Sampling Based Motion Planning with Reachable Volumes: Application to Manipulators and Closed Chain Systems

Troy McMahon[1], Shawna Thomas[1], Nancy M. Amato[1]

Abstract— **Reachable volumes are a geometric representation of the regions the joints of a robot can reach. They can be used to generate constraint satisfying samples for problems including complicated linkage robots (e.g. closed chains and graspers). They can also be used to assist robot operators and to help in robot design. We show that reachable volumes have an $\mathcal{O}(1)$ complexity in unconstrained problems as well as in many constrained problems. We also show that reachable volumes can be computed in linear time and that reachable volume samples can be generated in linear time in problems without constraints.**

We experimentally validate reachable volume sampling, both with and without constraints on end effectors and/or internal joints. We show that reachable volume samples are less likely to be invalid due to self-collisions, making reachable volume sampling significantly more efficient for higher dimensional problems. We also show that these samples are easier to connect than others, resulting in better connected roadmaps. We demonstrate that our method can be applied to 262-dof, multi-loop, and tree-like linkages including combinations of planar, prismatic and spherical joints. In contrast, existing methods either cannot be used for these problems or do not produce good quality solutions.

I. Introduction

Constrained motion planning problems are problems where a set of constraints are placed on the motion of an object (robot). These constraints could require that the robot remain in contact with a surface or maintain a specific clearance. They could also require that certain joints of the robot remain in contact with each other (e.g., closed chains). Motion planning with constraints is applicable to parallel robotics [13], grasping and manipulation [7, 15], computational biology [16], and animation [5].

Randomized motion planning methods such as the graph-based PRM [6] and the tree-based RRT [8] have had a good deal of success solving traditional motion planning problems. Unfortunately, these methods are poorly suited for constrained problems where the probability of randomly generating a sample satisfying the constraints approaches zero [9]. Previous methods have developed specialized samplers which generate samples that satisfy constraints [2, 4, 17]. Such samplers can be used in combination with existing PRM-based methods to solve problems with constraints. However, these methods are either unable to handle high degree of freedom (dof) systems or are unsuited for systems with spherical or prismatic joints.

In [12] we presented the concept of *reachable volumes* which can be applied to linkages, closed chains and tree-like robots with prismatic, spherical and planar joints. In comparison, most previous methods assume a planar robot with 1D articulated joints. Unlike many previous methods (e.g., cyclic coordinate decent [19], reachable distances [17], and inverse kinematics [2]) which focus on end effector constraints, reachable volumes allow for constraints on internal joints as well as end effectors and for multiple constraints to be applied simultaneously. We showed that the reachable volume of a chain is equivalent to the Minkowski sum of the reachable volumes of its links, which gives us an efficient method for computing reachable volumes. We also presented a family of samplers that use reachable volumes to generate constraint-satisfying configurations. We showed that these samplers were linear in the geometric complexity of the reachable volumes.

In this paper, we show that reachable volumes have an $\mathcal{O}(1)$ complexity for unconstrained problems as well as for many constrained problems. This allows us to compute the Minkowski sum of two reachable volumes in constant time, which in turn allows us to generate reachable volume samples in linear time. For problems with more complex constraints, we present an $\mathcal{O}(|L|^2\ |S|$ Complexity$(S))$ method for generating samples (where S is the set of constraints and $|L|$ is the number of bodies in the robot).

We perform an extensive experimental validation of the reachable volume samplers presented in [12]. Our results show that reachable volume sampling produces more valid samples than existing methods, that reachable volume samples are easier to connect than other samples, and that reachable volume sampling is more efficient at solving high dimensional problems than existing methods. Additionally, our results show that the running time of the sampler is linear with respect to the number of bodies in the robot, which confirms the theoretical bound shown previously.

The main contributions of this work include:

- proof of $\mathcal{O}(1)$ complexity of reachable volumes in unconstrained systems and in many constrained systems,
- proof of linear time complexity of reachable volume sampler in problems without constraints and
- evaluation of reachable volume samplers over a wide variety of systems including linkages, closed chains and tree-like robots with as many as 262 dof.

This research supported in part by NSF awards CNS-0551685, CCF-0833199, CCF-0830753, IIS-0916053, IIS-0917266, EFRI-1240483, RI-1217991, by NIH NCI R25 CA090301-11, by Chevron, IBM, Intel, Oracle/Sun and by Award KUS-C1-016-04, made by King Abdullah University of Science and Technology (KAUST). [1]Parasol Lab., Dept. of Comp. Sci. and Eng., Texas A&M Univ., College Station, Texas, USA. {tmcmahon,sthomas,amato}@cse.tamu.edu

II. Related Work

We give an overview of previous methods that are applicable to motion planning systems with constraints.

A. Sampling-based methods

A number of sampling-based motion planning methods have been proposed. Two of the most widely used methods are Probabilistic Roadmaps (PRMs) [6] and Rapidly-Exploring Random Trees (RRTs) [8]. While they have been applied to a wide variety of problems, they both have been shown to be poorly suited for problems with spatial constraints because the probability of randomly sampling constraint-satisfying configurations approaches zero [9].

Consequently, PRMs and RRTs have been adapted for use in spatially constrained systems. *Gradient decent methods* push randomly generated configurations onto a constrained surface [9, 20]. They are capable of solving problems with single-loop, articulated joint, closed chains. PRM-MC combines PRMs and Monte Carlo methods to generate samples that satisfy closure constraints for single loop closed chains up to 100 links [3]. In [18, 14], Trinkle and Milgram develop a method that uses C-space analysis for path planning while ignoring self collisions. They show results for a set of planar parallel star-shaped manipulators. Alternative Task-space and Configuration-space Exploration (ATACE) for path planning with constrained manipulators uses a randomized gradient decent method for constrained manipulators [21]. They present results for a 9 dof manipulator robot with a set of end effector constraints. In [22] Zhang et. al. present a Monte Carlo method for generating closed chain samples. This method uses analytical inverse kinematics to ensure that the sub-loops of closed chain robots are sampled in an unbiased manner and is shown to be applicable to 2D chains, closed chains and protein molecules with over 200 degrees of freedom.

B. Applications of Minkowski sums to motion planning

There have been a number of previous applications of Minkowski sums to motion planning. For example, the M-Sum Planner [10] is a hybrid motion planner that generates random samples for the angular coordinates of the environment, denoted as C-slices because they represent a slice of C-space in which the angular coordinates are fixed. For each C-slice, they compute the Minkowski sum of the robot and the obstacles in the environment. They then sample along the boundary of the Minkowski sum and connect the samples. Finally, they sort the C-slices and connect the nodes in nearby C-slices to form a roadmap. This method generates samples faster than biased samplers and nearly as fast as uniform sampling.

C. Reachable Distance and Reachable Volumes

The reachable distance of an articulated linkage is the range of distances that its end effector can reach with respect to its base [17]. The reachable distance of the linkage is computed by recursively computing the reachable distances of subsets of the linkage. This method efficiently produces samples for linkages, single and multiple loop closed chains, and constrained problems but is limited to planar joints.

In [12] we introduced *reachable volumes* (RV) which are a higher dimensional generalization of reachable distances. This work addresses the problem of motion planning with linkage robots. Formally, a linkage robot consists of a set of links L that are connected by a set of joints J. This work allows for spherical, planar and prismatic joints as well as combinations of different joint types. In contrast, most previous work assumes a planar robot comprised of planar (and sometimes prismatic) joints. For constrained problems we assume a set of constraints S=(S_1,...,$S_{|J|}$), where S_j is a subset of the environment in which joint j must be located.

This work defines RV-space to be a space of the same dimensionality as workspace in which the origin is located at one of the joints or end effectors of the robot (referred to as the root). RV-space has no obstacles and does not take validity into consideration. The reachable volume of a joint/end-effector is the set of points in RV-space that it can reach while satisfying the constraints in S. The reachable volume of a chain is the reachable volume of its end effector. Examples of reachable volumes are shown in Figure 1.

The reachable volume of a chain of links C is equal to the Minkowski sum of the reachable volumes of the links:

$$\texttt{Reachable}(C) = \texttt{Reachable}(l_1) \bigoplus \texttt{Reachable}(l_2)$$

$$\bigoplus \ldots \bigoplus \texttt{Reachable}(l_N)$$

We can therefore compute the reachable volume of C by computing the Minkowski sum of the reachable volumes of the links in C. The reachable volumes of more complicated robots like closed chains and tree-like graspers can be computed by decomposing the robot into chains.

We presented a method for sampling linkages, closed chains and tree-like robots by recursively placing joints in their reachable volumes. This method can be used in combination with sampling based methods such as the PRM. Our methods for computing reachable volumes and generating samples require $\mathcal{O}(|L|)$ Minkowski sum operations and are $\mathcal{O}(|L|)$ in the complexity of these operations.

III. Computing Minkowski Sums of Reachable Volumes

In this section we show that reachable volumes have an $\mathcal{O}(1)$ complexity for unconstrained problems as well as for certain constrained problems. We show that this enables us to compute the Minkowski sums of two reachable volumes, which combined with the methods presented in [12] allows us to generate samples in linear time with respect to the number of joints in the robot.

We first observe that the reachable volume of a linkage can be represented by a maximum value which represents the farthest distance from the origin that the linkage can reach and a minimum distance which represents the closet point to the origin that the end effector can reach (0 if it can reach the origin). For single link linkages, both the minimum

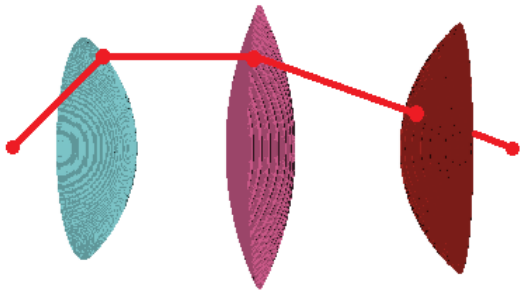

(a) The reachable volume of the first (teal), second (pink) and third (red) joints of a 4 link chain whose end effector is constrained.

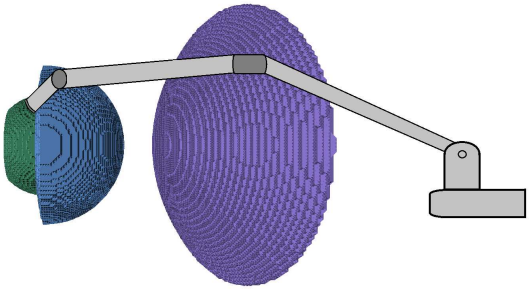

(b) The reachable volume of the knuckles (green), wrist (blue) and elbow (purple) of a robotic arm that is constrained to be grasping an object.

Fig. 1. Examples of reachable volumes.

and maximum values are equal to the length of the linkage. Using this representation, the reachable volume of a chain is the set of points whose distance from the origin is between these minimum and maximum values.

We next observe that for spherical, planar and (non-offset) prismatic joints, the reachable volume is the set of points between a specified minimum and maximum distance from the origin. These reachable volumes can be represented in constant space by storing the minimum and maximum distances. Consider a reachable volume $R1$ that is represented by the minimum value $R1_{min}$ and the maximum value $R1_{max}$ and a second reachable volume $R2$ that is represented by the minimum value $R2_{min}$ and the maximum value $R2_{max}$. The Minkowski sum of $R1$ and $R2$ can be represented by the the minimum value $(R1 \bigoplus R2)_{min}$ and the maximum value $(R1 \bigoplus R2)_{max}$ which are computed as follows:

$$(R1 \bigoplus R2)_{min} = \begin{cases} \max(R1_{min} - R2_{max}, 0) & \text{if } R1_{min} > R2_{min} \\ \max(R2_{min} - R1_{max}, 0) & \text{otherwise} \end{cases}$$

$$(R1 \bigoplus R2)_{max} = R1_{min} + R2_{min}$$

We observe that the Minkowski sum of $R1$ and $R2$ is also a reachable volume represented by a minimum and a maximum value. Inductively, we can conclude that the Minkowski sums of the reachable volumes of planar, prismatic and spherical joints will always be regions within a specified minimum and maximum distance from the origin. We also observe that the Minkowski sum of $R1$ and $R2$ can be computed in constant time regardless of how many joints and links are in the linkages that correspond to $R1$ and $R2$.

For constrained problems, the complexity of computing Minkowski sums depends on the geometry of the constraints S. To compute reachable volumes exactly, we must be able to compute Minkowski sums and intersections on the geometry. For problems where these computations are not feasible, we compute the reachable volume of the chain without constraints (using the method from [12]), then we separately compute the Minkowski sum of each constraint and the reachable volume of the portion of the chain C after the joint where the constraint S is applied:

$$RV(C, J, S) = RV_{0,|J|} \cap (S_1 \bigoplus RV_{1,|J|}) \cap \ldots$$

$$\cap (S_{|J|} \bigoplus RV_{|J|-1,|J|})$$

where $RV_{0,|J|}$ is the reachable volume of the chain C and $RV_{j,|J|}$ is the reachable volume of the portion of C after joint j (without constraints).

The result is a set of objects whose intersection is the reachable volume of the chain. Minkowski sum operations are commutative, so we can compute $RV_{j,|J|}$ first. Because $RV_{j,|J|}$ is the reachable volume of a linkage, its reachable volume is defined by 2 concentric spheres and can be computed as described above. The Minkowski sum of S_j and $RV_{j,|J|}$ is the Minkowski sum of S_j and the area between concentric circles, which can be computed in time proportional to the complexity of S_j. Computing reachable volumes using this method requires time of $\mathcal{O}(|J|\ |S|\ \text{Complexity}(S))$ time and $\mathcal{O}(|S|\ \text{Complexity}(S))$ space where $|J|$ is the number of joints, $|S|$ is the number of constraints and $\text{Complexity}(S)$ is the complexity of the constraints. Samples can therefore be generated in $\mathcal{O}(|J|^2\ |S|\ \text{Complexity}(S))$ time.

IV. Sampling in Intersection Reachable Volumes

The reachable volume sampler positions joints by placing them in their reachable volume given the joints that have already been sampled. This is done by placing the joint in the intersection of the reachable volumes of chains connecting it to joints that have already been sampled (we refer to these joints as *neighbors*). In this section we present a set of methods for generating samples in this intersection and discuss when each method is applicable.

The *intersection* method is applicable to reachable volumes that are the intersection of spheres. This method selects a random point along the circle formed by this intersection. This method is useful for sampling joints where two or more neighbors have already been sampled.

The *bounding patch* method is applicable to reachable volumes that are the intersection of a sphere-like reachable volume and a set of other reachable volumes. This method constructs a patch on the surface of the sphere that encompasses the intersection with the other reachable volumes. It then samples on this patch until it finds a joint that is in all of the other reachable volumes. This method is useful for joints where one neighbor has already been sampled.

The *bounding cube* method constructs an axis aligned bounding box around the reachable volume and then samples within this bounding box until it finds a sample that is in all of the reachable volumes. This method is used to sample joints where no neighbors have already been sampled.

We observe that each of these methods is complete in that they sample over a joint's entire reachable volume. Consequently, they can be used by the reachable volume samplers to provide probabilistically complete sampling. We also observe that the complexity of these methods is linear with respect to the number of reachable volumes involved.

V. Evaluation for Unconstrained Systems

We compare reachable volume sampling to uniform sampling [6] and an incremental sampling method, I-CD, which incrementally tests links along the chain starting at the base for collision before sampling the next link. I-CD detects invalid links as soon as they are sampled eliminating the need to sample the rest of the chain when collisions are found. For closed chains, we also compare to a Cyclic Coordinate Decent (CCD) sampler [19] which uses CCD to produce closed chain configurations. We demonstrate that our method can be applied to high dof linkages and closed chains, problems for which other methods are not suited. A more detailed description of I-CD and CCD is provided in [11] along with the full set of results.

A. Environments and robots studied

Walls. The walls environment (Figure 2(a)) is a commonly used benchmark which consists of 3 chambers separated by 2 walls. Holes in the walls allow the robot to travel between the chambers. In this environment we ran experiments using free flying chain linkages of varying dof (22–262).

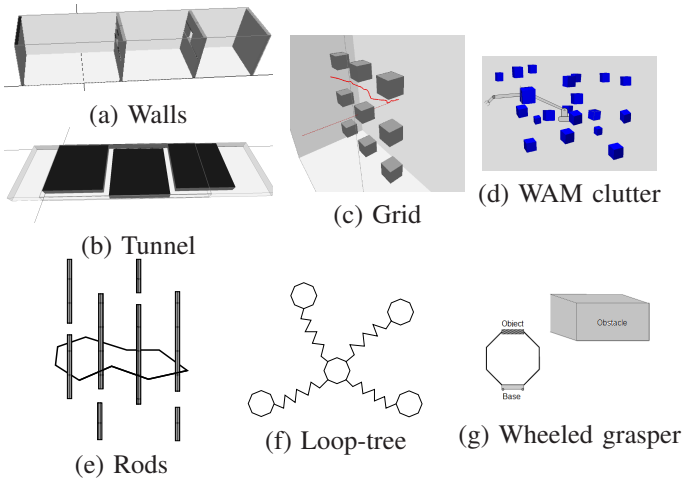

Fig. 2. Environments studied.

Tunnel. The tunnel environment (Figure 2(b)) is another commonly used benchmark which consists of 2 chambers connected by a long narrow tunnel. In this environment we ran experiments using free flying chains and closed chains of varying dof (22–262).

Grid. The grid environment (Figure 2(c)) consists of a set of cube obstacles arranged in a grid. Two types of robots are investigated: a 16 joint (32 dof) fixed-based linkage and two fixed-based tree-like robots comprising of an arm and a grasper formed by 2 subchains where one has 8 links in the arm and 4 links in each grasper yielding 32 dof and the other has 16 links in the arm and 8 links in each grasper yielding 64 dof.

WAM clutter. The WAM clutter (W-CL) environment (Figure 2(d)) consists of a Barrett WAM robotic arm surrounded by a clutter of obstacles. The WAM arm consists of a 6-dof arm with three graspers attached to it (total 15 dofs). This robot is interesting in that it includes both planar and spherical joints, demonstrating that our method is applicable to robots that include different types of joints.

Rods. The rods environment (Figure 2(e)) consists of 4 rods. Closed chains may enclose the rods and move onto different rods through breaks in them. In this environment we used 22 and 70 dof single loop closed chain.

Loop-tree robot. The loop-tree robot (Figure 2(f)) consists of an 8-link central loop with 4 8-link branches attached to it. At the end of each branch another 8-link loop is attached giving this robot a total of 5 loops and 160 dof. Experiments are run in a completely free environment.

Wheeled grasper. We finally study a wheeled robot with 2 graspers attached to it (Figure 2(g)). The graspers have spherical joints and need to transport an object under a low hanging environment, thus forming a closed chain. We study 19 dof and 67 dof variations of this environment.

B. Experimental Setup

In our experiments we construct roadmaps and study their quality and associated cost. For roadmap construction we use a PRM method with 2000 nodes and k-closest neighbor selection with a Scaled Euclidean distance metric and $k = 8$ for selecting node pairs to connect. In our linkages experiments we use a binary straight line local planner [6], while in our closed chain experiments we use a rotate-at-s local planner [1] with s = 0.5. All computation was performed on Brazos, a major computing cluster at Texas A&M University. The processing nodes consisted of quad-core Intel Xeon processors running at 2.5 Ghz, with 15 GB of RAM. All experiments had a maximum time allocation of 20 hours, and all results are averaged over 10 runs.

C. For linkages and tree-like robots

Figure 3 compares the performance of each method in generating and connecting roadmaps with 2000 samples for various environments with linkages and tree-like robots. Stars indicate methods that were unable to generate 2000 samples in the allotted 20 hours.

The sampler success rate (Figure 3(a)) is the proportion of samples that are valid (e.g., collision free). This indicates how efficient a method is at generating valid samples which can be used for roadmap construction. In lower dimensional problems, uniform sampling and I-CD have higher success rates than reachable volume sampling. However, as the

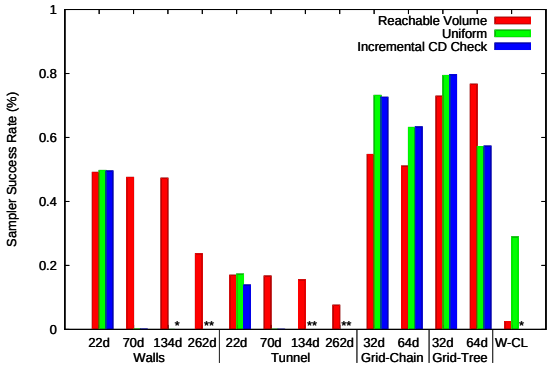

(a) Sampler Success Rate

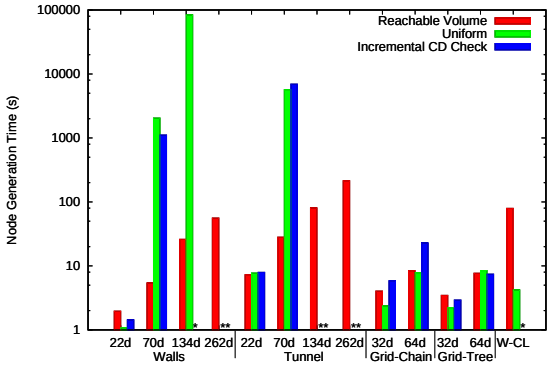

(b) Node GenerationTime

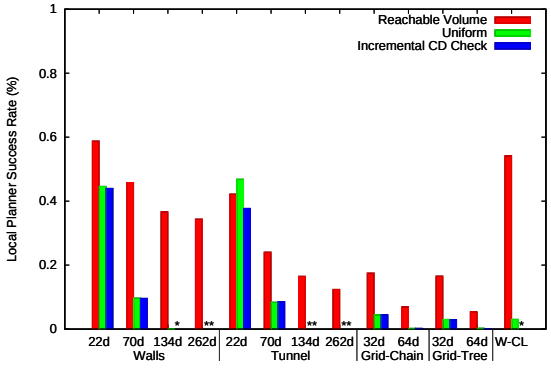

(c) Local Planner Success Rate

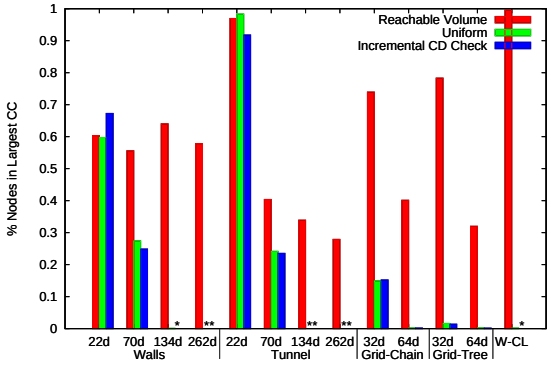

(d) % Nodes in Largest CC

Fig. 3. Experimental results for linkages and tree-like robots in various environments for 2000 samples. Stars indicate methods unable to generate samples in the allotted time. Note that (b) uses a log scale.

dof of the problem increases, reachable volume sampling outperforms the other methods and in some cases is the only method able to generate roadmaps in the allotted time. Interestingly, the success rate of reachable volume sampling does not significantly decrease with problem dimension.

Figure 3(b) provides the time required for each method to generate 2000 valid samples. We see that reachable volume sampling is slower than the others in lower dimensional problems such as the 22 dof linkages. However, in higher dimensional problems, the time is considerably better because reachable volume samples are less compact and thus less likely to have self-collisions which become more problematic as the robot complexity increases. In the highest dimensional problems shown, only reachable volume sampling was able to complete within the allotted time (20 hours).

Figure 3(c) shows the percentage of local planner calls that are successful. This directly determines the number of edges that can be added to the roadmap which in turn impacts how well connected it is. The local planner success rate for reachable volume sampling is consistently higher than the other methods which indicates that reachable volume sampling produces samples that are easier to connect. This is because the joint orientations of reachable volume samples are more uniform meaning that connecting them is less likely to result in self-collision. The performance difference in local planner success rate becomes more significant with increasing problem dimensionality. This trend is especially noticeable for tree-like robots where the local planner can fail because of collisions between the branches of the robot.

The size of a roadmap's largest connected component (CC) indicates how well connected it is. It also directly affects the number of different queries the roadmap can solve. Thus, roadmaps with larger percentages of samples in the largest CC are more desirable. Figure 3(d) shows that reachable volume sampling produce roadmaps with more samples in the largest CC than the other methods. This suggests that reachable volume sampling is doing a better job of finding connections between various areas of C-free, such as between the different chambers in the walls environment. This trend is particularly noticeable with the 70 dof linkage and the tree-like robots.

D. For closed chains

Figure 4 gives the performance of reachable volume sampling for robots containing closed chains. Again, 2000 valid, constraint-satisfying samples are created for each problem. Neither uniform sampling or I-CD are able to generate constraint-satisfying samples for any of the robots in the time allotted. Only reachable volume sampling can handle systems with spherical joints.

As expected, the sampler success rate decreases as problem complexity increases (Figure 4(a)), yet reachable volume sampling is still able to generate valid, constraint-satisfying samples for single loops up to 262 dof and complex robots like the loop-tree with 160 dof. This trend is echoed in the time required to generate such samples (Figure 4(b)).

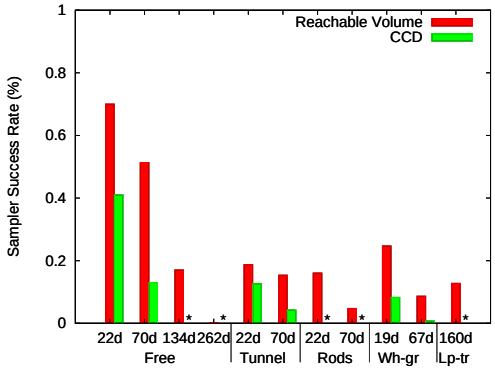

(a) Sampler Success Rate

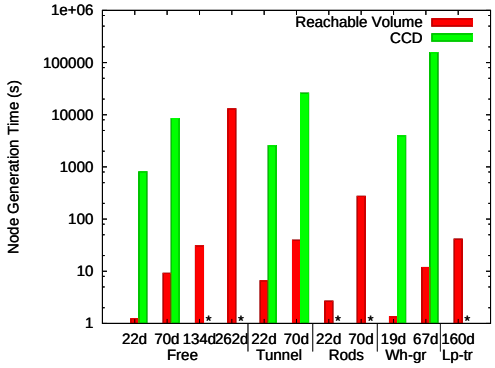

(b) Node GenerationTime

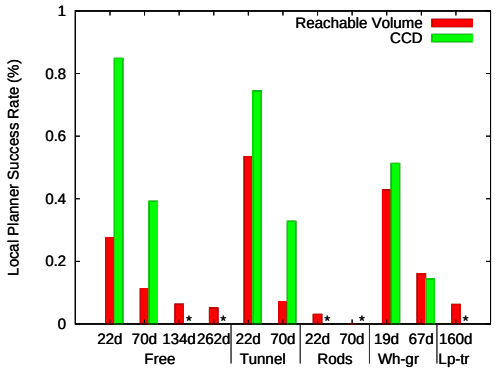

(c) Local Planner Success Rate

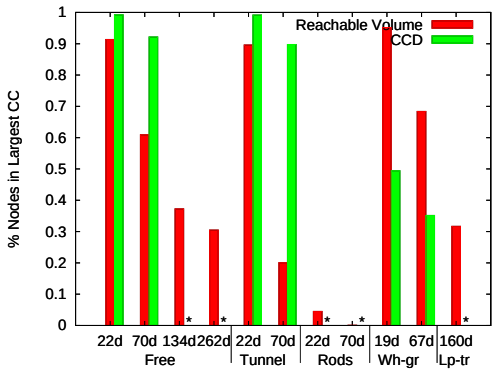

(d) % Nodes in Largest CC

Fig. 4. Reachable volume performance for closed chains in the following environments for 2000 samples: free, tunnel, rods, wheeled grasper (Wh-gr), and the loop-tree robot (Lp-tr). Uniform sampling and I-CD are not feasible for these robots. Note that plot (b) uses a log scale.

In comparison to CCD sampling, reachable volume sampling consistently produced samples that were more likely to be successful and required considerably less time to generate successful samples. In the rods and 134/262 dof free environments, the CCD sampler did not finish in the allotted time of 20 hours while the reachable volume sampler successfully generated samples and produced well connected roadmaps. Moreover reachable volume sampling can be applied to problems such as the loop-tree robot where CCD sampling is not applicable.

Figures 4(c) and 4(d) indicate that the reachable volume samples are successfully being connected to form significant connected components. The exception is the rods environment where there were very few successful connections which may point to the need for a more powerful local planner or a better distance metric to identify connectable samples. However, even in the loop-tree environment where the local planner success rate is low (0.13), we are still able to generate CCs which include nearly 1/3 of the samples.

CCD sampling did result in a higher local planner success than reachable volume sampling and in some environments it did produce roadmaps that were slightly more connected, however in most situations these differences would not justify the significantly higher running time of CCD.

VI. Evaluation for Constrained Systems

We evaluate the reachable volume sampler on a variety of environments with constraints.

A. Environments and robots studied

The **robot with cord** environment (Figure 5(a)) consists of a linkage robot with a cord attached to one of its joints. The robot's motion is constrained by the length of the cord, so that the distance between the joint and the base of the cord cannot exceed the length of the cord (light gray region). This scenario could be encountered when an industrial robot is operating a tool with an external power supply. We use 2 variations of this environment, one in which the robot consisted of 9 links (16 dof) with the cord attached to the 6th joint (cord-16) and another in which the robot consisted of 32 links (64 dof) with the cord attached to the 21st joint (cord-64). This environment demonstrates that our method can handle constraints on internal joints.

The **fixed base grasper** environment (Figure 5(b)) consists of a fixed base tree-like robot whose end effectors are constrained to be grasping one of the obstacles in the environment (green region). We include results for a 16-dof variations of this environment (gr-16) and a 32-dof variation (gr-34). This environment demonstrates that our method can be applied to grasping problems.

The **constrained closed chain** environment (Figure 5(c)) consists of a 22-dof closed chain (cc-22) or a 70-dof closed chain (cc-70). Our second environment set consists of closed chains ranging from 22-dof to 134-dof. Constraints are applied to 3 of the chain's joints so that these joints must always be within a small distance of each other.

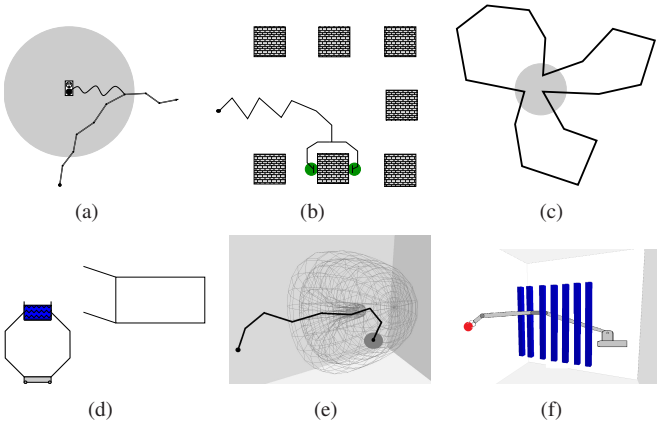

Fig. 5. Environments: (a) a fixed base arm with a cord attached to one of its joints. (b) A grasper with its end effectors constrained to positions where it is grasping one of the objects in the environment (green region). (c) A closed chain where 3 nodes are constrained to be near each other (shown in gray). (d) A wheeled grasper carrying a bucked that must remain level with the ground. (e) A linkage with its end effector fixed outside of the bug-trap and its end effector constrained to a region inside the bug-trap. (f) A WAM arm that must reach through a set of bars to grasp an object on the other side.

The **grasper with bucket** (Figure 5(d)) is a variation of the wheeled grasper environment in which the grasper is carrying a bucket that must remain level with the ground. We include results for a 22-dof variation of this environment (wb-22) and a 70-dof variation (wb-70).

The **bug-trap cleaner** (Figure 5(e)) environment is a variation of the bug-trap benchmark in which a fixed base robotic manipulator arm must clean out the bug-trap. The base of the arm is located outside of the bug-trap while the end effector is constrained to be inside. We perform experiments using 16-dof (bt-16) and 64-dof (bt-64) manipulators.

The **WAM bars** (w-b) environment (Figure 5(f)) consists of a Barrett WAM robotic arm which consists of a 6-dof arm with three graspers attached to it (total 15 dofs). The graspers are constrained to be grasping an object that is separated from the robot by a set of bars. The robot must reach through the bars to grasp the object.

B. Results

Our results show that reachable volume samples are more likely to be valid than samples produced by other methods (Figure 6(a)) and that the running time (Figure 6(b)) of reachable volume sampling is less than uniform sampling with filtering and significantly less than CCD. The sole exception is the bug-trap environment where uniform sampling with filtering has a lower running time; however the benefits in terms of connectivity and local planner success still make reachable volume sampling preferable in this environment.

We also show (Figure 6(c)) that the running time of the reachable volume sampler is relatively small compared to the overall roadmap construction time. The exception is the bug-trap environment where samples must be within the narrow passage of the bug-trap, resulting in a very low sampler success rate and a sampling time which dominates roadmap construction time regardless of the method used.

The largest CC size (Figure 6(d)) shows that reachable volume sampling consistently produces better connected roadmaps than the other methods. In many of the environments, reachable volume sampling achieved a largest CC size of close to 1, which indicates that these roadmaps are connected across the entire environment and that the problem has been solved. In comparison, CCD and uniform with filtering never had a largest CC size greater than 0.1, indicating they were unable to solve any of the environments.

In the grid and WAM environments, both uniform with filtering and CCD fail to generate roadmaps. This is interesting because it demonstrates one of the shortcomings of these methods. The grid and WAM environments both consists of tree-like graspers with constraints applied to their end effectors. While a method like CCD can converge to an end effector constraint, neither CCD nor uniform with filtering can ensure that the base of the fingers of the grasper is in a position where the other graspers can reach their associated constraints. These problems require a more powerful method, such as reachable volumes, which can position the base of the grasper in a position where all of the end effectors can reach their constraints.

VII. Conclusion

We demonstrate that reachable volume sampling can be applied to a wide variety of problems including high degree of freedom chains, tree-like linkages, closed chains, and combinations of these problems. We show results for constrained problems with as many as 70 dof and for unconstrained problems with as many as 262 dof. We show that reachable volume sampling produces more connectable samples faster than existing methods for constrained systems with spherical joints and with combinations of planar and spherical joints. In contrast, most previous methods either cannot be applied to these problems, do not produce quality solutions or have a significantly higher running time.

We show that reachable volume have a constant complexity in unconstrained problems as well as in many constrained problems. This allows us to perform Minkowski sum operations in constant time and generate samples in linear time.

References

[1] N. M. Amato, O. B. Bayazit, L. K. Dale, C. V. Jones, and D. Vallejo. Choosing good distance metrics and local planners for probabilistic roadmap methods. *IEEE Trans. Robot. Automat.*, 16(4):442–447, August 2000.

[2] L. Han and N. M. Amato. A kinematics-based probabilistic roadmap method for closed chain systems. In *New Directions in Algorithmic and Computational Robotics*, pages 233–246. A. K. Peters, Boston, MA, 2001. book contains the proceedings of the International Workshop on the Algorithmic Foundations of Robotics (WAFR), Hanover, NH, 2000.

[3] Li Han. Hybrid probabilistic roadmap — Monte Carlo motion planning for closed chain systems with spherical joints. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 920–926, New Orleans, LA, 2004.

[4] Li Han and Lee Rudolph. A unified geometric approach for inverse kinematics of a spatial chain with spherical joints. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 4420–4427, Roma, Italy, 2007.

[5] Marcelo Kallmann, Amaury Aubel, Tolga Abaci, and Daniel Thalmann. Planning collision-free reaching motion for interactive object manipulation and grasping. *Computer Graphics Forum*, 22(3):313–322, September 2003.

[6] L. E. Kavraki, P. Švestka, J. C. Latombe, and M. H. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Trans. Robot. Automat.*, 12(4):566–580, August 1996.

[7] Keith Kotay, Daniela Rus, Marsette Vona, and Craig McGray. The self-reconfiguring robotic molecule: Design and control algorithms. In Pankaj K. Agarwal, Lydia E. Kavraki, and Matthew T. Mason, editors, *Robotics: The Algorithmic Perspective*, pages 375–386. A. K. Peters, Boston, MA, 1998. book contains the proceedings of the International Workshop on the Algorithmic Foundations of Robotics (WAFR), Houston, TX, 1998.

[8] S. M. LaValle and J. J. Kuffner. Rapidly-exploring random trees: Progress and prospects. In *New Directions in Algorithmic and Computational Robotics*, pages 293–308. A. K. Peters, 2001. book contains the proceedings of the International Workshop on the Algorithmic Foundations of Robotics (WAFR), Hanover, NH, 2000.

[9] S.M. LaValle, J.H. Yakey, and L.E. Kavraki. A probabilistic roadmap approach for systems with closed kinematic chains. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 1671–1676, Detroit, MI, 1999.

[10] J.-M. Lien. Hybrid motion planning using minkowski sums. In *Proceedings of the Robotics: Science and Systems Confrence (RSS)*, June 2008.

[11] T. McMahon, S. Thomas, and N. M. Amato. Sampling based motion planning with reachable volumes. Technical Report TR13-008, Department of Computer Science, Texas A&M University, College Station Tx., 2013.

[12] T. McMahon, S. Thomas, and N. M. Amato. Sampling based motion planning with reachable volumes: Theoretical foundations. In *IEEE Int. Conf. Robot. Autom. (ICRA)*, 2014.

[13] Jean-Pierre Merlet. Still a long way to go on the road for parallel mechanisms. In *ASME Bienneal Mech. Rob. Conf.*, Montreal, Canada, 2002.

[14] R. James Milgram and Jeff C. Trinkle. The geometry of configuration spaces for closed chains in two and three dimensions. *Homology, Homotopy Appl.*, 6(1):237–267, 2004.

[15] A. Nguyen, L. J. Guibas, and M. Yim. Controlled module density helps reconfiguration planning. In *New Directions in Algorithmic and Computational Robotics*, pages 23–36. A. K. Peters, Boston, MA, 2001. book contains the proceedings of the International Workshop on the Algorithmic Foundations of Robotics (WAFR), Hanover, NH, 2000.

[16] Amit P. Singh, Jean-Claude Latombe, and Douglas L. Brutlag. A motion planning approach to flexible ligand binding. In *Int. Conf. on Intelligent Systems for Molecular Biology (ISMB)*, pages 252–261, 1999.

[17] X. Tang, S. Thomas, P. Coleman, and N. M. Amato. Reachable distance space: Efficient sampling-based planning for spacially constrained systems. *Int. J. Robot. Res.*, 29(7):916–934, 2010.

[18] Jeff C. Trinkle and R. James Milgram. Complete path planning for closed kinematic chains with spherical joints. *Int. J. Robot. Res.*, 21 (9):773–789, 2002.

[19] Li-Chun Tommy Wang and Chih Cheng Chen. A combined optimization method for solving the inverse kinematics problem of mechanical manipulators. *IEEE Trans. Robot. Automat.*, 7(4):489–499, 1991.

[20] Jeffery H. Yakey, Steven M. LaValle, and Lydia E. Kavraki. Randomized path planning for linkages with closed kinematic chains. *IEEE Trans. Robot. Automat.*, 17(6):951–958, 2001.

[21] Zhenwang Yao and K. Gupta. Path planning with general end-effector constraints: using task space to guide configuration space search. In *Proc. IEEE Int. Conf. Intel. Rob. Syst. (IROS)*, pages 1875–1880, Edmonton, Alberta, Canada, 2005.

[22] Y. Zhang, K. Hauser, and J. Luo. Unbiased, scalable sampling of closed kinematic chains. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2013.

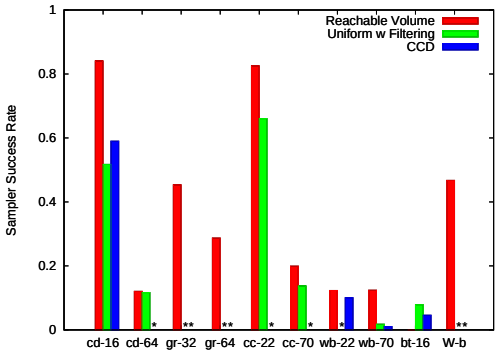

(a) Sampler Success Rate

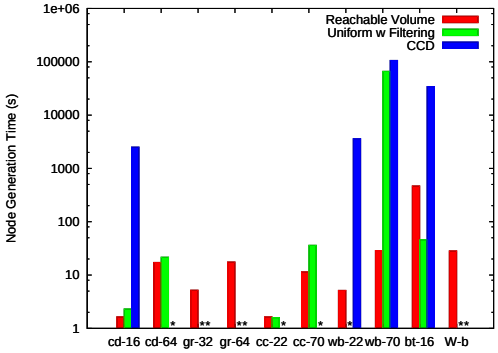

(b) Node Generation Time

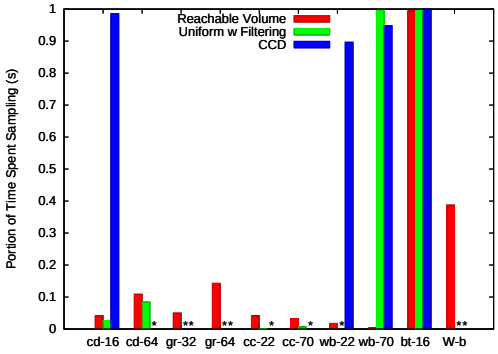

(c) Portion of construction time spent generating nodes

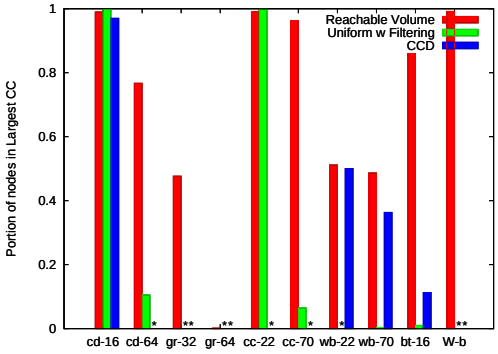

(d) % Nodes in Largest CC

Fig. 6. Experimental results for linkages and tree-like robots in various environments for 2000 samples. Stars indicate methods unable to generate samples in the allotted time or were not applicable. Note that (b) uses a log scale.