
Multi-Modal Motion Planning in Non-Expansive Spaces

Kris Hauser and Jean-Claude Latombe

Department of Computer Science, Stanford University,
{khauser, latombe}@cs.stanford.edu

Abstract: The motion planning problems encountered in manipulation and legged locomotion have a distinctive multi-modal structure, where the space of feasible configurations consists of overlapping submanifolds of non-uniform dimensionality. These spaces do not possess expansiveness, a property that characterizes whether planning queries can be solved with traditional sample-based planners. We present a new sample-based multi-modal planning algorithm and analyze its completeness properties. In particular, it converges quickly when each mode is expansive relative to the submanifold in which it is embedded. We also present a variant that has the same convergence properties, but works better for problems with a large number of modes. These algorithms are demonstrated in a legged locomotion planner.

1 Introduction

Probabilistic roadmap (PRM) planners (Chapter 7 of [4]) are state-of-the-art approaches for motion planning in high-dimensional configuration spaces under geometrically complex feasibility constraints. They approximate the connectivity of the feasible set \mathcal{F} using a network of randomly sampled configurations connected by straight line paths. It is widely known that PRMs can be slow when \mathcal{F} has poor *expansiveness* [11], or, informally, contains “narrow passages” [9]. In certain *non-expansive* spaces, where \mathcal{F} consists of overlapping submanifolds of varying dimensionality, the narrow passages are arbitrarily thin, and PRMs do not work at all.

This type of *multi-modal* structure characterizes a range of motion planning problems found in manipulation and legged locomotion. Here, each submanifold in \mathcal{F} is identified by a *mode*, a set of fixed contact points maintained between the robot and its environment. The planner then chooses a discrete sequence of modes (a sequence of contacts to make and break), as well as continuous single-mode paths through them (the joint space motions to achieve those changes of contacts).

In problems where each mode is low-dimensional, such as planar manipulation of multiple objects [1, 15], the primary challenge lies in the combinato-

rial complexity of mode sequencing. But in problems with high-dimensional modes, the geometric complexity of single-mode planning poses an additional challenge. Complete planning is intractable, so PRMs are often used for single-mode planning. This approach has solved several specific problems in manipulation with grasps and regrasps [14, 16] and legged locomotion [2, 8]. But what happens to the overall planner’s reliability when it is based on a large number of unreliable single-mode PRM queries (thousands or more)? PRMs cannot report that no path exists, so when a single-mode query fails, we cannot tell if the query is truly infeasible, the goal configuration was picked badly, or the PRM planner just needed more time. So far, little attention has been paid to the theoretical performance guarantees of these algorithms when applied to general multi-modal problems.

This paper presents MULTI-MODAL-PRM, a general-purpose multi-modal planning algorithm for problems with a finite number of modes, and investigates its theoretical completeness properties. MULTI-MODAL-PRM builds a PRM across modes by sampling configurations in \mathcal{F} and in the transitions between modes. We prove that, like PRMs, MULTI-MODAL-PRM will eventually find a feasible path if one exists, and convergence is fast as long as each mode is favorably expansive when restricted to its embedded submanifold.

We also present a more practical variant, INCREMENTAL-MMPRM, which searches for a small candidate subset of modes which are likely to contain a solution path, and then restricts MULTI-MODAL-PRM to these modes. INCREMENTAL-MMPRM has the same asymptotic completeness properties as MULTI-MODAL-PRM, but can be substantially faster for problems with a large number of modes (which is common). We demonstrate the application of INCREMENTAL-MMPRM in a legged locomotion planner [8], showing that it is indeed reliable.

2 Multi-Modal Planning

This section defines multi-modal problems, explains how to use PRMs in single- and multi-modal planning, and summarizes the pitfalls that have made many existing planners incomplete.

2.1 PRMs and Non-Expansive Spaces

PRM planners approximate the connectivity of \mathcal{F} , the feasible subset of a robot’s configuration space, using a roadmap of configurations (referred to as milestones) connected by simple paths (usually straight-lines). The concept of expansiveness was introduced to characterize how quickly a roadmap converges to an accurate representation [11]. So this paper can be self-contained, we review the basic algorithm and its properties in the Appendix.

Formally, \mathcal{F} is *expansive* if there exist constants $\epsilon, \alpha, \beta > 0$ such that \mathcal{F} is $(\epsilon, \alpha, \beta)$ -expansive (see Appendix). Otherwise, \mathcal{F} is non-expansive. In




Fig. 1: (a) A non-expansive, multi-modal feasible space \mathcal{F} . (b) \mathcal{F} can be decomposed into three modes that are individually expansive.

all expansive spaces, the probability that a PRM fails to solve a planning query decreases to zero exponentially as the roadmap grows. However, the convergence rate can become arbitrarily slow as ϵ , α , and β approach zero. An expansive space is poorly expansive (i.e., ϵ , α , and β are low) if it contains narrow passages. However, a smoothed analysis shows that narrow passages are unstable with respect to small perturbations of the input geometry, and are therefore unlikely to occur (except by design) [3].

By contrast, some non-expansive spaces do not exist by chance, but rather for structural reasons. If \mathcal{F} contains a cusp, it is non-expansive, but PRMs might still work well everywhere away from the cusp, since removing a tiny regions around the cusp makes \mathcal{F} expansive without changing its connectivity. But if \mathcal{F} contains regions of varying dimensionality, then PRM planners have probability zero of answering most queries. For example, \mathcal{F} may consist of 2D regions connected by 1D curves (Figure 1a). In these spaces, the PRM convergence bounds take the meaningless value of 1.

Varying dimensionality is *inherent* in the structure of multi-modal problems. In these problems, the submanifolds that form \mathcal{F} can be enumerated from the problem definition, e.g., by considering all possible combinations of contacts. However, their number may be huge.

2.2 Multi-Modal Problem Definition

The robotic system moves between a finite set of modes, Σ . Each mode $\sigma \in \Sigma$ is assigned a mode-specific feasible space \mathcal{F}_σ , and $\mathcal{F} = \bigcup_{\sigma \in \Sigma} \mathcal{F}_\sigma$ (Figure 1b). The feasibility constraints of a mode are divided into two classes.

- *Dimensionality-reducing* constraints, often represented as functional equalities $C_\sigma(q) = 0$. Define the submanifold \mathcal{C}_σ as the set of configurations that satisfy these constraints.
- *Volume-reducing* constraints, often represented as inequalities $D_\sigma(q) > 0$. These may cause \mathcal{F}_σ to be empty. Otherwise, \mathcal{F}_σ has the same dimension as \mathcal{C}_σ but lower volume (Figure 2a).




Fig. 2: (a) At a mode σ , motion is constrained to a subset \mathcal{F}_σ of a submanifold \mathcal{C}_σ with lower dimension than \mathcal{C} . (b) To move from configuration q at stance σ to q'' at an adjacent stance σ' , the motion must pass through a transition configuration q' .

For example, in legged locomotion, σ is a fixed set of footfalls. \mathcal{F} can be embedded in a configuration space \mathcal{C} , which consists of parameters for a free-floating robot base and the robot's joint angles. Enforcing contact at the footfalls imposes multiple closed-loop kinematic constraints, and these in turn define \mathcal{C}_σ (a submanifold of uniform lower dimension, except at singularities). Collision avoidance and stability constraints are volume-reducing, and restrict \mathcal{F}_σ to a subset of \mathcal{C}_σ .

2.3 Planning Between Two Modes

PRMs can be used for single-mode planning restricted to \mathcal{C}_σ . This requires adapting configuration sampling and path segment feasibility testing to the submanifold, since the sampler must have nonzero probability of generating a configuration in \mathcal{F}_σ , and a path segment on \mathcal{C}_σ may not be a straight line. Two approaches are parameterizing \mathcal{C}_σ with an atlas of charts [5], or using numerical methods to move configurations from the ambient space onto \mathcal{C}_σ [13]. PRM planners will plan quickly as long as \mathcal{F}_σ is expansive (restricted to \mathcal{C}_σ).

Suppose the robotic system is at configuration q at mode σ . To switch to σ' , we must plan a path in \mathcal{F}_σ that ends in a configuration q' in $\mathcal{F}_\sigma \cap \mathcal{F}_{\sigma'}$ (Figure 2b). The region $\mathcal{F}_\sigma \cap \mathcal{F}_{\sigma'}$ is called the *transition* between σ and σ' , and q' is called a *transition configuration*. Transitions are at least as constrained as σ or σ' because they must simultaneously satisfy the constraints of both stances. Hence, they often have zero volume relative to \mathcal{C}_σ or $\mathcal{C}_{\sigma'}$. Thus, transition configurations should be sampled explicitly from $\mathcal{C}_\sigma \cap \mathcal{C}_{\sigma'}$. Like single-mode sampling, transition sampling can be addressed using explicit parameterization or numerical techniques [6].

The existence of q' is a necessary condition for a single-mode path to connect q and σ' , and is also a good indication that a feasible path exists. Sampling q' is also typically faster than planning a single-mode path. These two observations are instrumental in the implementation of INCREMENTAL-

MMPRM. But the existence of q' is not sufficient for a path to exist, because q and q' might lie in different connected components of \mathcal{F}_σ .

2.4 Planning in Multiple Modes

Because PRMs usually work well for single-mode planning, multi-modal planning is usually addressed by combining several single-mode plans into a multi-modal plan. Given a discrete set of modes $\{\sigma_1, \dots, \sigma_m\}$, a multi-modal planner explores a *mode graph* \mathcal{G} . Each vertex in \mathcal{G} represents a mode, and an edge connects each pair of *adjacent* modes. We assume we are given a cheap computational *adjacency test*, which tests a necessary condition for the existence of a path moving between any pair of modes σ and σ' . For example, an adjacency test in legged locomotion tests whether σ' adds exactly one footfall within the reach of σ . The test should prune out as many unnecessary edges as possible to help reduce the size of \mathcal{G} .

Because \mathcal{G} tends to be extremely large, one natural strategy explores only a small part of \mathcal{G} . For example, a search-like algorithm incrementally builds a tree \mathcal{T} of configurations reachable from the start. \mathcal{T} is initialized to q_{start} . Each expansion step picks a configuration q at mode σ in \mathcal{T} , enumerates each adjacent mode σ' , then uses a PRM to plan a feasible single-mode path from q to a transition configuration q' in $\mathcal{F}_\sigma \cap \mathcal{F}_{\sigma'}$. If successful, q' is added to \mathcal{T} . This repeats until the goal is reached, and the single-mode paths leading to the goal are concatenated into a multi-modal path.

2.5 Completeness Challenges in Multi-Modal Planning

Special challenges arise when combining several single-mode PRM queries to solve a multi-modal problem. MULTI-MODAL-PRM addresses the following issues, any of which may cause a planner to fail to find a path.

Because any single-mode query might be infeasible, PRM planners must be terminated with failure after some cutoff time. Set the cutoff too low, and the planner may miss critical paths; too high, and the planner wastes time on infeasible queries. Prior work has usually just tried to avoid infeasible queries [2, 6, 12], allowing the cutoff to be set high. Another approach avoids cutoffs by interleaving computation among queries [6, 15, 14].

Because transitions $\mathcal{F}_\sigma \cap \mathcal{F}_{\sigma'}$ may have zero volume in \mathcal{F}_σ or $\mathcal{F}_{\sigma'}$, transition configurations should be sampled explicitly from $\mathcal{C}_\sigma \cap \mathcal{C}_{\sigma'}$. Furthermore, more than one configuration q' may need to be sampled in each transition $\mathcal{F}_\sigma \cap \mathcal{F}_{\sigma'}$, because a configuration may lie in a component that is disconnected in \mathcal{F} from the start configuration q , or one that is disconnected in \mathcal{F}' from a target configuration q'' .

3 Multi-Modal-PRM

This section presents the general MULTI-MODAL-PRM algorithm, and gives an overview of its theoretical completeness properties.

3.1 Algorithm

MULTI-MODAL-PRM builds PRMs across all modes, connecting them at explicitly sampled transition configurations (Figure 3). Suppose there are m modes, $\Sigma = \{\sigma_1, \dots, \sigma_m\}$. For each $\sigma_i \in \Sigma$ maintain a roadmap \mathcal{R}_{σ_i} of \mathcal{F}_{σ_i} . Define the sampler $\text{SAMPLE-MODE}(\sigma_i)$ as follows. Uniformly sample a configuration q from \mathcal{C}_{σ_i} . If q is in \mathcal{F}_{σ_i} , q is returned; if not, SAMPLE-MODE returns failure. Similarly, define $\text{SAMPLE-TRANS}(\sigma_i, \sigma_j)$ to rejection sample from $\mathcal{F}_{\sigma_i} \cap \mathcal{F}_{\sigma_j}$. MULTI-MODAL-PRM is defined as follows:

MULTI-MODAL-PRM(q_{start}, q_{goal}, N)

Add q_{start} and q_{goal} as milestones to the roadmaps corresponding to their modes (σ_{start} and σ_{goal}).

Repeat N times:

1. For each mode σ_i , sample a configuration q using $\text{SAMPLE-MODE}(\sigma_i)$. If it succeeds, add q to \mathcal{R}_{σ_i} as a new *milestone*, and connect it to existing milestones.
2. For each pair of adjacent modes σ_i and σ_j , sample a configuration q using $\text{SAMPLE-TRANS}(\sigma_i, \sigma_j)$. If it succeeds, add q to \mathcal{R}_{σ_i} and \mathcal{R}_{σ_j} , and connect it to all visible milestones in $\mathcal{R}_{\sigma_i} \mathcal{R}_{\sigma_j}$.

Build an aggregate roadmap \mathcal{R} by connecting roadmaps at matching transition configurations.

If q_{start} and q_{goal} are connected by a path in \mathcal{R} , terminate with success. Otherwise, return failure.

3.2 Summary of Theoretical Results

Section 4 proves that, under certain conditions, the probability that MULTI-MODAL-PRM is incorrect (returns failure when a feasible path actually exists) is less than ce^{-dN} , where c and d are positive constants and N is the number of iterations. This means that as more time is spent planning, the probability of failure decreases quickly to zero. The constants c and d do not explicitly depend on the dimensionality of the configuration space, or the total number of modes. Furthermore, since a constant number of samples ($m + n$, where n is the number of adjacencies) are drawn per iteration, MULTI-MODAL-PRM also converges exponentially in the total number of samples drawn.

This bound holds if the following conditions are met:

1. The set of modes $\Sigma = \{\sigma_1, \dots, \sigma_m\}$ is finite.
2. If \mathcal{F}_{σ_i} is nonempty, then it is expansive.
3. If \mathcal{F}_{σ_i} is nonempty, $\text{SAMPLE-MODE}(\sigma_i)$ succeeds with nonzero probability.
4. If $\mathcal{F}_{\sigma_i} \cap \mathcal{F}_{\sigma_j}$ is nonempty, $\text{SAMPLE-TRANS}(\sigma_i, \sigma_j)$ samples each connected component of $\mathcal{F}_{\sigma_i} \cap \mathcal{F}_{\sigma_j}$ with nonzero probability.

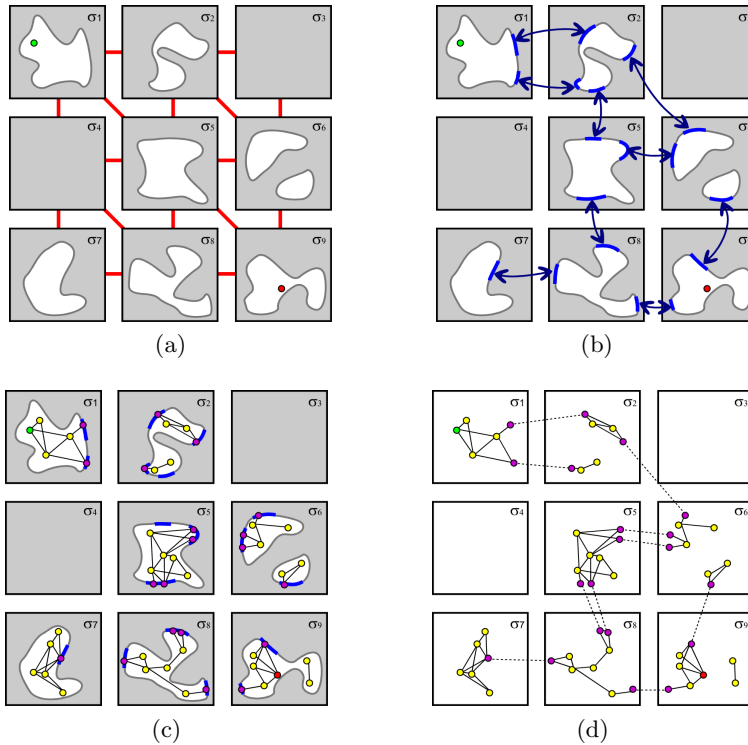


Fig. 3: Illustrating MULTI-MODAL-PRM on an abstract example problem. (a) A mode graph with nine modes, with adjacent modes connected by lines. (b) The transition regions, with arrows indicating how they map between modes. (c) Building roadmaps. Light dots are milestones sampled from modes, dark ones are sampled from transitions. (d) The aggregate roadmap. Transition configurations connected by dashed lines are identified.

4 Proof of Completeness Properties

This section proves that MULTI-MODAL-PRM is probabilistically complete and exponentially convergent in the number of iterations. The proof shows that three processes exponentially converge: 1) BASIC-PRM under rejection sampling; 2) roadmaps connecting transition components; and 3) roadmaps along any sequence of modes which contains a feasible multi-step path.

4.1 Exponential Convergence

We say a process is *exponentially convergent* in N if the probability of failure is less than ae^{-bN} for positive constants a and b . A useful *composition principle* allows us to avoid stating the coefficients a and b , which become cumbersome. If two quantities are subject to exponentially decreasing upper bounds, their sums and products are themselves subject to exponentially decreasing bounds.

The product is trivial, and $a_1e^{-b_1N} + a_2e^{-b_2N}$ is upper bounded by ae^{-bN} for $a = (a_1 + a_2)$ and $b = \max b_1, b_2$. So if two processes are exponentially convergent in N , then the probability that both of them succeed, or either one succeeds, also converges exponentially in N .

4.2 PRMs converge under rejection sampling

The first lemma states that PRMs converge exponentially, not just in the number of milestones (as was proven in [11]), but also in the number of rejection samples drawn. Let p be the probability that a random sample from \mathcal{C} lies in \mathcal{F} . From N configurations randomly sampled from \mathcal{C} , let the milestones \mathcal{M} be those that are feasible.

Lemma 1. *If \mathcal{F} is expansive and $p > 0$, a roadmap \mathcal{R} constructed from \mathcal{M} connects two query configurations q_1 and q_2 with probability exponentially convergent in N .*

Proof. Since the configuration samples are independent, the size of \mathcal{M} is binomially distributed. Hoeffding's inequality gives an upper bound to the probability that \mathcal{M} has n or fewer milestones:

$$Pr(|\mathcal{M}| \leq n) \leq \exp(-2(Np - n)^2/N).$$

If $|\mathcal{M}| \leq n$, the probability of failure is at most 1. On the other hand, if $|\mathcal{M}| > n$, the probability that \mathcal{R} fails to connect q_1 and q_2 is less than ce^{-dn} , where c and d are the constants in Theorem 2 (see Appendix). Since these events are mutually exclusive, have the overall probability of failure ν

$$\begin{aligned} \nu &\leq Pr(|\mathcal{M}| \leq n) \cdot 1 + Pr(|\mathcal{M}| > n)c \exp(-dn) \\ &\leq \exp(-Np^2/2) + c \exp(-Ndp/2) \end{aligned} \tag{1}$$

where we have set $n = pN/2$. Using the composition principle, this bound is exponentially decreasing in N as desired.

4.3 Convergence of paths connecting transition components

Here we consider the probability of finding a path in \mathcal{F} between arbitrary connected subsets A and B , by building a roadmap using N rejection samples in A , B , and \mathcal{F} .

Let \mathcal{M} and p be defined as in Lemma 1. Suppose we rejection sample A and B respectively by drawing samples uniformly from supersets A' and B' , with probability of success p_A and p_B . From N configurations sampled from A' , let the milestones \mathcal{M}_A be the configurations in A . Define \mathcal{M}_B similarly.

Lemma 2. *Suppose \mathcal{F} is expansive, and p , p_A , and p_B are nonzero. Let \mathcal{R} be the roadmap constructed from all milestones \mathcal{M} , \mathcal{M}_A , and \mathcal{M}_B . If A and B are in the same connected component in \mathcal{F} , then the probability that \mathcal{R} contains a path between A and B is exponentially convergent in N .*

Proof. View any pair of milestones q_A in \mathcal{M}_A and q_B in \mathcal{M}_B as PRM query configurations. Then \mathcal{R} contains a path between A and B when (event X) \mathcal{M}_A is nonempty, (event Y) \mathcal{M}_B is nonempty, and (event Z) the roadmap formed from \mathcal{M} connects q_A and q_B .

The probability that N rejection samples from A' fails to find a milestone in A is at most $(1 - p_A)^N \leq e^{-Np_A}$. So event X is exponentially convergent. The same holds for Y by symmetry. Finally, event Z is exponentially convergent in N by Lemma 1 and since q_A and q_B are in the same connected component. The lemma holds by composition of X , Y , and Z .

4.4 Convergence of Multi-Modal-PRM

Lemma 3. *Let the assumptions at the end of Section 3 hold. Between any two configurations, if any feasible multi-step path exists, there is some feasible path that makes a finite number of mode switches.*

Proof. Expansiveness implies ϵ -goodness, which implies that each connected component of the feasible space has volume at least $\epsilon > 0$. Let ϵ_0 be such that each mode is ϵ_0 -good. Then, each mode can only contain $1/\epsilon_0$ components, with m/ϵ_0 components overall.

Theorem 1. *Let the assumptions at the end of Section 3 hold. If q_s and q_g can be connected with a feasible multi-step path, then the probability that MULTI-MODAL-PRM finds a path is exponentially convergent in the number of iterations N .*

Proof. If q_s and q_g can be connected with a feasible path, there is a feasible path with a finite number of mode switches. Let $y(t)$ be such a path. Suppose the path travels through mode σ_k , starting at transition connected component T_1 and ending at T_2 . SAMPLE-MODE and SAMPLE-TRANS have properties allowing Lemma 2 to be applied to roadmap \mathcal{R}_{σ_k} . Therefore, the probability that \mathcal{R}_{σ_k} connects T_1 and T_2 exponentially converges to 1. The theorem follows from repeating this argument for all modes along the path, and using the composition principle.

Dimensionality and the total number of modes m do not explicitly affect the coefficients in the convergence bound (although the total cost per iteration is at least linear in m). The modes' expansiveness measures and the parameters p and p' have a straightforward effect on the convergence rate: when expansiveness or the parameters increase, the bound moves closer to zero. The bound also moves closer to zero if fewer modes are needed to reach the goal, or if the goal can be reached via multiple paths.

5 An Incremental Variant

Even executing a few iterations of MULTI-MODAL-PRM is impractical if the number of modes is large. Typical legged locomotion queries have over a bil-

lion modes, but only tens or hundreds of steps are needed to go anywhere within the range of visual sensing. The INCREMENTAL-MMPRM variant uses heuristics to produce a small subset of modes that are likely to contain a path to the goal. Limiting planning to these modes make planning much faster. But heuristics are not always right, so INCREMENTAL-MMPRM is designed to gracefully degrade back to MULTI-MODAL-PRM if necessary.

5.1 Overview

INCREMENTAL-MMPRM alternates between *refinement* and *expansion*. At each round r , it restricts itself to building roadmaps over a candidate set of modes Σ_r . We set Σ_0 to the empty set, and all single-mode roadmaps (except the start and goal) are initially empty. It repeats the following for rounds $r = 1, \dots, N$:

1. *Expansion*. Add new modes to Σ_{r-1} to produce the next candidate mode set Σ_r .
2. *Refinement*. Incrementally build roadmaps in Σ_r by performing n_r mode and transition samples (e.g., one iteration of MULTI-MODAL-PRM).

The performance of INCREMENTAL-MMPRM depends mainly on the expansion heuristic. The heuristic does not affect asymptotic convergence, as long as the candidate mode set grows until it cannot expand any further (at which point INCREMENTAL-MMPRM behaves exactly like MULTI-MODAL-PRM). But practically, it has a large impact on running time. The heuristic below can be applied to any multi-modal planning problem, and improves running time by orders of magnitude.

5.2 Expansion: Search Among Feasible Transitions

For some systems, any set of modes leading to the goal (such as those found with heuristic search) could be a reasonable choice of candidate modes. But in many systems, *most* modes and transitions are infeasible, so this approach would cause the planner to waste most of its time building roadmaps in infeasible modes. If the expansion step produces candidate modes that are likely to contain a feasible path, overall planning speed will be improved, even if expansion incurs additional computational expense.

Search among feasible transitions (SAFT) uses the existence of a feasible transition as a good indication that a feasible path exists as well (as in [2]). SAFT incrementally builds a mode graph \mathcal{G} , but without expanding an edge until it samples a feasible transition configuration. To avoid missing transitions with low volume, SAFT interleaves transition sampling between modes (as in [6]). It maintains a list \mathcal{A} of “active” transitions. Each transition T in \mathcal{A} has an associated priority $p(T)$, which decreased as more time is spent sampling T . The full algorithm is as follows:

SAFT-INIT (performed only once at round 0)

Add the start mode σ_{start} to \mathcal{G} . Initialize \mathcal{A} to contain all transitions out of σ_{start} .

SAFT-EXPAND(r) (used on expansion round r)

Repeat the following:

1. Remove a transition T from \mathcal{A} with maximum priority. Suppose T is a transition from σ to σ' . Try to sample a configuration in $\mathcal{F}_\sigma \cap \mathcal{F}_{\sigma'}$.
2. On failure, reduce the priority $p(T)$ and reinsert T into \mathcal{A} .
3. On success, add σ' to the mode graph \mathcal{G} . For each transitions T' out of σ' , add T' to \mathcal{A} with initial priority $p(T')$.

Repeat until \mathcal{G} contains a sequence of modes (not already existing in Σ_{r-1}) connecting σ_{start} to σ_{goal} . Add these modes to Σ_{r-1} to produce Σ_r .

As in [6], SAFT-EXPAND may find paths more quickly if the initial $p(T)$ estimates the probability that T is nonempty. Other heuristics, such as a bias toward easier steps, could also be incorporated into $p(T)$ [8].

5.3 Refinement: Strategies to Improve Connectivity

The running time of INCREMENTAL-MMPRM is also substantially affected by the refinement strategy. First, if the sample count parameter n_r is too high, the planner might waste time on a candidate set of modes that contains no feasible path; too low, and the planner will expand the candidate mode set unnecessarily. We set n_r proportional to the number of modes, with some minor tuning of the proportionality constant.

Second, given a fixed n_r , the planner should allocate single-mode planning computations to maximize its chance of finding a feasible path. If a mode’s roadmap is already highly connected, then more samples are unlikely to improve connectivity, so we bias sampling toward modes with poorly connected roadmaps. Furthermore, if the aggregate roadmap \mathcal{R} already contains paths connecting two modes separated by σ , additional planning in σ is unlikely to improve connectivity. Thus, we bias sampling toward modes that could potentially connect large components of \mathcal{R} .

Also, rather than use BASIC-PRM for single-mode planning, we use the SBL variant, which is much faster in practice [17]. SBL grows trees rooted from transition configurations, and delays checking the feasibility of straight line paths.

6 Experiments in a Legged Locomotion Planner

These algorithms are compared in a legged locomotion planner designed for rocky, steep terrain, and applied to NASA’s ATHLETE robot, a six-legged




Fig. 4: Three terrains for testing the locomotion planner.

lunar vehicle with 36 revolute joints. See [8] for details of the problem specifications and implementation. In previous work, we used a two stage algorithm that worked well for a four-limbed robot [2] and a bipedal robot [7]. The algorithm is, essentially, INCREMENTAL-MMPRM with a different refinement stage. Using a SAFT-like method, the first (exploration-like) stage produces a sequence of footsteps to take, and transition configurations for each step. The second (refinement-like) stage plans single-mode paths to connect the transition configurations. On ATHLETE, this second stage fails often even on seemingly simple problems, greatly increasing the number of times that the planner returns to exploration. This is costly, since each exploration usually takes several minutes. Here, we test how INCREMENTAL-MMPRM improves reliability of the refinement stage, assuming a good candidate set of modes has already been provided.

We used the three test terrains of Figure 4: flat ground (Flat), a smoothly undulating terrain (Hills), and a stair step with a height of 0.5 times the diameter of ATHLETE’s chassis (Step). In each terrain, the planner first samples several hundred candidate footfalls on the terrain surface. Each mode σ is defined as a fixed set of simultaneously reached footfalls. SAFT was called once to find a sequence of modes Σ_1 reaching the goal, where each subsequent mode adds or removes a footfall. We compare three refinement methods: 1) sampling one feasible configuration in each transition, and calling SBL to connect them with single-mode paths (Single trans); 2) the basic MULTI-MODAL-PRM algorithm (I-MMPRM-Basic); and 3) MULTI-MODAL-PRM using the connection strategy outlined in Section 5.3 (I-MMPRM-Connect). We tested each method 10 times on identical Σ_1 but with different random seeds, and terminating the planner if no path was found after 30,000 total samples.

Table 1 reports the results. The single-transition method fails in every case. The methods based on MULTI-MODAL-PRM successfully find a path within the iteration limit nearly every time. I-MMPRM-Connect is faster than I-MMPRM-Basic, especially when the steps have varying degrees of difficulty. On flat ground, Σ_1 contained one particularly difficult step (depicted in Figure 4), and hence I-MMPRM-Connect was over four times faster.

The failure of the single-transition method is consistent with the explanation that, by accident, transition configurations were sampled in disconnected components, which of course cannot be connected with single-mode paths.

Problem	Method	% successful	Time (s)
Flat	Single trans	0%	501 (88.3)
	I-MMPRM-Basic	90%	409 (180)
	I-MMPRM-Connect	100%	106 (36.6)
Hills	Single trans	0%	149 (29.5)
	I-MMPRM-Basic	100%	181 (103)
	I-MMPRM-Connect	100%	99.5 (28.3)
Step	Single trans	0%	258 (36.9)
	I-MMPRM-Basic	100%	507 (197)
	I-MMPRM-Connect	100%	347 (138)

Table 1: Locomotion planning statistics on various problems with a fixed mode sequence, averaged over 10 runs. Standard deviations in parentheses.

The MULTI-MODAL-PRM-based methods succeed by sampling multiple configurations in each transition (typically around 10 before finding a solution).

In similar experiments with a bipedal humanoid robot [7], the single-transition method worked fairly consistently, with an 80-90% success rate. The methods based on MULTI-MODAL-PRM were again consistently successful, but needed very few configurations per transition (typically 1 or 2 before finding a solution). This modest reliability increase came at a moderate computational overhead; I-MMPRM-Connect averaged about twice as long as Single-Trans at finding a feasible path.

7 Conclusion

This paper presented MULTI-MODAL-PRM, a new sample-based multi-modal planning algorithm for problems with a finite number of modes. We proved that MULTI-MODAL-PRM has probability of failure exponentially converging to 0 in the number of samples drawn, if the feasible space of each mode is expansive relative to its embedded submanifold. We also presented a variant, INCREMENTAL-MMPRM, that restricts planning to an incrementally growing set of candidate modes, and is orders of magnitude faster in problems with a large number of modes. We demonstrated the reliability of these techniques in experiments in a locomotion planner.

When applied to the ATHLETE six-legged robot, MULTI-MODAL-PRM dramatically improves the planner’s reliability over a simpler incomplete method. When applied to a humanoid biped, MULTI-MODAL-PRM modestly improves reliability with moderate overhead. These results suggest that disconnected feasible spaces occurred more frequently in ATHLETE than in the biped. Future work might aim to discover the causes of this phenomenon. We suspect characteristics of ATHLETE’s kinematics (e.g., singularities) or its larger number of legs in contact might be contributing factors.

Most systems with contact are posed as having a continuous (uncountably infinite) number of modes, which are then discretized for planning [8, 15, 16].

The completeness results in this paper hold only if the set of discretized modes contains a solution path. If no such path exists, the modes may have been discretized poorly. In future work, we hope to investigate the completeness and convergence rate of multi-modal planners that sample modes from continuous sets of modes.

Appendix: Probabilistic Roadmaps and Expansiveness

This appendix reviews the basic PRM algorithm and its theoretical completeness properties in expansive spaces.

A.1 Basic Algorithm

PRM planners address the problem of connecting two configurations q_{start} and q_{goal} in the feasible space \mathcal{F} , a subset of configuration space \mathcal{C} . Though it is prohibitively expensive to compute an exact representation of \mathcal{F} , feasibility tests are usually cheap. So to approximate the connectivity of \mathcal{F} , PRM planners build a *roadmap* \mathcal{R} , a network of feasible configurations (called *milestones*) connected with straight-line segments. A basic algorithm operates as follows:

BASIC-PRM(q_{start}, q_{goal}, n)
 Add q_{start} and q_{goal} to \mathcal{R} as milestones.
 Repeat n times:

1. Sample a configuration q uniformly from \mathcal{C} , and test its feasibility. Repeat until a feasible sample is found.
2. Add q to \mathcal{R} as a new milestone. Connect it to nearby milestones q' in \mathcal{R} if the line segment between q and q' lies in \mathcal{F} .

If \mathcal{R} contains a path between q_{start} and q_{goal} , return the path.
 Otherwise, return ‘failure’.

If

a PRM planner produces a path successfully, the path is guaranteed to be feasible, but if it fails, then we cannot tell whether no path exists or the cutoff n was set too low.

A.2 Performance in Expansive Spaces

BASIC-PRM and several variants have been shown to be probabilistically complete, that is, the probability of incorrectly returning failure approaches 0 as n increases. One particularly strong completeness theorem proves that PRMs converge exponentially, given that \mathcal{F} is *expansive* [11].

The notion of expansiveness expresses the difficulty of constructing a roadmap that captures the connectivity of \mathcal{F} . The success of PRMs in high




Fig. 5: A poorly expansive space. (a) The visibility set $\mathcal{V}(q)$ in region S_2 , of a point q in S_1 . (b) $\text{LOOKOUT}_\beta(S_1)$ is the set of points that see at least a β fraction of S_2 .

dimensional spaces is partially explained by the fact that expansiveness is not explicitly dependent on the dimensionality of \mathcal{F} . Let $\mu(S)$ measure the volume of any subset $S \subseteq \mathcal{F}$ (with $\mu(\mathcal{F})$ finite), and let $\mathcal{V}(q)$ be the set of all points that can be connected to q with a straight line in \mathcal{F} . The *lookout* set of a subset S of \mathcal{F} is defined as the subset of S that can “see” a substantial portion of the complement of S (see Figure 5). Formally, given a constant $\beta \in (0, 1]$ and a subset S of a connected component \mathcal{F}' in \mathcal{F} , define

$$\text{LOOKOUT}_\beta(S) = \{q \in S \mid \mu(\mathcal{V}(q) \setminus S) \geq \beta \mu(\mathcal{F}' \setminus S)\}$$

For constants $\epsilon, \alpha, \beta \in (0, 1]$, \mathcal{F} is said to be $(\epsilon, \alpha, \beta)$ -*expansive* if:

1. For all $q \in \mathcal{F}$, $\mu(\mathcal{V}(q)) \geq \epsilon \mu(\mathcal{F})$.
2. For any connected subset S , $\mu(\text{LOOKOUT}_\beta(S)) \geq \alpha \mu(S)$.

The first property is known as ϵ -goodness, and states that each configuration “sees” a significant fraction of \mathcal{F} . The second property can be interpreted as follows. View S as the visibility set of a single roadmap component \mathcal{R}' . Let \mathcal{F}' be the component of feasible space in which S lies. Then, with significant probability (at least $\alpha \mu(S)$), a random configuration will simultaneously connect to \mathcal{R}' and significantly reduce the fraction of \mathcal{F}' not visible to \mathcal{R}' (by at least β).

The primary convergence result of [11] can be restated as follows:

Theorem 2. *If \mathcal{F} is $(\epsilon, \alpha, \beta)$ -expansive, then the probability that a roadmap of n uniformly, independently sampled milestones fails to connect q_{start} and q_{goal} is no more than ce^{-dn} for some positive constants c and d .*

The constants c and d are simple functions of ϵ , α , and β . If \mathcal{F} is favorably expansive (ϵ , α , and β are high), the bound is close to zero, and BASIC-PRM will find a path between q_{start} and q_{goal} relatively quickly. If, on the other

hand, \mathcal{F} is poorly expansive (ϵ , α , and β are low), then PRM performance might be poor for certain query configurations q_{start} and q_{goal} . A complementary theorem proven in [10] states that a PRM planner will succeed with arbitrarily low probability for any fixed n in spaces with small α and β .

References

1. R. Alami, J.-P. Laumond, and T. Siméon. Two manipulation planning algorithms. In K. Goldberg, D. Halperin, J.-C. Latombe, and R. Wilson, editors, *Alg. Found. Rob.*, pages 109–125. A K Peters, Wellesley, MA, 1995.
2. T. Bretl, S. Lall, J.-C. Latombe, and S. Rock. Multi-step motion planning for free-climbing robots. In *WAFR*, Zeist, Netherlands, 2004.
3. S. Chaudhuri and V. Koltun. Smoothed analysis of probabilistic roadmaps. In *Fourth SIAM Conf. of Analytic Algorithms and Comp. Geometry*, 2007.
4. 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.
5. J. Cortés and T. Siméon. Sampling-based motion planning under kinematic loop-closure constraints. In *WAFR*, Zeist, Netherlands, 2004.
6. K. Hauser, T. Bretl, and J.-C. Latombe. Learning-assisted multi-step planning. In *IEEE Int. Conf. Rob. Aut.*, Barcelona, Spain, 2005.
7. K. Hauser, T. Bretl, and J.-C. Latombe. Non-gaited humanoid locomotion planning. In *IEEE Humanoids*, Tsukuba, Japan, 2005.
8. K. Hauser, T. Bretl, J.-C. Latombe, and B. Wilcox. Motion planning for a six-legged lunar robot. In *WAFR*, New York, NY, 2006.
9. D. Hsu, L. E. Kavraki, J.-C. Latombe, R. Motwani, and S. Sorkin. On finding narrow passages with probabilistic roadmap planners. In *WAFR*, pages 141–153, Natick, MA, 1998.
10. D. Hsu, J. Latombe, and H. Kurniawati. On the probabilistic foundations of probabilistic roadmap planning. *Int. J. Rob. Res.*, 25(7):627–643, 2006.
11. D. Hsu, J.-C. Latombe, and R. Motwani. Path planning in expansive configuration spaces. In *IEEE Int. Conf. Rob. Aut.*, pages 2219–2226, 1997.
12. Y. Koga and J.-C. Latombe. On multi-arm manipulation planning. In *IEEE Int. Conf. Rob. Aut.*, pages 945–952, San Diego, CA, 1994.
13. S. M. LaValle, J. H. Yakey, and L. E. Kavraki. A probabilistic roadmap approach for systems with closed kinematic chains. In *IEEE Int. Conf. Rob. Aut.*, Detroit, MI, 1999.
14. C. L. Nielsen and L. E. Kavraki. A two level fuzzy prm for manipulation planning. In *IEEE/RSJ Int. Conf. Int. Rob. Sys.*, pages 1716–1721, 2000.
15. D. Nieuwenhuisen, A. F. van der Stappen, and M. H. Overmars. An effective framework for path planning amidst movable obstacles. In *WAFR*, New York, NY, 2006.
16. A. Sahbani, J. Cortés, and T. Siméon. A probabilistic algorithm for manipulation planning under continuous grasps and placements. In *IEEE/RSJ Int. Conf. Int. Rob. Sys.*, pages 1560–1565, Lausanne, Switzerland, 2002.
17. G. Sánchez and J.-C. Latombe. On delaying collision checking in PRM planning: Application to multi-robot coordination. *Int. J. of Rob. Res.*, 21(1):5–26, 2002.