# Reduced Complexity Multi-Scale Path-Planning on Probabilistic Maps 

Florian Hauer ${ }^{1}$ Panagiotis Tsiotras ${ }^{2}$


#### Abstract

We present several modifications to the previously proposed MSPP algorithm that can speed-up its execution considerably. The MSPP algorithm leverages a multiscale representation of the environment in $n$ dimensions. The information of the environment is stored in a tree data structure representing a recursive dyadic partitioning of the search space. The information used by the algorithm is the probability that a node in the tree corresponds to an obstacle in the search space. Such trees are often created from mainstream perception algorithms, and correspond to quadtrees and octrees in two and three dimensions, respectively. We first present a new method to compute the graph neighbors in order to reduce the complexity of each iteration, from $O\left(|V|^{2}\right)$ to $O(|V| \log |V|)$. We then show how to delay expensive intermediate computations until we know that new information will be required, hence saving time by not operating on information that is never used during the search. Finally, we present a way to remove the very expensive need to calculate a full multi-scale map with the use of sampling and derive an theoretical upperbound of the probability of failure as a function of the number of samples.


## I. Introduction

The path-planning problem of an agent moving in a $n$ dimensional environment full of obstacles is considered. It is assumed, without loss of generality, that the search space is contained within an $n$-dimensional hypercube. Each dimension of this hypercube is divided by successive dyadic partitioning. This partitioning generates a tree data structure that contains all the environmental data, organized hierarchically with increasing resolution at each successive depth. In the 2-D and the 3-D cases this multiresolution data structure is commonly known as a quadtree and an octree, respectively. The information stored in each node of the tree is the probability that the volume of the search space corresponding to this node is occupied by an obstacle. The problem is then to find an obstacle-free path between a start point and a goal point, or to report failure if no obstacle-free path exists.

Path-planning algorithms rely on perception algorithms to map the environment and localize the agent on the map. Commonly used data structures to represent perceived environments include multiresolution representations resulting from many common perception algorithms [3], [10]. The use of hierarchical multiresolution data structures is motivated by several observations: First, information collected about the environment is not uniform, as each sensor has its own range, resolution and noise properties. The information used by perception algorithms is then naturally multiscale; estimation

[^0]and inference is often used to extract the best information out of noisy measurements, leading to a probabilistic representation of the environment [10]. Second, on-board computational resources might be limited, thus not allowing an agent to systematically use all perceived information. Furthermore, precise information about the environment far away from the agent might not be valid, or may even be irrelevant, if the robot is far away from the obstacle. A multiscale representation of the collected data allows to choose the resolution for each region of the space as needed. For planning purposes, for instance, local information is typically important over the short-term (e.g., obstacle avoidance), while far away information affects only long-term objectives such as reaching a goal or exploring the environment.

Several approaches have been used in the past to incorporate multiscale information during planning. Bottomup approaches use the information at the finest resolution and then combine it in increasingly coarser resolutions. Top-down approaches solve the path-planning problem at the coarser resolution and then progressively increase the resolution of the solution [6], [8]. Using both approaches can lead to fast optimal algorithms, as shown in [7]. But the preprocessing of the data required by this approach is too expensive (in terms of time and memory) for online applications. Another approach consists in using the information at different resolutions at the same time. This idea is explored in [2], where areas near the current vehicle are represented accurately, while farther-away areas are coarsely-encoded by using a transformation on the wavelet coefficients. The approach is shown to be complete and very fast.

A similar approach is used to create a local map in [1], but quadtrees are used instead of wavelets and only local planning is done. Other algorithms have been developed using multi-resolution maps, but they are often applied to a given non-uniform grid, without using the information at different resolution scales for the same region of the search space [4], [9]. More recently, the MSPP algorithm [5] extended the work of [2] to $n$ dimensions in a reformulation using $2^{n}$ trees instead of wavelets. The notion of $\varepsilon$-obstacles guaranteeing completeness for any value of the threshold $\varepsilon$ was also introduced in the same paper.

In this paper, we propose several modifications to the MSPP algorithm introduced in [5] to accelerate computations and extend the range of potential applications. Neighbor checking, a bottleneck operation of the MSPP algorithm, is reworked to reduce its complexity from $O\left(|V|^{2}\right)$ to $O(|V| \log |V|)$. The order in which operations are executed is also modified in order to minimize the computations. The full reduced graph is not computed, only its nodes are computed; the edges are calculated on-the-fly, which saves
computations; since most edges are never used by the search algorithm. Finally, we introduce a way to work without a multi-scale map, but instead we use a predicate to determine whether a point of the search space is an obstacle or not. This modification allows us to remove the expensive map creation step. It also saves memory, since the full multiscale map does not need to be known in advance. The trade-off here is giving up completeness for the sake of increased runtime performance and memory reduction, as well as an extension of the range of potential application where accurate information about the environment is not known and is collected incrementally via sampling.

## II. Notation and Previous Work

In this section, we briefly introduce the notation used in the rest of the paper. We refer the reader to [5] for more details on the original MSPP algorithm and the notation used throughout the rest of the paper.

## A. Multiresolution World Representation

The environment $\mathcal{W} \subset \mathbb{R}^{d}$ is assumed to be a $d$ dimensional grid world. Without loss of generality, we assume that each elementary cell of the grid world is a unit hypercube and there exists an integer $\ell>0$ such that the world is contained within a hypercube of side length $2^{\ell}$. The world $\mathcal{W}$ is encoded as a tree $\mathcal{T}=(\mathcal{N}, \mathcal{R})$ representing the multi-scale information, with $\mathcal{N}$ being the set of nodes and $\mathcal{R}$ being the set of edges describing their relations. Nodes of $\mathcal{T}$ are represented by two indices, $k$ and $p$, corresponding to the depth of the node $n_{k, p}$ in the tree and the position of the center of the node in the search space. A node $n_{k, p}$ represents a hypercube in the search space $\mathcal{W}$ centered at $p$ and of size $2^{k}$, and is denoted by $H\left(n_{k, p}\right)$. The children of the node $n_{k, p}$ are denoted by $n_{k-1, q_{i}}, i \in\left[1,2^{d}\right]$ where $q_{i}=p+2^{k-2} e_{i}$ and where $e_{i}$ is each of the $2^{d}$ ( $d$-dimensional) vectors generated by $[ \pm 1, \pm 1, \ldots, \pm 1]$. A node is called a leaf of $\mathcal{T}$ if it has no children.

The information $V\left(n_{k, p}\right)$ contained in each node $n_{k, p} \in$ $\mathcal{N}$ is the probability of the existence of an obstacle in the space represented by the node, computed by

$$
\begin{equation*}
V\left(n_{k, p}\right)=\frac{\text { Volume of obstacles in } H\left(n_{k, p}\right)}{\text { Volume of } H\left(n_{k, p}\right)} \tag{1}
\end{equation*}
$$

## B. The Path-Planning Problem

Two nodes in $\mathcal{T}$ are neighbors if their corresponding hypercubes share a hyperface, specifically, their intersection is a hypercube of dimension $d-1$. A necessary and sufficient condition for two nodes $n_{k_{1}, p_{1}}$ and $n_{k_{2}, p_{2}}$ to be neighbors is that both of the following two conditions are satisfied:

- The expression

$$
\begin{equation*}
\left\|p_{1}-p_{2}\right\|_{\infty}=2^{k_{1}-1}+2^{k_{2}-1} \tag{2}
\end{equation*}
$$

holds,

- There exists a unique $i \in[1, d]$, such that

$$
\begin{equation*}
\left|\left(p_{1}-p_{2}\right)_{i}\right|=2^{k_{1}-1}+2^{k_{2}-1} \tag{3}
\end{equation*}
$$

where $\left(p_{1}-p_{2}\right)_{i}$ is the $i^{\text {th }}$ component of the vector.

In the case of a 2-D or a 3-D environment with a uniform grid, conditions (2) and (3) imply 4-connectivity or 6-connectivity, respectively.

We define a path $\pi=\left(n_{k_{1}, p_{1}}, n_{k_{2}, p_{2}}, \ldots, n_{k_{N}, p_{N}}\right)$ in $\mathcal{T}$ to be a sequence of nodes $n_{k_{i}, p_{i}} \in \mathcal{N}$, each at corresponding position $p_{i}$ and depth $\ell-k_{i}$, such that two consecutive nodes of the sequence are neighbors. A path is called a finest information path (FIP) if all its nodes are leafs of $\mathcal{T}$. Leaf nodes represent the best resolution, and hence the finer information contained in the tree $\mathcal{T}$.

Due to noisy measurements, and to prevent overconfidence and numerical issues, perception algorithms will not allow $V\left(n_{k, p}\right)$ to reach $100 \%$. We therefore introduce the notion of $\varepsilon$-obstacles. Specifically, given $\varepsilon \in[0,1)$, a node $n_{k, p} \in \mathcal{N}$ is an $\varepsilon$-obstacle if

$$
\begin{equation*}
V\left(n_{k, p}\right) \geq 1-2^{-d k} \varepsilon \tag{4}
\end{equation*}
$$

A path $\pi$ is $\varepsilon$-feasible if none of its nodes are $\varepsilon$-obstacles.
Given the representation of $\mathcal{W}$ encoded in the tree $\mathcal{T}$, the problem is to find an $\varepsilon$-feasible FIP between two nodes in the tree, $n_{\text {start }}$, representing the starting node, and $n_{\text {goal }}$, representing the goal node, and to report failure if no such path exists.

## C. The MSPP Algorithm

The MSPP algorithm is a backtracking algorithm that iteratively builds a solution from the starting node $n_{\text {start }}$ until the goal node $n_{\text {goal }}$ is reached. At each iteration $i$, a local representation $\mathcal{G}_{i}$ of the environment, called the reduced graph, is computed and the best path to the goal on this graph is used as a heuristic to decide which direction to follow. The current candidate solution built by the algorithm at iteration $i$ is denoted by $\pi_{\mathrm{start}}^{i}$. Therefore, $\pi_{\mathrm{start}}^{i}$ is an $\varepsilon$-feasible FIP from $n_{\text {start }}$ to $n_{k_{i}, p_{i}}$, the node reached by the algorithm at iteration $i$. The best path from $n_{k_{i}, p_{i}}$ to the goal on the reduced graph $\mathcal{G}_{i}$ is denoted by $\pi_{i}^{\text {goal }}$. The reduced graph $\mathcal{G}_{i}$ is computed by first identifying its vertices (corresponding to nodes of $\mathcal{T}$ ) via a top-down exploration of $\mathcal{T}$. The first element of $\pi_{i}^{\text {goal }}$ is used to build the global solution $\pi_{\text {start }}^{i}$. If the path $\pi_{i}^{\text {goal }}$ does not exist, the algorithm backtracks.

The main lines of the MSPP algorithm are shown in Algorithm 1] Refer to [5] for full details of the MSPP algorithm.

## III. Fast Neighbor Computation - MSPP-FN

In the original MSPP algorithm [5], neighbors for the reduced graph are computed by testing whether every pair of vertices satisfies the neighborhood properties. As shown in [5], this step is the bottleneck during each iteration, having complexity $O\left(|V|^{2}\right)$, where $|V|$ is the number of vertices. A new way to compute neighbors by reducing the complexity from $O\left(|V|^{2}\right)$ to $O(|V| \log |V|)$ proceeds as follows; for each vertex, neighbor candidates are generated, and the existence of these candidate vertices of $\mathcal{G}_{i}$ is verified. The details are given next. MSPP-FN will be used in the rest of the paper to refer to the MSPP algorithm using the fast neighbor computation.

```
Algorithm 1: The MSPP Algorithm - Simplified for
clarity
    Data: Tree \(\mathcal{T}\), Start node \(n_{\text {start }}\), Goal node \(n_{\text {goal }}\)
    Result: \(\varepsilon\)-feasible FIP from \(n_{\text {start }}\) to \(n_{\text {goal }}\) or failure
    \(i \leftarrow 0, n_{k_{i}, p_{i}} \leftarrow n_{\text {start }}, \pi_{\text {start }}^{0} \leftarrow\left[n_{k_{i}, p_{i}}\right]\);
    while Goal not found AND no failure do
        \(\left(\tilde{\mathcal{G}}_{i}, v_{\text {start }, i}, v_{\text {goal }, i}\right) \leftarrow\) ReducedGraph \(\left(T, n_{k_{i}, p_{i}}\right)\);
        \(\pi_{i}^{\text {goal }} \leftarrow\) ShortestPath \(\left(\tilde{\mathcal{G}}_{i}, v_{\text {start }, i}, v_{\text {goal }, \mathrm{i}}\right)\);
        if \(\operatorname{exists}\left(\pi_{i}^{\text {goal }}\right)\) then
            \(n_{k_{i+1}, p_{i+1}} \leftarrow\) firstElement \(\left(\pi_{i}^{\text {goal }}\right)\);
            \(\pi_{\text {start }}^{i+1} \leftarrow\left[\begin{array}{ll}\pi_{\text {start }}^{i} & n_{k_{i+1}, p_{i+1}}\end{array}\right]\);
        else
            backtrack;
        \(i \leftarrow i+1 ;\)
```


## A. Tree Data Structure For Vertices

In order to do fast searches over the vertices of the reduced graph, we keep them in a tree structure. Let $\mathcal{T}_{i}$ define this tree structure. As the original tree $\mathcal{T}$ is traversed to select nodes for $\mathcal{G}_{i}, \mathcal{T}_{i}$ is constructed by copying every element of $\mathcal{T}$ traversed by the selection process, except for $\varepsilon$-obstacles. $\mathcal{T}_{i}$ is then a tree with the same structure as $\mathcal{T}$, but its branches are shorter. In other words, $\mathcal{T}_{i}$ is a pruned version of $\mathcal{T}$ whose leave nodes are the vertices of $\mathcal{G}_{i}$.

Note that, for implementation, $\mathcal{T}_{i}$ does not change significantly between two consecutive iterations, so it is computationally cheaper to modify $\mathcal{T}_{i-1}$ than to create a new data structure at each iteration. Memory allocation is the most expensive operation when creating new nodes. Modifying $\mathcal{T}_{i-1}$ allows to only have to allocate memory for nodes of $\mathcal{T}_{i}$ that did not exist in $\mathcal{T}_{i-1}$. The copy process is then modified to add nodes only if they do not already exist, and remove excessive nodes when reaching a node corresponding to a vertex of $\mathcal{G}_{i}$. The pseudo-code is given in Function 1. The vertex list is also removed since the information is already contained in $\mathcal{T}_{i}$. The function GetRGFastNeighbor is called with the root of $\mathcal{T}$, the root of $\mathcal{T}_{i}$ (created during initialization) and the current node $n_{k_{i}, p_{i}}$.

## B. Same Size Neighbors

Generating neighbors is easy when the nodes have the same size, so we will consider this case first. Given a node $n_{k, p}$, we want to find all its neighbors having the same size that correspond to vertices of $\mathcal{G}_{i}$. Same size implies the same depth in $\mathcal{T}$, so every neighbor will have the same depth index $k$. Also, the neighbor conditions and the fact that the nodes are centered on a grid, imply that only one dimension of the position vector can be changed at a time, that is, the neighbors' positions $p_{\text {nhb }, i}$ can only be

$$
p_{\mathrm{nhb}, i}=p+2^{k} b_{i}, 1 \leq i \leq 2 d
$$

with

$$
b_{i}= \begin{cases}d_{i} & \text { if } i \leq d \\ -d_{i} & \text { otherwise }\end{cases}
$$

```
Function 1: GetRGFastNeighbor ()
    Data: Node \(n_{k, p}\) (in \(\mathcal{T}\) ), Node \(t_{k, p}\) (in \(\mathcal{T}_{i}\) ), Current
            node \(n_{k_{i}, p_{i}}\)
    if \(\left(\left\|p-p_{i}\right\|_{2}-\frac{\sqrt{d}}{2} 2^{k_{i}} \geq \alpha 2^{k}\right.\) OR \(\left.\operatorname{isLeaf}\left(n_{k, p}\right)\right)\) AND
    doesNotContainPath \(\left(n_{k, p}\right)\) then
            if \(n_{k, p}\) is not a \(\varepsilon\)-obstacle then
            Remove all descendants of \(t_{k, p}\);
        else
            Remove \(t_{k, p}\) and its descendants;
    else
        foreach \((m, q)\) index of children of \((k, p)\) do
            if \(t_{m, q}\) does not exist then
                Create \(t_{m, q}\) child of \(t_{k, p}\);
            GetRGFastNeighbor \(\left(n_{m, q}, t_{m, q}, n_{k_{i}, p_{i}}\right)\);
```



Fig. 1: Generating neighbors for the construction of $\mathcal{G}_{i}$. Same size neighbors case: $H\left(n_{k, p}\right)$ is the center square and the $p_{\mathrm{nhb}, i}$ are the generated position candidates for the neighbors.
where $d_{i}$ is the $i^{t h}$ vector of the standard basis of $\mathbb{R}^{d}$. If $p_{\mathrm{nhb}, i}$ is within the bounds of the search space, $n_{k, p_{\mathrm{nhb}, i}}$ is in $\mathcal{T}_{i}$ and $n_{k, p_{\mathrm{nhb}, i}}$ is a leaf of $\mathcal{T}_{i}$, then $n_{k, p_{\mathrm{nhb}, i}}$ is a valid neighbor of $n_{k, p}$. Figure 1 shows in the center the hypercube corresponding to $n_{k, p}$ and the neighbor candidates around it. The red arrows represent the vectors $2^{k} b_{i}$.

Searching the tree $\mathcal{T}_{i}$ can be done on average in $O(\log |V|)$, and the number of candidates to check is $2 d$.

## C. Larger Neighbors

Consider now the case of finding the larger neighbors of $n_{k, p}$. The previous result can still be used, but it will generate points inside larger neighbors instead of their positions. The search through the tree works as follows. It starts with the root of the tree, representing the entire environment, as the current node. As long as the current node has children (recall that our data structure assumes that they either all exist or none of them exists), the child whose hybercube contains the searched point $p_{\text {nhb, } i}$ is selected as the current node. That is, at each step, the search process selects the node at the next level of resolution whose hypercube contains $p_{\text {nhb }, i}$. The search stops if either the current node is at $p_{\text {nhb }, i}$ or the current node does not have children. At the end of the process, the current node is a neighbor of $n_{k, p}$ and if it is


Fig. 2: Generating neighbors for the construction of $\mathcal{G}_{i}$. Larger neighbors case: $H\left(n_{k, p}\right)$ is the smaller square around $p$ and $p_{\mathrm{nhb}, 1}$ is the first generated position candidate for the neighbors. The dashed squares represents the hypercubes corresponding to nodes visited during the search for $p_{\text {nhb, } 1}$ in the tree.
a leaf, then it is also a vertex of $\mathcal{G}_{i}$. A larger node could contain $p$ and then not be a neighbor, but since $n_{k, p}$ exists, that node would have children and the search process will never stop in such a situation.

Figure 2 shows, in dashed lines, the last four nodes that would be explored while searching for $p_{\text {nhb }, 1}$. If $n_{k, p_{\mathrm{nhb}, 1}}$ does not exist, the algorithm will stop at one of its ancestors, which will be a neighbor of $n_{k, p}$. Note that the search cannot stop at the largest ancestor shown, since it contains $n_{k, p}$, so all the children exist.

## D. Smaller Neighbors

The last case to consider is when there are smaller neighbors of $n_{k, p}$. The search for $p_{\mathrm{nhb}, i}$ in $\mathcal{T}_{i}$ will return a node that is not a leaf. In this case, the exploration of


Fig. 3: Generating neighbors for the construction of $\mathcal{G}_{i}$. Smaller neighbors case: $H\left(n_{k, p}\right)$ is the left square and $p_{\text {nhb, } 1}$ is the first generated position candidate for the neighbors. The larger dashed square is $H\left(n_{k, p_{\mathrm{nhb}, 1}}\right)$ and the blue squares correspond to the descendant of $n_{k, p_{\mathrm{nhb}, 1}}$ that are neighbors of $n_{k, p}$.
children of $p_{\text {nhb }, i}$ can lead to the neighbors. Note that $p_{\mathrm{nhb}, i}$ was generated by moving in the direction $b_{i}$, but since the neighbors are smaller, the move was too large, and hence
neighbors of $n_{k, p}$ are leaf nodes, descendant of $p_{\text {nhb,i }}$ in the direction $-b_{i}$.

Figure 3 shows what happens for smaller neighbors. The larger dashed square is the hypercube corresponding to the candidate neighbor $p_{\text {nhb, } 1}$, but that node is not a leaf, so it is not a vertex of $\mathcal{G}_{i}$. Exploring its children (until leaf nodes) in the direction $-b_{1}=-d_{1}$, will lead to all its descendants that are neighbor with $n_{k, p}$, and in $\mathcal{G}_{i}$, since they will be leaf nodes. The neighbors are drawn in blue in Figure 3 .

## E. Computing All Neighbors in $\mathcal{G}_{i}$

When looking for all neighbors, nodes can be treated in any order, in particular, from smallest to largest. For the smallest nodes, all neighbors are larger. If all smaller nodes have been treated before, for a given node $n_{k, p}$, the smaller neighbors will already have been found and the only information missing is the larger nodes. All neighbor pairs can then be found by looking for larger neighbors for each node ordered from the smallest to the largest. Finding larger and same size neighbors is done in $O(\log |V|)$ for each of the $|V|$ nodes, so finding all neighboring pairs in $\mathcal{G}_{i}$ is then be done in $O(|V| \log |V|)$.

## F. Computing All Neighbors of a Given Node $n_{k, p}$

Finding all neighbors of a given node $n_{k, p}$ can be done using the pseudo-code in Function 2 For each direction $b_{i}$, we compute the candidate neighbor position $p_{\mathrm{nhb}, i}$ and search in $\mathcal{T}_{i}$ for the corresponding node. If the node is a leaf, it means that a larger or same size neighbor has been found; otherwise, the leaf descendants, in the direction $-b_{i}$, of the node found are smaller size neighbors.

```
Function 2: findNeighbors()
    Data: Node \(n_{k, p}\)
    neighbors= \(\emptyset\);
    foreach \(i\) in \([1,2 d]\) do
        \(n_{m, q}=\operatorname{find}\left(\mathcal{T}_{i}, p_{\mathrm{nhb}, \mathrm{i}}\right)\);
        if \(\operatorname{isLeaf}\left(n_{m, q}\right)\) then
            neighbor \(s=\) neighbor \(s \cup n_{m, q}\);
        else
            addLeafInDir( \(n_{m, q},-b_{i}\), neighbors);
    return neighbors;
```

```
Function 3: addLeafInDir()
    Data: Node \(n_{k, p}\), Direction b, List neighbors
    foreach \(i\) in \(\left[1,2^{d}\right]\) do
        if \(b^{T} e_{i}>0\) then
            \(n=\operatorname{child}\left(n_{k, p}, i\right)\);
            if isLeaf( \(n\) ) then
                neighbor \(s=n e i g h b o r s \cup n\);
            else
                addLeafInDir( \(n, b\), neighbors \()\);
```


## IV. Multi-Scale Path Planning Without Full Information Map - MSPP-S

Although in 2D or 3D geometric workspaces, the multiscale map is often the result of perception algorithms, this is not always the case. When the search space is the configuration space and it is different from the geometric workspace, computing the multi-scale map might be very expensive, as it requires to analyze every single cell of the map. Furthermore, in some cases, we may only have access to a predicate of whether a point of the search space is an obstacle or not. A robotic arm, for example, is usually parameterized by the position of each joint; given a configuration, the spatial position of each link can be computed, and self-collision or collision with obstacles is checked in the geometric workspace. It is assumed in this section that we have such a predicate, say isObstacle $(s)$, that informs us if a point $s$ of the search space is an obstacle.

In the proposed approach, sampling is used to estimate the obstacle probabilities of the nodes in $\mathcal{G}_{i}$. Since we are using an estimate instead of the exact node probability values, completeness of the algorithm is not ensured. Note, however, that if a large enough number of samples is drawn, the estimated probabilities will be close to their actual values, and loss of completeness is very unlikely.

Similarly to the original MSPP algorithm, the proposed algorithm, MSPP-S (as MSPP with sampling), decomposes the space using a grid, which has fine resolution near the current position, and the resolution becomes increasingly coarser farther away. An empty tree data structure $\mathcal{T}_{i}$ is created to represent that grid. It is empty in the sense that it does not have any information about the obstacles, it is a pure geometric partition of the search space. For each node $n_{k, p}$ of the partition, the predicate can be used for a given number $N_{\text {samples }}$ of random points drawn in the search space corresponding to the node. An estimate of the probability of obstacles can then be calculated from those results and used to fill up the tree $\mathcal{T}_{i}$. The value of the node is approximated by

$$
\hat{V}\left(n_{k, p}\right)=\frac{\text { Number of obstacles sampled }}{N_{\text {samples }}} .
$$

Similarly to the data structure used for neighbor checking, it is less costly to modify the data structure from the previous iteration than to recreate a new data-structure at each iteration. Moreover, in that case, some information will already exist in the data structure, and sampling only needs to be done for the newly added nodes.

Note that the data structure created is similar to the one created for neighbor checking, hence the same notation $\mathcal{T}_{i}$. Since only the structure matters for neighbor checking, and not the actual values, this data structure can also be used for the neighbor checking step.

The pseudo-code for the vertex selection is given in Function 4 and the edges can be computed as described in Section III-E

## V. Minimal Reduced Graph Construction

Constructing $\mathcal{G}_{i}$ can be costly and only part of the information might be used at each iteration to solve for the

```
Function 4: GetRGVerticesWithSampling()
    Data: Node \(t_{k, p}\) (in \(\mathcal{T}_{i}\) ), Current node \(n_{k_{i}, p_{i}}\)
    if \(\left(\left\|p-p_{i}\right\|_{2}-\frac{\sqrt{d}}{2} 2^{k_{i}} \geq \alpha 2^{k}\right.\) AND
    doesNotContainPath \(\left(t_{k, p}\right)\) then
            Remove all descendants of \(t_{k, p}\);
            if \(n_{k, p}\) has not been sampled yet then
                    Sample \(N_{\text {samples }}\) in \(H\left(n_{k, p}\right)\);
                Estimate \(\hat{V}\left(n_{k, p}\right)\);
    else
        foreach \((m, q)\) index of children of \((k, p)\) do
            if \(t_{m, q}\) does not exist then
                Create \(t_{m, q}\) child of \(t_{k, p}\);
            GetRGVerticesWithSampling \(\left(t_{m, q}, n_{k_{i}, p_{i}}\right)\);
```

shortest path. In this work, it is assumed that the planning problem on $\mathcal{G}_{i}$ is solved using the $\mathrm{A}^{*}$ algorithm although this is not restrictive. In the $\mathrm{A}^{*}$ algorithm, nodes are kept in a priority queue, called $O P E N$, ordered by $f$-values, where $f=g+h$ with $g$ the cost-to-go and $h$ an admissible heuristic to the goal. While $O P E N$ has elements, the first element is removed and for each of the neighbors, if they have not been closed yet, the $g$-value is updated, and it is added to the $O P E N$ priority queue. The algorithm stops when the first element of the $O P E N$ priority queue is the goal.

In the $\mathrm{A}^{*}$ algorithm, knowing the neighbors of a node is only useful when that node is taken out of the $O P E N$ queue. Similarly, the obstacle probability is only needed to calculate the $g$-value of a node.
By delaying those calculations until the necessary information is required, improvement in execution speed is expected. The following changes allow to save computations in the new algorithm:

- the ReducedGraph function only computes the nodes of the reduced graph
- during the $A^{*}$ algorithm, neighbors of a node are computed when the node is selected from the OPEN priority queue to be explored. If sampling is being used, sampling is only made the first time the $g$-value is calculated.
At the end, the algorithm will only have calculated the neighbors for the nodes in the $C L O S E$ list, and estimated the obstacle probability for nodes in $O P E N \cup C L O S E$. In the worst case, the $\mathrm{A}^{*}$ algorithm will explore every vertex and every edge, so all neighbors will be calculated and all nodes will be sampled similarly to the naïve case. But, in general, the number of neighbors calculated and the number of node sampled will be largely reduced compared to the naïve case. This is confirmed by the numerical examples in the next section.


## VI. Probabilistic bounds of MSPP-S

As the estimate $\hat{V}\left(n_{k, p}\right)$ is used instead of the actual obstacle probability $V\left(n_{k, p}\right)$, a bad estimate could lead to missing solutions and losing completeness of the algorithm. In particular, if $\hat{V}\left(n_{k, p}\right)$ overestimates $V\left(n_{k, p}\right)$, the node
$n_{k, p}$ might wrongly be evaluated as a $\varepsilon$-obstacle which would prevent the algorithm from finding any path passing through $n_{k, p}$, and potentially the only solution, thus breaking the completeness of the algorithm.

In this section, we derive an analytic worst case bound for the probability of failure of the MSPP-S algorithm.

We assume that there is an underlying grid and that each unit cell is either free space or obstacle.

## A. Definitions

To deal with the probability of misevaluating the obstacle probability of a cell, we redefine the notion of obstacles with a threshold $\gamma$ and the event of wrongly evaluating a node as an obstacle.

Definition 1. Let $\varepsilon, \gamma>0$. A node $n_{k, p}$ is $a \varepsilon, \gamma$-obstacle if

$$
\begin{equation*}
\hat{V}\left(n_{k, p}\right) \geq 1-2^{-d k} \varepsilon+\gamma \tag{5}
\end{equation*}
$$

Definition 2. Let $\varepsilon, \gamma>0 . M\left(n_{k, p}\right)$ is the event that the node $n_{k, p}$ is a $\varepsilon, \gamma$-obstacle and is not a $\varepsilon$-obstacle.
B. Bounds on $P\left(M\left(n_{k, p}\right)\right)$

Proposition 1. Let $\varepsilon, \gamma>0$ and $n$ the number of sampled points in a node $n_{k, p}$. Then

$$
\begin{equation*}
P\left(M\left(n_{k, p}\right)\right) \leq \exp \left(-2 \gamma^{2} n\right) \tag{6}
\end{equation*}
$$

Proof. Suppose $M\left(n_{k, p}\right)$ is true. Then 4 and 5 are verified, then from (5)-(4), we get

$$
\begin{equation*}
\hat{V}\left(n_{k, p}\right)-V\left(n_{k, p}\right) \geq \gamma \tag{7}
\end{equation*}
$$

Suppose that $n_{k, p}$ is composed of $N$ unit cells, including $N_{o}$ obstacles. Let $\mu=N_{o} / N$, then $V\left(n_{k, p}\right)=N_{o} / N=\mu$. Uniformly sampling a random point in $n_{k, p}$ is similar to uniformly picking one of the $N$ unit cells, that is, an obstacle is picked with probability $\frac{N_{o}}{N}=\mu$. Let $x_{i}$ be the random variable associated with the $i^{\text {th }}$ obstacle test. $x_{i}$ takes value 1 for obstacles and 0 for free space, that is $x_{i}$ is a Bernouilli trial. Let $n_{o}=\sum_{i=1}^{n} x_{i} . n_{o}$ is the number of successes in $n$ Bernoulli trials, so $n_{o}$ follows a binomial distribution. Using (7), changing variables and using Hoeffding's inequality (111) to $\boxed{12})$, we get that

$$
\begin{align*}
& P\left(M\left(n_{k, p}\right)\right)  \tag{8}\\
\leq & P\left(\hat{V}\left(n_{k, p}\right)-V\left(n_{k, p}\right) \geq \gamma\right)  \tag{9}\\
= & P\left(\frac{n_{o}}{n}-\mu \geq \gamma\right)  \tag{10}\\
= & P\left(n_{o} \geq n(\gamma+\mu)\right)  \tag{11}\\
\leq & \exp \left(-2 \gamma^{2} n\right) . \tag{12}
\end{align*}
$$

Proposition 2. Let a node $n_{k, p}$. If $k \geq k_{\max }=\left\lceil\frac{1}{d} \log _{2} \frac{\varepsilon}{\gamma}\right\rceil$, then $n_{k, p}$ cannot be a $\varepsilon, \gamma$-obstacle and $M\left(n_{k, p}\right)$ never happens.

Proof.

$$
\begin{align*}
& k \geq k_{\max }=\left\lceil\frac{1}{d} \log _{2} \frac{\varepsilon}{\gamma}\right\rceil \geq \frac{1}{d} \log _{2} \frac{\varepsilon}{\gamma}  \tag{13}\\
\Rightarrow & 2^{d k} \geq \frac{\varepsilon}{\gamma}  \tag{14}\\
\Rightarrow & \gamma \geq \varepsilon 2^{-d k}  \tag{15}\\
\Rightarrow & 1-2^{-d k} \varepsilon+\gamma \geq 1 \tag{16}
\end{align*}
$$

Proposition 3. Let a node $n_{k, p}$. If $k \leq k_{\min }=\left\lfloor\frac{1}{d} \log _{2} n\right\rfloor$, then computing $V\left(n_{k, p}\right)$ is less expensive than computing $\hat{V}\left(n_{k, p}\right)$. So $\varepsilon$-obstacles can be used and $M\left(n_{k, p}\right)$ never happens.
Proof. The cost of computing $V\left(n_{k, p}\right)$ is $2^{d k}$. The cost of computing $\hat{V}\left(n_{k, p}\right)$ is $n$.

$$
\begin{align*}
& k \leq k_{\min }=\left\lfloor\frac{1}{d} \log _{2} n\right\rfloor \leq \frac{1}{d} \log _{2} n  \tag{17}\\
\Rightarrow & 2^{d k} \leq n \tag{18}
\end{align*}
$$

## C. Probability of failure of MSPP-S

We assume here the MSPP-S algorithm uses $\varepsilon, \gamma$-obstacles when $k>k_{\min }$ and $\varepsilon$-obstacles otherwise. We also assume that the algorithm evaluates $\hat{V}\left(n_{k, p}\right)$ at most once for each node, it is evaluated when the value is needed and then the value is kept in memory.
Proposition 4. Given $\varepsilon, \gamma, n>0$, an upperbound on the probability of failure of MSPP-S is given by

$$
\begin{equation*}
P(\text { failure }) \leq 1-\left(1-\exp \left(-2 \gamma^{2} n\right)\right)^{n b_{o c c}} \tag{19}
\end{equation*}
$$

where

$$
\begin{equation*}
n b_{o c c}=\frac{2^{d\left(\ell-k_{\min }\right)}-2^{d\left(\ell-k_{\max }+1\right)}}{2^{d}-1} \tag{20}
\end{equation*}
$$

Proof. $M\left(n_{k, p}\right)$ can happen for every node such that $k_{\min }<$ $k<k_{\max }$, that is, the maximum number of occurences $n b_{o c c}$ of $M\left(n_{k, p}\right)$ is the number of nodes verifying $k_{\text {min }}<k<$ $k_{\text {max }}$,

$$
\begin{align*}
n b_{o c c} & =\sum_{k_{\min }<k<k_{\max }} 2^{d(\ell-k)}  \tag{21}\\
& =\frac{2^{d\left(\ell-k_{\min }\right)}-2^{d\left(\ell-k_{\max }+1\right)}}{2^{d}-1} \tag{22}
\end{align*}
$$

If $M\left(n_{k, p}\right)$ never happens, the algorithm will never discard potential paths, and stay complete, that is,

$$
\begin{align*}
P(\text { success }) & \geq P\left(M\left(n_{k, p}\right) \text { never happens }\right)  \tag{23}\\
& =P\left(\neg M\left(n_{k, p}\right)\right)^{n b_{o c c}} . \tag{24}
\end{align*}
$$

Hence,

$$
\begin{equation*}
P(\text { failure }) \leq 1-\left(1-\exp \left(-2 \gamma^{2} n\right)\right)^{n b_{o c c}} \tag{25}
\end{equation*}
$$



Fig. 4: Bound on the probability of failure with the parameters $\ell=5, d=1, \varepsilon=90 \%, \gamma=0.35 \%$ and $Z=2$.

## D. Upperbound when multiple independent solutions exist

Suppose there exists $Z$ independant solutions to the planning problem, that is, the search space can de partitionned in $Z$ regions in which there exists at least one solution. Suppose for simplificity that all regions have the same size, $1 / Z^{\text {th }}$ of the original space.
The probability of failing in one region is then bounded by
$P($ failure in one region $) \leq 1-\left(1-\exp \left(-2 \gamma^{2} n\right)\right)^{n b_{o c c} / Z}$.
The probability of the algorithm failing is smaller that the algorithm failing in each regions simultaneously since there might exist solutions crossing regions, so

$$
\begin{equation*}
P(\text { failure }) \leq\left(1-\left(1-\exp \left(-2 \gamma^{2} n\right)\right)^{n b_{o c c} / Z}\right)^{Z} \tag{27}
\end{equation*}
$$

Fig 4 shows the upperbound on the probability of failure as a function of the number of sampled points $n$ for a given set of parameters of the algorithm. For small $n$, the probability is very close to 1 since we have a very poor precision in the estimate of the obstacle probabilities. As $n$ grows, the probabilty for every single cell gets better hence the probability of failure decreases. The grow of $n$ also increases the value of $k_{\text {min }}$, thus creating drops in the maximum number of occurences of $M\left(n_{k, p}\right)$ and in the probability of failure. As $k_{\min }$ reaches $k_{\max }-1$ the maximum number of occurences of $M\left(n_{k, p}\right)$ goes to 0 and the probability of failure of the algorithm then becomes 0 .

The upperbounds derived are very conservative in the sense that:

- in most cases, only part of the nodes $n_{k, p}$ with $k_{\text {min }}<$ $k<k_{\max }$ are evaluated;
- if $M\left(n_{k, p}\right)$ happens for a node that is not part of the solution, the solution will still be found by the algorithm;
- if multiple solutions exist but are not independent, the probability is still largely reduced, but not exponentially;
- in typical environement, large areas of free space exists, thus multiplying the number of possible solutions and largely reducing the probability of failure.

In practise, failure of the MSPP-S algorithm has not been observed.

## VII. Results

## A. Comparison in Random Environemnts

In this section, we compare the original MSPP algorithm against the proposed extensions and also against the $\mathrm{A}^{*}$ run on a uniform grid. MSPP-FN refers to the variant with the new neighbor test and MSPP-S refer to the variant using sampling (it also uses fast neighbors). Obstacle maps were randomly generated and then used to solve path-planning problems via these four algorithms. The problem was solved for dimensions ranging from 2 to 5 with a tree depth of 5 , that is, for search spaces ranging from $2^{2 \times 5}=1024$ to $2^{5 \times 5} \simeq 3 \times 10^{7}$.


Fig. 5: Comparison of the execution time of the $\mathrm{A}^{*}$, MSPP and MSPP-FN algorithms. The results are shown in logarithmic scale.

Figure 5 shows the average execution time (in log scale) of the MSPP, the MSPP-FN and the $\mathrm{A}^{*}$ algorithms on the randomly generated maps. In this figure, the time to create the map is not taken into account in order to compare the pure performance of the planning algorithms, that is, it is just the time to find a path on an already existing map.

For the smaller search spaces, we see very few differences between all the algorithms, as expected. As the dimension and the size of the search space grow however, the MSPP algorithm becomes much faster than the $\mathrm{A}^{*}$, by more than two orders of magnitude in dimension 5. By the same token, the MSPP-FN algorithm is even faster (by $50 \%$ ) over the baseline MSPP algorithm in dimension 5.

In Figure 6, the cost of creating the map is taken into account. This is done in order to compare the results of using the MSPP-S algorithm. Three algorithms are compared here, the $\mathrm{A}^{*}$ algorithm with the construction of the graph, the MSPP algorithm with the construction of the multi-scale map and the MSPP-S algorithm. Similarly to the previous case, on a small search space, there is little or no improvement. As the problem dimension increases, however, the improvement gets much better. The MSPP-S algorithm is three orders


Fig. 6: Comparison of the time to construct the map and run the $\mathrm{A}^{*}$ or MSPP-FN algorithm versus the time to run the MSPP-S algorithm for which a map does not need to be computed. The results are shown in logarithmic scale.
of magnitude faster than creating a map and using the $\mathrm{A}^{*}$ algorithm, and more than ten times faster than constructing a multi-scale map and using the original MSPP algorithm.

## B. Application to a Robot Arm



Fig. 7: Initial and final pose of the planning problem for the PR2 arm.

The planning algorithm was used to plan a trajectory for an arm of the PR2 robot. Planning was done in the configuration space using 4 joints of the arm. Figure 7 shows the initial configuration and the desired final configuration; the robot needs to move a book from the top shelf to the second shelf. The depth of the tree was set to 5 , creating a search space of size $2^{4 \times 5} \simeq 3 \times 10^{7}$.

The path-planning problem was solved three times to compare the variants of the algorithm: First, the multiscale map was built by exploring the entire search-space and the MSPP algorithm was used to find the solution. All algorithms were performed on the same desktop computer running Ubuntu Linux. Building the map was the most timeconsuming process. It takes on average 4 minutes and 52
seconds and solving the path-planning problem takes on average 47 seconds. Using the MSPP-FN algorithm on the same map, the problem was solved in 4 seconds on average.

## VIII. CONCLUSIONS

In this paper, we have introduced several modifications and extensions to the original MSPP algorithm, first presented in [5], to increase its computational efficiency. The resulting multi-scale path-planning algorithms, called MSPP-FN and MSPP-S offer several non-trivial improvements over the previous MSPP algorithm. First, the complexity of each iteration of the algorithm is reduced by changing the manner by which the adjacency relationships in the reduced graph are computed at each iteration. Second, the range of applications of the algorithm has been widened, by allowing the use of an obstacle predicate rather than accurate prior knowledge of a full information multi-scale map. This extension results in much fewer requirements in terms of memory allocation and theoretical bounds on the probability of failure were derived. Third, reordering the operations performed by the algorithm allows one to minimize computations by avoiding information that is not needed during execution. We have compared the original MSPP algorithm to the proposed MSPP-FN and MSPP-S algorithms and found runtime improvements by over $50 \%$. Both algorithms outperform $\mathrm{A}^{*}$ by more than two orders of magnitude.

## REFERENCES

[1] S. Behnke. Local multiresolution path planning. In Robocup 2003: Robot Soccer World Cup VII, pages 332-343. Springer, 2004.
[2] R. V. Cowlagi. Hierarchical Motion Planning for Autonomous Aerial and Terrestrial Vehicles. PhD thesis, Georgia Institute of Technology - School of Aerospace Engineering, 2011.
[3] F. Endres, J. Hess, J. Sturm, D. Cremers, and W. Burgard. 3-D Mapping With an RGB-D Camera. IEEE Transactions on Robotics, 30(1):177-187, Feb 2014.
[4] D. Ferguson and A. Stentz. Using interpolation to improve path planning: The Field D* algorithm. Journal of Field Robotics, 23(2):79101, 2006.
[5] F. Hauer, A. Kundu, J. M. Rehg, and P. Tsiotras. Multi-scale perception and path planning on probabilistic obstacle maps. In Robotics and Automation (ICRA), 2015 IEEE International Conference on, pages 4210-4215. IEEE, 2015.
[6] S. Kambhampati and L. S. Davis. Multiresolution path planning for mobile robots. IEEE Journal of Robotics and Automation, 2(3):135145, 1986.
[7] Y. Lu, X. Huo, and P. Tsiotras. Beamlet-like data processing for accelerated path-planning using multiscale information of the environment. In 49th IEEE Conference on Decision and Control (CDC), pages 3808-3813, Atlanta, Georgia USA, December 15-17, 2010.
[8] D. K. Pai and L.-M. Reissell. Multiresolution rough terrain motion planning. IEEE Transactions on Robotics and Automation, 14(1):1933, 1998.
[9] C. Petres, Y. Pailhas, P. Patron, Y. Petillot, J. Evans, and D. Lane. Path planning for autonomous underwater vehicles. IEEE Transactions on Robotics, 23(2):331-341, 2007.
[10] K. M. Wurm, A. Hornung, M. Bennewitz, C. Stachniss, and W. Burgard. OctoMap: A probabilistic, flexible, and compact 3D map representation for robotic systems. In ICRA 2010 Workshop on Best Practice in 3D Perception and Modeling for Mobile Manipulation, Anchorage, Alaska, May 3-8, 2010.


[^0]:    ${ }^{1} \mathrm{PhD}$ candidate, School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0150, Email:fhauer3@gatech.edu
    ${ }^{2}$ Professor, School of Aerospace Engineering and Institute for Robotics and Intelligent Machines, Georgia Institute of Technology, Atlanta, GA 30332-0150, Email:tsiotras@gatech.edu

    Support for this work has been provided by ARO MURI award W911NF-11-1-0046 and ONR award N00014-13-1-0563

