

# IMPROVING EFFICIENCY OF SAMPLING-BASED MOTION PLANNING VIA MESSAGE-PASSING MONTE CARLO

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Sampling-based motion planning methods, while effective in high-dimensional spaces, often suffer from inefficiencies due to irregular sampling distributions, leading to suboptimal exploration of the configuration space. In this paper, we propose an approach that enhances the efficiency of these methods by utilizing low-discrepancy distributions generated through Message-Passing Monte Carlo (MPMC). MPMC leverages Graph Neural Networks (GNNs) to generate point sets that uniformly cover the space, with uniformity assessed using the the  $\mathcal{L}_p$ -discrepancy measure, which quantifies the irregularity of sample distributions. By improving the uniformity of the point sets, our approach significantly reduces computational overhead and the number of samples required for solving motion planning problems. Experimental results demonstrate that our method outperforms traditional sampling techniques in terms of planning efficiency.

## 1 INTRODUCTION

Sampling-based motion planning is a key method for robot navigation in complex environments. From Probabilistic Roadmaps (PRMs) Kavraki et al. (1996) to Rapidly-exploring Random Trees (RRTs) LaValle & James J. Kuffner (2001), the focus of much of the research in this area has been on developing more efficient and optimal motion planning algorithms Janson et al. (2015); Karaman & Frazzoli (2011) that can handle high-dimensional spaces Orthey & Toussaint (2021; 2019), various types of constraints Karaman & Frazzoli (2013); Yerushova & LaValle (2009), and dynamic environments Phillips & Likhachev (2011); Otte & Frazzoli (2016).

Instead of refining motion planning algorithms, our work targets the fundamental building block underlying these approaches—the core sampling process itself. We introduce a novel sampling strategy based on Message-Passing Monte Carlo (MPMC) Rusch et al. (2024), a graph neural network architecture trained to generate low discrepancy point sets on unit hypercubes of arbitrary dimension. We expand the generic MPMC algorithm with the introduction of a novel training objective tailored to high-dimensional spaces, ensuring that generated points are optimally distributed and scalable.

This leads to a potent unbiased state sampling technique, that can be seamlessly integrated into any sampling-based planner, and that requires neither conditioning on the workspace description, nor a steer function, nor past examples, nor start and goal information. This approach provides strong theoretical guarantees and it also outperforms traditional techniques across various benchmarks, in environments of varying complexity and dimensionality.

Our contributions are stated as follows:

- We introduce the first application of MPMC neural network point set generation in motion planning, a significant novelty for unbiased sampling techniques.
- We propose a novel training objective tailored to high-dimensional configuration spaces, ensuring the generated points are well-suited for complex planning problems.
- We support our approach with rigorous theoretical justification, linking the training objective of MPMC point sets to a tighter upper bound on the distance from the optimal path.

- We establish superior planning efficiency against the current gold standard sampling approach in motion planning on a variety of PRM benchmarks, including challenging high-dimensional environments.
- We demonstrate the effectiveness of our sampling technique on a real-world UR5 robot arm, showing its potential for practical deployment in robotics.

Our MPMC-based sampling strategy offers a powerful and versatile alternative to traditional sampling methods. By targeting the sampling process directly, we open up new possibilities for improving the efficiency of a wide array of motion planning algorithms in both simulated and real-world robotic systems.

## 2 RELATED WORK

### 2.1 LEARNING FOR SAMPLING-BASED PLANNERS

The incorporation of machine learning into motion planning has opened up new ways to accelerate pathfinding by learning from past experiences. One common strategy involves storing and reusing previously computed paths or solutions. For example, methods such as path libraries Berenson et al. (2012), sparse roadmaps Coleman et al. (2014), and local obstacle roadmaps Lien (2010) allow a robot to retrieve and adapt previously successful solutions to new, but similar, planning problems. These approaches reduce the computation time by narrowing the search space using knowledge from past instances.

Another line of research enhances the sampling process by learning distributions that guide planners toward more promising regions of the configuration space. Some methods employ problem-invariant distributions Lehner & Albu-Schäffer (2018), while others adapt based on the workspace environment Chamzas et al. (2021). Deep learning further extends these concepts by learning from prior planning tasks, enabling distributions that condition on both the workspace and specific start-goal configurations of new problems Chamzas et al. (2022). These learned distributions effectively bias the sampling process, improving convergence rates and solution quality in new planning scenarios.

In contrast, our method departs from these past learning-based approaches by introducing a deterministic sampling strategy obtained via neural network training that does not depend on conditioning or prior knowledge.

### 2.2 LOW-DISCREPANCY CONSTRUCTIONS

Over the past century, numerous constructions of low-discrepancy point sets and sequences have been proposed. Most of the early constructions are deeply rooted in number theory and abstract algebra. A widely used building block of previous low-discrepancy constructions is the one-dimensional van der Corput sequence van der Corput (1935) in base  $b$ , which is generalized to a higher dimensional setting via the Halton sequence Halton (1960a). Each of the  $d$  coordinates in a Halton sequence corresponds to a distinct van der Corput sequence in base  $b$ , with the bases selected to be co-prime. Faure sequences, as introduced in Faure (1982), offer a similar construction to Halton sequences but incorporate permutations of the digits in base  $b$ .

Another broad class of low-discrepancy constructions are today known as digital  $(t, s)$ -sequences, first introduced in Niederreiter (1987). These constructions include the widely known Sobol' sequence Sobol' (1967a) which is constructed using tools from linear algebra involving primitive polynomials and well-chosen generating matrices defined over finite fields.

A further distinct approach, rooted in a different branch of number theory, emerged with Korobov's introduction of the good-lattice method Korobov (1963). This technique utilizes modular arithmetic and prime number properties to construct a structured, grid-like set of integration nodes. Since its introduction, lattice rules have undergone significant extensions and refinements. Comprehensive discussions and valuable reviews of these developments are available in Haber (1970); Sloan (1985); Nuyens (2014); Dick et al. (2022).

Several approaches have recently emerged targeting the construction of points for fixed dimension and number of points. In Doerr & De Rainville (2013), new low-discrepancy point sets were pro-

posed by optimizing permutations applied to the Halton sequence. Another approach, known as subset selection, was introduced to select  $k < N$  points from an  $N$ -element set that minimize the discrepancy. An exact algorithm for this selection was presented in Clément et al. (2022), while a swap-based heuristic approach was employed in Clément et al. (2024). Additionally, Clément et al. (2023) proposed a non-linear programming method to generate point sets with optimal star-discrepancy for fixed dimension and number of points. However, this approach faces significant computational challenges, limiting its practical application to finding optimal sets for only up to 21 points in two dimensions and 8 points in three dimensions. Message-Passing Monte Carlo (MPMC) falls into this class of constructions by optimizing the discrepancy of point sets that are fixed for a chosen dimension and number of points. However, MPMC significantly differs from all previous methods by its explicit use of machine learning. Moreover, it has been shown in Rusch et al. (2024) that MPMC generates point sets with significantly better distributional properties compared to any previous method, reaching optimal or near-optimal discrepancy.

### 3 METHODS

Message-Passing Monte Carlo (MPMC) is a machine learning approach designed to generate low-discrepancy point sets, which are the key for efficiently covering space in a uniform manner. MPMC leverages Graph Neural Networks (GNNs) and tools from Geometric Deep Learning to generate these point sets. The method focuses on the geometric properties needed to ensure uniformity and it is highly versatile for generating points across different dimensions. In this section we describe MPMC and its extensions for the motion planning application domain.

#### 3.1 MESSAGE-PASSING MONTE CARLO SAMPLING

Message-Passing Monte Carlo (MPMC) Rusch et al. (2024) leverages Graph Neural Networks (GNNs) to generate point sets that cover the space in a uniform manner. This uniformity can be assessed through measures of irregularity termed discrepancy. While there exist a plethora of different uniformity measures, we focus on the  $\mathcal{L}_p$ -discrepancy here. That is, given a set of points  $\{\mathbf{X}_i\}_{i=1}^N$  in the unit hypercube  $[0, 1]^d$  and  $p \geq 1$ , the  $\mathcal{L}_p$ -discrepancy is defined as,

$$\mathcal{L}_p^p(\{\mathbf{X}_i\}_{i=1}^N) := \int_{[0,1]^d} \left| \frac{\#\left(\{\mathbf{X}_i\}_{i=1}^N \cap [0, \mathbf{x}]\right)}{N} - \mu([0, \mathbf{x}]) \right|^p d\mathbf{x}, \quad (1)$$

where  $\#\left(\{\mathbf{X}_i\}_{i=1}^N \cap [0, \mathbf{x}]\right)$  counts how many points of  $\{\mathbf{X}_i\}_{i=1}^N$  fall inside the box  $[0, \mathbf{x}] = \prod_{i=1}^d [0, x_i]$  for  $\mathbf{x} = (x_1, \dots, x_d) \in [0, 1]^d$ , and  $\mu(\cdot)$  denotes the usual Lebesgue measure. A differentiable closed-form solution to the high-dimensional integral in (1) for the case of  $p = 2$ , known as Warnock’s formula Warnock (1972), can be computed in  $\mathcal{O}(N^2d)$ . Consequently, MPMC utilizes the  $\mathcal{L}_2$ -discrepancy as the objective function minimized during training to generate point sets with low discrepancy. The full MPMC model transforming random input points into low-discrepancy point sets via deep GNNs is depicted in Fig. 1.

##### 3.1.1 FAST MPMC TRAINING IN HIGHER DIMENSIONS

As described in Rusch et al. (2024), the  $\mathcal{L}_2$ -discrepancy fails to distinguish random from highly uniform points in high dimensions for moderate amount of points. Therefore, it was suggested in Rusch et al. (2024) to minimize the Hickernell  $\mathcal{L}_2$ -discrepancy  $D_{H,2}$  Hickernell (1998) in this case, that sums over all  $\mathcal{L}_2$ -discrepancies of  $k$ -dimensional projections, with  $1 \leq k \leq d$ . A straightforward implementation of  $D_{H,2}$  has a computational complexity of  $\mathcal{O}((2^d - 1)N^2d)$ , which would involve evaluating the  $\mathcal{L}_2$ -discrepancy for over a billion projections for  $d = 32$ , a relatively small number of points. To mitigate this issue, the authors in Rusch et al. (2024) proposed limiting the summation to a smaller, randomly selected subset of projections. In contrast, in this work we leverage the closed-form solution of the Hickernell  $\mathcal{L}_2$ -discrepancy from Joe (1997) that can be computed again

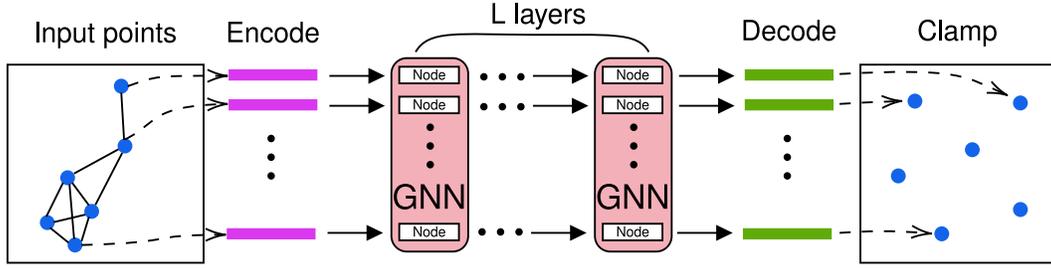


Figure 1: Schematic of the MPMC model reproduced from Rusch et al. (2024). First, random input points are encoded to a high dimensional representation. Second, the encoded representations are passed through a deep GNN, where the underlying computational graph is constructed based on nearest neighbors using the positions of the initial input points. Finally, the node-wise output representations of the final GNN layer are decoded and clamped back into the unit hypercube.

in  $\mathcal{O}(N^2d)$  instead of  $\mathcal{O}((2^d - 1)N^2d)$ . More concretely,

$$\begin{aligned}
 D_{H,2}^2(\{\mathbf{X}_i\}_{i=1}^N) &= \sum_{\emptyset \neq s \subseteq \{1, \dots, d\}} \mathcal{L}_2^2(\{\mathbf{X}_i^{(s)}\}_{i=1}^N) \\
 &= \left(\frac{4}{3}\right)^d - \frac{2}{N} \sum_{i=1}^N \prod_{k=1}^d \left(\frac{3}{2} - \frac{\mathbf{X}_{i,k}^2}{2}\right) + \frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N \prod_{k=1}^d [2 - \max(\mathbf{X}_{i,k}, \mathbf{X}_{j,k})],
 \end{aligned}$$

where  $\mathbf{X}_{i,k}$  is the  $k$ -th entry of  $\mathbf{X}_i$ , and  $\{\mathbf{X}_i^{(s)}\}_{i=1}^N$  is the projection of  $\{\mathbf{X}_i\}_{i=1}^N$  onto  $[0, 1]^{|s|}$ . This enables very fast training for generating MPMC points in higher dimensions that minimize the discrepancy over all projections.

### 3.2 MOTION PLANNING VIA PROBABILISTIC ROAD MAPS

In this work, the focus is on the efficiency of sampling-based motion planning in the sense of task success with respect to number of points sampled. Furthermore, for convenience in dealing with pre-trained sets of points of fixed number, we use a particular instantiation of the Probabilistic Roadmap (PRM) algorithm, which we refer to as ps-PRM (for pre-sampled PRM). This specific version of the PRM algorithm is described in Algorithm 1.

The algorithm begins by sampling a fixed number  $N$  of points from the entire space to form the set  $S$ . Next, we prune out milestones from the initial set  $S$  that do not fall in free space, resulting in a refined set of valid samples  $S_v$ , which is combined with the start and goal nodes

$\{\mathbf{X}_{\text{start}}, \mathbf{X}_{\text{goal}}\}$  to obtain the PRM graph vertex set  $V$ . For each node  $v$  in  $V$ , the algorithm identifies nearby nodes within a radius  $r_N$ , with the set of such points denoted  $P_{\text{near}}$ . For each neighbor  $p$  in  $P_{\text{near}}$ , the algorithm checks for collision-free paths between  $v$  and  $p$ . If the path is free of collisions, a bidirectional edge is added between  $v$  and  $p$  to the edge set  $E$ . Finally, the algorithm attempts to find a shortest path from  $\mathbf{X}_{\text{start}}$  to  $\mathbf{X}_{\text{goal}}$  and returns the path if one exists; otherwise, it indicates failure.

### 3.3 THEORETICAL GUARANTEE

Next, we outline a theoretical justification of minimizing discrepancy as a means to improving efficiency of sampling-based motion planning. To this end, we introduce another uniformity measure

---

#### Algorithm 1 ps-PRM Algorithm

---

- 1:  $S \leftarrow \text{Sample}(N)$
  - 2:  $S_v \leftarrow \text{PruneInvalid}(S)$
  - 3:  $V \leftarrow \{\mathbf{X}_{\text{start}}, \mathbf{X}_{\text{goal}}\} \cup S_v$
  - 4:  $E \leftarrow \emptyset$
  - 5: **for all**  $v \in V$  **do**
  - 6:      $P_{\text{near}} \leftarrow \text{Near}(V \setminus \{v\}, v, r_N)$
  - 7:     **for all**  $p \in P_{\text{near}}$  **do**
  - 8:         **if**  $\text{CollisionFree}(v, p)$  **then**
  - 9:              $E \leftarrow E \cup \{(v, p)\} \cup \{(p, v)\}$
  - 10:         **end if**
  - 11:     **end for**
  - 12: **end for**
  - 13: **return**  $\text{ShortestPath}(\mathbf{X}_{\text{start}}, \mathbf{X}_{\text{goal}}, V, E)$
-

216  
217  
218  
219  
220  
221  
222  
223  
224  
225  
226  
227  
228  
229  
230  
231  
232  
233  
234  
235  
236  
237  
238  
239  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
250  
251  
252  
253  
254  
255  
256  
257  
258  
259  
260  
261  
262  
263  
264  
265  
266  
267  
268  
269

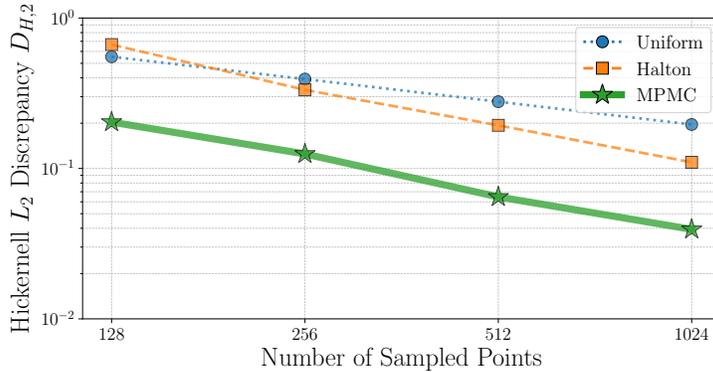


Figure 2: Hickernell  $\mathcal{L}_2$ -discrepancy in  $d = 10$  for Uniform, Halton and MPMC sampling.

known as dispersion, which is commonly used in assessing the efficiency of sampling-based motion planning. Given a point set  $\{\mathbf{X}_i^s\}_{i=1}^N$  in  $[0, 1]^d$  and  $p \geq 1$ , the  $l_p$ -dispersion is then defined as,

$$\mathcal{D}_p(\{\mathbf{X}_i^s\}_{i=1}^N) = \sup_{\mathbf{s} \in [0,1]^d} \min_{1 \leq i \leq N} \|\mathbf{s} - \mathbf{X}_i\|_p.$$

Dispersion is closely related to discrepancy through the following inequality established in Niederreiter (1992),

$$\mathcal{D}_\infty(\{\mathbf{X}_i^s\}_{i=1}^N) \leq \mathcal{L}_\infty(\{\mathbf{X}_i^s\}_{i=1}^N)^{1/d},$$

with the  $\mathcal{L}_\infty$ -discrepancy. Based on this and following Janson et al. (2018), one can provide a deterministic sampling guarantee. Namely, using the PRM planning algorithm, assuming some mild assumptions are satisfied (e.g., feasible path  $\delta$ -clearance), and choosing a radius  $r_N$  of the PRM algorithm based on  $N$  samples according to,

$$r_N = 2\alpha\sqrt{d}\mathcal{L}_\infty(\{\mathbf{X}_i^s\}_{i=1}^N)^{\frac{1}{d}},$$

for some  $\alpha > 1$  and  $r_N$  satisfying conditions corresponding to the path  $\delta$ -clearance, one can guarantee that the cost of the returned path is within a factor of  $\frac{1}{\alpha-1}$  of the optimal  $\delta$ -clear path. Clearly, a low discrepancy induces a guarantee of the resulting path to be close to the optimal  $\delta$ -clear path.

## 4 EXPERIMENTAL EVALUATION

We evaluate our motion planning solution in simulation, using a variety of environments ranging from 2D mazes to higher dimensional spaces and in a physical setting with a UR5 robot arm. We assess performance according to the number of sampled points and success rate in finding a path. Code for generating MPMC points can be found at <https://github.com/tk-rusch/MPMC>

### 4.1 SAMPLING METHODS

The performance of our method is evaluated against Uniform sampling and also against Halton Halton (1960b) and Sobol Sobol' (1967b) (for the 2-D maze setup only) sequences, two widely-used low-discrepancy Quasi-Monte Carlo methods. Sobol sequences are designed to efficiently cover spaces by minimizing gaps between sample points, making them ideal for complex problems. Halton sequences, on the other hand, generate points using prime number bases, providing effective uniform coverage. However, both Sobol and Halton sequences suffer from correlation issues in higher dimensions Rusch et al. (2024).

To obtain a statistical measure of performance across repeated runs, we randomize the selection of the MPMC points from the pool of trained point sets (batches can contain 8, 16 or 32 point sets based on the training instance). Similarly for the deterministic Halton and Sobol sequences, we initialize a new sequence where the previous one ended. Hence, when sampling point sets of size  $N$ , we take



Figure 3: Visualization of the 2-D maze, SE(3) and UR5 simulation experiments. 2D mazes of increasing difficulty are presented on the first row: level 1 (left), level 2 (middle) and level 3 (right). The start position is depicted in green and the end goal in orange. The SE(3) cubicle experiment is on the bottom left, with the inverted L shaped object required to navigate the maze, then through the passage to the other side of the board and back to reach the goal configuration (in the bottom right corner). The SE(3) Twistycool setup is in the middle of the bottom row, the red object starting on the left has to rotate its way through the passage in the wall to reach its goal. The 5-D UR5 robot experiment is on the bottom right, it shows the arm reaching into a green box while avoiding contact with the table or box.

the points indexed from  $(i - 1) N + 1$  to  $i * N$  for the  $i$ -th run, with  $i$  taking integer values from 1 to the number of runs fixed for the experiment.

To illustrate differences between the sampling techniques considered in the context of this work, we provide the Hickernell  $\mathcal{L}_2$ -discrepancies of point sets of sizes  $\{128, 256, 512, 1024\}$  in 10 dimensions in Figure 2. Halton sequences never surpass a 2 times advantage over Uniform sampling, and can even offer worse discrepancy for smaller  $N$ . MPMC points, on the other hand, reach substantially lower Hickernell  $\mathcal{L}_2$ -discrepancy, consistently close to 3 times lower than that of Halton and achieving up to a 5-fold improvement over Uniform sampling for larger point sets.

## 4.2 EXPERIMENTS

### 4.2.1 2-D MAZE

The initial experiments consist in using PRM to navigate a series of three 2-D mazes of increasing difficulty depicted in Figure 3. These are designed only to be increasingly difficult to solve for PRM with Uniform sampling before comparing to the other sampling techniques to avoid introducing any bias favoring a specific scheme. The maps are of size 640 by 480 and the agent is assumed to be a disk of radius 6 for collision checking. A single start and end goal is considered and 50 planning attempts are run by level, number of sampled points and sampling technique. The shortest path solver used is A\* Hart et al. (1968).

### 4.2.2 OMPL BENCHMARKS

All of the Uniform, Halton and MPMC sampling schemes are put to the test on experiments from the popular Open Motion Planning Library (OMPL) Moll et al. (2015). The base code PRM implementation is only modified to account for our fixed number pre-sampling. We ensure sufficient compute time to reach failure and run 50 iterations per sampling method, per number of points sampled and per scenario described below:



Figure 4: Demonstration of a motion plan on hardware. The goal is to reach into the box without making contact.

**SE(3) rigid body puzzles** The Special Euclidean Group in 3 Dimensions  $SE(3)$  is composed of a 3-D translational component and a 3-D rotational component. The latter is often conveniently sampled in quaternion space, with uniform sampling on the 4D unit sphere manifold. In keeping with standard practice, we evaluate samplers’ efficient coverage only on the translational 3-D Euclidean space component. Two such puzzles, Cubicle and Twistycool, are benchmarked from the OMPL. For the latter example, we have reduced the default bounds to allow this harder problem to be solved with the same list of  $n$  values used throughout. Both examples can be visualized in Figure 3

**$d$ -D hypercube corridor** On the  $d$ -dimensional hypercube, the valid region is defined such that there exists an index  $k$  with the following constraints: for all dimensions  $i < k$ , the  $i$ -th coordinate must be less than or equal to a threshold edge width  $\lambda$ , and for all dimensions  $i > k$ , it must be greater than or equal to  $1 - \lambda$ . This results in a valid subspace that resembles narrow passageways along the hypercube edges, leading from one corner of the hypercube to the diagonally opposite corner. For dimension  $d \in \{2, 3, 10\}$ , the value of the edge width is tuned to ensure the point sets available are of sufficient size to establish a meaningful comparison between samplers (respectively  $\lambda \in \{0.1, 0.2, 0.37\}$ ).

**10-D kinematic chain** The kinematic chain experiment features a robot with multiple links, operating in a 10-dimensional configuration space. The robot must navigate through an environment with obstacles represented as line segments. Initially, the robot’s first link is at zero radians, and subsequent links are arranged with a specific angular offset. The goal is to move the robot from this starting position to a target configuration where the first link aligns nearly with a desired orientation, while avoiding collisions with the environment and itself.

#### 4.2.3 UR5 ROBOT ARM

The UR5 robot arm is a popular collaborative robot with a 6 degree of freedom workspace. In this benchmark, our goal is to reach into a sideways box without contacting the table or the box. We forgo an end effector and thus we do not utilize the last DOF of the robot (which rotates the end effector and in this setting does not effect the workspace of the manipulator). Simulation, sampling, and planning is implemented in the Klamp’t software package. We compare our sampling method to sampling from the uniform distribution. For a visualization of the task in simulation see Fig. 3. We also demonstrate our results on real hardware.

## 5 RESULTS

### 5.1 2-D MAZE

The performance of the different sampling techniques for all three levels of the 2-D mazes are presented in Figure 5. On all three levels and for all fixed number of points, MPMC offers the best value per sampled point set size on all but two instances (outperformed by Halton by 2% on level 1 with  $N = 32$  and by Sobol by 8% on level 3 with  $N = 128$ ). Indeed, MPMC’s efficiency superiority is all the more established on the harder level, where it can reach close to 90% with only 256 points sampled, a feat that takes Halton and Sobol twice the number of points to reach, and a performance level Uniform sampling fails to reach even with 4 times that number.

378  
 379  
 380  
 381  
 382  
 383  
 384  
 385  
 386  
 387  
 388  
 389  
 390  
 391  
 392  
 393  
 394  
 395  
 396  
 397  
 398  
 399  
 400  
 401  
 402  
 403  
 404  
 405  
 406  
 407  
 408  
 409  
 410  
 411  
 412  
 413  
 414  
 415  
 416  
 417  
 418  
 419  
 420  
 421  
 422  
 423  
 424  
 425  
 426  
 427  
 428  
 429  
 430  
 431

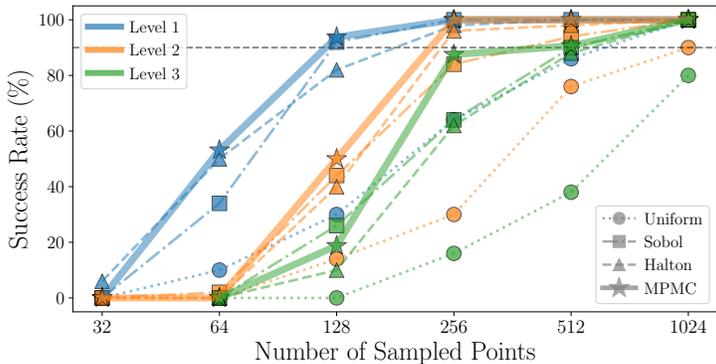


Figure 5: Success rates of PRM on the 3 levels of 2-D maze motion planning versus the number of points sampled. Results for each level are grouped by color and the data by sampler is identified by marker and line styles.

In fact, Uniform sampling is substantially less efficient as it requires sampling at least 4 times the number of points of low discrepancy methods to reach the 90% mark on all levels. Elsewhere, Sobol and Halton sequences seem to offer comparable performance overall, with relative ranking seemingly contingent on the level and number of points sampled. They often fail to meet the performance of MPMC with significant margins (e.g. 34% success for Halton on level 1 with  $N = 64$  against MPMC’s 54%, under 64% for both Halton and Sobol on level 3 with  $N = 256$  to MPMC’s 88%).

## 5.2 OMPL AND UR5 BENCHMARKS

The benchmarking success rates are provided in Table 1, along with the average number of valid points sampled by each scheme and its standard deviation.

On the SE(3) Cubicle and Twistycool experiments, MPMC globally achieves the best performance, although with somewhat marginal advantage. This tightness of margins can be linked to the importance of the SO(3) angle quaternion sampling which is done uniformly in all setups. Halton and Uniform schemes appear to be relatively more contingent on the scenario and are harder to clearly rank.

However, MPMC outperforms both Uniform and Halton samplers more consistently and significantly when used on all the samples’ components. Indeed, over the 2 and 3 dimensional hypercube examples, MPMC sampling solves around twice as many runs as its best competitor with 128 and 256 point regimes and maintains an advantage in the region of 50% with larger point sets.

The previous observation remains, to a solid extent, valid in higher dimensions. Although relative performance gains are in multiple cases well reduced, MPMC sampling maintains top-performing status. Furthermore, there remains instances where it offers large performance gains.

This is evident on the 10-D kinematic chain test, where MPMC reaches 24% success rate with 256 samples, compared to only 12% for Uniform and 14% for Halton. Similarly, a significant performance gap is established between MPMC (80%) and its counterparts (Halton 56% and Uniform 62%) with 512 points on the 10-D hypercube benchmark. In the realistic scenario of the UR5 planning task (5-D), MPMC consistently surpasses Uniform and Halton sampling (3 to 8-fold higher success rates). We also demonstrate that our approach is readily deployed to real hardware (Figure 4).

## 6 DISCUSSION

This work introduces the application of MPMC neural network point set generation in motion planning, offering a novel, unbiased approach to sampling. Our proposed method includes a custom training objective specifically designed for high-dimensional configuration spaces, ensuring that the

Table 1: Benchmark results on the OMPL suite and UR5 robot arm comparing performance in terms of success rate (SR in %) versus  $N$  the number of points sampled  $\{128, 256, 512, 1024\}$ . Also provided are the mean and standard deviation of the number of valid milestones in each experiment (denoted  $|V|$ )

Experiment	Space	Sampler	128		256		512		1024	
			SR (%)	$ V $	SR (%)	$ V $	SR (%)	$ V $	SR (%)	$ V $
Cubicle	SE(3)	Uniform	0	$56.4 \pm 5.0$	0	$112.9 \pm 7.2$	2	$221.8 \pm 13.6$	26	$412.6 \pm 60.8$
		Halton	0	$57.4 \pm 3.4$	0	$112.3 \pm 4.8$	4	$221.1 \pm 8.9$	36	$382.4 \pm 98.9$
		MPMC	0	$58.8 \pm 3.3$	0	$112.3 \pm 3.6$	10	$221.2 \pm 14.5$	46	$367.4 \pm 100.5$
Twistycool	SE(3)	Uniform	8	$89.2 \pm 14.7$	4	$180.4 \pm 22.3$	26	$332.7 \pm 71.6$	70	$485.0 \pm 221.7$
		Halton	6	$89.9 \pm 10.8$	16	$171.4 \pm 30.2$	20	$336.0 \pm 63.0$	56	$584.1 \pm 193.4$
		MPMC	10	$89.5 \pm 11.5$	8	$175.2 \pm 27.6$	30	$327.7 \pm 67.1$	72	$499.3 \pm 212.7$
Hypercube	$\mathbb{R}^2$	Uniform	24	$15.8 \pm 13.5$	24	$19.1 \pm 24.1$	44	$56.4 \pm 50.0$	60	$104.1 \pm 85.6$
		Halton	24	$7.9 \pm 11.1$	36	$19.2 \pm 23.5$	36	$47.0 \pm 48.9$	36	$67.5 \pm 90.7$
		MPMC	48	$15.4 \pm 12.4$	60	$31.6 \pm 24.4$	68	$69.3 \pm 45.5$	84	$161.0 \pm 71.0$
Hypercube	$\mathbb{R}^3$	Uniform	24	$7.9 \pm 7.5$	34	$17.1 \pm 14.9$	44	$29.5 \pm 27.4$	54	$64.1 \pm 51.5$
		Halton	20	$5.3 \pm 6.5$	32	$15.0 \pm 14.2$	34	$27.6 \pm 27.3$	28	$42.4 \pm 52.3$
		MPMC	56	$10.3 \pm 6.3$	64	$22.5 \pm 12.7$	76	$40.6 \pm 22.6$	96	$103.1 \pm 20.1$
Hypercube	$\mathbb{R}^{10}$	Uniform	2	$7.0 \pm 2.0$	12	$9.7 \pm 2.5$	62	$17.7 \pm 4.3$	88	$33.9 \pm 5.5$
		Halton	2	$6.6 \pm 1.8$	12	$10.4 \pm 2.3$	56	$17.7 \pm 2.9$	98	$32.6 \pm 3.4$
		MPMC	0	$6.3 \pm 1.7$	16	$10.7 \pm 2.4$	80	$18.9 \pm 1.9$	100	$31.1 \pm 2.9$
Kinematic Chain	$\mathbb{R}^{10}$	Uniform	4	$30.2 \pm 4.8$	12	$57.5 \pm 9.6$	22	$103.5 \pm 26.3$	54	$172.5 \pm 64.9$
		Halton	2	$30.0 \pm 4.7$	14	$55.5 \pm 9.1$	22	$103.7 \pm 21.4$	62	$156.0 \pm 64.5$
		MPMC	6	$31.5 \pm 4.3$	24	$51.8 \pm 7.5$	26	$102.3 \pm 21.5$	60	$161.7 \pm 71.9$
UR5	$\mathbb{R}^5$	Uniform	0	$47.6 \pm 5.8$	10	$92.2 \pm 7.8$	6	$184.3 \pm 11.5$	12	$368.6 \pm 16.1$
		Halton	2	$49.4 \pm 6.9$	4	$93.1 \pm 7.5$	14	$185.8 \pm 10.8$	18	$369.2 \pm 15.6$
		MPMC	6	$49.1 \pm 3.0$	31	$97.4 \pm 3.5$	75	$181.4 \pm 3.8$	81	$360.6 \pm 4.3$

generated points are optimized for complex planning tasks. Also, we provide theoretical proof that the MPMC point sets offer a tighter upper bound on the distance from the optimal path.

Through extensive experiments on PRM benchmarks, we demonstrate that our method significantly improves planning efficiency over commonly used sampling methods, especially in high-dimensional and challenging environments. Furthermore, we validate the real-world applicability of our technique by successfully implementing it on a UR5 robot arm, highlighting its potential for deployment in practical robotic systems.

One limitation of our approach is the need to retrain the MPMC model for each specific number of points  $N$  and dimensionality  $d$  to ensure optimal performance. This retraining requirement can be computationally intensive, particularly for applications where the configuration space frequently changes. Future work will focus on refining the neural network architecture and training procedure to achieve generalization across both  $N$  and  $d$ , reducing the need for retraining and making the method more flexible across different planning scenarios.

Moreover, while we have demonstrated the benefits of MPMC within the PRM framework, future research will explore the potential advantages of integrating this technique into more sophisticated sampling-based planners, potentially extending its impact across a broader class of planning algorithms. Finally, an exciting avenue for future exploration lies in adapting our technique to compute-critical applications, such as acrobatic flight or high-speed driving, where planners like the Model Predictive Path Integral (MPPI) Williams et al. (2017) require real-time performance and could greatly benefit from efficient, high-quality sampling methods.

## REFERENCES

- Dmitry Berenson, Pieter Abbeel, and Ken Goldberg. A robot path planning framework that learns from experience. In *2012 IEEE International Conference on Robotics and Automation*, pp. 3671–3678, 2012. doi: 10.1109/ICRA.2012.6224742.
- Constantinos Chamzas, Zachary Kingston, Carlos Quintero-Pena, Anshumali Shrivastava, and Lydia E. Kavraki. Learning sampling distributions using local 3d workspace decompositions for motion planning in high dimensions. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, May 2021. doi: 10.1109/icra48506.2021.9561104. URL <http://dx.doi.org/10.1109/ICRA48506.2021.9561104>.

- 486 Constantinos Chamzas, Aedan Cullen, Anshumali Shrivastava, and Lydia E. Kavraki. Learning to  
487 retrieve relevant experiences for motion planning, 2022. URL [https://arxiv.org/abs/  
488 2204.08550](https://arxiv.org/abs/2204.08550).
- 489 François Clément, Carola Doerr, and Luís Paquete. Star discrepancy subset selection: problem  
490 formulation and efficient approaches for low dimensions. *J. Complexity*, 70:Paper No. 101645,  
491 34, 2022. ISSN 0885-064X,1090-2708.
- 492 François Clément, Carola Doerr, and Luís Paquete. Heuristic approaches to obtain low-discrepancy  
493 point sets via subset selection. *J. Complexity*, 83:Paper No. 101852, 2024. ISSN 0885-064X,1090-  
494 2708.
- 495 François Clément, Carola Doerr, Kathrin Klamroth, and Luís Paquete. Constructing optimal  $\mathcal{L}_\infty$   
496 star discrepancy sets. *Preprint*, 2023. <https://arxiv.org/abs/2311.17463>.
- 497 David Coleman, Ioan A. Sucan, Mark Moll, Kei Okada, and Nikolaus Correll. Experience-based  
498 planning with sparse roadmap spanners, 2014. URL [https://arxiv.org/abs/1410.  
499 1950](https://arxiv.org/abs/1410.1950).
- 500 Josef Dick, Peter Kritzer, and Friedrich Pillichshammer. *Constructions of Lattice Rules*, pp. 95–139.  
501 Springer International Publishing, Cham, 2022.
- 502 Carola Doerr and François-Michel De Rainville. Constructing low star discrepancy point sets with  
503 genetic algorithms. In *Proceedings of the 15th Annual Conference on Genetic and Evolutionary  
504 Computation*, pp. 789–796, New York, NY, USA, 2013. Association for Computing Machinery.  
505 ISBN 9781450319638.
- 506 Henri Faure. Discrepance de suites associées à un système de numération (en dimension  $s$ ). *Acta  
507 Arithmetica*, 41:337–351, 1982.
- 508 S. Haber. Numerical evaluation of multiple integrals. *SIAM Review*, (12):481–526, 1970.
- 509 John H. Halton. On the efficiency of certain quasi-random sequences of points in evaluating multi-  
510 dimensional integrals. *Numer. Math.*, 2:84–90, 1960a.
- 511 John H. Halton. On the efficiency of certain quasi-random sequences of points in evaluating  
512 multi-dimensional integrals. *Numerische Mathematik*, 2:84–90, 1960b. URL [https://api.  
513 semanticscholar.org/CorpusID:122173470](https://api.semanticscholar.org/CorpusID:122173470).
- 514 Peter Hart, Nils Nilsson, and Bertram Raphael. A formal basis for the heuristic determination of  
515 minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2):100–107,  
516 1968. doi: 10.1109/tssc.1968.300136. URL [https://doi.org/10.1109/tssc.1968.  
517 300136](https://doi.org/10.1109/tssc.1968.300136).
- 518 Fred Hickernell. A generalized discrepancy and quadrature error bound. *Mathematics of computa-  
519 tion*, 67(221):299–322, 1998.
- 520 Lucas Janson, Edward Schmerling, Ashley Clark, and Marco Pavone. Fast marching tree: a fast  
521 marching sampling-based method for optimal motion planning in many dimensions, 2015. URL  
522 <https://arxiv.org/abs/1306.3532>.
- 523 Lucas Janson, Brian Ichter, and Marco Pavone. Deterministic sampling-based motion planning:  
524 Optimality, complexity, and performance. *The International Journal of Robotics Research*, 37(1):  
525 46–61, 2018.
- 526 Stephen Joe. *Formulas for the computation of the weighted L2 discrepancy*. Citeseer, 1997.
- 527 Sertac Karaman and Emilio Frazzoli. Sampling-based algorithms for optimal motion planning.  
528 *CoRR*, abs/1105.1186, 2011. URL <http://arxiv.org/abs/1105.1186>.
- 529 Sertac Karaman and Emilio Frazzoli. Sampling-based optimal motion planning for non-holonomic  
530 dynamical systems. *2013 IEEE International Conference on Robotics and Automation*, pp. 5041–  
531 5047, 2013. URL <https://api.semanticscholar.org/CorpusID:9370488>.

- 540 L.E. Kavraki, P. Svestka, J.-C. Latombe, and M.H. Overmars. Probabilistic roadmaps for path plan-  
541 ning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*,  
542 12(4):566–580, 1996. doi: 10.1109/70.508439.
- 543 N.M. Korobov. Number-theoretic methods of approximate analysis. *Fizmatgiz, Moscow*, 1963. In  
544 Russian.
- 545 Steven M. LaValle and Jr. James J. Kuffner. Randomized kinodynamic planning. *The International*  
546 *Journal of Robotics Research*, 20(5):378–400, 2001. doi: 10.1177/02783640122067453. URL  
547 <https://doi.org/10.1177/02783640122067453>.
- 548 Peter Lehner and Alin Albu-Schäffer. The repetition roadmap for repetitive constrained motion  
549 planning. *IEEE Robotics and Automation Letters*, 3(4):3884–3891, 2018. doi: 10.1109/LRA.  
550 2018.2856925.
- 551 Jyh-Ming Lien. *Planning Motion in Environments with Similar Obstacles*, pp. 89–96. 07 2010.  
552 ISBN 9780262289801. doi: 10.7551/mitpress/8727.003.0013.
- 553 Mark Moll, Ioan A. Şucan, and Lydia E. Kavraki. Benchmarking motion planning algorithms: An  
554 extensible infrastructure for analysis and visualization. *IEEE Robotics & Automation Magazine*,  
555 22(3):96–102, September 2015. doi: 10.1109/MRA.2015.2448276.
- 556 Harald Niederreiter. *Random number generation and quasi-Monte Carlo methods*. SIAM, 1992.
- 557 Harold Niederreiter. Point sets and sequences with small discrepancy. *Monatshefte für Mathematik*,  
558 104(4):273–337, 1987.
- 559 Dirk Nuyens. *The construction of good lattice rules and polynomial lattice rules*, pp. 223–256. De  
560 Gruyter, Berlin, Boston, 2014.
- 561 Andreas Orthey and Marc Toussaint. Rapidly-exploring quotient-space trees: Motion planning using  
562 sequential simplifications, 2019. URL <https://arxiv.org/abs/1906.01350>.
- 563 Andreas Orthey and Marc Toussaint. Sparse multilevel roadmaps for high-dimensional robotic mo-  
564 tion planning. In *2021 IEEE International Conference on Robotics and Automation (ICRA 2021)*,  
565 pp. 7851–7857, Piscataway, NJ, 2021. IEEE. doi: 10.1109/ICRA48506.2021.9562053.
- 566 Michael W. Otte and Emilio Frazzoli. Rrtx: Asymptotically optimal single-query sampling-based  
567 motion planning with quick replanning. *The International Journal of Robotics Research*, 35:797  
568 – 822, 2016. URL <https://api.semanticscholar.org/CorpusID:20802196>.
- 569 Mike Phillips and Maxim Likhachev. Sipp: Safe interval path planning for dynamic environments.  
570 *2011 IEEE International Conference on Robotics and Automation*, pp. 5628–5635, 2011. URL  
571 <https://api.semanticscholar.org/CorpusID:16210040>.
- 572 T Konstantin Rusch, Nathan Kirk, Michael M Bronstein, Christiane Lemieux, and Daniela Rus.  
573 Message-passing monte carlo: Generating low-discrepancy point sets via graph neural networks.  
574 *Proceedings of the National Academy of Sciences*, 121(40):e2409913121, 2024.
- 575 Ian H. Sloan. Lattice methods for multiple integration. *Journal of Computational and Applied*  
576 *Mathematics*, 12-13:131–143, 1985. ISSN 0377-0427.
- 577 I.M Sobol’. On the distribution of points in a cube and the approximate evaluation of integrals. *USSR*  
578 *Computational Mathematics and Mathematical Physics*, 7(4):86–112, 1967a. ISSN 0041-5553.
- 579 I.M Sobol’. On the distribution of points in a cube and the approximate evaluation of in-  
580 tegrals. *USSR Computational Mathematics and Mathematical Physics*, 7(4):86–112, 1967b.  
581 ISSN 0041-5553. doi: [https://doi.org/10.1016/0041-5553\(67\)90144-9](https://doi.org/10.1016/0041-5553(67)90144-9). URL <https://www.sciencedirect.com/science/article/pii/0041555367901449>.
- 582 J.G. van der Corput. Verteilungsfunktionen i–ii. *Proc. Akad. Amsterdam*, 38:813–821, 1058–1066,  
583 1935.
- 584 Tony T Warnock. Computational investigations of low-discrepancy point sets. In *Applications of*  
585 *number theory to numerical analysis*, pp. 319–343. Elsevier, 1972.

594 Grady Williams, Andrew Aldrich, and Evangelos A. Theodorou. Model predictive path integral  
595 control: From theory to parallel computation. *Journal of Guidance, Control, and Dynamics*,  
596 40(2):344–357, 2017. doi: 10.2514/1.G001921. URL [https://doi.org/10.2514/1.](https://doi.org/10.2514/1.G001921)  
597 G001921.

598  
599 Anna Yershova and Steven M. LaValle. Motion planning for highly constrained spaces. In  
600 Krzysztof R. Kozłowski (ed.), *Robot Motion and Control 2009*, pp. 297–306, London, 2009.  
601 Springer London. ISBN 978-1-84882-985-5.

602  
603  
604  
605  
606  
607  
608  
609  
610  
611  
612  
613  
614  
615  
616  
617  
618  
619  
620  
621  
622  
623  
624  
625  
626  
627  
628  
629  
630  
631  
632  
633  
634  
635  
636  
637  
638  
639  
640  
641  
642  
643  
644  
645  
646  
647