problem. The stake

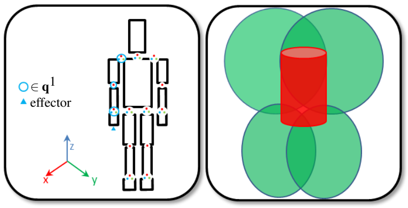

Fig. 3 Left: Robot in a rest configuration. The right arm is denoted as the limb R^1. Each colored dot represents a degree of freedom around an axis. Right: Volumes of the robot. The red geometry denotes W^0 and must remain collision-free. The green spheres are the W^k.

of this part is to select an efficient sequence, in particular by reducing the number of contact in the sequence. To create contacts in an efficient manner, we consider each limb as a manipulator attached to the root, and select the most relevant contact from a database of precomputed configurations (Figure 2). The relevance is defined as the contribution to the quasi static balance of the robot, and as the contribution to the motion of the root. This task efficiency is measured based on the Extended FORce Transmission ratio (EFORT), that we proposed in [19]. Details are presented in Section 3.

1.5 Notation conventions and definitions

A vector $\mathbf{x}$ is denoted with a bold lower case letter. A matrix $\mathbf{A}$ is denoted with a bold upper case letter. A set C is denoted with a upper case italic letter. Scalar variables and functions are denoted with lower case italic letters, such as r or $f(\mathbf{x})$.

A robot is a kinematic chain R, comprising $n+6$ degrees of freedom (DOFs). R is composed of l limbs $R^k, 1 \leq k \leq l$, attached to a root. It is described by a configuration $\mathbf{q} \in SE(3) \times \mathbb{R}^n$. We define some relevant projections of $\mathbf{q}$:

- $\mathbf{q}^k$ denotes the configuration (a vector of joint values) of the limb R^k (Fig. 3);
- $\mathbf{q}^{\bar{k}}$ denotes the vector of joint values of R **not** related to R^k. We define for convenience $\mathbf{q} = \mathbf{q}^k \oplus \mathbf{q}^{\bar{k}}$;
- $\mathbf{q}^0 \in SE(3)$ denotes the position and orientation of the root of the robot R.

The environment O is defined as the union of the obstacles O_i it contains. The volume encompassing the trunk of the robot is denoted W^0 (Fig. 3-right: central cylinder). The range of motion of a limb R^k is denoted W^k (Fig. 3-right: the four ellipses).

$$W^k = \left\{ \mathbf{x} \in \mathbb{R}^3 : \exists \mathbf{q}^k, \mathbf{p}^k(\mathbf{q}^k) = \mathbf{x} \right\} \tag{1}$$

where $\mathbf{p}^k$ denotes the end-effector position of R^k (translation only) for $\mathbf{q}^0$ the null displacement. We also define $W = \bigcup_{k=1}^{l} W^k$. Finally, we define $W^k(\mathbf{q}^0), 0 < k \leq l$ as the volume W^k translated and rotated by the rigid displacement $\mathbf{q}^0$.

2 Computation of a guide trajectory in C_{reach} (Stage 1)

We consider the problem of computing a guide trajectory $\mathbf{q}^0(t) : [0,1] \longrightarrow SE(3)$ for the geometrical root of a multiped robot, connecting start and goal configurations. As said in previous section, the stake is to enforce that any configuration $\mathbf{q}^0$ of the guide is truly feasible, i.e. can be mapped into a balanced configuration in contact. We denote by $C_{contact} \subset SE(3) \times \mathbb{R}^n$ the contact submanifold of the robot. We say that a root placement $\mathbf{q}^0$ is *truly feasible* if there exists $\mathbf{q}^{\overline{0}}$ such that

$$\mathbf{q}^0 \oplus \mathbf{q}^{\overline{0}} \in C_{contact}$$

The set of all truly-feasible root placements is denoted by C_{reach}. By extension, a trajectory $\mathbf{q}^0(t)$ is truly feasible if $\forall t \in [0,1], \mathbf{q}^0(t) \in C_{reach}$. We denote by π a function that associates to any truly-feasible $\mathbf{q}^0$ a configuration of $C_{contact}$ (π might not be uniquely defined: it is the reciprocal of a projector).

$$\pi : \quad \mathbf{q}^0 \in SE(3) \longrightarrow \mathbf{q}^0 \oplus \mathbf{q}^{\overline{0}} \longrightarrow C_{contact} \tag{2}$$

For a two-stage acyclic contact planner to be exact and complete, we need the combination of two conditions on a guide trajectory generator: all the generated trajectories must be *truly feasible* (sufficient condition); the guide planner must be complete, i.e. it must be able to discover any truly feasible trajectory (necessary condition)[2].

2.1 Conditions for true feasibility

By default, the true feasibility implies a constructive demonstration by exhibiting a valid $\mathbf{q}^{\overline{0}}$. This is the approach chosen by [4]. However, computing a valid $\mathbf{q}^{\overline{0}}$ is costly in practice. In this section we rather define a necessary condition and a sufficient condition for true feasibility that do not require this explicit computation.

True feasibility: necessary condition

For a contact to be possible, a volume $O_i \in O$ necessarily intersects with the range of motion $W(\mathbf{q}^0)$ (Fig. 2–1). Furthermore, if $\mathbf{q}^0$ is truly feasible, then the trunk of the robot $W^0(\mathbf{q}^0)$ is necessarily not colliding with the environment O.

Therefore we can approximate C_{reach} with a set $C_{reach} \subset C^1_{reach}$ with the reachability condition defined as:

$$C^1_{reach} = \{\mathbf{q}^0 : W(\mathbf{q}^0) \cap O \neq \emptyset \wedge W^0(\mathbf{q}^0) \cap O = \emptyset\} \tag{3}$$

---

[2] The proof is immediate, using a contradiction approach.

It is straighforward to prove that $C_{reach} \subset C^1_{reach}$ (by construction of the included set). This inclusion is very important: it directly implies that any motion-planning algorithm with a guaranty of completeness in $SE(3) \times \mathbb{R}^n$ is complete in $C_{reach} \times \mathbb{R}^n$. This is a strong reduction of the search space, which can be directly applied to any existing method.

The condition C^1_{reach} is only necessary which means that one such root placement might not be truly feasible: in practice it is not guaranteed to find a valid sequence of contacts for every guide trajectory in C^1_{reach}.

Sufficient condition for true feasibility

In building C^1_{reach}, we directly considered the including hull of the root body W^0, obtaining a necessary condition. A trivial sufficient condition for true feasibility can be constructed as a variation of it, by replacing W^0 with a bounding volume B^{max} encompassing the whole robot in a given pose, except for the effector surfaces to be in contact[3]. We denote by C^∞_{reach} the set of root placement corresponding to the sufficient condition. Of course $C^\infty_{reach} \subset C_{reach}$.

In general, the inclusion is strict, which means that we lose the completeness of the two-stage contact planner (i.e. the planner is not able to discover a trajectory for which the root remains inside $C_{reach} \setminus C^\infty_{reach}$). However, the sufficient condition is interesting since we can guarantee that any such trajectories lead to a valid sequence of contact.

2.2 *Reachability: a compromise condition*

The sufficient condition is not interesting in practice since it leads the solver to lose too many interesting trajectories. However, the necessary condition is not perfect neither, since the first stage of the planner would stop on a necessary guide that is not truly feasible in practice. It might be possible to find a shape B that is necessary and sufficient; however, it seems intuitively very unlikely in general (we were not able to find a counter-example). The construction of a shape $W^0 \subset B \subset B^{max}$ leading to necessary and sufficient condition (or the proof of its inexistence) is out of the scope of this work.

However between W^0 and B^{max}, a trade-off can be found between a necessary and a sufficient condition. We define W^0_s as the volume W^0 subject to a scaling transformation by a factor $s \in \mathbb{R}^+$. We then consider the spaces C^s_{reach}

---

[3] This condition is trivial in the sense that the resulting W has a zero measure. Consequently the likelihood of randomly shooting a configuration with W in contact is null. For the need of the proof, the trivial sufficient condition is enough. In practice,the construction of a non-trivial including shape W^0 was possible for all the robot structures we considered, even if we cannot yet systematize the construction procedure.

$$C^{s}_{reach} = \{\mathbf{q}^0 : W(\mathbf{q}^0) \cap O \neq \emptyset \wedge W^0_s(\mathbf{q}^0) \cap O = \emptyset\} \qquad (4)$$

The higher s is, the closer the reachability condition is to being sufficient. When $s = 1$, the reachability condition is equivalent to the necessary condition. In practice, it is easy to adjust s to keep most of the interesting guides without introducing incorrect guides. The chosen s can then be statistically validated, as discussed in Section 4.

2.3 Computing the guide trajectory in C^{s}_{reach} with RB-PRM.

Once a value of s has been fixed, any sampling-based motion planner can be used to plan a trajectory in C^{s}_{reach}. We have chosen to implement RB-PRM as a variation of Visibility-PRM [14], which usually leads to a smaller set of nodes than classical PRM planners. The associated drawback is that the paths returned by the planner might not be the shortest ones, which is typically not an issue in highly cluttered environments.

To sample more efficiently configurations of C^{s}_{reach}, we bias the sampling process to generate near obstacles configurations, similarly to [1] to generate configurations in narrow passages. First, a configuration is randomly set to a random point on the surface of one obstacle. The configurations are then translated and rotated randomly until the reachability condition is satisfied.

3 From a guide trajectory to a discrete sequence of contact configurations (Stage 2)

As an input of this stage, we consider a truly feasible guide trajectory $\mathbf{q}^0(t) : [0,1] \longrightarrow SE(3)$ for the root of the robot R. We now consider the second problem of computing a trajectory $\mathbf{q}^{\overline{0}}(t)$ of the limb of the robot. From the true feasibility, we know that such a trajectory exists. On the opposite to previous works [8, 4], the stake here is not to find such a trajectory but rather to select one with particular properties. In particular, we show here how to build a contact sequence with a small number of contact variations and good-efficiency and naturalness of the postures.

More precisely, any projector reciprocal π introduced in (2) can be used to expand $\mathbf{q}^0(t)$ into a whole-body trajectory. We propose here a particular construction of π leading to interesting contact sequences.

3.1 Extension of the guide trajectory

The guide trajectory $\mathbf{q}^0(t)$ is first discretized into a sequence of j key placements:

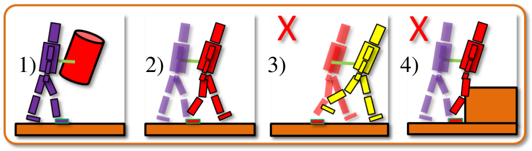

Fig. 4 Contacts are maintained unless their position is too far, or the environment prevents it.

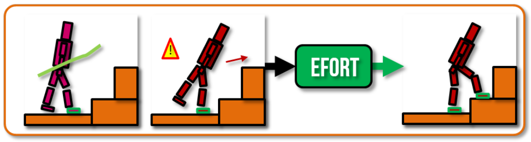

Fig. 5 Contacts are generated when the configuration is not balanced.

$$\mathbf{Q}^0 = [\mathbf{q}_0^0; \mathbf{q}_i^0; \ldots, \mathbf{q}_{j-1}^0]$$

where $\mathbf{q}_0^0$ and $\mathbf{q}_{j-1}^0$ respectively correspond to the start and goal configurations. To ensure continuity in the contact transition phases, we rewrite π under the following recursive form for any $0 < i < j$:

$$\pi : \begin{cases} \mathbf{Q}^0 \longrightarrow C_{contact} \\ \mathbf{q}_i^0 \longrightarrow \mathbf{q}_i^0 \oplus g(\mathbf{q}_{i-1}, \mathbf{q}_i^0) \end{cases}$$

We initialize the recurrence with $\pi(\mathbf{q}_0^0) = \mathbf{q}_0$ the initial configuration of the robot. The function g is defined independently by g^k for each limb R^k. In defining g^k, two aspects must be considered. Is the limb R^k in contact? And which criteria is it optimizing?

Maintaining a contact: If possible, a limb in contact at time $i-1$ remains in contact at i. The contact is broken if a direct projection of the current limb configuration to the nearest contact fails (due to joint limits or collision). Independently of whether R^k is in contact or not, the limb configurations are selected to be the closest to the previous configuration (by minimizing $\|\mathbf{p}^k(\mathbf{q_i}) - \mathbf{p}^k(\mathbf{q_{i-1}})\|$).

Once a first candidate configuration is taken for all limbs, the quasi-static balance is tested by checking that the weight wrench is in the contact friction cone, using the geometric approach described in [16]. If the balance is not obtained, new contact are randomly generated using the following procedure.

Creating a contact: Consider a configuration where some limbs are in contact, some are free and the quasi-static balance is not enforced. A new contact configuration is obtained by projecting a random configuration to the nearest contacts. The sampling is biased by selecting configurations so that the limbs previously in contact are already close to contact and at least one additional limb is close to the environment. The random configuration is then projected onto the contact by iteratively minimizing the distance between the end effectors and the environment, under the constraint of respecting the joint limits and avoiding the collision.

3.2 Contribution to the global movement: the EFORT criteria

EFORT criterion: If only relying on the random sampling to select new contacts, the planner would produce inefficient postures. The resulting contact sequence is then poorly efficient and is very unnatural. Moreover, the limbs are not efficiently configured and are not able to efficiently follow the general movement: contacts break very frequently.

When creating additional contacts, we therefore propose to select particular configurations that allow to exert a force compatible with the direction of motion. This task efficiency is measured with the Extended FORce Transmission ratio (EFORT) [19]. The measure of EFORT is given by

$$\alpha_{EFORT}(\mathbf{q}^k, \mathbf{m}) = [\mathbf{m}^T(\mathbf{J}\mathbf{J}^T)\mathbf{m}]^{-\frac{1}{2}}(\nu_0 \mathbf{n}^T \mathbf{m}) \tag{5}$$

where $\mathbf{J}$ is the Jacobian matrix of the limb R^k in configuration $\mathbf{q}^k$; ν_0 is the friction coefficient of the contact surface; $\mathbf{n}$ is the normal of the contact surface; and $\mathbf{m} = -(\mathbf{q}^0_{i+1} - \mathbf{q}^0_i)$ is the direction opposite to the motion. The first part of the equation optimizes the ratio between the joint torque and the resulting force applied along $\mathbf{m}$. The second part minimizes the odds of slipping while applying a force in the direction $\mathbf{m}$.

Optimization at creation: In practice, a database of configurations is stored for each limb, which can be considered as manipulator arm. The database is implemented as an octree data structure, indexed by the end-effector positions of the configurations (and additionally storing $\mathbf{J}$). Upon request, the octree returns a set of configurations which are close to contact (Fig. 2-3). These candidates are sorted based on their task efficiency, given by α_{EFORT}. The first candidate in this list which satisfies the balance criterion and is collision free is selected and is projected on the contact surface using an inverse kinematics solver.

4 Results

The main strength of our planner is that it efficiently works for arbitrary robot shapes. We first validate this aspect by producing a large variety of movements with three very different robots (humanoid, insectoid, dexterous hand) in five challenging scenarios. Two evaluations of the method are provided: qualitatively, by displaying the naturalness of the contact sequence in the companion video; and quantitatively by statistically measuring the validity of the compromise condition (Sec. 4.2) and the performances of the algorithms (Sec. 4.3).

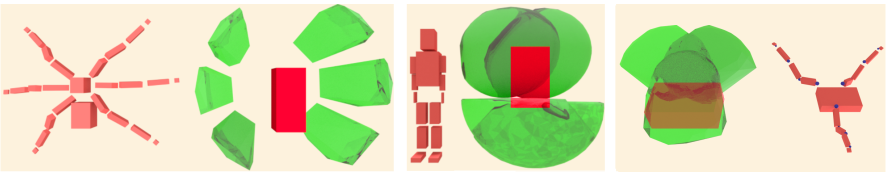

Fig. 6 Robots and associated volumes: in red W_s^0; in green the range of motion of each limb.

4.1 Robot models and scenarios

Fig. 6 describes the robot used in the experiments. The humanoid robot has four limbs, each with 7 DOFs. It has a total of 34 DOFs. The insectoid robot has six limbs, each with 8 DOFs, and a total of 54 DOFs. The hand has three fingers, each with 6 DOFs and a total of 24 DOFs.

In all the scenarios considered, the formulation of the problem is always the same: a start and goal configuration are provided as an input of the scenario, and the framework outputs a sequence of statically balanced contact configurations connecting the start and goal configurations. A companion video available at `http://youtu.be/LmLAHgGQJGA` (anonymous link) displays the complete contact sequence obtained in all these scenarios. The video only renders the contact configuration (not the interpolation between contacts, which is out of the scope of the paper).

Truck egress (humanoid and insectoid): The robot must leave a truck the doors of which are blocked: it has to crawl through the front window. Figure 7 presents the sequence of contacts obtained for both robots: RB-PRM can find solutions in highly cluttered environments with narrow passages.

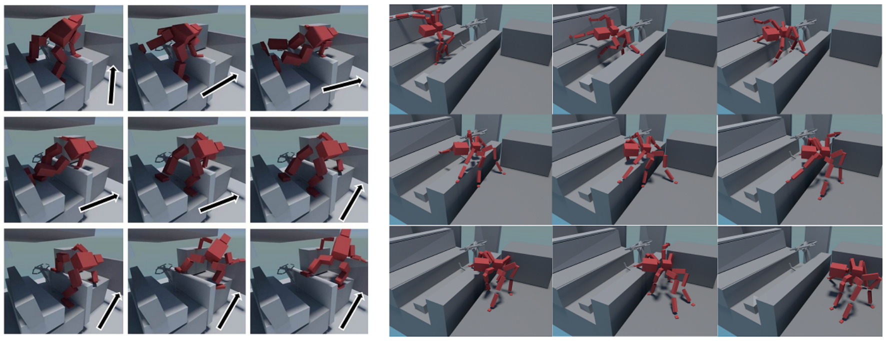

Fig. 7 The computed contact sequences for the truck egress scenario. Only selected postures are shown for the insect.

Climbing (humanoid and insectoid): The robot has to climb on a wall with several grasps disposed along it. In this scenario, we give stronger conditions for the sampled root placements: We require that more than one range of motion W^k collide with obstacles of the environment. Fig. 8 presents the contact sequence obtained for both robots. In this case, the generation of the contact for the insect failed (due to the compromise criterion C^s_{reach}).

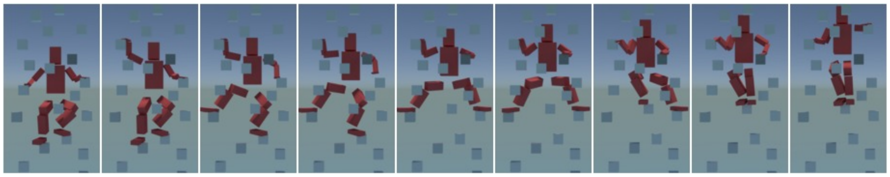

Fig. 8 The computed sequence for the climbing scenario.

Manipulation of a pen (3-finger hand): This scenario is proposed to illustrate the genericity of our approach: we consider a manipulation task for a robotic hand and use our contact planner to compute a guide trajectory for the fingers, considered as effectors (Figure 9). Although we do not address the hard issue of accounting for rolling motions, the planner is able to compute the shown sequences, this in less than 5 seconds.

Other scenarios (humanoids): The standing-up scenario (already presented in Fig. 1) is a setup taken from [8]: it corresponds to a long narrow passage in the configuration space. In the crouching scenario, demonstrated in the companion video, the character automatically goes from a standing to a crouching position to crawl under an obstacle.

4.2 Parametrization of the reachability condition

To find the appropriate C^s_{reach} in which to sample the guide trajectory, we computed the rejection rate for various values of s for each robot in the most cluttered truck scenario. For a given value of s, 10^6 root positions and orientations are computed in C^s_{reach}. In each case we try to generate a collision free contact configuration, with a database comprising $N = 10^5$ sample configurations for each limb. The rejection rate is the ratio of between the number of failures and the number of trials. We can see that this rate decreases as s grows, as shown in Figure 10. For the humanoid, we chose $s = 2.2$, and for the insect, $s = 2.8$.

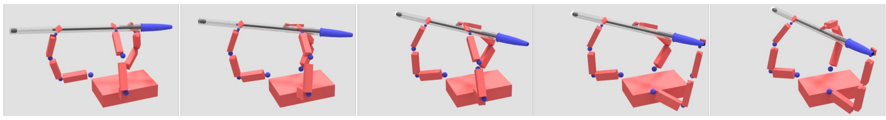

Fig. 9 Contact sequence found for a pen manipulation in a zero gravity environment.

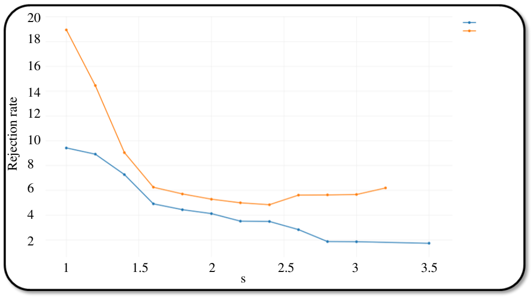

Fig. 10 Rejection rates for the humanoid (orange) and the insectoid (blue) depending on *s* for the truck scenario.

4.3 Performance

The number of samples used for generating the contacts of each limb is 10000. Table 1 presents the average time (in seconds) spent in the different phases of the planner, for each phase and each scenario, as well as the number of contact phases of the sequence.

We observe that many contacts are required for the insect, which can be explained by its restricted range of motion. The time spent generating the navigation graph can be approximated to one minute. The time spent in generating the graph of the climbing scenario, despite the relatively open environment, is explained by the additional restrictions imposed on the reachability condition. The difficulty to find a balanced configuration essentially influences the time spend generating the contacts.

The number of contact in the sequence gives a rough estimation of its duration in seconds. Except for the humanoids crawling out of the trucks, all the contact gener-

	Generate RB-PRM (offline)	Generating the contact sequence	Number of contact states
Truck egress (humanoid)	73	15	10
Truck egress (insectoid)	70	23	48
Climbing (humanoid)	25	5	15
Climbing (insectoid)	21	27	51
Crouching (humanoid)	5	6	22

Table 1 **Average time** (in seconds) spent in RB-PRM generation, and the online generation of the contact sequence.

ation are real-time: given a computed RB-PRM, a legged figure or a dexterous hand can generate its motion in real time. Aside from the quality of the generated trajectories observed in the video, this computation time is a major practical achievement.

5 Discussion and future work

In this paper we consider acyclic contact planning in cluttered environments, formulated as two sub problems that we address sequentially: $\mathscr{P}_1$: the computation of a guide trajectory for the root of the robot that can be extended ; $\mathscr{P}_2$: the computation of a discrete sequence of contacts along this trajectory. Our contribution to $\mathscr{P}_1$is a generic characterization of the properties that the guide trajectory must respect, in particular to enforce the completeness of the acyclic contact planner. We introduced a low dimensional space C_{reach} which can be mapped into the contact submanifold of the robot, approximated and efficiently sampled by our Reachability-Based planner. Our contribution to $\mathscr{P}_2$is an pragmatic contact generation scheme which can take into account criteria to enforce interesting properties on the generated contacts (such as robustness, energy efficiency or naturalness). One such criterion, EFORT, is used to demonstrate the method and optimizes a force exertion compatible with the direction of motion.

Aside from the theoretical contributions, we demonstrated that the results of our method outperform the previous approaches, allowing a very interesting compromise on three criteria which are hard to conciliate: genericity, performance, and quality of the solution. **Regarding genericity**, the reachability condition, coupled with an approach based on limb decomposition, allows the method to address arbitrary multiped robots. The only pre-requisite is the specification of the volumes W^0 which can be rigorously adjusted from a statistical analysis. **Regarding performance**, our framework outperforms existing methods in addressing either $\mathscr{P}_1$or $\mathscr{P}_2$, leading to computation cost close to real-time. **Regarding the quality of the trajectories**, a parametrization of the reachability condition allows us to compute relevant trajectories in all the scenarios presented, with low rejection rates. As for [4], failures can still occur, due to the compromise criterion used in computing the guide trajectory.

Future work will focus on a more accurate formulation of C_{reach} to address this limitation. Additionally, we will work on theoretically characterizing the benefits of using the EFORT criterion when generating the sequence (our understanding is yet only intuitive), and quantifying the reward it term of energetic efficiency or naturalness. Generating the complete motions, by interpolating between the computed contact sequences, is our logic next objective. This is a mandatory steps before going on the real robot. A reasonable approach is to search the optimal trajectory connecting two postures of the sequence, under the constraint of maintaining stability. In particular, we will consider optimization schemes based on reduced dynamic models, possibly extending the preliminary work of [2].

References

1. Amato, N.M., Burchan, O., Lucia, B., Dale, K., Jones, C., Vallejo, D.: Choosing good distance metrics and local planners for probabilistic roadmap methods. In: Proc. IEEE Int. Conf. Robot. Autom. (ICRA) (1998)
2. Audren, H., Vaillant, J., Kheddar, A., Escande, A., 1, K.K., Yoshida, E.: Model preview control in multi-contact motion–application to a humanoid robot. In: IEEE/RSJ Int. Conf. on Int. Robots and Syst. (IROS) (2014)
3. Baudouin, L., Perrin, N., Moulard, T., Lamiraux, F., Stasse, O., Yoshida, E.: Real-time replanning using 3D environment for humanoid robot. In: IEEE-RAS Int. Conf. on Humanoid Robots (Humanoid'11). Bled, Slovenia (2011)
4. Bouyarmane, K., Escande, a., Lamiraux, F., Kheddar, a.: Potential field guide for humanoid multicontacts acyclic motion planning. 2009 IEEE Int. Conf. on Robot. and Auto pp. 1165–1170 (2009). DOI 10.1109/ROBOT.2009.5152353
5. Bouyarmane, K., Kheddar, A.: Multi-Contact Stances Planning for Multiple Agents. In: ICRA'11: Int. Conf. on Robot. and Auto. Shanghai Int. Conf. Center, Shanghai, Chine (2011)
6. Bretl, T., Rock, S., Latombe, J.C., Kennedy, B., Aghazarian, H.: Free-climbing with a multi-use robot. In: M.H.A. Jr., O. Khatib (eds.) ISER, *Springer Tracts in Advanced Robot.*, vol. 21, pp. 449–458. Springer (2004)
7. Deits, R., Tedrake, R.: Footstep planning on uneven terrain with mixed-integer convex optimization. In: 14th IEEE-RAS Int. Conf. on Humanoid Robots, Humanoids 2014, Madrid, Spain, November 18-20, 2014 (2014)
8. Escande, A., Kheddar, A., Miossec, S., Garsault, S.: Planning Support Contact-Points for Acyclic Motions and Experiments on HRP-2. In: O. Khatib, V. Kumar, G.J. Pappas (eds.) ISER, *Springer Tracts in Advanced Robot.*, vol. 54, pp. 293–302. Springer (2008)
9. Hauser, K., Bretl, T., Harada, K., Latombe, J.C.: Using motion primitives in probabilistic sample-based planning for humanoid robots. In: S. Akella, N.M. Amato, W.H. Huang, B. Mishra (eds.) WAFR, *Springer Tracts in Advanced Robot.*, vol. 47, pp. 507–522. Springer (2006)
10. Kajita, S., Kanehiro, F., Kaneko, K., Fujiwara, K., Harada, K., Yokoi, K., Hirukawa, H.: Biped Walking Pattern Generation by using Preview Control of Zero-Moment Point. In: IEEE Int. Conf. Robot. and Auto (ICRA). Taipei, Taiwan (2003)
11. Kovar, L., Gleicher, M., Pighin, F.: Motion graphs. In: ACM Trans. on Graphics, vol. 21. ACM, New York, NY, USA (2002)
12. Mordatch, I., Lowrey, K., Todorov, E.: Ensemble-CIO: Full-Body Dynamic Motion Planning that Transfers to Physical Humanoids (2015)
13. Mordatch, I., Todorov, E., Popović, Z.: Discovery of complex behaviors through contact-invariant optimization. ACM Trans. on Graphics **31**(4) (2012)
14. Nissoux, C., Siméon, T., Laumond, J.: Visibility based probabilistic roadmaps. In: Proceedings 1999 IEEE/RSJ Int. Conf. on Int. Robots and Syst. Hum. and Environment Friendly Robots with High Int. and Emot. Quot., October 17-21,1999, Hyundai Hotel, Kyongju, Korea (1999)
15. Pettré, J., Laumond, J.P., Siméon, T.: A 2-stages locomotion planner for digital actors. In: Proceedings of the 2003 ACM SIGGRAPH/Eurographics symposium on Computer animation, SCA '03. Eurographics Association (2003)
16. Qiu, Z., Escande, A., Micaelli, A., Robert, T.: Human motions analysis and simulation based on a general criterion of stability. In: Int. Symposium on Digital Human Modeling (2011)
17. Siméon, T., Laumond, J., Cortes, J., Sahbani, A.: Manipulation planning with probabilistic roadmaps. Int. Journal of Robot. Research **23**(7-8) (2004)
18. Stilman, M.: Global Manipulation Planning in Robot Joint Space With Task Constraints. IEEE Trans. on Robot. **26**(3) (2010)
19. Tonneau, S., Pettré, J., Multon, F.: Using task efficient contact configurations to animate creatures in arbitrary environments. Computers & Graphics **45**(0) (2014)