On Improving the Clearance for Robots in High-Dimensional Configuration Spaces

Roland Geraerts

Mark H. Overmars

institute of information and computing sciences, utrecht university

technical report UU-CS-2005-024

www.cs.uu.nl

On Improving the Clearance for Robots in High-Dimensional Configuration Spaces*

Roland Geraerts Mark H. Overmars
Institute of Information and Computing Sciences
Utrecht University, the Netherlands
Email: {roland,markov}@cs.uu.nl.

Abstract

In robot motion planning, many algorithms have been proposed that create a path for a robot in an environment with obstacles. Since it can be hard to create such paths, most algorithms are aimed at finding *a* solution. However, for most applications the paths must be of a good quality as well. That is, paths should preferably be short, be smooth, and should preserve some clearance from the obstacles. In this paper, we focus on improving the clearance of paths. Existing methods only extract high clearance paths for rigid, translating bodies. We present an algorithm that improves the clearance along paths of a broader range of robots which may reside in arbitrary high-dimensional configuration spaces. Examples include planar, free-flying and articulated robots.

1 Introduction

Motion planning is one of the fundamental problems in robotics. The motion planning problem can be defined as finding a path between a start and goal placement of a robot in an environment with obstacles. The last decade, efficient algorithms have been devised to tackle this problem. They are successfully applied in fields such as mobile robots, manipulation planning, CAD systems, virtual environments, protein folding and human robot planning. See e.g. the books of Latombe [1] and Lavalle [2] for many interesting results.

One important aspect of motion planning is the quality of the path. The quality is often measured in terms of length and clearance. Many applications require a short path since redundant motions will take longer to execute. In practical situations, the path often has to keep some minimum amount of clearance to the obstacles because it can be difficult to measure and control the precise position of a robot. Traveling along a path with a certain amount of minimum clearance reduces the chances of collisions due to these uncertainties. In a nuclear power plant for example, it may be desirable to minimize the risk of heat or radioactive contamination [3]. In this type of environment, calculating a safe path (off-line) is more significant than calculating a low quality path on-line. Clearance is used for other purposes as well: in [4] for instance, it is used to guide multiple units in a virtual environment.

*This research was supported by the Dutch Organization for Scientific Research (N.W.O.). This research was also supported by the IST Programme of the EU as a Shared-cost RTD (FET Open) Project under Contract No IST-2001-39250 (MOVIE - Motion Planning in Virtual Environments).

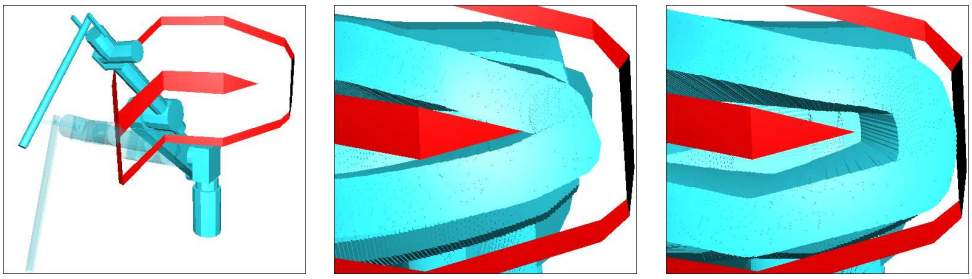

(a) Articulated robot (b) Initial path (c) Optimized path

Figure 1: Improving the clearance along a path traversed by an articulated robot with six degrees of freedom. Figure (a) shows the start and (blurred) goal placement of the robot arm. In (b), we zoom in on the initial path. The swept volume of the significant parts of the robot is shown. As can be seen, the clearance along this path is very small. Our new algorithm successfully increases the clearance along the path which is visualized in (c).

Indeed, enlarging the clearance is important in various applications. It is though far from trivial how to do this. Previous work limited their efforts to rigid, translating bodies. In this paper, we present a new algorithm that improves the clearance along paths for a broader range of robots including planar, free-flying and articulated robots, see Figure 1.

1.1 Related work

Many motion planning algorithms create a roadmap (or graph) which represents collision-free motions that can be made by the moving object in an environment with obstacles. From this graph a path is obtained by a Dijkstra's shortest path query. Since these calculations can be performed off-line, we refer to this stage as preprocessing. The paths usually are optimized in a post-processing stage.

In [5], an augmented version of Dijkstra's algorithm is used to extract a path based on other criteria than length. The minimum clearance along the path is maximized by incorporating a higher cost for edges that represent a small amount of clearance. Unfortunately, the paths rarely provide optimal solutions because they are restricted to the randomly generated nodes in the roadmap. Even if the nodes are placed on the medial axis [6], the edges will in general not lie on the medial axis, and hence the extracted paths do not have a guaranteed amount of clearance.

Another category of algorithms create a path along the Generalized Voronoi Diagram (GVD). The GVD (or medial axis) for a robot with n degrees of freedom (DOFs) is defined as the collection of k-dimensional geometric features ($0 \leq k < n$) which are in general $(n+1-k)$-equidistant to the obstacles. As an example, consider Figure 2(a) that shows the outline of a solid bounding box and a part of the medial axis for a translating robot. The medial axis of this robot consists of a collection of surfaces, curves and points. The surfaces are defined by the locus of 2-equidistant closest points to the bounding box. The curves have 3-equidistant closest points and the points have 4-equidistant closest points to the bounding box. These features are connected if the free space in which the robot operates is also connected [7]. Hence, the GVD is a complete representation for motion planning purposes. Most importantly, paths on the GVD have appealing properties such as large clearance from obstacles.

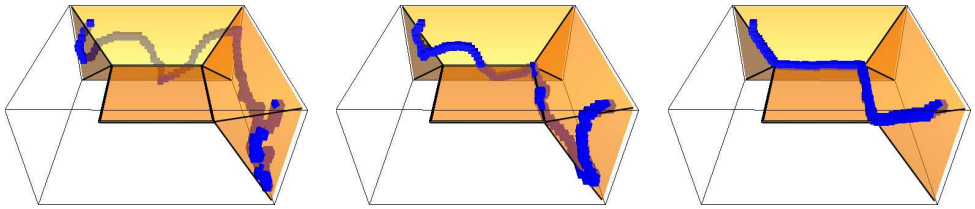

(a) Initial path (b) Retracted path (c) Optimal path

Figure 2: Retraction of a path in an environment that only consists of a solid bounding box. A part of the medial of this box is shown. Figure (a) shows the initial path. This path is retracted to the medial axis by the algorithm from [11] in (b). Figure (c) shows a path having the optimal amount of clearance. This path was obtained by our new algorithm.

Unfortunately, an exact computation of the GVD is not practical for problems involving many degrees of freedom (DOFs) and many obstacles as this requires the expensive and intricate computation of the configuration space obstacles. Therefore, the GVD is approximated in practice.

In [8], the GVD is approximated by applying spatial subdivision and isosurface extraction techniques. Although the calculations are easy, robust and can be generalized to higher dimensions, the technique only works for disjoint convex sites and consumes an exponential amount of memory which makes this technique impractical for problems involving many DOFs. Another approach incrementally constructs the GVD by finding the maximal inscribed disks in a two dimensional discretized workspace [9]. While this algorithm is also extensible to handle higher-dimensional problems, it suffers from the same drawback as the preceding algorithm. In [10], a technique is described that exploits the fast computation of a GVD using graphics hardware for motion planning in complex static and dynamic environments. The technique is though limited to a three-dimensional workspace for rigid translating robots.

The above preprocessing methods concentrate on creating a data structure from which paths can be extracted. In [11], the authors proposed a post-processing algorithm that adds clearance to a path by retracting it to the medial axis of the workspace. A path Π is divided into n configurations. Such a configuration corresponds to a particular placement of the robot in the environment. Each configuration is retracted to the medial axis (except for the start and goal configuration) as follows. First, the closest point (on the obstacles) between the robot and the obstacles is calculated. Then the robot is iteratively moved away from this point until the robot has 2-equidistant nearest points to the obstacles.

This method provides a fast and accurate retraction of possibly non-convex rigid bodies to the medial axis in two- and three-dimensional spaces. Unfortunately, as the retraction is performed by a series of translations of the robot, the method is not suitable for increasing the clearance along a path traversed by an articulated robot or a free-flying robot for which the rotational DOFs are more important than the translational ones. In addition, the method will in general not produce a maximal clearance path because the retraction is completed when the configurations on the path are placed somewhere on the medial axis. A configuration could have a higher clearance if it was retracted to a *ridge* of the medial axis. We define a ridge as the locus of points in configuration space $\mathcal{C}$ that represents a locally maximum clearance. See Figure 2 for an example. The crooked path from Figure 2(a) was retracted to the medial axis by the algorithm from [11]. Figure 2(b) shows that the retracted configurations sway on the

medial axis surfaces. In Figure 2(c), the path is retracted toward the ridges, resulting in an optimal clearance path.

In conclusion, there exists no efficient algorithm that can compute high-clearance paths for robots residing in high-dimensional configuration spaces. We will propose an algorithm that efficiently retracts a path toward the ridges of the medial axis for robots with an arbitrary number of DOFs. This algorithm may be applied in high-cost environments such as a factory in which a manipulator arm operates.

1.2 Paper outline

This paper is organized as follows: in section 2, we will introduce our retraction algorithm. We elaborate on the algorithmic details in section 3 and show that the algorithm is correct. Then, we apply the algorithm on problems involving planar, free-flying and 6-DOF articulated robots in section 4. Finally, we conclude in section 5 that our method is successfully able to improve the clearance along paths traversed by a broad range of robots.

2 Retraction Algorithm

The retraction algorithm tries to iteratively increase the clearance of the configurations on the path by moving them in a direction for which the clearance is higher, see Algorithm 1. We define clearance as the Euclidean distance between the pair of closest points on the robot and obstacles.

Our problem is as now follows. Convert a given path Π into a path Π' such that for each $\pi'_i \in \Pi'$ the clearance is either locally maximal or is larger than a given constant c_{min}. Our solution consists of a number of iterations. In each iteration, we choose a random direction dir. Then we try to move each configuration π_i in the chosen direction, i.e. $\pi'_i \leftarrow \pi_i \oplus dir$. If the clearance of π'_i is larger than the clearance of π_i, then π_i is replaced by π'_i. We stop retracting the path when the path has c_{min} clearance or when the *average* clearance does not improve anymore. (See Definition 1 which shows how to compute the average clearance.)

By updating the configurations, the path is forced to stretch and shrink during the retraction which causes the following two problems. First, the distance between two adjacent configurations in the path can become *larger* than the maximum step size. This happens for example when the path is pushed away from the obstacles. If this occurs, we insert some extra configurations in between them. Second, multiple configurations can be mapped to a small region by which the distance between two non-adjacent configurations is *smaller* than the step size. This occurs for example when pieces of the path are traversed twice. As we will see in the following section, these can easily be removed.

An impression of a retraction is visualized in Figure 3. This figure shows an initial and a final path traversed by a square robot in a simple two-dimensional workspace. The lines between them are the guided random walks of the configurations. We call these walks *guided* because a configuration is updated only if its clearance increases. Note that at some places, extra configurations were inserted while at other places configurations were removed. After 40 iterations, the initial path was successfully retracted to the medial axis, resulting in a large clearance path. While this example shows a retraction for a robot with only two DOFs, the retraction can also be applied to robots with more DOFs such as an articulated robot with six joints.

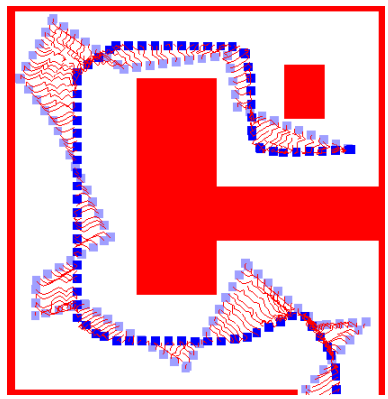

Figure 3: An impression of the retraction algorithm. The algorithm retracts the initial path (traversed by a square robot) to the medial axis. For each configuration on the path, the guided random walk (small curve) is drawn.

Algorithm 1 RetractPathConfigurationSpace(path Π)

Require: $\forall i : d(\pi_i, \pi_{i+1}) \leq step$

```
loop
    dir ← RandomDirection(step)
    for each π_i in Π do
        π' ← π_i ⊕ dir
        if Clearance(π'_i) > Clearance(π_i) then
            π_i ← π'
    ValidatePath(Π)
```

3 Algorithmic Details

A path consists of a series of configurations. We require that the distance between each pair of adjacent configurations is at most *step*. Hence, we need a distance metric which will be described below. The clearance of the configurations is improved by moving them in a random direction. We will show how to compute such a direction. Furthermore, we will show how to move the robot in this direction. After an iteration of the algorithm, the distance between two adjacent configurations can change. To keep a valid path, we will show how to insert and delete appropriate configurations. Finally, we show that the algorithm indeed retracts a path toward the ridges of the medial axis.

3.1 Distance metric

The importance of choosing a good distance metric is discussed in [12]. Such a metric often incorporates weights which can be used to control the relative importance of the DOFs of the robot. We distinguish three types of DOFS: translation, rotation$_1$ (rotation about the x-, y-, or z-axis) and rotation$_3$ (rotation in SO3). For example, a free-flying robot can be described by three translational DOFs and one rotational$_3$ DOF, and an articulated robot with six joints may be described by six rotational$_1$ DOFs.

We calculate the distance between two configurations q and r by summing the weighted partial distances for each DOF $0 \leq i < n$ that describes the configurations, i.e.

$$d(q,r) = \sqrt{\sum_{i=0}^{n-1} [w_i d(q_i, r_i)]^2}.$$

The calculation of distance $d(q_i, r_i)$ is dependent on the type of the DOF. For translation, we set $d(t', t'')$ to $|t' - t''|$. We split the calculation for a rotational$_1$ DOF in two parts: if the range is smaller than 2π, which occurs often for revolution joints in manipulator arms, we take the same distance measure as for a translational DOF. If the rotational DOF is periodic, i.e. the orientation at 0 radians equals the orientation at 2π radians, we take the smallest angle. More formally, we set $d(r', r'')$ to $|r' - r''|$ if r is not periodic, otherwise $d(r', r'') = \min\{|r' - r''|, r' - r'' + 2\pi, r'' - r' + 2\pi\}$. We use unit quaternions to represent rotations in 3D. The distance between two quaternions r' and r'' can be calculated by taking the shortest angle over a 4D-sphere, i.e. $d(r', r'') = \min\{2\arccos(r' \cdot r''), 2\pi - 2\arccos(r' \cdot r'')\}$. We do not use three Euler angles to represent an orientation because this representation has several drawbacks quaternions do not have, such as the 'gimbal lock' and the difficulty to define proper methods for sampling, interpolation and distance metric [13].

3.2 Random direction vector

We need to create a random direction q' such that the distance from configuration q to $q \oplus q'$ equals *stepsize*, see Algorithm 2. We set *stepsize* to $\frac{2}{3}$*step* as this is needed in the proof that the algorithm is correct. The direction q' is composed of elements from $\{q'_{tra}, q'_{rot_1}, q'_{rot_3}\}$ such that $\sqrt{\sum_{i=0}^{n-1} p_i^2} = stepsize$, where $p_i = w_i d(q_i, q_i \oplus q'_i)$. We choose a random value rnd_i between 0 and 1 for each DOF i of q'. Then we set the partial distance p_i for DOF i to $d(q_i, q_i \oplus q'_i) = \frac{rnd_i \, * \, stepsize}{||rnd * w||}$. We now know how much each DOF contributes to the total distance. We set the translational and rotational$_1$ components of q'_i to $\pm p_i$. The calculation

of the rotational$_3$ component is more complicated. We represent this component as a random 3D unit axis $a = (a_x, a_y, z_z)$ and an angle of revolution θ about that axis. (Since a revolution of more than π radians makes no sense, we constrain θ to $0 \leq \theta \leq \pi$.) This representation can easily be converted to a quaternion, i.e. $q'_{rot_3} = (a_x \sin(\theta/2), a_y \sin(\theta/2), a_z \sin(\theta/2), \cos(\theta/2))$. If we set θ to p_i, then the distance between quaternion q_{rot_3} and $q'_{rot_3} * q_{rot_3}$ equals p_i.

Theorem 1 (Direction vector)**.** *The above construction of the random direction vector q' ensures that $d(q, q \oplus q') = stepsize$.*

Proof: Let $\sqrt{\sum_{i=0}^{n-1} p_i^2} = stepsize$. We compose each value of the partial distance p_i of a random value rnd_i between 0 and 1, multiplied with the corresponding weight factor w_i, i.e. $p_i = rnd_i w_i$. By multiplying the term $rnd_i w_i$ by $\frac{stepsize}{||rnd*w||}$ we obtain $p_i = \frac{rnd_i w_i}{||rnd*w||} stepsize$. Now holds that $\sqrt{\sum_{i=0}^{i=n-1} [\frac{rnd_i w_i}{||rnd*w||} stepsize]^2} = stepsize$. So $w_i d(q_i, q_i + q'_i) = \frac{rnd_i w_i}{||rnd*w||} stepsize$. Hence the partial distance for DOF i is $d(q_i, q_i + q'_i) = \frac{rnd_i * w_i}{||rnd*w||} stepsize / w_i = \frac{rnd_i * stepsize}{||rnd*w||}$.

In other words, the value for DOF i in the direction vector q' has to be set to $\frac{rnd_i * stepsize}{||rnd*w||}$. In line 6, as $d(t, t + t') = |t - (t + t')| = |t'| = \pm p_i$, we set the value for a translational DOF to $\pm p_i$. The same argument holds for choosing the value for the rotational$_1$ DOF in line 8.

Finally, we have to proof that the distance between quaternion q and $q' * q$ equals p_i. ■

Theorem 2 (Direction vector of a quaternion)**.** *Let q and q' be two quaternions. The quaternion q' represents a random (unit) axis and an angle of revolution θ about that axis. If the range of θ is set to $0 \leq \theta \leq \pi$, then the distance $d(q, q' * q) = p_i$.*

Proof: The distance between two quaternions q and r is defined as $d(q, r) = \min\{2 \arccos(q \cdot r), 2\pi - 2 \arccos(q \cdot r)\}$. Since θ is positive, the dot product $q \cdot r$ is also positive (and is between 0 and 1). As a consequence, the arccos of the dot product will be between 0 and π. Hence, the distance between q and r equals to $d(q, r) = 2 \arccos(q \cdot r)$. This distance should be equal to the partial distance p_i:

$$d(q, r) = 2 \arccos(q \cdot r) = p_i$$

Let $q = q' * r$ and $r = (r_x, r_y, r_z, r_w)$. The quaternion q' represents a rotation by θ around a unit axis $a = (a_x, a_y, z_z)$: $q' = (a_x \sin(\theta/2), a_y \sin(\theta/2), a_z \sin(\theta/2), \cos(\theta/2))$. Then we get

$$2 \arccos((q' * r) \cdot r) = p_i$$

By substitution we get

$$2 \arccos(((r_x^2 + r_y^2 + r_z^2 + r_w^2) \cos(\frac{\theta}{2})) = p_i$$

The length of a quaternion that represents a rotation is always equal to 1. Hence, $r \cdot r = r_x^2 + r_y^2 + r_z^2 + r_w^2 = 1$.

By substitution we get

$$\theta = p_i$$

■

Besides choosing a random vector, we need to add a direction q' to configuration q. For translational and rotational$_1$ DOFs, we can add up the values. If rotational$_1$ DOF is periodic, we have to make sure that the value remains in the range between 0 and 2π. For the rotational$_3$ DOF, q'_{rot_3} has to be multiplied by q_{rot_3}.

Algorithm 2 RandomDirection(float *step*, weight vector w_i)

Require: $stepsize \leftarrow \frac{2}{3} * step$
for all $i : 0 \leq i < |\text{DOFs}|$ **do**
  $rnd_i \leftarrow \text{Random}(0, 1)$
for all $i : 0 \leq i < |\text{DOFs}|$ **do**
  $p_i \leftarrow rnd_i * stepsize / ||rnd * w||$
  if $\text{DOF}_i = \text{translation}$ **then**
    $q'_{i(tra)} \leftarrow \pm p_i$
  else if $\text{DOF}_i = \text{rotation}_1$ **then**
    $q'_{i(rot_1)} \leftarrow \pm p_i$
  else if $\text{DOF}_i = \text{rotation}_3$ **then**
    $\text{axis}_{xyz} \leftarrow (\text{Random}(-1, 1), \text{Random}(-1, 1), \text{Random}(-1, 1))$
    $\text{axis} \leftarrow \frac{1}{||axis||} * \text{axis}$
    $q'_{i(rot_3)} \leftarrow \text{Quaterion}(\text{axis}_x * \sin(p_i), \text{axis}_y * \sin(p_i), \text{axis}_z * \sin(p_i), \cos(p_i))$
return q'

3.3 Path validation

As a path is forced to stretch and shrink during the retraction, this path may not be valid anymore after an iteration of Algorithm 1. A path Π is valid if $\forall i : d(\pi_i, \pi_{i+1}) \leq step$. In this section we will show how to construct a new valid path.

First, we make a copy of Π (which is the path before updates were taken place), i.e. $\Pi' \leftarrow \Pi$. After one iteration, the configurations of Π' may have been updated. We construct a new path Π'' which will contain all configurations from Π' and new configurations to assure that Π'' will be valid.

For all i, we check if $d(\pi'_i, \pi'_{i+1}) > step$. This may occur when a configuration π'_i is moved. In general, π'_i can be left unchanged or moved in direction *dir*. There are two cases in which this distance will be larger than *step*. We handle them as follows.

In the first case, π'_i is left unchanged while π'_{i+1} is updated. First we append π'_i to Π''. Then, we either append π_{i+1} or $\pi = \text{Interpolate}(\pi_i, \pi'_{i+1}, 0.5)$ to Π''. We append the configuration that has the largest clearance. A similar procedure can be performed if π'_i is updated while π'_{i+1} is left unchanged.

In the second case, both π'_i and π'_{i+1} are updated. Again, we first we append π'_i to Π''. Then, we either append both π_i and π_{i+1} or $\pi = \text{Interpolate}(\pi'_i, \pi'_{i+1}, 0.5)$ to Π''. If the clearance of π is larger than the minimum clearance of π_i and π_{i+1}, we append π, and vice versa.

During the construction of Π'', configurations may be interpolated. For translational and rotational_1 DOFs, we can use linear interpolation. The interpolation between two quaternions can be performed by spherical linear interpolation (SLERP), see [13] for implementation issues.

We mentioned that a path can shrink during the retraction. As a consequence, multiple configurations may be mapped to a small region. We remove the configurations π_i for which holds that $d(\pi_{i-1}, \pi_{i+1}) \leq step$. After the deletion of π_i, π_{i-1} will be adjacent to π_{i+1}.

Theorem 3 (Path validation)**.** *The algorithm will keep the distance between each two adjacent configurations below (or equal to) step.*

Proof: In the first case, π'_i is left unchanged while π'_{i+1} is updated. In any case, it holds that

$d(\pi'_i, \pi_{i+1}) \leq step$, $d(\pi_{i+1}, \pi'_{i+1}) \leq step$ and $d(\pi'_i, \pi) = d(\pi, \pi'_{i+1}) \leq 2 * step * \frac{2}{3} * 0.5 < step$. Hence, this construction assures that the distance between the last two appended configurations is at most *step*. A similar proof can be given if π'_i is updated while π'_{i+1} is left unchanged.

In the second case, both π'_i and π'_{i+1} are updated. If π is appended, the distance between the new configurations in Π'' can be $d(\pi'_i, \pi) = d(\pi, \pi'_{i+1}) \leq 3 * step * \frac{2}{3} * 0.5 \leq step$. In the other case, $d(\pi'_i, \pi_i) \leq step$, $d(\pi_i, \pi_{i+1}) \leq step$ and $d(\pi_{i+1}, \pi'_{i+1}) \leq step$. As the maximum distance will be at most step, path Π'' will be valid.

We mentioned that a configuration π_i is removed if $d(\pi_{i-1}, \pi_{i+1}) \leq step$. After the deletion of π_i, configuration π_{i-1} will be adjacent to configuration π_{i+1}. By construction, the distance between them will be at most *step*.

In conclusion, the algorithm will keep the distance between each two adjacent configurations below (or equal to) *step*. Hence, the resulting path Π'' will be valid. ■

3.4 Retraction toward the ridges of the medial axis

We will now show that the algorithm retracts a path toward the *ridges* of the *medial axis*. We define the medial axis as the locus of points in $\mathcal{C}$-space having at least two equidistant closest points on the boundary of the $\mathcal{C}$-space obstacles. We define a ridge as the locus of points in $\mathcal{C}$-space that represent a locally maximum clearance.

In each iteration of the algorithm, we only update a configuration if its clearance increases. Such an update can lead to an insertion and a deletion of configurations. If a configuration π is inserted, then the clearance of π will be equal to or higher than the clearance of the configuration before it was updated. If a configuration is deleted, then it will not play a role anymore. Hence, the clearance along the path never decreases.

If we would not pin down the start and goal configurations of the path then the following would happen. As we move the configurations with a particular step size that is larger than zero, there is a moment for which they will have obtained a (locally) maximal amount of clearance. Hence, the configurations would eventually be placed on the ridges.

Unfortunately, as the start and goal configurations *are* pinned down, not all configurations will be placed on the ridges. However, these configurations will eventually obtain a high clearance as only improvements are allowed.

4 Experiments

In this section, we will investigate how much the two retraction algorithms – retraction in the workspace (Algorithm from [11]: $\mathcal{W}$-retraction) and retraction in the configuration space (Algorithm 1: $\mathcal{C}$-retraction) – can improve the clearance along six paths. We integrated these algorithms in a single motion planning system called SAMPLE (System for Advanced Motion PLanning Experiments), which we implemented in Visual C++ under Windows XP. All experiments were run on a 2.66GHz Pentium 4 processor with 1 GB internal memory. We used Solid 3.5 for collision checking [14].

	Dimensions of the bounding box	
	environment	**robot**
Planar	100×100	1×1
Simple corridor	$40 \times 11 \times 30$	$0.2 \times 0.2 \times 0.75$
Corridor	$40 \times 17 \times 40$	$5 \times 1 \times 5$
Wrench	$160 \times 160 \times 160$	$68 \times 24 \times 8$
Hole[1]	$40 \times 40 \times 40$	$5 \times 5 \times 10$
Manipulator	$10 \times 10 \times 10$	variable

Table 1: The axis-aligned bounding boxes of the environments and robots

4.1 Experimental setup

We considered the environments and their corresponding paths depicted in Figure 4. They have the following properties (see Table 1 for their dimensions):

Planar: This simple two-dimensional environment contains a path traversed by a square robot that can only translate in the plane. As the robot has two translational DOFs, a retraction in the workspace will result in a path having the optimal amount of clearance. The results will show if a retraction in the $\mathcal{C}$-space is competitive.

Simple corridor: This simple three-dimensional environment with lots of free space features a path traversed by a small free-flying cylinder. Both algorithms will introduce an extra amount of clearance as they both move the robot to middle of the corridors. However, the $\mathcal{C}$-retraction algorithm should outperform the $\mathcal{W}$-retraction algorithm as also rotational DOFs are considered.

Corridor: The environment consists of a winding corridor that forces a free-flying elbow shaped robot to rotate. As there is little room between the walls of the corridor and the robot, it may be hard to increase the clearance along the path.

Wrench: This environment features a relative large free-flying object (wrench) in a workspace that consists of thirteen crossing beams. The wrench is rather constrained at the start and goal positions. We expect that the $\mathcal{W}$-retraction algorithm will be outperformed by the $\mathcal{C}$-retraction algorithm as the rotational DOFs are of major concern in this environment.

Hole: The free-flying robot, which has six DOFs (three translational DOFs and a rotation_3 DOF), consists of four legs and must rotate in a complicated way to get through the hole. Only where the path goes through the hole, the clearance is small. Hence, the improvement of the minimum amount of clearance along the path would reveal the power of the $\mathcal{C}$-retraction algorithm.

Manipulator: The articulated robot has six rotational DOFs and operates in a constrained environment. The clearance along the path is very small and cannot be improved by the $\mathcal{W}$-retraction algorithm because it cannot handle rotational DOFs. Again, an increase in the minimum and average amount of clearance along the path will show the potential of the $\mathcal{C}$-retraction algorithm.

We subdivided each path in sequential configurations such that the distance between each two adjacent configurations is at most some predetermined step size. The step sizes for the paths can be found in Table 2. The distance metric from section 3.1 can use different weights

[1] The dimensions of the hole are $5 \times 5 \times 0.5$. The legs of the robot are 1 thick.

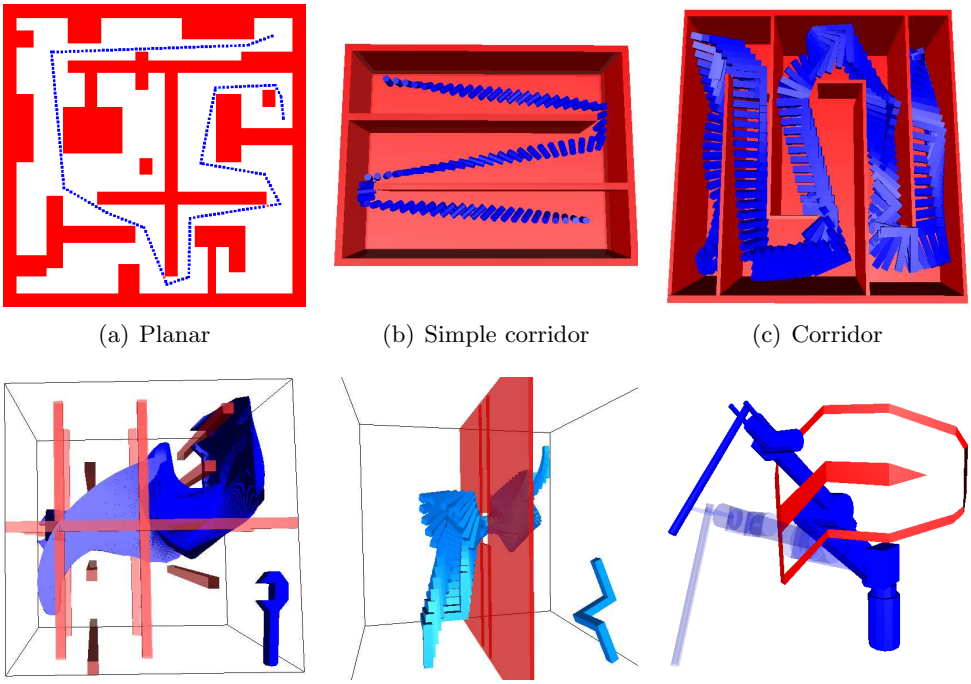

(a) Planar (b) Simple corridor (c) Corridor

(d) Wrench (e) Hole (f) Manipulator

Figure 4: The six test environments and their corresponding paths

	step size
Planar	1.0
Simple corridor	0.4
Corridor	0.7
Wrench	3.0
Hole	1.0
Manipulator	0.1

Table 2: The step sizes for the robots

	Type of DOF of the robot		
	translational	**rotational$_1$**	**rotational$_3$**
Planar	1, 1		
Simple corridor	1, 1, 1		3
Corridor	1, 1, 1		7
Wrench	6, 6, 6		30
Hole	1, 1, 1		11
Manipulator		6, 6, 6, 2, 2, 2	

Table 3: The weights for each DOF of the robots

for the DOFs of a robot. These are enumerated in Table 3.

We recorded the minimum, maximum and *average* clearance along the paths. The average clearance gives an indication of the amount of free space in which a path can be moved without colliding with the obstacles:

Definition 1 (Average clearance along a path)**.** *Let* $\Pi = \pi_0, \ldots, \pi_{n-1}$ *be the configurations along a path* Π *such that* $\forall i : d(\pi_i, \pi_{i+1}) \leq$ *stepsize. Then the average clearance equals to* $\frac{1}{n}\sum_{i=0}^{n-1} Clearance(\pi_i)$.

As the $\mathcal{C}$-retraction algorithm is non-deterministic, we run this algorithm 100 times for each experiment and report the averages.

4.2 Experimental results

The results of the experiments are stated in Table 4 and visualized in Figure 5.

Planar: A retraction in the workspace results in a path having the optimal amount of clearance. The statistics show that the $\mathcal{C}$-retraction technique reaches these optimal values. However, for robots that have two translational DOFs, we recommend the workspace technique as this technique is faster.

Simple corridor: Indeed, a large amount of clearance was introduced by the retraction algorithms. At the expense of five extra seconds of computing time, the $\mathcal{C}$-retraction technique doubled the minimum amount of clearance and increased the average clearance with 39% with respect to the $\mathcal{W}$-retraction technique.

Corridor: While there is little room between the walls of the corridor and the robot, the techniques were still able to increase the clearance along the path. Again, the $\mathcal{C}$-retraction technique outperforms the $\mathcal{W}$-retraction technique but this takes much more computation time.

Wrench: The algorithms needed relatively much time as the environment was larger compared to the other ones. Both algorithms were successful in increasing the clearance. The $\mathcal{C}$-retraction technique performed slightly better with respect to increasing the average and maximum clearance. However, the $\mathcal{W}$-retraction technique was 6% better with respect to the minimum clearance.

Hole: The $\mathcal{W}$-retraction technique doubled the amount of minimum and average clearance along the path. The $\mathcal{C}$-retraction technique outperforms the $\mathcal{W}$-retraction technique because all DOFs were taken into account. The minimum amount of clearance along the path was further improved with 33%.

Manipulator: The minimum clearance along the initial path was very close to zero. The $\mathcal{C}$-retraction technique successfully introduced some clearance along the path. Although there is little room for the manipulator to move, the algorithm was able to double the average clearance along the path. This extra clearance may be crucial for high-cost environments to guarantee safety. For clarity, we only visualized a part of the sweep volume of the manipulator in Figure 5.

The running times indicate that improving the clearance along paths may be too slow to be applied online. However, in applications where safety is important the running times are not that crucial. For example, due to the difficulty of measuring and controlling the precise position of a manipulator arm, the arm can be damaged if it moves in the vicinity of obstacles. Improving the clearance at the cost of a few minutes of calculation time can prevent damage to the robot and its environment.

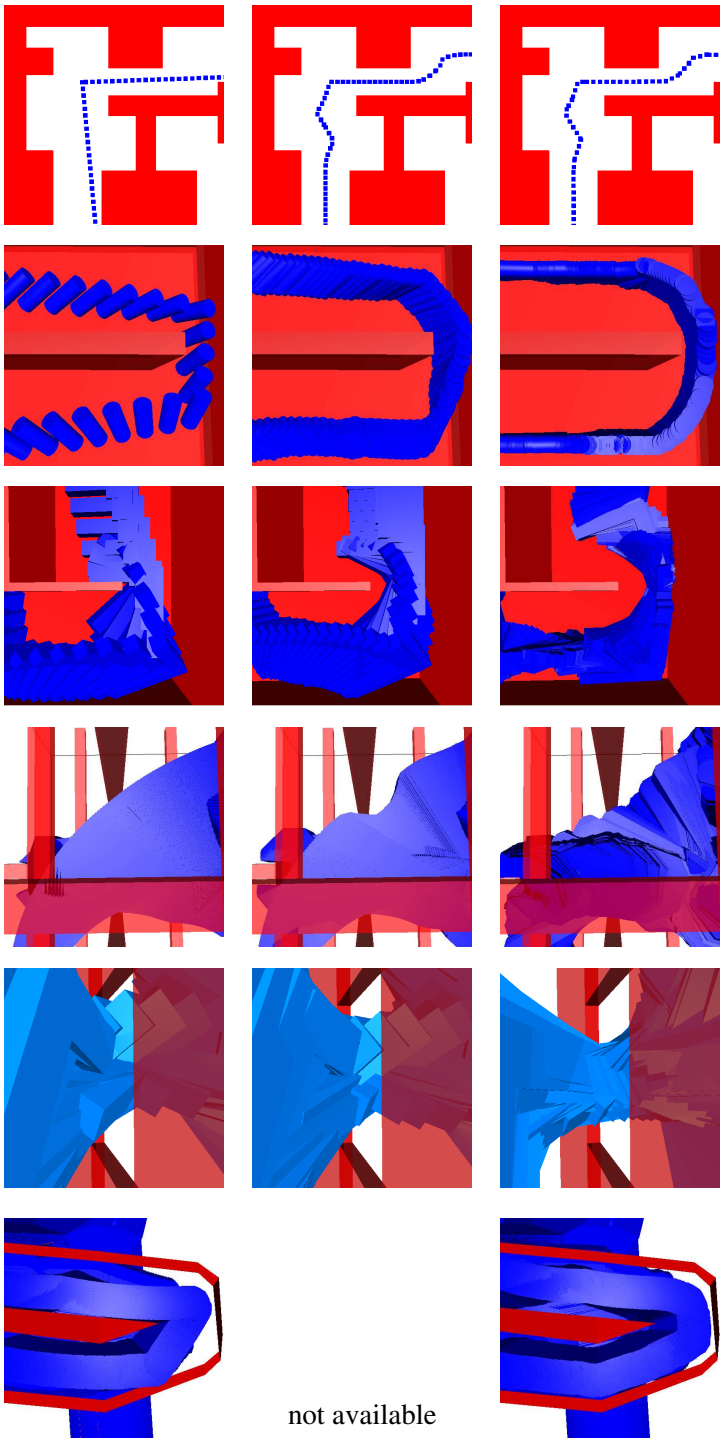

(a) Initial paths (b) $\mathcal{W}$-retraction (c) $\mathcal{C}$-retraction

Figure 5: A close up of the paths in the six environments. The pictures in the left column show pieces of the initial paths. The paths in the middle column are the result of the $\mathcal{W}$-retraction technique while the paths in the right have been created by one particular run of the $\mathcal{C}$-retraction technique.

Planar	Clearance			Time
	min	**avg**	**max**	**s**
Initial path	0.00	2.47	7.15	0.0
$\mathcal{W}$**-retraction**	1.79	4.48	8.32	0.8
$\mathcal{C}$**-retraction**	1.79	4.49	8.32	9.4

Simple corridor	Clearance			Time
	min	**avg**	**max**	**s**
Initial path	0.16	1.91	3.83	0.0
$\mathcal{W}$**-retraction**	0.62	2.62	3.96	0.7
$\mathcal{C}$**-retraction**	1.21	3.64	4.25	6.0

Corridor	Clearance			Time
	min	**avg**	**max**	**s**
Initial path	0.01	0.59	2.44	0.0
$\mathcal{W}$**-retraction**	0.22	1.15	3.22	1.0
$\mathcal{C}$**-retraction**	0.27	1.87	4.57	27.6

Wrench	Clearance			Time
	min	**avg**	**max**	**s**
Initial path	0.00	4.17	11.32	0.0
$\mathcal{W}$**-retraction**	2.11	7.12	12.38	12.4
$\mathcal{C}$**-retraction**	1.99	7.83	15.03	373.8

Hole	Clearance			Time
	min	**avg**	**max**	**s**
Initial path	0.28	1.81	5.97	0.0
$\mathcal{W}$**-retraction**	0.79	3.08	6.85	0.6
$\mathcal{C}$**-retraction**	1.05	3.44	7.24	12.7

Manipulator	Clearance			Time
	min	**avg**	**max**	**s**
Initial path	0.00	0.14	0.35	0.0
$\mathcal{W}$**-retraction**	n.a.	n.a.	n.a.	n.a.
$\mathcal{C}$**-retraction**	0.05	0.29	0.43	26.8

Table 4: Clearance statistics for the six environments. A larger clearance statistic indicates a better result. These statistics are averages over 100 independent runs.

5 Conclusions and Future Work

We presented a new algorithm that improves the clearance along paths traversed by a broad range of robots which may reside in high-dimensional configuration spaces. The algorithm retracts paths of planar, free-flying and articulated robots toward the ridges of the medial axis of the $\mathcal{C}$-space by a series of guided random walks. The algorithm improved upon existing algorithms since higher amounts of clearance for a larger diversity of robots are obtained.

A drawback of the algorithm is that it cannot be applied on paths in real-time. We expect that the running times can be dramatically decreased by incorporating learning techniques which will be a topic of future research. However, when on-line behavior is wanted, a complete graph can be retracted to the medial axis in the preprocessing phase. We showed that a path can be extracted from this graph in real-time.

We believe that this approach will enhance the quality of motion planners.

References

[1] J.-C. Latombe, *Robot Motion Planning.* Kluwer, 1991.

[2] S. LaValle, *Planning Algorithms.* http://msl.cs.uiuc.edu/planning, 2005.

[3] V. Boor, A. Kamphuis, C. Geem, E. Schmitzberger, and J. Bouchet, "Formalisation of path quality," Delivery MOLOG project, 2000.

[4] A. Kamphuis and M. Overmars, "Finding paths for coherent groups using clearance," in *Eurographics/ACM SIGGRAPH Symposium on Computer Animation*, 2004, pp. 19–28.

[5] J. Kim, R. Pearce, and N. Amato, "Extracting optimal paths from roadmaps for motion planning," in *IEEE International Conference on Robotics and Automation*, 2003, pp. 2424–2429.

[6] J.-M. Lien, S. Thomas, and N. Amato, "A general framework for sampling on the medial axis of the free space," in *IEEE International Conference on Robotics and Automation*, 2003, pp. 4439–4444.

[7] H. Choset and J. Burdick, "Sensor-based exploration: The hierarchical generalized Voronoi graph," *International Journal of Robotics Research*, vol. 19, no. 2, pp. 96–125, 2000.

[8] J. Vleugels and M. Overmars, "Approximating Voronoi diagrams of convex sites in any dimension," *International Journal of Computational Geometry & Applications*, vol. 8, pp. 201–221, 1998.

[9] E. Masehianand, M. Admin-Naseri, and S. Khadem, "Online motion planning using incremental construction of medial axis," in *IEEE International Conference on Robotics and Automation*, 2003, pp. 2928–2933.

[10] K. Hoff, T. Culver, J. Keyser, M. Lin, and D. Manocha, "Interactive motion planning using hardware-accelerated computation of generalized Voronoi diagrams," in *IEEE International Conference on Robotics and Automation*, 2000, pp. 2931–2937.

[11] R. Geraerts and M. Overmars, "Clearance based path optimization for motion planning," in *IEEE International Conference on Robotics and Automation*, 2004, pp. 2386–2392.

[12] N. Amato, O. Bayazit, L. Dagle, C. Jones, and D. Vallejo, "Choosing good distance metrics and local planners for probabilistic roadmap methods," in *IEEE International Conference on Robotics and Automation*, 1998, pp. 630–637.

[13] J. Kuffner, "Effective sampling and distance metrics for 3D rigid body path planning," in *IEEE International Conference on Robotics and Automation*, 2004, pp. 3993–3998.

[14] G. van den Bergen, *Collision Detection in Interactive 3D Environments*. Morgan Kaufmann, 2003.