

Shortcutting Revisited: Improved Sampling and Irreplaceable Optimality

Emil F. Awad¹, Hannah Erickson¹, Alexander J. LaValle¹, Nicoletta Prencipe², Başak Sakçak^{1,3}, Kalle G. Timperi¹, and Steven M. LaValle¹

¹ Faculty of Information Technology and Electrical Engineering
University of Oulu, Finland

² CNRS, Institut de Recherche Mathématique de Rennes,
University of Rennes, France

³ Dept. of Advanced Computing Sciences, Maastricht University, the Netherlands

Abstract. This paper addresses the problem of optimizing collision-free paths that are typically produced by sampling-based motion planning algorithms. The well-known method of shortcutting attempts to optimize a given path by iteratively picking a pair of points along it at random and checking whether it can be replaced by a straight line or geodesic while avoiding collisions. A novel concept of *irreplaceably optimal path* is introduced, and convergence to such paths is studied. Irreplaceable optimality is a general notion that applies well beyond the current setting, and it is distinct from other concepts such as local optimality or optimality within a homotopy class. Improvements over prior shortcutting methods are also proposed, including skipping redundant optimizations and replacing uniform random sampling with sample sequences that provide deterministic convergence guarantees. Experiments were done for both 2D illustrative cases and multilink robots with up to 20 degrees of freedom.

Keywords: path optimization, sampling-based motion planning, geodesics, functional analysis, shortest paths, motion and path planning

1 Introduction

Jagged solution paths are a well-known nuisance of using sampling-based motion planning (SBMP) algorithms. The collision-free paths produced by standard algorithms such as RRT-Connect [12] or PRMs [11] technically solve the problem of connecting initial and goal configurations with the paths traversing the free subset of the configuration space, but the paths are typically unusable in raw form. Asymptotically optimal SBMP approaches, such as RRT*[10] or BIT*[5], address this issue by gradually improving the entire search graph of collision free paths. Although in theory the methods yield the shortest possible paths as the running time tends to infinity, it is common in practice to quickly post-process and optimize the first computed path so that it becomes more usable in practice, thereby avoiding the burden of improving all paths in the search graph.

Shortcutting is a preferred method of improving path quality due to its efficiency and implementation ease [1,6,9]. In each iteration, the method picks two

points at random along the path and attempts to replace the path between them by a *geodesic* (straight line). If the geodesic path segment passes collision checking, then the path is updated. Iterations continue for a fixed number of steps or until a termination condition is met that has tracked how much the path has been improving in recent iterations.

Shortcutting can be seen as a two-dimensional sampling operation: by selecting an initial and a final point along a path and connecting them with a geodesic. Sampling these two points is independent of the dimension of the configuration space in which the path lies, but checking whether its collision-free is not. We found that discarding possible shortcuts when they are already a part of the path can be done without checking if its collision-free, which considerably reduces the running time.

Alternatively, not considered here, shortcutting methods are occasionally combined with standard path planning in a hybrid approach [8,16]. There are, of course, many other ways to improve paths incrementally, such as energy-based methods [20,21] or functional optimization [2], and the goal may be some other desirable property like smoothness rather than shorter path lengths [19]. Notably, in [4], gradient descent among an increasing set of linear constraints was used to gradually shorten the length of a path. The authors of [4] noted that plain shortcutting leads to initial improvements more quickly than their gradient-based method, but that the gradient method provided more benefit during the final optimizations.

The current paper is motivated by two basic questions:

1. Is randomly picking interval endpoints the best way to sample?
2. What does shortcutting actually converge to?

In spite of the widespread use of shortcutting for several decades, relatively little is known about these seemingly obvious questions. This is surprising considering the popularity of SBMP and the need to quickly improve their computed paths.

To address the first question, we formally define the problem of sampling intervals in $[0, 1]$, the normalized domain of the path $\tau : [0, 1] \rightarrow \mathcal{C}_{\text{free}}$. We consider both random and deterministic sampling methods, and progress measures such as dispersion, much in the same way they were considered for PRM and grid-search comparisons in [14]. We introduce several alternatives to the standard approach of randomly picking interval endpoints, and experimentally evaluate them on computed examples in both 2D for illustrative purposes, and in dimensions up to 100, to establish their generality. Most importantly, we introduce interval sampling methods that are far superior to the usual approach. Many open questions remain in this context.

Addressing the second question is surprisingly complicated. In many topologies it is evident that repeated shortcutting will not converge to anything, and even when the sequence converges in relatively simple cases like \mathbb{R}^d , the limit path may have counter-intuitive properties. The upshot is that denseness of the shortcutting sequence alone seems not to guarantee convergence, although requiring stronger uniformity seems to improve the situation. We introduce a new notion called *irreplaceable optimality*. This is a necessary condition of optimal

paths, similar to Pontryagin’s maximum principle [3,15], but expressed in terms of shortest paths without differential constraints and with global obstacles that cause nonconvexity of $\mathcal{C}_{\text{free}}$. This new notion of optimality is neither equivalent to global optimality nor local optimality within a homotopy class (shortcutting permits jumping into a ‘better’ class). We establish that if the sampling is syndetically dense then repeated shortcutting of a path in \mathbb{R}^d results in either an irreplaceably optimal path or fails to be irreplaceably optimal in a specific manner, to which we refer as boundary-optimizable.

2 Problem Formulation and Definitions

In broad strokes, we take as input a valid, collision-free path leading from a start to goal configuration, and return a shorter path between the start and goal by recursively choosing two points on the path and attempting to replace the subpath between them with a geodesic.

2.1 Properties of the Configuration Space

Let \mathcal{C} be the *configuration space* of a robot. In the simplest case, $\mathcal{C} \subset \mathbb{R}^d$ with $n \geq 2$, but generally it may be any topological manifold. The obstacles are represented by an open subset \mathcal{O} of the robot workspace \mathcal{W} , which can be a subset of \mathbb{R}^2 or \mathbb{R}^3 . Let $\mathcal{R}(q) \subseteq \mathcal{W}$ be the space occupied by the robot at configuration $q \in \mathcal{C}$. The configurations that result in collisions give the set $\mathcal{C}_{\text{obs}} := \{q \in \mathcal{C} \mid \mathcal{R}(q) \cap \mathcal{O} \neq \emptyset\}$. We may assume that \mathcal{C}_{obs} is an open set.¹ The *free configuration space* $\mathcal{C}_{\text{free}} := \mathcal{C} \setminus \mathcal{C}_{\text{obs}}$ is then the closed set of configurations in which the robot does not intersect with any obstacles.

A *path* $\sigma : [0, 1] \rightarrow \mathcal{C}$ connecting a start and a goal configuration $q_0 = \sigma(0) \in \mathcal{C}$ and $q_1 = \sigma(1) \in \mathcal{C}$, respectively, is a continuous mapping to the robot configuration space. A path is *collision-free* if $\sigma(\tau) \in \mathcal{C}_{\text{free}}$ for all $\tau \in [0, 1]$. We will denote with $\sigma|_{[s_1, s_2]}$ the restriction of the path σ to the interval $[s_1, s_2] \subseteq [0, 1]$. We denote by \mathcal{P} the set of paths in \mathcal{C} . In addition, we denote by $[a, b]$ the straight line segment connecting two points $a, b \in \mathcal{C}$.

We consider configuration spaces that are topological manifolds equipped with a distance metric $d : \mathcal{C} \times \mathcal{C} \rightarrow [0, \infty)$ compatible with the manifold topology. The length of a path is defined as

$$\ell(\sigma) = \sup_{\{\tau_0=0 < \tau_1 < \dots < \tau_n=1; n \in \mathbb{N}\}} \sum_{i=1}^n d(\sigma(\tau_i), \sigma(\tau_{i-1})), \quad (1)$$

in which the supremum is taken over all finite partitions of $[0, 1]$. For a path restricted to the interval $[s_1, s_2]$, the supremum in (1) is taken over all finite partitions of $[s_1, s_2]$ and we will denote its length by $\ell(\sigma|_{[s_1, s_2]})$. Without loss of generality, we will assume that all the paths are parametrized using a *constant speed parametrization* which means $\ell(\sigma|_{[a, b]}) = (b - a)\ell(\sigma)$ for $0 \leq a \leq b \leq 1$.

¹ This is the case, for example, if the sets $\mathcal{R}(q)$ depend continuously on q in the Hausdorff metric.

Definition 1 (interval space). *The interval space is the set $\mathcal{I} = \{(s_1, s_2) \in [0, 1]^2 \mid s_1 < s_2\}$.*

A path σ is a *geodesic* if $d(\sigma(t_i), \sigma(t_j)) = |t_i - t_j|d(\sigma(0), \sigma(1))$ for all $t_i, t_j \in [0, 1]$. Suppose the distance metric is length-induced, that is, it satisfies $d(q_0, q_1) = \inf\{\ell(\sigma) \mid \sigma(0) = q_0 \text{ and } \sigma(1) = q_1\}$. Then, a geodesic exists and it is a shortest path connecting two configurations. However, in general, we cannot guarantee uniqueness. Uniqueness is guaranteed, for example, for geodesics that are shortest paths in Euclidean spaces where the distance metric is given by the Euclidean norm. In this paper, we focus on geodesics that are length-induced such that they are the shortest paths connecting two points in \mathcal{C} .

To avoid problems with non-rectifiable paths, that is, paths with infinite length, and problems with chattering, we assume that $\partial\mathcal{C}_{\text{obs}}$ is piecewise real-analytic. Furthermore, we assume that \mathcal{C}_{obs} is a *collared space*, which means that there exists some embedding (homeomorphism) $f : \partial\mathcal{C}_{\text{obs}} \times [0, 1] \rightarrow \mathcal{C}_{\text{free}}$ with $f(q, 0) = q$. The latter assumption is needed to avoid having collision-free paths that only have zero clearance.

2.2 Shortcut Definitions

We gather here the different definitions and assumptions that have been considered regarding the proof of convergence for repeated shortcutting. We will focus on convergence in \mathbb{R}^d under the Euclidean norm. Geodesics in \mathbb{R}^d are straight line segments. For any two points $x_1, x_2 \in \mathbb{R}^d$, we denote by $[x_1, x_2]$ the line segment connecting x_1 to x_2 . Definitions 2 and 3 are illustrated in Figure 1.

Definition 2 (candidate shortcut, shortcut). *Let $\sigma : [0, 1] \rightarrow \mathcal{C}$ be a path, let $s = (s_0, s_1) \in \mathcal{I}$, and let $\gamma_s : [0, 1] \rightarrow \mathcal{C}$ be a geodesic with $\gamma_s(t) = \sigma(s_t)$ for $t \in \{0, 1\}$. Then γ_s is*

- (1) a candidate shortcut, if $\sigma|_{[s_0, s_1]}$ is not a geodesic, and
- (2) a shortcut, if it is a candidate shortcut and $\gamma_s([0, 1]) \subset \mathcal{C}_{\text{free}}$.

Denote by $B_\varepsilon(A)$ the open ε -neighbourhood around the set A . Also, denote by \overline{A} and $\text{int}(A)$ the closure and interior of the set A , respectively.

Definition 3 (candidate shortcut point, shortcut point, extremal shortcut point, limit shortcut point). *Let $\sigma : [0, 1] \rightarrow \mathcal{C}$ be a path, let $s = (s_0, s_1) \in \mathcal{I}$, and let $\gamma_s : [0, 1] \rightarrow \mathcal{C}$ be a geodesic with $\gamma_s(t) = \sigma(s_t)$ for $t \in \{0, 1\}$. Then*

- (1) s is a candidate shortcut point, if γ_s is a candidate shortcut,
- (2) s is a shortcut point, if γ_s is a shortcut,
- (3) s is an extremal shortcut point, if γ_s is a shortcut and $B_\varepsilon(\gamma_s([0, 1])) \not\subset \mathcal{C}_{\text{free}}$ for all $\varepsilon > 0$, and
- (4) s is a limit shortcut point, if $s \in \overline{\text{int}(S_\sigma)}$,

in which S_σ denotes the set of shortcut points of σ .




Fig. 1. (top left) A path σ from q_0 to q_1 that navigates around obstacles. **(bottom left)** The pair (e, f) is not a candidate shortcut point because the geodesic between $\sigma(e)$ and $\sigma(f)$ is already part of the path. (a, h) is a candidate shortcut point but not a shortcut point because the geodesic between $\sigma(a)$ and $\sigma(h)$ is blocked. (d, g) is a shortcut point but not a limit shortcut point because even though the geodesic between $\sigma(d)$ and $\sigma(g)$ is unblocked, the tri-tangent ensures that there is no open set in the interval space that has (d, g) as a limit point. (b, c) is a limit shortcut point. **(right)** A depiction of the interval space for σ . Grey regions indicate where the geodesic is blocked by an obstacle. Blue regions are areas where the geodesic is already in σ . The blue dots indicate the geodesics that make up σ . The points (a, h) , (d, g) , and (e, f) are shown. The set S consists of limit shortcut points and contains (b, c) .

We now define the notion of *irreplaceably optimal paths* which intuitively captures the desired property of a limit path obtained via shortcutting. It turns out however, that it is theoretically possible for a shortcutting sequence not to converge to such a path. In such cases, there exist ways to shortcut the limit path after the infinite shortcutting sequence has converged, although these shortcuts cannot be made during the actual shortcutting phase (after a finite number of steps). An illustration of this situation is given in Figure 4 in Section 3 below.

Definition 4 (irreplaceably optimal path). A path $\sigma \in \mathcal{P}$ is irreplaceably optimal if σ has no limit shortcut points, i.e., $\text{int}(S_\sigma) = \emptyset$.

Definition 5 (boundary-optimizable path). A path $\sigma \in \mathcal{P}$ is boundary-optimizable at $s \in \mathcal{I}$, if $s \in \text{int}(S_\sigma)$ and s is an extremal shortcut point.

We now define *replacement*, the process by which paths are iteratively shortened by shortcutting.

Definition 6 (replacement). In configuration spaces where geodesics between two points are unique, the replacement $R(\sigma, s)$ of a path σ at a shortcut point




Fig. 2. Effect of the choice of initial shortcuts to the limiting path. **(a)** An initial path (blue) and two possible shortcuts (green, yellow) are shown. **(b)** Taking either of the shortcuts in (a) makes it impossible to converge to the optimal path (red).

$s = (s_1, s_2) \in \mathcal{I}$ is defined as the substitution of $\sigma|_{[s_1, s_2]}$ with the geodesic between $\sigma(s_1)$ and $\sigma(s_2)$ and reparametrization to maintain constant speed.

Note that in particular σ and $R(\sigma, s)$ have the same start and end points, and $\ell(R(\sigma, s)) < \ell(\sigma)$.

Given a sequence of points in the interval space, the process of repeated shortcutting proceeds as per Algorithm 1. This yields a sequence of paths that becomes monotonically shorter.

Algorithm 1 Generation of the shortcutting sequence.

```

1:  $i = 0, j = 0, (a_n)_{n=0}^\infty \in \mathcal{I}$ 
2: loop
3:   if  $a_i$  is a candidate shortcut point of  $\sigma_j$  then
4:     if  $a_i$  is a shortcut point of  $\sigma_j$  then           ▷ Checks if  $\gamma_{a_i}$  is collision-free
5:        $\sigma_{j+1} \leftarrow R(\sigma_j, a_i)$                  ▷ Replacement function
6:       yield  $\sigma_{j+1}$                                    ▷ Incrementally create the sequence of paths
7:        $j \leftarrow j + 1$ 
8:    $i \leftarrow i + 1$ 

```

3 Convergence in Euclidean Space

Shortening an initial path via iterative shortcutting will generally not lead to the shortest possible collision-free path from the start configuration to the goal configuration. Each shortcut restricts the set of possible paths that can be obtained at the limit (see Figure 2). We conjecture that every shortcutting sequence converges uniformly in the limit. A uniformly convergent subsequence is guaranteed by the Arzelà-Ascoli Theorem [22] owing to the fact that any shortcutting sequence is uniformly equicontinuous in the convex hull of the initial path.

We consider convergence under two properties of the sampling sequence.

Definition 7 (non-negligent random sampling). Let $(X_n)_{n=0}^\infty$ be a sequence of random variables, each taking values in \mathcal{I} . We say that $(X_n)_{n=0}^\infty$ is non-negligent in \mathcal{I} if for every open set $A \subset \mathcal{I}$ there exists some $p(A) > 0$ with probability $\mathbb{P}(X_n \in A) \geq p(A)$ for all $n \in \mathbb{N}$.

Definition 8 (syndetically dense sequence). A sequence $(a_n)_{n=0}^\infty \in \mathcal{I}$ is syndetically dense, if for every open set $A \in \mathcal{I}$ there exists some $r(A) \in \mathbb{N}$ such that every finite subsequence of consecutive entries $(a_k, a_{k+1}, \dots, a_m)$ contains at least one point in A whenever $m - k \geq r(A)$.

The Halton sequence is syndetically dense (see Section 4 and the appendix). Note that a sequence that is dense in \mathcal{I} may not be syndetically dense.




Fig. 3. For any $\varepsilon > 0$, the shortest curve α that connects $\sigma_k(t_1)$ and $\sigma_k(t_2)$ lying distance $L_{k+1} := \|\sigma_k(t_1) - \sigma_k(t_2)\|$ apart, while exiting the open tube $B_\varepsilon(\Gamma_{k+1}(t))$, must travel through some point z equidistant from these points, and has length $L_\alpha = (L_{k+1}^2 + 4\varepsilon^2)^{1/2}$. Then any curve φ of fixed length $\ell(\varphi) = L_k$ with the same start and end points and maximal sup-norm relative to $\widehat{\sigma}_{k+1}$ must satisfy $\varepsilon_\varphi := \|\varphi - \widehat{\sigma}_{k+1}\|_\infty = \frac{1}{2}(L_k^2 - L_{k+1}^2)^{1/2}$. In particular this is true for the restriction of σ_k onto $[t_1, t_2]$.

Lemma 9. Let $(\sigma_n)_{n \in \mathbb{N}} \subset \mathbb{R}^d$ be a sequence of paths, generated by iteratively choosing shortcuts in \mathcal{I} . Then $\lim_{n \rightarrow \infty} \|\sigma_{n+1} - \sigma_n\|_\infty = 0$.

Proof. Consider two consecutive paths σ_k and σ_{k+1} and let $s = (s_1, s_2) \in \mathcal{I}$ be the shortcut point for which $\sigma_{k+1} = R(\sigma_k, s)$. For $n \in \{k, k+1\}$, let $\Gamma_n(s) := \sigma_n([s_1, s_2])$ so that $\Gamma_{k+1}(s)$ is the straight line segment corresponding to the shortcut. $D_{k+1} := L_k^2 - L_{k+1}^2$. Let $\varphi : [s_1, s_2] \rightarrow \mathbb{R}^d$ be a curve of length L_k , satisfying $\varphi(s_i) = \sigma_k(s_i)$ for $i \in \{1, 2\}$, and denote by $\widehat{\sigma}_{k+1}$ the restriction $\sigma_{k+1}|_{[s_1, s_2]}$. Then φ maximizes the supremum norm $\varepsilon_\varphi := \|\varphi - \widehat{\sigma}_{k+1}\|_\infty$ precisely when the point $z = \varphi(\frac{s_1+s_2}{2})$ is equidistant from the two endpoints, i.e. $\|z - \sigma_k(s_1)\| = \|z - \sigma_k(s_2)\|$, and in addition satisfies $\text{dist}(z, \Gamma_{k+1}(s)) = \frac{1}{2}(L_k^2 - L_{k+1}^2)^{1/2} = \varepsilon_\varphi$, see Figure 3. Thus, $\varepsilon_{k+1} \leq \frac{1}{2}(L_k^2 - L_{k+1}^2)^{1/2} \rightarrow 0$ as $k \rightarrow \infty$ and the result follows. ■

Denote by $d_{\mathcal{H}}(A, B)$ the Hausdorff distance between $A, B \subset \mathbb{R}^d$.

Lemma 10. Let $(\sigma_n)_{n \in \mathbb{N}} \subset \mathbb{R}^d$ be a sequence of paths, generated by iteratively choosing shortcuts in \mathcal{I} . For every $\varepsilon > 0$ there exists some $\delta(\varepsilon) > 0$ with the property that for all $q \in \mathbb{N}$ and all $\delta < \delta(\varepsilon)$ there exists some $K(q) \in \mathbb{N}$ such that the interval $I_s := [s - \delta, s + \delta]$ satisfies $d_{\mathcal{H}}(\sigma_n(s), \sigma_{n+m}(I_s)) < \varepsilon$ for all $s \in (\delta, 1 - \delta)$, all $m \in \{1, \dots, q\}$ and $n \geq K(q)$.

Proof. Let $\varepsilon > 0$. Since the lengths $\ell(\Gamma_n)$ of the curves $\Gamma_n = \sigma_n([0, 1])$ form a monotonically decreasing sequence, and each σ_n is parametrized with a constant-speed parametrisation on $[0, 1]$, it follows that there exists some $N \in \mathbb{N}$ and $\delta(\varepsilon) > 0$ for which $\|\sigma_n(t) - \sigma_n(s)\| < \varepsilon/2$ whenever $|t - s| < \delta(\varepsilon)$ and $n \geq N$ (uniform equicontinuity of the family $\{\sigma_n\}$).

Now, let $q \in \mathbb{N}$ and $\delta < \delta(\varepsilon)$. Lemma 9 implies that there exists some $K(q) \geq N$ for which $\|\sigma_{n+1} - \sigma_n\|_\infty < \varepsilon/2q$ for all $n \geq K(q)$. Then for all $s \in (\delta, 1 - \delta)$ and $t \in [s - \delta, s + \delta]$ and $n \geq K(q)$ and $p \in \{1, \dots, q\}$ we have

$$\begin{aligned} \|\sigma_{n+p}(t) - \sigma_n(s)\| &\leq \sum_{j=1}^p \|\sigma_{n+j}(t) - \sigma_{n+j-1}(t)\| + \|\sigma_n(t) - \sigma_n(s)\| \\ &< q(\varepsilon/2q) + (\varepsilon/2) = \varepsilon. \end{aligned} \quad \blacksquare$$

In Lemma 11 and Proposition 12 below, we assume that $(\sigma_n)_{n \in \mathbb{N}} \subset \mathbb{R}^d$ is obtained via a sequence of shortcuts $(s^n)_{n \in \mathbb{N}}$ that are generated either by non-negligent random sampling or by syndetically dense sampling in \mathcal{I} .

Lemma 11. *Assume that a subsequence $(\sigma_{n_k})_{k \in \mathbb{N}}$ converges uniformly to some path σ and that $s = (s_1, s_2)$ is a non-extremal shortcut point for σ . Then for any $\varepsilon > 0$ there exists a subsequence of shortcuts $(s^{m_k})_{k \in \mathbb{N}}$ for which $\sigma_{m_k}(s_i^{m_k}) \in B_\varepsilon(\sigma(s_i))$ for all $k \in \mathbb{N}$ and $i = 1, 2$. In the case of probabilistic sampling, the result holds almost surely.*

Proof. Since s is a non-extremal shortcut point, we may assume that $\varepsilon > 0$ satisfies $B_\varepsilon([\sigma(s_1), \sigma(s_2)]) \subset \mathcal{C}_{\text{free}}$. Choose some $\delta_1 > 0$ for which $[s_i - \delta_1, s_i + \delta_1] \subset [0, 1]$ for $i = 1, 2$. Due to uniform equicontinuity, there exists some $N_1 \in \mathbb{N}$ and $\delta_2 \leq \delta_1$ for which $\|\sigma_n(t) - \sigma_n(s_i)\| < \varepsilon/2$ whenever $|t - s_i| < \delta_2$ and $n \geq N_1$ and $i = 1, 2$. Since the subsequence $(\sigma_{n_k})_{k \in \mathbb{N}}$ converges, there exists some $N_2 \geq N_1$ for which $\sigma_{n_k}(s_i) \in B_{\varepsilon/2}(\sigma(s_i))$ for all $k \geq N_2$ and $i = 1, 2$. Define $\delta_0 := \min\{\delta_2, \delta(\varepsilon/2)\}$, where $\delta(\varepsilon/2)$ is given by Lemma 10, and let $I_i := [s_i - \delta_0, s_i + \delta_0]$ for $i = 1, 2$. It follows that for each $q \in \mathbb{N}$ there exists some $K(q) \geq N_2$ for which $\sigma_{n_k+p}(I_i) \subset B_\varepsilon(\sigma(s_i))$ for all $k \geq K(q)$ and $p \in \{1, \dots, q\}$.

For a syndetically sampled shortcut sequence, let q be the maximum waiting time before a sample from $I = I_1 \times I_2$ is chosen. It then follows from the above that infinitely many shortcut points s^{m_k} satisfy $\sigma_{m_k}(s_i^{m_k}) \in B_\varepsilon(\sigma(s_i))$, in which $m_k = n_k + p_k$ with $p_k \in \{1, \dots, q\}$ for $k \geq K(q)$. If non-negligent random sampling is used, we have $\mathbb{P}(s^n \in I) = p(I) > 0$, independently for each index $n \in \mathbb{N}$. Then the probability of always missing I from some index M onwards is $\prod_{n \geq M} (1 - p(I)) = 0$, from which the result follows. \blacksquare

For any pair of points $x, y \in \mathbb{R}^d$, we denote by $[x, y]$ the straight line segment connecting x to y . For a point x and set A , we write $\text{dist}(x, A) := \inf_{a \in A} \|x - a\|$ to denote the distance of x from A .

Proposition 12. *Assume that a subsequence $(\sigma_{n_k})_{k \in \mathbb{N}}$ converges uniformly to some path σ . Then σ is either irreplaceably optimal so that $\text{int}(S_\sigma) = \emptyset$, or else σ is boundary-optimizable at every $s \in \text{int}(S_\sigma)$.*

Proof. Assume contrary to the claim that σ is not irreplaceably optimal (which implies $\text{int}(S_\sigma) \neq \emptyset$) and that there exists some $s = (s_1, s_2) \in \text{int}(S_\sigma)$ which is not an extremal shortcut point. Since $\sigma \neq R(\sigma, s)$, there exists some $p \in (s_1, s_2)$ for which $\sigma(p) \notin [\sigma(s_1), \sigma(s_2)]$. Let $D := \text{dist}(\sigma(p), [\sigma(s_1), \sigma(s_2)])$ and let $L_s := \|\sigma(s_2) - \sigma(s_1)\|$. It can be verified (using the Pythagorean theorem) that the shortest path γ connecting $\sigma(s_1)$ to $\sigma(s_2)$ and travelling through $\sigma(p)$ has length $\ell(\gamma) \geq rL_s > L_s$ in which $r := ((2DL_s^{-1})^2 + 1)^{1/2} > 1$. We aim to show that the above would imply the contradiction that the path can be shortened indefinitely by an amount bounded away from zero. Since s is not an extremal shortcut point, there exists some $\varepsilon_1 > 0$ for which $B_{\varepsilon_1}([\sigma(s_1), \sigma(s_2)]) \subset \mathcal{C}_{\text{free}}$. There also exists some $\varepsilon_2 \leq \varepsilon_1$ with the property that for any combination of points $z(t) \in B_{\varepsilon_2}(\sigma(t))$ with $t \in \{s_1, s_2, p\}$, the combined length

$$L_z = L(z(s_1), z(s_2), z(p)) := \ell([z(s_1), z(p)]) + \ell([z(p), z(s_2)])$$

satisfies $|L_z - \ell(\gamma)| < \frac{1}{4}(r-1)L_s$, while $L_z^* = L^*(z(s_1), z(s_2)) := \ell([z(s_1), z(s_2)])$ satisfies $|L_z^* - L_s| < \frac{1}{4}(r-1)L_s$. These imply $L_z > L_s + \frac{3}{4}(r-1)L_s$ and $L_z^* < L_s + \frac{1}{4}(r-1)L_s$. Lemma 11 guarantees the existence of a subsequence $(s^{m_k})_{k \in \mathbb{N}}$ for which $\sigma_{m_k}(s_1^{m_k}) \in B_{\varepsilon_2}(\sigma(s_1))$ and $\sigma_{m_k}(s_2^{m_k}) \in B_{\varepsilon_2}(\sigma(s_2))$ for all $k \in \mathbb{N}$. On the other hand, for every index k there exists some $u(k) \in \mathbb{N}$ for which $n_{u(k)} > m_k$ and $\sigma_{n_{u(k)}}(p) \in B_{\varepsilon_2}(\sigma(p))$ and $\sigma_{n_{u(k)}}(s_i) \in B_{\varepsilon_2}(\sigma(s_i))$ for $i = 1, 2$.

Due to the above estimates, the lengths of the paths $\Gamma_s^n := \sigma_n([s_1, s_2])$ satisfy $\ell(\Gamma_s^{n_{u(k)}}) - \ell(\Gamma_s^{m_{k+1}}) > \frac{1}{2}(r-1)L_s$ for all $k \in \mathbb{N}$. Thus, the length of the path is reduced infinitely many times by a fixed amount in the parameter interval (s_1, s_2) . Since the length of the entire path σ_n is monotonically decreasing in any fixed parameter interval, this implies that the length of the limit path satisfies $\ell(\sigma) = -\infty$, which is a contradiction. ■

Figure 4 illustrates a situation in which non-negligent random sampling or syndetically dense sampling converges to a path that is boundary-optimizable and not irreplaceably optimal due to extremal shortcut points. Note that a small perturbation of the obstacle boundaries would make the extremal shortcut points disappear. We conjecture that the existence of extremal shortcut points is a degenerate property of a shortcutting problem, and can always be removed in this way. The proof of this result remains for future work.

4 Particular Sampling Methods

This section explains sampling methods used in the experiments in Section 5. Note that all of these methods use a dense sampling sequence to generate the shortcutting sequence of Algorithm 1.

Checking whether a candidate shortcut is collision-free (step 4 of Algorithm 1) is the most computational costly step of the shortcutting method. Therefore, given that in practice shortcutting methods are mostly done on polygonal paths expressed as a sequence of configurations, it is easy to discard a point in the sampling sequence that is not even a candidate shortcut. For that we check whether the geodesic corresponding to the point (interval) of the sampling




Fig. 4. (left) An environment and an initial path that does not converge to an irreplaceably optimal path. The tips of the triangular obstacles and the top of the trapezoidal obstacle are all collinear. Only non-limit shortcut points are capable of causing the path to go across the top of the trapezoid. **(right)** The path that the initial path converges to after repeated shortcutting by a syndetically dense or non-negligent random sampling sequence. Shortcuts that go from segment $[a, b]$ to $[c, d]$ are limit shortcut points in this path, but they are also extremal shortcut points. The dotted line is a path that results from taking an extremal shortcut.

sequence lies in the same edge; see point (e, f) in Figure 1. Therefore, in step 3 of Algorithm 1 we first check whether the point of the sequence is a candidate shortcut, before checking whether it is collision free. This additional step reduces the number of collision-checks the shortcutting methods perform, which makes them run considerably faster.

Random shortcutting Here, we choose the sequence $(a_n)_{n \in \mathbb{N}}$ of Algorithm 1, to be a sequence of uniformly random points chosen from the interval \mathcal{I} . For the coming experiments, we use Python’s random generator from Numpy. Note that even though the $(a_1, a_2) \in \mathcal{I}$ are generated randomly, the distance between these two random points is not uniform. Its distance $|s_1 - s_2|$ will have a higher probability of being smaller. Therefore, random shortcutting is biased toward smaller shortcuts. This could be fixed by identifying interval endpoints and treating the problem as sampling an interval on a circle [13]; this is not considered here. The convergence is ultimately probabilistic.

Halton shortcutting Shortcutting achieves most of its optimization in the very first iterations, and it is usually implemented using a random number generator. If we instead use a Halton sequence, we obtain lower dispersion and discrepancy than with random sampling [17]. This helps to visit every point in \mathcal{I} more uniformly than standard uniform randomness. It also provides deterministic convergence, as opposed to probabilistic. For the Halton shortcutting method, let $(a_n)_{n \in \mathbb{N}}$ in Algorithm 1 be a sequence of Halton 2D numbers [7] in base 2 and 3. This sequence was generated using the Halton function from Python’s SciPy library, which scrambles the Halton numbers as per [18].

Sliding + Halton shortcutting After experimenting with different sequences and environments, we found that the strategy of fixing the initial distance between s_1 and s_2 , sliding the pair across the interval $[0, 1]$ starting from $s_1 = 0$




Fig. 5. Sliding numbers $a_i = (s_1^i, s_2^i)$. Here, $s_2^k - s_1^k = s_2^{k+1} - s_1^{k+1} = \lambda_i$, for some $i = 0, 1, \dots, n-1$. From a_k to a_{k+1} we slide with $s_j^{k+1} = s_j^k + \lambda_i(1 - \omega_i)$, for $j = 1, 2$.

until $s_2 \geq 1$, and then halving the interval length and repeating the process, consistently yielded better results in the initial iterations than the other sequences we tested. To construct these sliding numbers, we chose two finite sequences of numbers λ_i and ω_i of the same length n . The sequences λ_i and ω_i are called *lengths* and *overlaps*, respectively. For $i = 0, 1, \dots, n-1$, we have that $0 < \lambda_i \leq 1$ and $0 \leq \omega_i < 1$. Given λ_i and ω_i , Algorithm 2 describes the generation of the sliding numbers. Figure 5 illustrates how they *slide* through the interval $[0, 1]$. These sliding numbers are not dense in \mathcal{I} (see Figure 6).

The method first uses a fixed set of sliding numbers and then switches to a randomized Halton sequence. For the experiments we first attempt to shortcut using 247 sliding numbers before switching to a randomized Halton sequence. The switch occurs if either the path length has not lowered over the last 15 sliding numbers or if the limit of 247 points has been reached. These sliding points are generated using Algorithm 2, with $\lambda_i = 2^{-i}$ and $\omega_i = 1/2$, for $i = 0, 1, \dots, 6$. Figure 6 shows the resulting 247 sliding numbers. These parameters were chosen based on empirical evidence after many experiments over many different environments and initial paths.

Algorithm 2

Sliding($(\lambda_0, \dots, \lambda_{n-1})$,
 $(\omega_0, \dots, \omega_{n-1})$)

- 1: $L \leftarrow$ empty list
 - 2: **for** $i = 0$ to $n - 1$ **do**
 - 3: $\alpha \leftarrow \lambda_i(1 - \omega_i)$
 - 4: $s_1 \leftarrow 0$
 - 5: $s_2 \leftarrow \lambda_i$
 - 6: **while** $s_2 < 1$ **do**
 - 7: $L.append((s_1, s_2))$
 - 8: $s_1 \leftarrow s_1 + \alpha$
 - 9: $s_2 \leftarrow s_2 + \alpha$
 - 10: $L.append((s_1, 1))$
 - 11: **return** L
-




Fig. 6. The 247 sliding numbers used for the *slide + Halton* shortcutting method

5 Shortcutting Experiments

Experiments were performed for examples with configuration spaces of 2, 5, and 20 dimensions. Hundreds of runs were made over a variety of initial paths and environments. The performance criterion we use is the number of collision detection checks each method performs to reduce the path length. We choose




Fig. 7. Environment 1, initial path 1, shortcutting methods comparison. (a) Initial path and converging path for slide + Halton method. (b) For 100 runs of the shortcutting methods, we plot the mean cost with a shaded area representing one standard deviation.

this criterion because collision checking dominates the running time, especially in higher dimensions, and it is independent of the implementation. One collision check corresponds to checking if a candidate shortcut is collision-free. Some experiments are highlighted here and more appear in the appendix.

5.1 2D Point Robots

We first show our shortcutting methods for a point translating in a 2D environment, resulting in a 2D configuration space. We use two different environments, each with two different unoptimized initial paths. Each initial path was generated with the bidirectional RRT algorithm [12].

Environment 1 For maze-like environments slide + Halton gives the best results. In Figure 7(b), note that random and Halton are very close together, and Halton has a slightly lower mean. Note that slide + Halton almost converges by just using the 247 sliding points, whereas random and Halton need thousands of collision checks to get near the slide + Halton result and their lengths decay very gradually. All methods seem to converge to a stable cost.

Environment 2 This kind of environment tends to have the most variability among the shortcutting methods. This is because there are many narrow spaces to shortcut through. In Figure 8, we have the most extreme of these cases. All methods get stuck in one homotopy class, until they eventually find a shortcut that changes their homotopy class and converge to the same path. In both Figures 8 and 9, the mean of random and Halton are similar, but the variability of Halton drops more quickly. In Figure 8(b), all methods converge similarly, whereas in Figure 9(b), slide + Halton converges faster. Note that in Figures 8(a) and 9(a) we have the same environment, same initial and goal configurations, and the resulting optimized paths and convergences rate are different. This is because the shortcutting process is very dependent on the given initial path.

Different irreplaceably optimal paths Figure 10 shows an environment in which starting from the same initial path, the same method converges to three




Fig. 8. Environment 2, initial path 1, shortcutting methods comparison. (a) Initial path and converging path for slide + Halton method. (b) For 100 runs of the shortcutting methods, we plot the mean cost with a shaded area representing one standard deviation.




Fig. 9. Environment 2, initial path 2, shortcutting methods comparison. (a) Initial path and converging path for slide + Halton method. (b) For 100 runs of the shortcutting methods, we plot the mean cost with a shaded area representing one standard deviation.

different paths based on the particular sample sequence. If shortcutting methods are randomized, then we often get different results by rerunning the same method. This behaviour is shown in the video: <https://youtu.be/K500a1Ex9kA>. It is easier to see why this happens in two dimensions than in higher ones, as explained in Figure 2. Each of the resulting paths in Figure 10(a) is an irreplaceably optimal path, but in Figure 10(b) we can see that the shortest of these paths is chosen less frequently. Also, slide + Halton always converges to one path; this happens because the first sliding points are fixed, and in this case they reduce the possibility of converging to another path.

5.2 Higher-dimensional Multilink Robots

Here we consider articulated bodies, modeled as a multilink arm formed of line segments, with fixed base point, in a 2D environment with obstacles; see Figure 11. The configuration space is an d -dimensional torus $\mathcal{C} = (S^1)^d = T^d$, in which each component corresponds to S^1 , the range of each joint angle. Therefore, to evaluate the path length of these n -link robots, we will use the standard flat

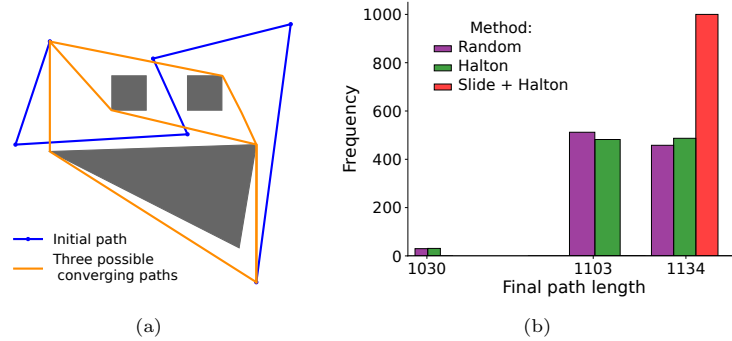


Fig. 10. Same initial path with different possible converging paths.

Riemannian metric on the n -torus. Link-obstacle pairs are checked for collisions, but not link-link pairs. For each robot, we attempted shortcutting with initial paths that were generated using the bidirectional RRT algorithm. More experiments are shown in the appendix. You can also watch animations of the resulting paths in the video: <https://youtu.be/iz3P0gfrQrw>.




Fig. 11. Two different multilink arm robots in different environments. (a) 5-link arm robot. (b) 20-link arm robot.

We observed similar results as from the 2D experiments for the first iterations, but these higher dimensional experiments seem to converge to many more irreplaceably optimal paths, which was unexpected. Therefore, we focus on showing the frequency of these different resulting paths rather than showing the early convergence rates of the methods. We will also show an example of the most frequent converging path of the Slide + Halton method.

In Figures 12(a) and 13(a) we show the two unoptimized initial paths that we used to test the shortcutting methods for a 5-link and 20-link arm robot, respectively. Figure 14 shows the results. Even though in both results, they appear to converge to more than 10 different paths, we cannot assume that we




Fig. 12. (a) 5-link arm robot initial path with length of 19.49. (b) Slide + Halton most frequent converging path with length of 9.69. Each snapshot is a node of the path.




Fig. 13. (a) 20-link arm robot initial path with length of 36.55. (b) Slide + Halton most frequent converging cost: 18.03. Each snapshot is a node of the resulting path.

have converged to an irreplaceably optimal path. It could be that we have already settled on an irreplaceably optimal path, or it could be the case that it need a many more iterations to converge. In practice, the most frequently occurring of these paths seem unlikely to be further optimized by shortcutting, as shown in Figures 12(b) and 13(b). The robot seems to barely graze the obstacles upon path execution, indicating that it is quite optimized (it is easier to see why obstacles are grazed in 2D configuration spaces). Also, in Figure 14, random and Halton methods seem to have the same frequency of converging paths, whereas slide + Halton is restricted to a subset of these paths and has different frequencies.

6 Conclusions

This paper focused on sequences of sample points that were “open-loop” meaning that the sample points were not based on any properties of the existing path or the environment. A sampling sequence that takes feedback about the environment or the path into account may produce superior paths, perhaps by taking early shortcuts that are likely to open up new regions of the configuration space

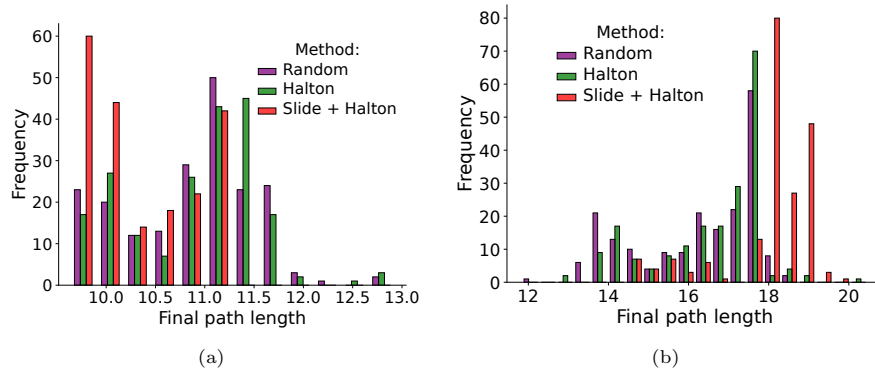


Fig. 14. Histograms (a) and (b) are the results of shortcutting methods for the initial path in Figures 12(a) and 13(a), respectively.

for later shortcuts, by avoiding shortcuts likely to converge to a very suboptimal final path (like in Figure 2), or by taking shortcuts that cause potentially advantageous environment features to be incorporated into the path (e.g. inflections, reflex vertices, etc.). Feedback might also be able to provide an estimate on the size of some of the unblocked areas of the interval space, which could in turn indicate how fine a sampling is likely to be necessary.

We found that when we use a randomized shortcutting method, there might be many different possible converging paths with different costs. Determining the best path accessible via shortcutting given a particular environment and starting path is still an open question. One potential direction would be to avoid shortcuts that converge to a poor irreplacably optimal path could be helped by determining a data structure that keeps a record of the interim paths generated by the shortcuts so that portions of the shortcutting sequence can be unwound and other sequences can be explored. This would be particularly helpful in environments like the one pictured in Figure 10 where the best possible irreplacably optimal path has a very narrow sampling window compared to the longer irreplacably optimal path. Incorporating a hybridization process [16] among the intermediate paths may also be useful.

A logical next step is to extend the concept of irreplacable optimality to non-holonomic and kinodynamic settings, with the intention of minimizing something other than path length (for example, minimizing the time required to traverse the path). For example, shortcutting in the context of kinodynamic planning using time-optimal controls was explored in [13].

Acknowledgments This work was supported by a European Research Council Advanced Grant (ERC AdG, ILLUSIVE: Foundations of Perception Engineering, 101020977), Academy of Finland (BANG! 363637). The authors would like to thank Vadim Weinstein for useful discussions while revising this paper.

References

1. S. Berchtold and B. Glavina. A scalable optimizer for automatically generated manipulator motions. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'94)*, volume 3, pages 1796–1802. IEEE, 1994.
2. J. T. Betts. Survey of numerical methods for trajectory optimization. *Journal of Guidance, Control, and Dynamics*, 21(2):193–207, March-April 1998.
3. V.G. Boltyanski, R.V. Gamkrelidze, E.F. Mishchenko, and L.S. Pontryagin. The maximum principle in the theory of optimal processes of control. *IFAC Proceedings Volumes*, 1(1):464–469, 1960.
4. M. Campana, F. Lamiroux, and J-P. Laumond. A gradient-based path optimization method for motion planning. *Advanced Robotics*, 30(17-18):1126–1144, 2016.
5. J.D. Gammell, S.S. Srinivasa, and T.D. Barfoot. Batch informed trees (BIT*): Sampling-based optimal planning via the heuristically guided search of implicit random geometric graphs. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 3067–3074, 2015.
6. R. Geraerts and M.H. Overmars. Creating high-quality paths for motion planning. *The international journal of robotics research*, 26(8):845–863, 2007.
7. J. H. Halton. On the efficiency of certain quasi-random sequences of points in evaluating multi-dimensional integrals. *Numerische Mathematik*, 2:84–90, 1960.
8. K. Hauser and V. Ng-Thow-Hing. Fast smoothing of manipulator trajectories using optimal bounded-acceleration shortcuts. In *2010 IEEE international conference on robotics and automation*, pages 2493–2498. IEEE, 2010.
9. P. Huang, Y. Shaoul, and Y. Li. Benchmarking Shortcutting Techniques for Multi-Robot-Arm Motion Planning. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2025.
10. S. Karaman and E. Frazzoli. Sampling-based algorithms for optimal motion planning. *The International Journal of Robotics Research*, 30(7):846–894, 2011.
11. L.E. Kavraki, P. Svestka, J.-C. Latombe, and M.H. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*, 12(4):566–580, 1996.
12. J.J. Kuffner and S.M. LaValle. RRT-connect: An efficient approach to single-query path planning. In *IEEE International Conference on Robotics and Automation (ICRA)*, volume 2, pages 995–1001, 2000.
13. A.J. La Valle, B. Sakcak, and S.M. LaValle. Bang-bang boosting of RRTs. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2869–2876, 2023.
14. S.M. LaValle, M.S. Branicky, and S.R. Lindemann. On the relationship between classical grid search and probabilistic roadmaps. *International Journal of Robotics Research*, 23(7/8):673–692, July/August 2004.
15. D. Liberzon. *Calculus of variations and optimal control theory: a concise introduction*. Princeton University Press, 2012.
16. R. Luna, I.A. Şucan, M. Moll, and L.E. Kavraki. Anytime solution optimization for sampling-based motion planning. In *2013 IEEE international conference on robotics and automation*, pages 5068–5074. IEEE, 2013.
17. H. Niederreiter. *Random Number Generation and Quasi-Monte-Carlo Methods*. Society for Industrial and Applied Mathematics, Philadelphia, 1992.
18. Art B. Owen. A randomized halton algorithm in r, 2017.
19. J. Pan, L. Zhang, and D. Manocha. Collision-free and smooth trajectory computation in cluttered environments. *The International Journal of Robotics Research*, 31(10):1155–1175, 2012.

20. S. Quinlan and O. Khatib. Elastic bands: connecting path planning and control. In *IEEE International Conference on Robotics and Automation (ICRA)*, volume 2, pages 802–807, 1993.
21. N. Ratliff, M. Zucker, J.A. Bagnell, and S. Srinivasa. CHOMP: Gradient optimization techniques for efficient motion planning. In *2009 IEEE International Conference on Robotics and Automation*, pages 489–494. IEEE, 2009.
22. W. Rudin. *Real and Complex Analysis*. McGraw Hill, 1986.