

Spanning-Tree Based Coverage of Continuous Areas by a Mobile Robot

Yoav Gabriely and Elon Rimon

Dept. of Mechanical Engineering
Technion, Israel Institute of Technology
yoav@robots.technion.ac.il, elon@robby.technion.ac.il

December 1, 1999

Abstract

This paper considers the problem of covering a continuous planar area by a square-shaped tool attached to a mobile robot. Using a tool-based approximation of the work-area, we present an algorithm that covers every point of the approximate area for tasks such as floor cleaning, lawn mowing, and field demining. The algorithm, called Spanning Tree Covering (STC), subdivides the work-area into disjoint cells corresponding to the square-shaped tool, then follows a spanning tree of the graph induced by the cells, while covering every point precisely once. We present and analyze three versions of the STC algorithm. The first version is off-line, where the robot has perfect apriori knowledge of its environment. The off-line STC algorithm computes an optimal covering path in linear time $O(N)$, where N is the number of cells comprising the area. The second version of STC is on-line, where the robot uses its on-board sensors to detect obstacles and construct a spanning tree of the environment while covering the work-area. The on-line STC algorithm completes an optimal covering path in time $O(N)$, but requires $O(N)$ memory for its implementation. The third version of STC is “ant”-like. In this version, too, the robot has no apriori knowledge of the environment, but it may leave pheromone-like markers during the coverage process. The ant-like STC algorithm runs in time $O(N)$, and requires only $O(1)$ memory. Finally we present simulation results of the three STC algorithms, demonstrating their effectiveness in cases where the tool size is significantly smaller than the work-area characteristic dimension.

1 Introduction

The *mobile robot covering problem* can be formulated as follows. Let a tool of a specific planar shape be attached to a mobile robot, and let \mathcal{A} be a continuous planar work-area bounded by obstacles. Then the mobile robot has to move the tool along a path such that every point of \mathcal{A} is covered by the tool along the path. The covering problem is currently receiving considerable attention for several reasons. First, sensor-based coverage by mobile robots seems amenable to the geometric planning techniques developed for the sensor-based robot navigation problem, a problem which recently received considerable attention [3, 18]. A second reason is current interest in competitive algorithms¹ for autonomous systems that operate with incomplete information [2, 10]. The optimal covering problem can be formulated as a generalization of the Traveling Salesperson Problem (TSP) for a continuous domain, and thus is NP hard [1]. It is therefore natural to seek competitive algorithms for the covering problem. Finally, recent work on ant-like coverage [23, 22] has spurred considerable interest in the possibility of achieving complex coverage behavior by bounded-resource agents who leave pheromone-like markers in the environment.

Besides its theoretical interest, the mobile robot covering problem has several important applications. One such application area is automatic floor cleaning and coating in facilities such as supermarkets [7] and train stations [24]. Other application areas are lawn mowing [14], hazardous waste cleaning [8], and field demining [12]. Furthermore, although we focus on mobile robot coverage, the same methods apply to other robotic coverage tasks, such as car painting and wall finishing during house construction.

¹An algorithm is *competitive* if its solution to every problem instance is a constant times the optimal solution to the problem with full information available.

Previous efforts on the covering problem have focused mainly on the *discrete* version of the problem, where an agent has to visit all nodes of an unknown graph by traversing its edges. Examples are the discrete covering algorithms of Dudek et al. [6] and Deng and Mirzaian [5], who achieve competitive graph coverage by using pebble markers that can be placed and later picked up by the searching agent. More recently, Wagner et al. [23, 22] have proposed discrete coverage algorithms that use pheromone-like markers to guide the coverage process and to coordinate cooperation between several agents. The *continuous* covering problem, where a robot must sweep a tool over a continuous area, has been studied by Ntafos [13]. Much like in our approach, Ntafos imposes a tool-based grid approximation over the work-area. Then he subdivides the grid into rectangular sub-grids that can be covered optimally, and computes a TSP path that visits all sub-grids. Ntafos considers only simple grids with no internal holes, and his algorithm covers the grid within 33% of its optimal covering path. In contrast, we consider general grids, we make no attempt at subdividing the grid into sub-grids, and our algorithm produces an optimal covering path. Furthermore, Ntafos’s algorithm is strictly off-line while ours has an on-line version. The continuous coverage problem has also been studied by Choset and Pignon [4] and by Pirzadeh and Snyder [15]. However, these works lack any formal analysis of the quality of the resulting coverage path. Finally, we note that map building or area exploration is distinct from the covering problem considered here. In the map building problem an autonomous agent has to completely scan an unknown environment using its sensors [9, 13, 17, 21]. Map building can thus be interpreted as achieving sensory coverage of an unknown environment. In contrast, area coverage requires physical sweeping of a tool over every point of a given work-area.

In this paper we make the following simplifying assumptions on the mobile robot covering problem. First, we assume that the tool is a square of size D . Second, the robot is allowed to move the tool only in directions orthogonal to the tool’s four sides, without rotating the tool during this motion. Third, we subdivide the work-area into square cells of size $2D$, and discard cells which are partially covered by obstacles (Figure 1). As later discussed in analysis and simulations, the quality of the approximation depends on the tool size D being significantly smaller than the work-area characteristic dimension. The algorithm, called *Spanning Tree Covering* (STC), follows a spanning tree of the graph induced by the cells, while covering every point precisely once. The STC algorithm can be interpreted in two ways. First, from a graph-theoretic viewpoint, the algorithm strives to embed the largest possible Hamiltonian cycle² in the given work-area. The resulting Hamiltonian cycle is automatically an optimal covering path, since the tool covers every cell (and hence every point in the underlying continuous area) precisely once. It should be noted that the problem of constructing a Hamiltonian cycle in a general grid-like graph is NP hard [11]. We show that for covering purposes, it is advantageous to first construct a finer grid in $O(N)$ steps, then construct a Hamiltonian cycle for this grid in additional $O(N)$ steps, where N is the total number of grid cells. Second, from a robot motion-planning viewpoint, we decompose the environment into simple cells and use a spanning tree as an adjacency graph for these cell. However, unlike previous covering paradigms, we carefully weave the traversal of the adjacency graph with the covering of individual cells. This integration of inter-cell traversal with intra-cell coverage is a novel approach that has not appeared in the literature before.

²A *Hamiltonian cycle* is a closed path in a graph which visits every node of the graph precisely once.

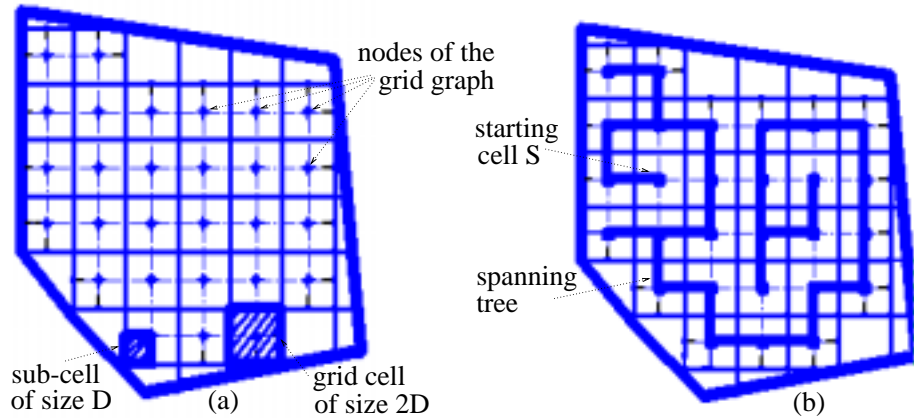


Figure 1: (a) Grid approximation of a given work-area. (b) A spanning tree for the grid.

In the next two sections we present and analyze three versions of the STC algorithm. The first version is off-line, where the robot has perfect apriori knowledge of its environment. The off-line STC algorithm computes an optimal covering path for the approximate work-area in linear time $O(N)$. The second version of STC is on-line, where the robot uses its on-board sensors to detect obstacles and construct a spanning tree of the environment while covering the work-area. The on-line STC algorithm completes an optimal covering path in time $O(N)$. However, it requires $O(N)$ memory, and this requirement limits the size of the work-areas that can be covered by the on-line algorithm. The third version of STC is “ant”-like. In this version, too, the robot has no apriori knowledge of the environment, but it may leave pheromone-like markers during the coverage process. The ant-like STC algorithm completes an optimal covering path in time $O(N)$, and requires only $O(1)$ memory. The analysis section also contains a fourth version of the STC algorithm, which is an adaptation of the on-line STC to an algorithm that achieves *exact* coverage of the work-area. The latter algorithm is amenable to competitive analysis, and we show that it achieves complete coverage with a competitive ratio of $1 + 3(\Pi D/A)$, where D is the tool size, Π is a quantity related to the obstacles’ perimeter, and A the total area. In practice $\Pi D \ll A$, and STC typically achieves an almost optimal coverage of the work-area. Finally we present simulation results of the three STC algorithms, demonstrating their effectiveness in cases where the tool size is significantly smaller than the work-area characteristic dimension.

2 The STC Algorithm

We describe three versions of the *Spanning Tree Covering* (STC) algorithm. In all three versions, we assume that the tool attached to the mobile robot is a square of size D . We also assume that the robot is allowed to move the tool only in the directions orthogonal to the tool’s four sides, without rotating the tool during this motion. We begin with the off-line version of STC.




Figure 2: An execution example of the off-line STC algorithm.

2.1 The Off-Line STC Algorithm

The input to the off-line STC algorithm is a geometrical description of a bounded planar environment populated by piecewise-smooth and stationary obstacles. The algorithm first subdivides the work-area into cells of size $2D$, and discards cells which are partially covered by obstacles. We define a graph structure, $G(V, E)$, by defining as nodes V the center points of each cell, and as edges E the line segments connecting centers of adjacent cells (Figure 1(a)). The algorithm next constructs a spanning tree for G (Figure 1(b)), and uses this tree to generate a covering path as follows.

Off-line STC Algorithm:

Input: A geometrical description of the environment, converted to a $2D$ -size grid with a graph structure G as described above. A starting cell S .

Pre-processing: Starting from S , construct a *spanning tree* for G using either DFS or PRIM. Subdivide every $2D$ -size cell into four identical *sub-cells* of size D .

Covering action: Starting at a sub-cell of S , move between neighboring sub-cells (each being identical to the tool's shape) along a path which *circumnavigates the spanning-tree along a counterclockwise direction*. Halt when the starting sub-cell is encountered again.

An execution example is illustrated in Figure 2. It can be seen that during the covering action the robot moves through sub-cells that lie on the right-side of an edge or a leaf-node of the spanning tree. Since a tree is contractible to a point, the circumnavigation of the spanning tree generates a simple closed path that brings the robot back to the starting sub-cell. The figure possesses two unrealistic features which were added for clarity. The tool size D is shown unrealistically large with respect to the work-area size, and the robot's path is shown curved while it is rectilinear according to the algorithm. We note that the off-line STC algorithm can assign edge weights and compute a minimal spanning tree for the weighted grid-graph using algorithms such as PRIM [16]. We can use this feature to control the spanning tree and hence the covering path geometry. For example, suitable edge weighting can generate covering paths that tend to scan the work-area along a particular coordinate direction, or covering paths that strive to follow obstacle perimeters. These possibilities are illustrated below.

2.2 The On-Line STC Algorithm

In the on-line version of STC the robot has no prior knowledge about the environment, except that obstacles are stationary. Rather, the robot must use its on-board sensors to detect obstacles and plan its covering path accordingly. We assume that the robot has position and orientation sensors, allowing it to locally recognize the $2D$ -size cells comprising the work-area. We also assume a range sensor, capable of identifying obstacles in the four cell neighboring the robot's current cell (Figure 3(a)). The practically important issues of sensor selection, sensor measurement errors, and sensor fusion are not considered here. Rather, we assume ideal sensors that provide perfect readings. In the on-line STC algorithm, the robot incrementally constructs a spanning tree for the grid representing the work-area. During the spanning tree construction, the robot subdivides every cell it encounters into four identical sub-cells of size D , each being identical to the tool size. The covering tool follows a sub-cell path which circumnavigates the incrementally constructed spanning tree, until the entire work-area grid is covered. A description of the algorithm follows.

On-line STC Algorithm:

Sensors: A position and orientation sensor. An obstacle detection sensor.

Input: A starting cell S , but no apriori knowledge of the environment.

Recursive function: A function $STC(w, x)$, where x is the current cell and w the previous cell along the spanning tree.

Initialization: Call $STC(Null, S)$, where S is the starting cell.

$STC(w, x)$:

1. Mark the current cell x as an *old* cell.
2. While x has a *new* obstacle-free neighbor:
 - 2.1 Scan for first new neighbor of x in counterclockwise order, starting with parent cell w .
Call this neighbor y .
 - 2.2 Construct a spanning-tree edge from x to y .
 - 2.3 Move to a sub-cell of y as described below.
 - 2.4 Execute $STC(x, y)$.End of while loop.
3. If $x \neq S$, move back from x to a sub-cell of w as described below.
4. Return. (End of $STC(w, x)$.)

We now discuss several details of the algorithm. First note that the robot runs a DFS algorithm during the spanning tree construction. The counterclockwise scanning of neighbors specified in step 2 is crucial for ensuring that the robot circumnavigates the incrementally constructed spanning tree in counterclockwise order (Figure 3(a)). The robot may scan the neighboring cells in clockwise order, but then it would circumnavigate the spanning tree in clockwise order. Second, in step 2.2 the robot is located in a sub-cell of x and has to move into a new cell y . By construction there is already a spanning-tree edge from x to y . The robot moves from its current sub-cell in x to a sub-cell of y by following the right-side of the spanning tree edges, measured with respect to the robot's direction of motion (Figure 3(b)). It is shown below that the robot covers the sub-cells of x in counterclockwise order. Hence no matter where the cell y is located relative to x , the robot always has a sub-cell path from




Figure 3: (a) Counterclockwise scanning of four neighbors. (b) A move from x to a new cell y . (c) A return from x to a parent cell w .

x to y which follows the right-side of the spanning tree edges. When the robot returns to a parent cell w , it again moves through sub-cells that lie on the right-side of the spanning-tree edge connecting x with w , as shown in Figure 3(c). Finally, the PRIM algorithm is not useful here, since PRIM requires discontinuous jumps between nodes while in our problem the robot must physically follow a continuous covering path.

2.3 The Ant-Line STC Algorithm

The ant-like version of STC is also an on-line algorithm, and it uses the same sensors described above. However, now the robot has the ability to leave markers in the sub-cells it covers, using color, odor, heat, or pebble markers [19, 20]. We additionally assume the availability of a detection device, capable of inspecting the sub-cell markers in the current cell and its four immediate neighbors. Using this detection device, the robot identifies a neighboring cell as *new* when its four sub-cells are all unmarked.

Much like the on-line STC algorithm, here too the robot uses DFS to incrementally construct a spanning tree for the work-area grid. However, rather than store the spanning tree in its memory, the robot uses the sub-cell markers to identify parent cells along the spanning tree as follows. By construction, the robot covers and marks the sub-cells of a given cell in counterclockwise order. Hence as long as the four sub-cells of the current cell are not all marked, the robot need only scan the sub-cells in counterclockwise order and identify the transition between unmarked and marked sub-cells. The two sub-cells bordering this transition have an exterior edge in common. As depicted in Figure 4, this edge is necessarily the boundary edge between the current cell and its parent cell in the spanning tree. Moreover, the robot has no need to identify a parent cell from a cell whose four sub-cells are already marked, since once the robot covers the four sub-cells of a given cell it never returns to this cell again. A description of the algorithm follows.

Ant-like STC Algorithm:

Marking device: A device that marks the sub-cell currently being covered.

Sensors: A position and orientation sensor. An obstacle detection sensor. A marker detector.

Input: A starting cell S , but no apriori knowledge of the environment.

Find parent function: A function $Parent(x)$ which returns the parent of the current cell x




Figure 4: Parent detection using sub-cell markers.

in the spanning tree, using the technique described above.

Find new neighbor function: A function $Neighbor(w, x)$ which accepts as inputs the current cell x and its parent cell w , and returns the first *new* obstacle-free neighbor of x in counter-clockwise order, starting with w . Returns *Null* if x has no new neighbor.

Initialization: Set $x = S$.

While S has unmarked sub-cells:

1. Determine the parent cell of x , $w = Parent(x)$.
2. Determine a *new* neighbor of x , $y = Neighbor(w, x)$.
3. If y is not *Null*:
 - 3.1 Move to a sub-cell of y , while marking the sub-cells of x being covered.
 - 3.2 Set $x=y$.
4. Else (y is *Null*):
 - 4.1 Return from x to a sub-cell of w , while marking the sub-cells of x being covered.
 - 4.2 Set $x = w$.

End of while loop.

We now discuss several details of the algorithm. In the function $Parent(x)$, the parent of the starting cell S is not well defined. In this case $Parent(x)$ identifies the transition between unmarked and marked sub-cells of S in counterclockwise order, and returns the neighbor cell adjacent to this transition. In steps 3.1 and 4.1, the move between cells is executed by following the right-side of the spanning tree edges, as depicted in Figure 3. However, in the ant-like STC the robot does not explicitly construct the spanning tree. Rather, it computes the two spanning-tree edges required for its local motion. The edge from the current cell x to its parent cell is computed with $Parent(x)$. The edge from x to a new neighbor (step 3.1), or from a neighbor back to x (step 4.1), is computed from knowledge of the neighbor. Finally consider the sub-cell marking specified in the algorithm. As discussed above, $Parent(x)$ can identify the parent cell of x only when not all four sub-cells of x are marked. Hence *the robot marks a sub-cell only when the covering tool leaves this sub-cell*. In particular, when the covering tool returns to a parent cell w in step 4.1, the robot marks the sub-cells of x being covered during the return motion, but not the sub-cell of w into which the covering tool has entered. This feature allows the robot to identify the parent cell of w in the next step of the algorithm.

3 Algorithm Analysis

In this section we first discuss several properties of the three STC algorithms, establishing their correctness and deriving performance bounds. Then we present an adaptation of the on-line STC algorithm which achieves exact coverage and is amenable to competitive analysis. Additional results concerning the quality of the grid approximation appear in the appendix.

3.1 Properties of the STC Algorithm

For purposes of analysis, we define a *repetitive coverage* of a point p as the situation where p is being covered by the robot's tool, then exposed, and later covered again. We first establish that the three STC algorithms do not generate any repetitive coverage.

Lemma 3.1 (non-repetitive coverage). *The three STC algorithms: off-line STC, on-line STC, and ant-like STC, do not repetitively cover any work-area point.*

Proof: In the three algorithms, the robot guides the tool from one sub-cell to an adjacent sub-cell using translational motion in the direction perpendicular to the edge common to both sub-cells. Since each sub-cell is identical to the tool size, the motion between adjacent sub-cells involves no repetitive coverage of either sub-cell. To complete the proof, we have to show that no sub-cell is visited more than once by the covering tool. In the three algorithms, the robot moves the tool through a sequence of sub-cells that circumnavigates the spanning tree (Figure 2). If we consider the spanning-tree edges as having some non-zero thickness, the curve which circumnavigates the spanning tree is a simple closed curve. By construction, any two spanning-tree edges which do not emanate from the same node lie at least two sub-cells apart (Figure 1). Hence the curve circumnavigating the spanning-tree cannot cross the same sub-cell twice. Consequently the tool which traces this curve cannot visit any sub-cell twice. This, together with the non-repetitive covering of adjacent sub-cells, implies the result. \square

The following proposition considers the covering pattern of individual cells.

Proposition 3.2. *The three STC algorithms cover the sub-cells of any particular cell in a counterclockwise (but not necessarily contiguous) order.*

Proof: In all three algorithms, the robot moves the covering tool along a sub-cell path which locally follows the right-hand side of the spanning-tree, measured with respect to the tool's direction of motion. If a cell x has no new neighbors, the center of x is a leaf-node of the spanning tree, and the tool circumnavigates this node in counterclockwise order as depicted in Figure 5(a). If x has new neighbors, the robot selects the *first* neighbor in counterclockwise order, starting with the parent cell of x . Let y denote the first neighbor and w denote the parent cell. The covering tool next follows the right-side of the spanning tree to a sub-cell of x which allows it to exit into y . (This sub-cell is determined by the spanning-tree edge from x to y .) As depicted in Figure 5(b)-(d), motion along the right-side of the spanning tree implies that the tool moves through sub-cells of x in counterclockwise order.

When the covering tool returns from y to x , it follows the same spanning-tree edge that lead it from x to y . Hence the tool re-enters x through a sub-cell which is adjacent to the sub-cell from which it left x . Let x' and x'' denote the sub-cells of x associated with

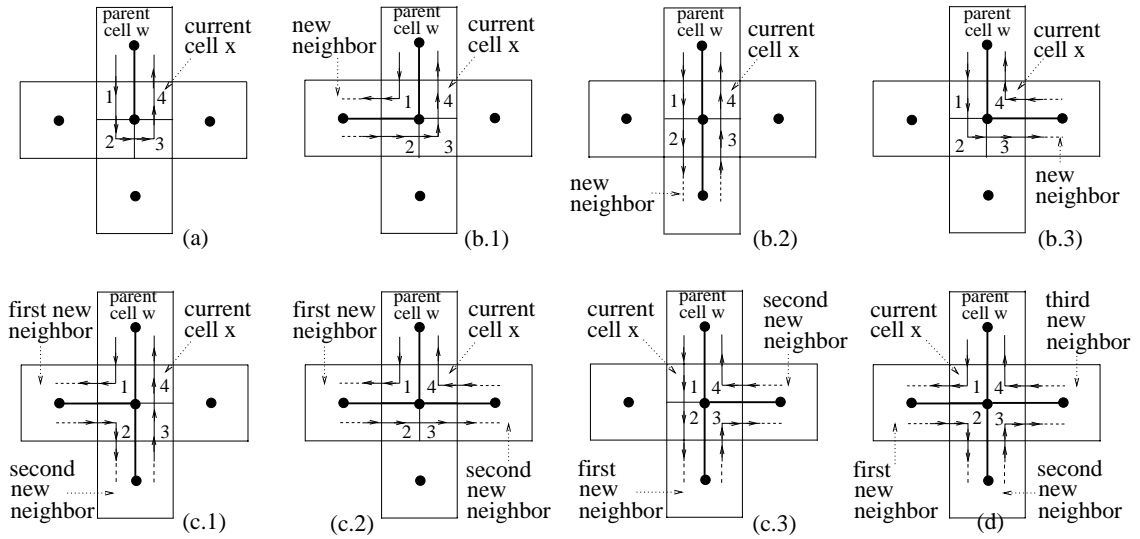


Figure 5: The sub-cell paths associated with (a) a cell with no new neighbors, (b)-(d) a cell with a new neighbor in three possible locations.

exiting from x into y and returning from y to x . We now show that x'' is adjacent to x' in counterclockwise order. There are two cases to consider. In the first case x' is the first sub-cell of x into which the tool entered from w . In this case x'' cannot lie before x' , since the sub-cell before x' borders the parent cell w , and this would imply that $y = w$. In the second case x' is not the first sub-cell of x . In this case the sub-cell before x' is already covered. It follows that x'' cannot lie before x' , since according to Lemma 3.1 no sub-cell is covered twice. The covering tool thus repeatedly enters and returns from new neighbors of x in counterclockwise order, until it reaches the last sub-cell of x . The last sub-cell of x lies on the right-side of the spanning-tree edge that leads from x to w , and the covering tool returns from this sub-cell to w . \square

The proposition has two important implications. First, it ensures that the covering tool can always access a new neighbor of the current cell x through empty sub-cells of x which follow the right-hand side of the spanning tree. Figure 5 illustrates these sub-cell paths for the four possible locations of a target neighbor. Second, the proposition asserts that all cells are covered by the *same* counterclockwise pattern. This feature is crucial for the ant-like STC, which relies on a fixed coverage pattern in order to uniquely identify the parent cell using sub-cell markers. The next lemma considers the completeness of the three STC algorithms.

Lemma 3.3 (complete coverage). *The three STC algorithms: off-line STC, on-line STC, and ant-like STC, cover every cell which is accessible from the starting cell S .*

Proof: Every obstacle-free cell defines a node in the work-area graph, and every pair of adjacent cells shares an edge in this graph. Hence the component of the work-area grid accessible from S is a connected graph. The three STC algorithms run DFS (possibly PRIM in the off-line STC), which constructs a spanning tree that reaches *every* cell in the grid representing the accessible work-area. These cells are partitioned into sub-cells. By construction, every sub-cell touches the spanning tree either at a point (a leaf node), or along a segment (an edge-segment emanating from a node). Since every sub-cell touches

the spanning tree, the sub-cell path generated by circumnavigating the spanning tree passes through *every* sub-cell of the cells accessible from S . Since each cell is completely covered once its four sub-cells are visited by the covering tool, all cells accessible from S are covered. \square

The following theorem summarizes the coverage properties of the three STC algorithms.

Theorem 1. *The three STC algorithms: off-line STC, on-line STC, and ant-like STC, completely cover any connected grid representing a continuous work-area. Moreover, the covering is optimal in the sense that there is no repetitive coverage of any point in the continuous area underlying the grid.*

The theorem follows from Lemma 3.1 and Lemma 3.3. The next theorem lists several performance bounds for the three STC algorithm. For simplicity, we assume that the off-line STC runs DFS rather than the PRIM algorithm.

Theorem 2. *Let N be the number of cells in the grid representing the accessible work area. Then the three STC algorithms cover this area in $O(N)$ steps, using a covering path of total length $L = 4ND$ where D is the tool size. Furthermore, the on-line STC requires $O(N)$ memory, while the ant-like STC requires $O(1)$ memory.*

Let us sketch the derivation of these bounds. First, the bound $O(N)$ on the number of covering steps reflects the way all three STC algorithms operate. In each step a new sub-cell is being covered. Since there are $4N$ sub-cells in an N -cell grid, all three algorithms cover the grid in $4N$ steps. Next consider the covering path total length. The covering path is a rectilinear curve, in each of whose segments the tool traverses a distance D between adjacent sub-cells. Since there are $4N$ sub-cells, the total length of the covering path is $4ND$. The $O(N)$ memory requirement of the on-line STC algorithm corresponds to a worst case scenario, where the algorithm makes N recursive calls into new cells before completing the coverage of any cell. The $O(N)$ memory requirement imposes a strict limitation on the size of the work-areas that can be covered by a bounded-memory robot. However, this limitation can be alleviated by allowing the robot to operate in environments whose total area is estimated in advance. By installing memory according to the expected work-areas size, the on-line STC algorithm can cover these work-areas irrespective of their particular geometry. In contrast, the ant-like STC algorithm uses sub-cell markers to locally identify the spanning tree. The ant-like STC stores only the following two items, which require $O(1)$ memory. The first is the coverage state of the starting cell S , which allows the algorithm to determine completion of coverage. (This item can be eliminated in a more sophisticated formulation of the algorithm.) The second is the identity of the spanning-tree edge along which the covering tool entered the current cell. This knowledge allows the robot to uniquely identify the four sub-cells of the current cell, which in turn allows the robot to properly scan the neighbor cells.

3.2 Competitive Analysis

We now adapt the on-line STC algorithm to an algorithm called *exact STC*, which achieves exact coverage. Then we subject the exact STC algorithm to a competitive analysis, which

compares the length of the covering path generated by the algorithm to the optimal covering path with full information available. First we make several simplifying assumptions. We assume that the tool size D is sufficiently small so that the connectivity of the work-area is preserved by the grid approximation. According to Lemma A.1 in the appendix, this assumption is automatically satisfied when the minimal gap between obstacles in the environment is at least $2R$, where $R = 2\sqrt{2}D$. We also assume that every obstacle in the environment occupies at least one cell which borders the work-area grid. This assumption is also satisfied when the minimal gap between obstacles is at least $2R$. The ensuing analysis holds true for environments with narrower gaps between obstacles, but we wish to avoid the overhead incurred by discussing such environments.

In order to achieve exact work-area coverage, we use Lemma A.2 in the appendix. According to this lemma, any point of the work-area not included in the grid approximation lies at a distance of at most R from one of the obstacle boundaries. The robot thus augments the grid coverage with a sweep of a thickness- R neighborhood about the obstacle boundaries. Our third assumption is that the covering tool can freely rotate while following the contour of an obstacle boundary. Since the tool size is D and $R = 2\sqrt{2}D$, the freely rotating tool can sweep a thickness- R neighborhood about an obstacle using $2\sqrt{2}$ passes.

The exact STC algorithm simply runs the on-line STC algorithm with the following modification. Each time the current cell x has a neighbor cell occupied by an obstacle, the robot checks if this is a new obstacle. If it is a new obstacle, the robot halts the grid coverage and sweeps a thickness- R neighborhood about the obstacle. During this sweep, the robot marks the obstacle boundary as being already covered. The robot completes the sweep at the cell x , and from this cell it resumes the grid coverage according to the on-line STC.

The covering path generated by the exact STC algorithm consists of segments of two types. Segments which correspond to the grid coverage, and segments associated with the sweep around obstacle boundaries. The total length of the segments of the first type is $4ND$ according to Theorem 2, where N is the number of grid cells. The total length of the segments of the second type is as follows. The coverage of a thickness- R neighborhood about an obstacle requires *three* passes about the obstacle, at distances of $D/2$, $3D/2$, and $5D/2$ from the obstacle's boundary. It can be verified that the length of the curve at a distance $3D/2$ is the average of the length of the other two curves. Hence the total length of the segments of the second type is 3Π , where Π is the total length of the curves at a distance $3D/2$ from the obstacles. Our next task is to derive a lower bound on the length of the optimal covering path. In order to simplify the derivation, we assume that the covering tool moves along a piecewise linear path, maintaining a fixed orientation along each linear segment.

Lemma 3.4 (optimality bound). *Let a continuous and connected work-area have a total area A . Then the length L of any piecewise-linear covering path of the work-area satisfies $L \geq A/D_{\max}$, where D_{\max} is the tool's maximal sweeping width.*

Proof: Let L_i be the length of the i^{th} segment in the piecewise-linear covering path, and let D_i be the tool's covering width during this segment. Then for a complete work-area coverage $A \leq \sum_i L_i D_i \leq D_{\max} \sum_i L_i$. The covering path total length is given by $L = \sum_i L_i$. Hence any piecewise-linear covering path satisfies $L \geq A/D_{\max}$. \square




Figure 6: An adversarial environment for the exact STC algorithm.

In our case we assume that the covering tool always use a sweeping width of D . The total length of the covering path generated by the exact STC algorithm is $4ND + 3\Pi \leq A/D + 3\Pi$, since $4ND^2 \leq A$. It follows that the *competitive ratio* of the exact STC algorithm is:

$$\frac{A/D + 3\Pi}{A/D} = 1 + 3\frac{\Pi D}{A}.$$

In practice $\Pi D \ll A$, and the exact STC algorithm typically achieves an almost optimal coverage of the work-area. However, the exact STC algorithm may perform poorly in certain environments. Figure 6 depicts a very long corridor of length H , whose width is twice the tool size D . Both ends of the corridor enter small rooms whose size is negligible with respect to the corridor's size. The work-area grid is aligned with the corridor and enters both rooms. The exact STC algorithm covers the grid using an optimal path of length $2H$, where we have neglected the covering of the two small rooms. The globally optimal covering path also has a length of $2H$. However, the exact STC algorithm additionally sweeps a thickness- R neighborhood about the outer walls. Since the algorithm requires a complete sweep around the outer walls, the covering tool must pass through the corridor three times in each direction, or a total of six times, in order to complete three loops around the outer walls. Neglecting again the two rooms, the total length of the six passes through the corridor is $6H$. The total length of the covering path generated by the exact STC is $8H$, and the competitive ratio for this environment is $8H/2H = 4$. The poor performance of the exact STC algorithm in this environment is due to its inefficient sweep of a thickness- R neighborhood about the obstacles. It is reasonable to expect that more sophisticated boundary following methods would yield a tighter competitive ratio in such environments.

4 Simulation Results

In this section we present simulation results of the three STC algorithms. We begin with an example of the off-line STC algorithm which demonstrates the possibility of using edge weighting to influence the covering pattern. Then we consider an example of the on-line and ant-like STC algorithms. Finally we consider an office environment in which the tool size D is an order-of-magnitude smaller than a door's width. In this case the grid closely approximates the work-area, and the STC algorithms achieve a 95% coverage of the work-area.

Our first example executes the off-line STC algorithm on the environment shown in Figure 7. The off-line STC algorithm can use PRIM rather than DFS to pre-compute a

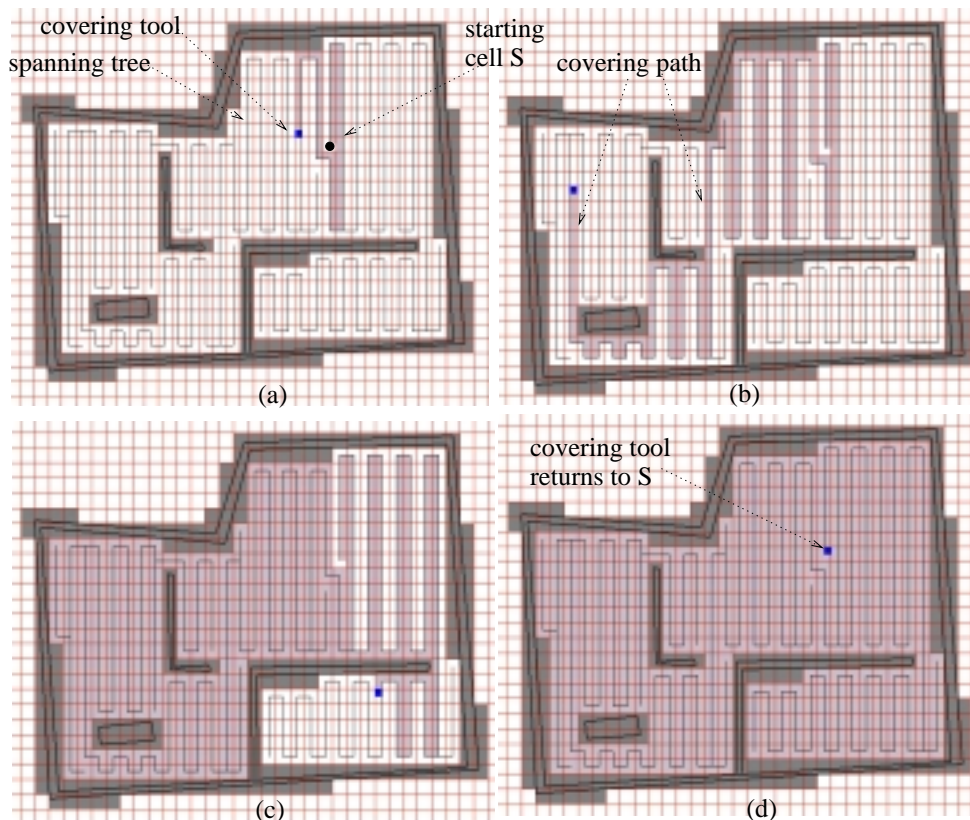


Figure 7: Four covering stages of the off-line STC algorithm using edge weighting.

spanning tree for the work-area. The PRIM algorithm admits arbitrary edge weighting, and here we assign edge weights that give preference to vertical edges over horizontal edges. The resulting spanning-tree edges tend to be vertically aligned, and the covering tool consequently scans the work-area along vertical strips. (For comparison, the spanning tree generated for the same environment by DFS appears in Figure 8(f).) Figure 7 shows four stages of the covering process. Figure 7(a) shows the spanning tree and the initial motion of the covering tool along the right-side of the spanning tree edges. Figure 7(b) shows the covering tool approaching the left edge of the work-area, using vertical covering sweeps. Figure 7(c) shows the covering tool entering the lower-right room. Finally, Figure 7(d) shows the completion of coverage, where the covering tool has completed the circumnavigation of the spanning tree and returns to the starting cell.

The next example, shown in Figure 8, runs the on-line STC and ant-like STC algorithms on the same environment. The two algorithms generate identical covering paths, and the covering stages shown in Figure 8 represent the execution of both algorithms. However, the two algorithms operate internally in a completely different way. The on-line STC maintains a data structure of the incrementally constructed spanning tree, while the ant-like STC uses sub-cell markers to locally identify the spanning tree. In Figure 8(a) the covering tool follows the right-side of the spanning tree edges along a sub-cell path which spirals outward from the starting cell S . In Figure 8(b) the covering tool has reached the outer boundary, and it follows this boundary until it reaches the entrance to the lower-right room. In Figure 8(c) the covering tool spirals inward to the center of the lower-right room, and in Figure 8(d) it




Figure 8: Six covering stages of the on-line and ant-like STC algorithms.

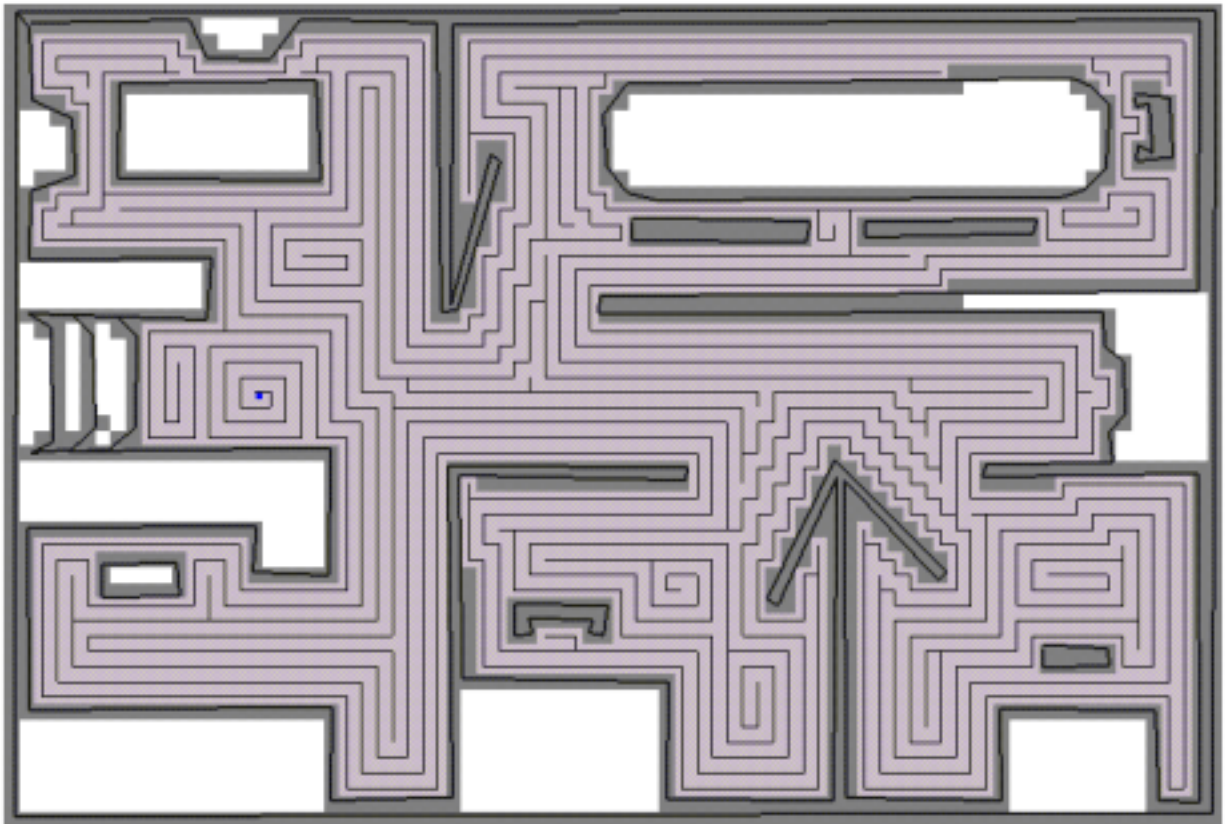


Figure 9: Coverage of an office environment using a tool of size $1/10$ of a door's width.

spirals outward along the complementary spiral, until it reaches the room's entrance. At this stage the lower-right room is completely covered. The covering tool next proceeds to the left portion of the environment. This portion is populated by two internal obstacles which divide the area into three cavities. Each of these cavities is covered by the same double-spiral pattern, and the resulting coverage is shown in Figure 8(e). Finally, the covering tool spirals inward to the starting cell S , resulting in a complete coverage of the work-area grid. The spanning tree constructed incrementally by the on-line STC algorithm appears in Figure 8(f). The same spanning tree is locally traced by the ant-like STC algorithm using sub-cell markers. It should be noted that the covering path is optimal, in the sense that there is no repeated coverage of any point in the area underlying the grid.

The last example, shown in Figure 9, attempts to provide a realistic office environment of a startup company making robotic vacuum-cleaners for office and home environments. This example consists of several rooms and offices which are populated by pieces of furniture and other items. The tool size D is selected to be an order-of-magnitude smaller than a door's width. The result of executing the on-line STC and ant-like STC algorithms on this environment appears in Figure 9, where the covered area is 92% of the total area. Note that the portions of the work-area not covered by the STC algorithms are concentrated around the obstacle perimeters. Hence it is possible to augment the grid coverage with a sweep around the obstacle boundaries as described above, in order to achieve exact coverage of the

entire work-area.

A Properties of the Work-Area Grid

The STC algorithm subdivides the work-area into $2D$ -size cells and discards cells which are occupied by obstacles. In this appendix we characterize the quality of this approximation in terms of the work-area parameters. Given a starting cell S , the *accessible work-area* is the connected component of the work-area which contains the cell S and is accessible by the mobile robot. The following lemma characterizes the connectivity of the grid approximation.

Lemma A.1. *The grid approximation of the accessible work-area contains all cells which can be accessed from the starting cell S by a path whose minimal distance from the obstacles is at least $R = 2\sqrt{2}D$, where D is the tool size.*

The bound is obtained by considering two cells arranged diagonally with a vertex in common. If these two cells are occupied by obstacles, they may disconnect the work-area grid as illustrated in Figure ???. We can in principal compute the minimum passage between obstacles in the work-area, then use the lemma to determine a tool size D which guarantees a grid approximation that preserves the connectivity of the accessible work-area. The next lemma characterizes the thickness of the area which is left out by the grid approximation.

Lemma A.2. *Any point of the accessible work-area not included in the grid approximation lies at a distance of at most R from one of the obstacle boundaries, where $R = 2\sqrt{2}D$.*

The lemma follows from the fact that the diameter of a $2D$ -size cell is $2R$. Hence all points which lie in a cell occupied by obstacles are at most R away from one of the obstacles. The lemma provides a justification for the fourth version of the STC algorithm, which achieves exact coverage by augmenting the grid coverage with a sweep of an area of thickness R about the obstacle boundaries. The last lemma gives a formula for the total area not included in the work-area grid. First we introduce some notation. The total area of the given work-area is denoted A , while the total area of the grid approximation is denoted \hat{A} . In the case of a polygonal environment, convex and concave obstacle vertices are defined as follows. At a convex vertex an obstacle protrudes into the free space, while at a concave vertex the free space protrudes into an obstacle. In the case of an environment bounded by smooth obstacles, we parametrize the obstacle boundaries by arc-length, using the curve $\alpha(s)$.

Lemma A.3. *If the work-area is bounded by polygonal obstacles, the total area not included in the grid approximation is bounded by*

$$A - \hat{A} \leq \Delta R - \frac{1}{2}R^2 \left(\sum_i (\pi - \alpha_i) - \sum_j \cot\left(\frac{\beta_j}{2}\right) \right),$$

where $R = 2\sqrt{2}D$, Δ is the obstacles' total perimeter, α_i are the angles of the concave obstacle vertices, and β_j are the angles of the convex obstacle vertices.

If the work-area is bounded by smooth obstacles, the total area not included in the grid approximation is bounded by

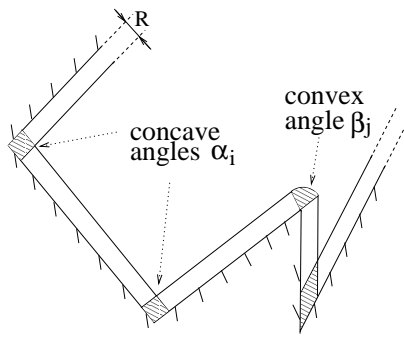


Figure 10: The computation of $A - \hat{A}$ in a polygonal environment.

$$A - \hat{A} \leq R \int_0^\Delta \|\alpha'(s)\| ds + \frac{1}{2} R^2 \int_0^\Delta \frac{\kappa(s)}{\|\alpha'(s)\|^2} ds,$$

where $\alpha(s)$ parametrizes the obstacles' boundary, $\alpha'(s)$ is the unit tangent at $\alpha(s)$, and $\kappa(s)$ is the curvature of the obstacles' boundary at $\alpha(s)$, for $0 \leq s \leq \Delta$.

The derivation of the bounds for $A - \hat{A}$ appears in Ref. [?]. Figure 10 illustrates the derivation for a polygonal work-area. According to Lemma A.2, points which are not included in the grid approximation are at a distance of at most R from the obstacles. The term ΔR is simply the area of a rectangle of thickness R , whose length Δ is obtained by unfolding the obstacles' boundary into a continuous straight line. The negative term $(R^2/2) \sum_i (\pi - \alpha_i)$ (since $0 \leq \alpha_i \leq \pi$) is the area at the concave obstacle vertices which is counted twice by the rectangle ΔR . The positive term $(R^2/2) \sum_j \cot(\beta_j/2)$ is the area of the circular sectors at the convex obstacle vertices, which has to be added to the rectangle ΔR . The bounds for $A - \hat{A}$ are quadratic in D , and in principal can be used to determine a tool size which would guarantee coverage within a desired area tolerance, as illustrated in Figure ??.

References

- [1] E. M. Arkin and R. Hassin. Approximation algorithms for the geometric covering salesman problem. *Discrete Applied Mathematics*, 55:197–218, 1994.
- [2] A. Blum, P. Raghavan, and B. Schieber. Navigating in unfamiliar terrain. In *STOC*, pages 494–504, 1991.
- [3] H. Choset and J. W. Burdick. Sensor based planning, part ii: Incremental construction of the generalized voronoi graph. In *IEEE Int. Conference on Robotics and Automation*, pages 1643–1649, 1995.
- [4] H. Choset and P. Pignon. Coverage path planning: The boustrophedon decomposition. In *Int. Conf. on Field and Service Robotics*, Canberra, Australia, December 1997.
- [5] X. Deng and A. Mirzaian. Competitive robot mapping with homogeneous markers. *IEEE Trans. on Robotics and Automation*, 12(4):532–542, 1996.

- [6] G. Dudek, M. Jenkin, E. Milios, and D. Wilkes. Robotic exploration as graph construction. *IEEE Trans. on Robotics and Automation*, 7(6):859–865, 1991.
- [7] H. Enders, W. Feiten, and G. Lawitzky. Field test of navigation system: Autonomous cleaning in supermarkets. In *IEEE Int. Conf. on Robotics and Automation*, pages 1779–1781, 1998.
- [8] S. Hedberg. Robots cleaning up hazardous waste. *AI Expert*, pages 20–24, May 1995.
- [9] S. Hert, S. Tiwari, and V. J. Lumelsky. A terrain covering algorithm for an auv. *Autonomous Robots*, 3:91–119, 1996.
- [10] C. Icking and R. Klein. Competitive strategies for autonomous systems. Technical Report CS:175, University of Hagen, Germany, 1995.
- [11] A. Itai, C. Papadimitriou, and J. Szwarcfiter. Hamiltonian paths in grid graphs. *SIAM J. of Computation*, 11:676–686, 1982.
- [12] J. D. Nicoud and M. K. Habib. The pemex autonomous demining robot: Perception and navigation strategies. In *IEEE/RSJ Int. Conf. on Intelligent Robot Systems*, pages 1:419–424, 1995.
- [13] S. Ntafos. Watchman routes under limited visibility. *Computational Geometry: Theory and Applications*, 1:149–170, 1992.
- [14] U. Peles and S. Abramson. *RoboMow by Friendly Machines*. www.friendlymachines.com, Israel, 1998.
- [15] A. Pirzadeh and W. Snyder. A unified solution to coverage and search in explored and unexplored terrains using indirect control. In *IEEE Int. Conf. on Robotics and Automation*, pages 2113–2119, 1990.
- [16] R. C. Prim. Shortest connection networks and some generalizations. *Bell System Tech. J.*, 36:1389–1401, 1957.
- [17] N. S. V. Rao and S. S. Iyengar. Autonomous robot navigation in unknown terrains: visibility graph based methods. *IEEE transactions on Systems, Man and Cybernetics*, 20(6):1443–1449, 1990.
- [18] E. Rimon. Construction of c-space roadmaps from local sensory data: What should the sensors look for? *Algorithmica*, 17(4):357–379, 1997.
- [19] R. A. Russell. Laying and sensing odor markings as a strategy for assisting mobile robot navigation tasks. *IEEE Robotics and Automation Magazine*, 2:3–9, Sep. 1995.
- [20] R. A. Russell. Heat trails as short-lived navigational markers for mobile robots. In *IEEE Int. Conf. on Robotics and Automation*, pages 3534–3539, 1997.

- [21] C. J. Taylor and D. J. Kriegman. Vision based motion planning and exploration algorithms for mobile robots. In *Workshop on Algorithmic Foundations of Robotics*, pages 69–83. A K Peters, 1995.
- [22] I. A. Wagner, M. Lindenbaum, and A. M. Bruckstei. Distributed covering by ant-robots using evaporating traces. *IEEE Transactions on Robotics and Automation*, 15(5):918–933, 1999.
- [23] I. A. Wagner, M. Lindenbaum, and A. M. Bruckstein. Efficiently searching a graph by a smell-oriented vertex process. *Annals of Mathematics and Artificial Intelligence*, 24:211–223, 1998.
- [24] H. Yaguchi. Robot introduction to cleaning work in the east japan railway co. *Adv. Robotics*, 10 (4):403–414, 1996.