On the Connectivity of Motion Spaces for Biologically-Inspired Legged Robots

Kris Hauser
School of Informatics and Computing
Indiana University
Bloomington, IN 47405, USA
hauserk@indiana.edu

Abstract— **This paper discusses recent analyses of the shape of motion spaces for legged robots in quasistatic equilibrium at a given set of contacts. Two robots were examined: HRP-2, a biologically inspired humanoid designed by AIST Japan, and ATHLETE, a six-legged, non-biologically inspired robot designed by NASA. Probabilistic roadmap techniques are used to compute the approximate connectivity of several hundred motion spaces. Results suggest that motion spaces for HRP-2 are connected, while those of ATHLETE are frequently disconnected. This signifies that when planning long-range locomotion, HRP-2 only needs to reason about its footsteps, while ATHLETE must reason about its pose as well. This makes planning much more difficult for ATHLETE, a claim that is supported by experience both in human teleoperation and in motion planning algorithms. If similar findings hold for other biological mechanisms, then it can be argued that evolution has selected for mechanisms that minimize the cognitive load needed for locomotion, which could form the foundation of a strong argument for biologically-inspired design.**

I. Introduction

The motion of a robot or biological system can be represented as a trajectory in a continuous motion space, whose shape is determined by certain operational constraints (e.g. maintaining balance, avoiding collision). The geometric complexity of a system's motion space is directly related to the amount of information needed for it to locomote. For example, a fish in open water or an insect flying in air have simple, relatively unconstrained motion spaces that permit free movement in any direction. By contrast, a creature with a complex motion space must use more sophisticated motion strategies, which necessitates more sophisticated development in cognition, perception, and reflexes. Likewise, robots with complex motion spaces require more sophisticated planning, sensing, and control strategies.

This paper hypothesizes that *natural organisms have simple motion spaces*. Evolution may select for organism morphologies that induce simple motion spaces, because they will need fewer cognitive resources. Furthermore, such spaces may also be found in robots inspired by nature, which would provide a strong argument for biologically-inspired robot design.

This paper specifically considers the connectivity of motion spaces encountered by legged robots. Singly-connected spaces are easier to reason about than disconnected spaces. If a space is singly connected, then the robot only needs to test if the endpoint is feasible: if so, then some path is guaranteed to exist. If a motion space is multiply connected, the robot must also test whether a feasible path exists.

(a) (b)

Fig. 1. The (a) HRP-2 humanoid robot and the (a) ATHLETE six-legged lunar robot investigated in this study.

This work addresses two robot platforms, the biologically-inspired humanoid HRP-2, and the non-biologically-inspired six-legged robot ATHLETE (Figure 1). Prior experience suggested that it is more difficult to operate ATHLETE in unstructured environments than HRP-2, either under human teleoperation or with the use of motion planning algorithms. One possible explanation is that ATHLETE encounters disconnected motion spaces more often than HRP-2. To test this hypothesis, we study the motion space of each robot when its contacts against the terrain are fixed. The motion space is defined as the set of configurations such that the robot does not collide with itself or the environment, and can stand with quasistatic equilibrium.

Fixed-contact motion spaces are significant for legged robots in the following way. If they are singly-connected, then the robot can plan multi-step locomotion by finding a sequence of feasible foot placements. This requires advance reasoning in the 3D space of foot placements and orientations. On the other hand, if the spaces are multiply-connected, then it must reason about foot placements *and* configurations. If the robot makes an incorrect pose when taking step A, then it may be unable to take step B, while a different pose would have enabled taking step B. This pitfall can be extrapolated so that the required lookahead is arbitrarily large. Thus, the presence of multiply-connected spaces necessitates reasoning in a space with 3+d dimensions, where d is the number of degrees-of-freedom of the system.

The experiments in this paper use probabilistic roadmap techniques to approximate motion space connectivity. Experiments over a wide range of terrains suggest that ATHLETE may encounter disconnected spaces often, even on flat terrain. All tested motion spaces for HRP-2 are connected. In future work we hope to perform further experiments on a larger set of robots, to examine whether the hypothesis holds for other biological and biologically-inspired systems. We also hope to better understand the properties of mechanisms that affect connectivity, in order to assist in robot design.

II. Legged Robot Platforms

HRP-2 (Humanoid Robotics Platform), developed by AIST Japan, is a 1.5 m tall, 58 kg bipedal humanoid robot [10]. It has six joints in each of its arms and legs, two joints in its waist, and two joints in its neck. Its human-like form was designed for use in human environments like homes, offices, and construction sites.

ATHLETE (All-Terrain Hex-Limbed Extra-Terrestrial Explorer) is a large, 850 kg, six-legged lunar vehicle developed by NASA's Jet Propulsion Laboratory [12]. Each of the six legs is nearly 2 m long and has six revolute joints, ending in a wheel. On irregular or steep terrain, the wheels may be braked so it can walk like a standard legged robot. Its legs were designed primarily for mechanical considerations, in particular, to fold compactly into the chassis and to have conveniently computed inverse kinematic solutions.

In prior work, we developed a motion planner for HRP-2 that worked well in practice for a wide variety of terrains, including rough terrain, ladders, and large stair steps [5]. This planner implicitly assumes that fixed-contact motion spaces are likely to be connected. To our surprise, the same planner failed quite often for ATHLETE, even on flat terrain. We suspected disconnected motion spaces to be the culprit, which led us to develop more sophisticated and far more reliable planning techniques that account for disconnected spaces [8].

Anecdotal reports from field trials also suggest that it is difficult for humans to teleoperate ATHLETE to walk over even slightly uneven terrain. The non-human form of ATHLETE may be nonintuitive to operate, but we also suspect that disconnected spaces significantly contribute to the difficulty. If humans are used to operating in connected spaces, then they will find reasoning about ATHLETE's motion to be very challenging.

III. Motion Spaces

The motion of a robot is a trajectory in the high dimensional configuration space $\mathcal{C}$. A configuration q is described by 6 parameters for the position and orientation of the chassis, and a parameter for each joint angle (resulting in 34 parameters for HRP-2 and 42 for ATHLETE). We will study the set of configurations at which the robot can stand stably, and support itself against gravity at fixed footfalls.

A. Motion Spaces at a Fixed Stance

Let σ be a fixed set of footfalls, which we call a *stance*. Here a footfall refers simultaneously to the foot that is placed on the terrain as well as the location at which it is placed (and may also refer to hand contacts for HRP-2).

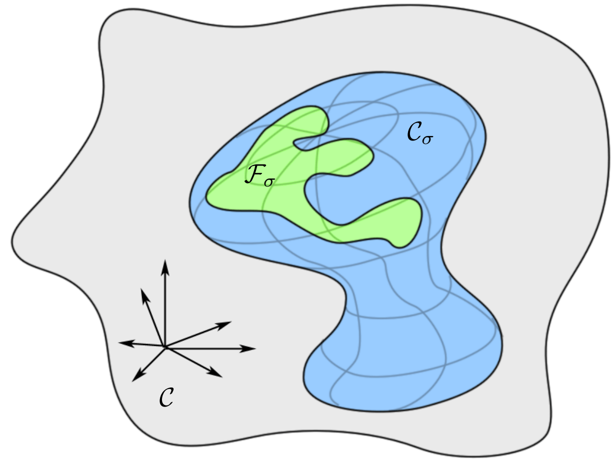

Fig. 2. At a mode σ, motion is constrained to a subset $\mathcal{F}_\sigma$ of a submanifold $\mathcal{C}_\sigma$ with lower dimension than $\mathcal{C}$.

The feasible space $\mathcal{F}_\sigma \subset \mathcal{C}$ is defined as the set of configurations that satisfy the following feasibility conditions, which will be described in more detail in Section IV-C:

1) *Contact.* The robot makes contact with the terrain at the stance σ.
2) *Collision.* The robot does not collide with itself or the environment, except at the specified contacts in σ.
3) *Equilibrium.* The robot is under quasistatic equilibrium, that is, there exists a set of joint torques and forces at the footfalls in σ that exactly counteracts gravity and keeps the robot still. Known friction is assumed at each contact.

Under these conditions, $\mathcal{F}_\sigma$ has a geometrically complex shape. The contact condition imposes one or more closed chain constraints, and restricts $\mathcal{F}_\sigma$ to lie on a lower dimensional submanifold $\mathcal{C}_\sigma$ of $\mathcal{C}$ (as illustrated in Figure 2). The collision and stability conditions reduce the volume, but not the dimensionality of $\mathcal{F}_\sigma$.

B. Multi-Step Motions

Suppose the robot stands with configuration q at a fixed stance σ, and considers taking a step to place a new footfall f against the environment. The robot will arrive at the stance $\sigma' = \sigma \cup \{f\}$. So, the robot must move along a path within $\mathcal{F}_\sigma$ between q and some *transition* configuration $q' \in \mathcal{F}_\sigma \cap \mathcal{F}_{\sigma'}$. The same condition must be achieved if the robot wishes to remove a footfall f to reach the stance $\sigma' = \sigma \setminus \{f\}$.

The transition region $\mathcal{F}_\sigma \cap \mathcal{F}_{\sigma'}$ is an important bottleneck in legged locomotion. It consists of configurations where the robot touches the footfall at f but can stand without applying any force at f. Thus, it is more constrained than either of $\mathcal{F}_\sigma$ or $\mathcal{F}_{\sigma'}$.

C. The Role of Feasible Space Connectivity

Figure 4 illustrates some of the pitfalls associated with multiply-connected feasible spaces $\mathcal{F}_\sigma$. In particular, if

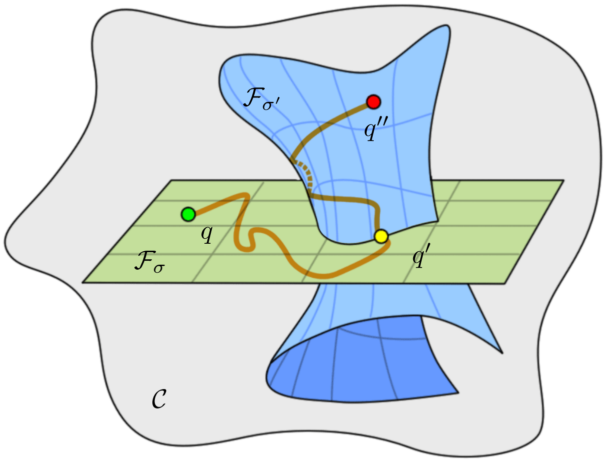

Fig. 3. To move from configuration q at stance σ to q'' at an adjacent stance σ', the motion must pass through a transition configuration q'.

spaces are known to be singly connected, one can infer the existence of a feasible path from σ to σ' as long as the transition $\mathcal{F}_\sigma \cap \mathcal{F}_{\sigma'}$ is non-empty. Thus, step feasibility is a function of stance only.

On the other hand, this conclusion cannot be drawn if multiply-connected spaces exist. Rather, the robot pose affects whether it can take a future step, because it may currently lie in the wrong connected component. Correcting this situation requires backtracking to a different stance and returning to stance σ at a different pose. Intuitively, it would seem that humans and other legged creatures rarely need to execute such maneuvers.

IV. Computing Approximate Connectivity With Roadmaps

Exact queries on $\mathcal{F}_\sigma$, such as querying the number of connected components, or querying whether two points are connected, are very difficult to compute for robots with many degrees of freedom. But the feasibility conditions can be tested quickly for points and curves in $\mathcal{C}$, which means that probabilistic roadmap techniques can be employed efficiently. We use these roadmaps to compute approximate properties of $\mathcal{F}_\sigma$.

A. Overview of Probabilistic Roadmaps

Probabilistic roadmaps (PRMs) are state-of-the-art techniques for motion planning in high-dimensional configuration spaces (see Chapter 7 of [2]). They approximate the shape of a feasible space using a roadmap, which is a network of configurations sampled at random from the feasible set (called *milestones*), and each pair of milestones is connected with an edge if a path segment between them (typically taken to be a straight line) is also feasible.

A PRM planner requires implementations of three subroutines: configuration sampling, configuration feasibility testing, and path segment feasibility testing. The implementations we use for legged robot feasible spaces are described below. Further details can be found in [7].

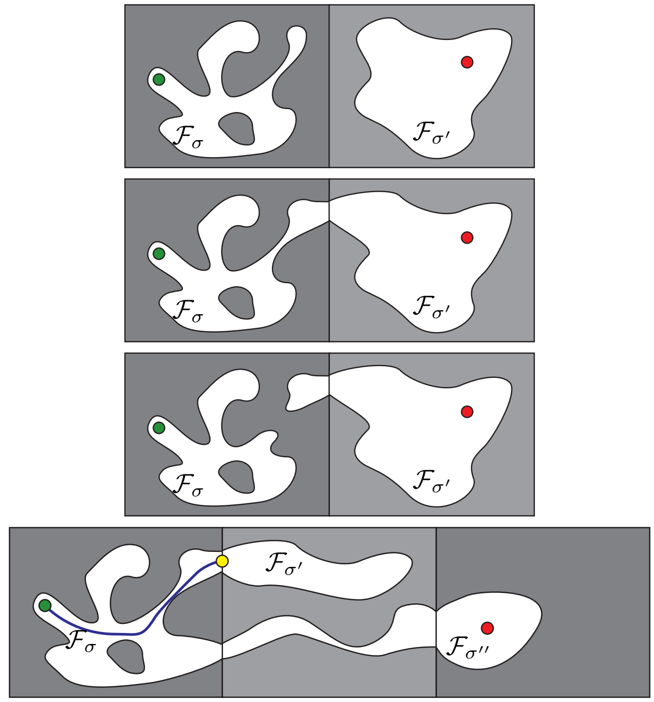

Fig. 4. If the feasible spaces of two stances σ and σ' do not intersect, there is no feasible motion from q_1 to q_2. If $\mathcal{F}_\sigma$ and $\mathcal{F}_{\sigma'}$ are both connected and intersect, then a feasible path definitively exists. If they intersect, but $\mathcal{F}_\sigma$ is multiply-connected, then a feasible path may or may not exist. If $\mathcal{F}_{\sigma'}$ is multiply-connected, then a motion that takes q_1 to $\mathcal{F}_{\sigma'}$ may end up in the wrong component of $\mathcal{F}_{\sigma'}$.

B. Configuration Sampling

To sample configurations in $\mathcal{F}_\sigma$ we rejection sample by generating a configuration q in $\mathcal{C}$, solving an inverse kinematics (IK) problem so that the robot meets the footfalls in $\mathcal{C}_\sigma$, then rejecting the configuration if it is infeasible (with the tests described in the following section). The IK step is necessary because a random sample from $\mathcal{C}$ has zero probability of satisfying the contact constraints.

To solve the IK problem, we first use analytic IK routines that solve for a subset of joint angles that minimize the robot-contact distance, and then use a numerical Newton-Raphson IK solver that optimizes over all joint variables [7]. Experiments suggest that this generates feasible configurations faster than either the analytic or numerical approaches alone. The rejection rate can also be decreased by incorporating the collision and equilibrium constraints as inequality constraints in the Newton-Raphson solver, as described in [6], and by sampling initial configurations from $\mathcal{C}$ with a smarter distribution, such as in [3], [7]. Overall, feasible configurations can be generated approximately at the rate of one every 10 ms – 1 s on a 2 GHz PC, depending on the stance.

C. Configuration Feasibility Tests

Testing for contact uses a simple forward kinematics computation. To efficiently test for collision, we use the Proximity Query Package (PQP), which uses techniques based on bounding volume hierarchies [4].

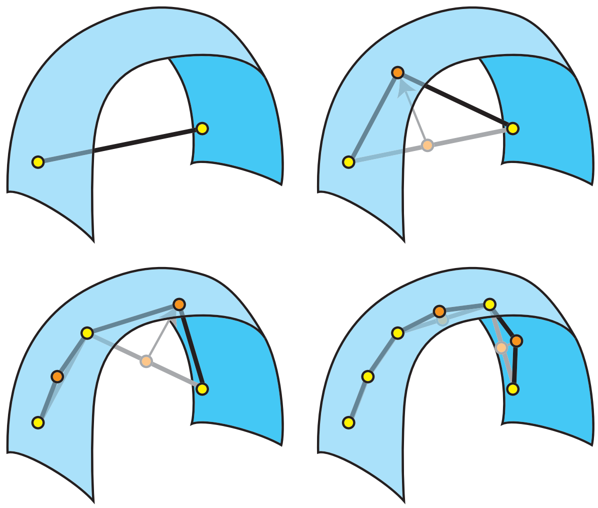

Fig. 5. Deforming a straight line path onto a submanifold. The path is recursively bisected until a tolerance is reached.

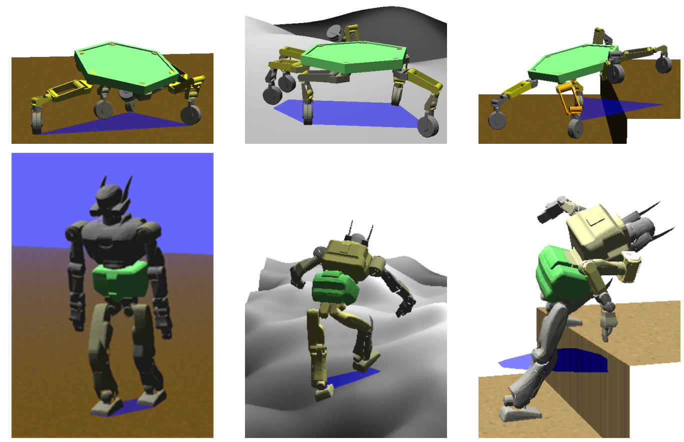

Fig. 6. The three test terrains: Flat, Hills, and Stair.

Now consider the equilibrium condition. On flat ground, this latter condition is nearly equivalent to 1) checking if the center of mass over the convex hull of the footfalls and 2) checking torque limits. On uneven ground, the conditions become significantly more complex, because contact and frictional forces must counteract each other as well as gravity. As a first approximation, the robot can be considered a rigid object with all joints fixed. If there are no valid contact forces that keep the rigid object stable, then the robot cannot be either. Given the force of gravity, and a polyhedral approximation to the friction cone, the equilibrium condition for rigid objects can be tested efficiently using a linear program [1]. A more refined stability test does not treat the robot as a rigid object, and considers the robot's bounded joint torques. This test also involves solving a linear program [6] and can be computed efficiently.

D. Path Segment Feasibility Tests

We use a recursive bisection technique to check the feasibility of straight-line paths between milestones q_1 and q_2. The midpoint q_{mid} between q_1 and q_2 has the highest probability of being infeasible, so it should be tested first. However, a straight-line path will not, in general, lie in $\mathcal{C}_\sigma$. Instead, we simultaneously deform the path onto $\mathcal{C}_\sigma$ and check its feasibility, as follows (Figure 5). Apply the Newton-Raphson method to the midpoint of q_1 and q_2 to produce q_{mid} on $\mathcal{C}_\sigma$. Then test if q_{mid} lies in $\mathcal{F}_\sigma$. If both steps succeed, recurse on both of the segments until a desired resolution ϵ has been reached; otherwise, the checker returns failure.

This method may still miss collisions that occur below the ϵ threshold. For our purposes, we consider it sufficient to take a small ϵ. If exact feasibility is desired, there exist bound-checking techniques that guaranteed exact checking of path segments [11].

E. Approximating Connectivity

Although PRMs are typically used for point-to-point motion planning, here we will use them to approximate the connectivity of feasible spaces. There are several theoretical results that state that the connectivity of a PRM will converge to the connectivity of the feasible space as more configurations are sampled. For example, it the feasible space satisfies an expansiveness condition (believed to widely hold in practice), then the probability that two points are connected in the space but not in the roadmap decreases exponentially to 0 as more configurations are sampled [9].

In the following experiments, we approximate the connectivity of $\mathcal{F}_\sigma$, by building a PRM $\mathcal{R}_\sigma$ of at least 1,000 configurations in $\mathcal{F}_\sigma$. This is done by sampling 10,000 configurations at random, testing their feasibility, and adding feasible configurations as milestones in $\mathcal{R}_\sigma$. The procedure is repeated until $\mathcal{R}_\sigma$ contains at least 1,000 milestones. Once this is complete, the path segments between all pairs of milestones are tested for feasibility, and are added to $\mathcal{R}_\sigma$ if feasible. In our experiments, the computation time for the overall procedure ranged between approximately 1 and 10 minutes per space.

We approximate the true connectivity $N(\sigma)$ of $\mathcal{F}_\sigma$ with the estimate $\hat{N}(\sigma)$ that counts the number of large connected components in $\mathcal{R}_\sigma$. We exclude small, spurious components that contain less than 1% of total milestones, which usually become aggregated into larger components when more samples are added. Errors in the estimate $\hat{N}(\sigma)$ occur when $\mathcal{F}_\sigma$ contains either tiny components or narrow passages, which are unstable features that are likely to disappear upon perturbation. Since these types of features are unstable, they are unlikely to be significantly utilized during motion and hence $\hat{N}(\sigma)$ is a reasonable approximation.

V. Experimental Results

Our experiments studied the feasible spaces of each robot platform on more than 100 stances drawn from three varied environments. These are depicted in Figure 6: flat ground (Flat), undulating terrain (Hills), and a stair-step (Stair).

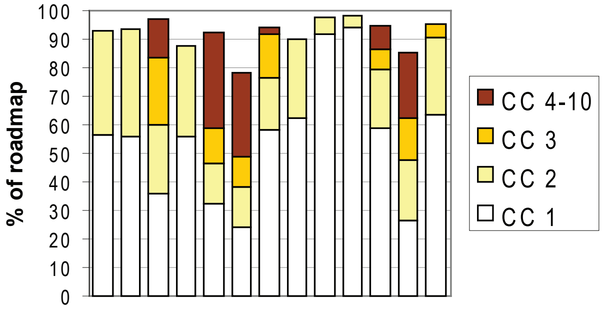

Fig. 7. Sizes of the ten largest connected components of the disconnected spaces in the Stair example.

We selected only stances with nonempty feasible spaces, as verified using random sampling.

For HRP-2, we found that each space consists of a single component. By contrast, 7/34, 9/46, and 13/37 feasible spaces for ATHLETE on the Flat, Hills, and Stair examples, respectively, have multiple components. Of those, more than half have three or more components. Figure 7 plots the component sizes for each disconnected space for the Stair example.

Upon further inspection, we found that each pair of components can usually be separated in a single joint — that is, there exists a joint k and a value θ such that all configurations q in one component have $q_k < \theta$ and those in the other component have $q_k > \theta$, where q_k is the kth joint angle of q. The joint is typically one of the two joints closest to the ground in a support leg. At this stage we do not have a clear explanation why this is the case.

VI. Conclusion

This paper presented a connectivity analysis of the fixed-contact motion spaces of two legged robots. Probabilistic roadmap techniques were used to approximate the connectivity of motion spaces under collision and quasistatic equilibrium constraints. From over 100 motion spaces tested over three simulated environments, all spaces for the humanoid robot HRP-2 were connected. In contrast, approximately 25% of the spaces for the non-biologically-inspired ATHLETE robot were disconnected. These results provide preliminary evidence for the hypothesis that biological organisms have evolved simple motion spaces, and these spaces are inherited by biologically-inspired robots.

Future work will perform further experiments on additional legged robots and organisms to examine whether the motion space hypothesis holds. Dynamic effects may be another promising avenue of study. Our preliminary analysis of the disconnections in ATHLETE's motion spaces suggest that some properties of ATHLETE's nonbiological morphology may be the culprit in causing the disconnections. These properties, however, remain elusive. Understanding the effects of mechanism design on motion spaces will help roboticists design robots that can reason about motion more efficiently, and are easier to operate.

Acknowledgment

This work was partially supported by a Siebel Fellowship. Methods used in this work were supported by and developed in conjunction with colleagues Tim Bretl, Kensuke Harada, Jean-Claude Latombe, and Brian Wilcox.

References

[1] T. Bretl, S. Lall, J.-C. Latombe, and S. Rock. Multi-step motion planning for free-climbing robots. In *WAFR*, Zeist, Netherlands, 2004.
[2] H. Choset, K. Lynch, S. Hutchinson, G. Kantor, W. Burgard, L. Kavraki, and S. Thrun. *Principles of Robot Motion: Theory, Algorithms, and Implementations*. MIT Press, 2005.
[3] 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 Int. Conf. Rob. Aut.*, Washington, D.C., 2002.
[4] S. Gottschalk, M. Lin, and D. Manocha. OBB-tree: A hierarchical structure for rapid interference detection. In *ACM SIGGRAPH*, pages 171–180, 1996.
[5] K. Hauser, T. Bretl, and J.-C. Latombe. Learning-assisted multi-step planning. In *IEEE Int. Conf. Rob. Aut.*, Barcelona, Spain, 2005.
[6] K. Hauser, T. Bretl, and J.-C. Latombe. Non-gaited humanoid locomotion planning. In *IEEE Humanoids*, Tsukuba, Japan, 2005.
[7] K. Hauser, T. Bretl, J.-C. Latombe, K. Harada, and B. Wilcox. Motion planning for legged robots on varied terrain. *Int. J. Rob. Res.*, 27(11-12):1325–1349, 2006.
[8] K. Hauser and J.-C. Latombe. Multi-modal motion planning in non-expansive spaces. In *Workshop on the Algorithmic Foundations of Robotics*, Guanajuato, Mexico, 2008.
[9] D. Hsu, J.-C. Latombe, and R. Motwani. Path planning in expansive spaces. *Int. J. Comp. Geometry and Applic.*, 9(4-5):495–512, 1999.
[10] K. Kaneko, F. Kanehiro, S. Kajita, H. Hirukawa, T. Kawasaki, M. Hirata, K. Akachi, and T. Isozumi. Humanoid robot HRP-2. In *IEEE Int. Conf. Rob. Aut.*, pages 1083–1090, New Orleans, LA, 2004.
[11] F. Schwarzer, M. Saha, and J.-C. Latombe. Exact collision checking of robot paths. In *Workshop on the Algorithmic Foundations of Robotics*, Nice, France, December 2002.
[12] B. H. Wilcox, T. Litwin, J. Biesiadecki, J. Matthews, M. Heverly, J. Morrison, J. Townsend, N. Ahmad, A. Sirota, and B. Cooper. Athlete: A cargo handling and manipulation robot for the moon. *J. Field Robot.*, 24(5):421–434, 2007.