
Piecemeal Graph Exploration by a Mobile Robot

(Extended Abstract) *

Baruch Awerbuch[†] Margrit Betke Ronald L. Rivest Mona Singh
Laboratory for Computer Science
Massachusetts Institute of Technology
Cambridge, MA 02139

Abstract

We study how a mobile robot can piecemeal learn an unknown environment. The robot's goal is to learn a complete map of its environment, while satisfying the constraint that it must return every so often to its starting position (for refueling, say). The environment is modelled as an arbitrary, undirected graph, which is initially unknown to the robot. We assume that the robot can distinguish vertices and edges that it has already explored.

We present a surprisingly efficient algorithm for piecemeal learning an unknown undirected graph $G = (V, E)$ in which the robot explores every vertex and edge in the graph by traversing at most $O(E + V^{1+o(1)})$ edges. This nearly linear algorithm improves on the best previous algorithm, in which the robot traverses at most $O(E + V^2)$ edges.

We also give an application of piecemeal learning to the problem of searching a graph for a "treasure".

1 Introduction

Environment learning and algorithmic motion planning for robots have recently become active research areas. The goal is to find efficient algorithms for a robot to learn about or navigate in its environment. Such algorithms are now useful in practice: there are working meal delivery robots in hospitals [15], and vehicles that navigate autonomously on highways [4]. More formal theoretical approaches to these problems have also been studied extensively (e.g., [20, 10, 22, 5]).

We study the problem of *piecemeal learning* of an unknown environment [8]. The robot's goal is to learn a complete map of its environment while satisfying the *piecemeal constraint* that learning must be done "a piece at a time," with the robot returning to the starting point s after each learning phase. Why might mobile robot exploration be done piecemeal? Robots may explore environments that are too risky or costly for humans: the inside of a volcano (e.g., CMU's Dante II robot) or the surface of Mars. The robot's hardware may be too expensive or fragile to stay long in dangerous conditions. Thus, it may be best to organize the learning into phases, allowing the robot to return to s before it breaks down or runs out of power. At the start position s , the robot can cool off, recharge, or drop off samples collected.

Approaches to modelling a robot's environment come from graph theory, computational geometry, on-line algorithms, and the theory of finite automata. The model used here was introduced by Betke, Rivest, and Singh [8]. The robot's task is to learn an unknown environment modelled as an undirected graph $G = (V, E)$ in a piecemeal manner. The robot's efficiency (or running time) is measured in terms of the number of edges traversed. The main difficulty in our work lies in designing efficient, but analyzable, robot exploration algorithms. We first give a simple algorithm that runs in $O(E + V^{1.5})$ time. We then improve this algorithm and give an almost linear time algorithm: it achieves $O(E + V^{1+o(1)})$ running time. The most efficient previously known algorithm has $O(E + V^2)$ running time.

A robot can explore grid-graphs with rectangular obstacles in a piecemeal manner in linear time, if the robot is

*We gratefully acknowledge support from NSF grant CCR-9310888, ARO grant DAAL03-86-K-0171, NSF grant 9217041-ASC, Air Force Contract TNDGAFOSR-86-0078, ARPA/Army contract DABT63-93-C-0038, NSF contract 9114440-CCR, DARPA contract N00014-J-92-1799, ARPA/ONR contract N00014-92-J-1310, the Siemens Corporation, and a special grant from IBM. The authors can be reached at baruch@blaze.cs.jhu.edu, margrit@lcs.mit.edu, rivest@theory.lcs.mit.edu, and mona@theory.lcs.mit.edu.

[†]Also at Johns Hopkins University, Baltimore, MD 21218.

Permission to make digital/hard copies of all or part of this material without fee is granted provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication and its date appear, and notice is given that copyright is by permission of the Association for Computing Machinery, Inc. (ACM). To copy otherwise, to republish, to post on servers or to redistribute to lists, requires specific permission and/or fee.
©OLT'95 Santa Cruz, CA USA © 1995 ACM 0-89723-5/95/0007..\$3.50

given a bound on the number of edges it may traverse in each learning phase (Betke, Rivest, and Singh [8]). We extend these results to show that the robot can learn *any* undirected graph piecemeal in almost linear time. It is open whether arbitrary, undirected graphs can be learned piecemeal in linear time.

The piecemeal constraint is most naturally satisfied by requiring the robot to explore in a near breadth-first manner, so that it is never much further away from the start vertex s than necessary to visit any unexplored vertex. In this manner, returns to s are efficient. Breadth-first search (BFS) on unknown graphs is also an important problem in its own right, with many applications. We consider one such application, *treasure hunting*, where the goal is to find a treasure (or a lost child, or a particular landmark) that is believed to be near s . If the robot knows that the treasure is close to its goal location, it should explore in a breadth-first manner from its current position.

BFS is a classic technique for searching graphs [19, 18, 11]. However, standard BFS is efficient only when the robot can efficiently switch or “teleport” from expanding one vertex to expanding another. In contrast, our model assumes a more natural scenario where the robot must *physically* move from one vertex to the next. We change the classical BFS model to a more difficult *teleport-free* exploration model, and give efficient *approximate* BFS algorithms where the robot does not move much further away from s than the distance from s to the unvisited vertex nearest to s .

The teleport-free BFS algorithms we first present never visit a vertex more than twice as far from s as the nearest unvisited vertex is from s . Our final teleport-free BFS algorithm satisfies the stronger condition that if the closest unvisited vertex to s is distance δ away, the robot is never more than $\delta + o(\delta)$ away from s .

For the treasure hunting problem, if the treasure is at a vertex that has shortest path distance δ_T away from s , then the robot traverses at most $O(E + V^{1+o(1)})$ edges, where E and V are the number of edges and vertices within radius $\Delta = \delta_T + o(\delta_T)$ from s . In contrast, we give an example to show that if the robot exactly satisfies the traditional BFS constraint (i.e., it cannot move further away from s than the unvisited vertex nearest to s), then it may traverse up to $O(E^2)$ edges. Our final treasure hunting algorithm is also a solution to the piecemeal learning problem.

Previous work

Many researchers have studied problems in environment learning and robot motion planning. Papadimitriou and Yanakakis [20] developed one of the first models for exploring unknown environments. They show how to find a shortest path in an unknown, undirected graph. Deng and Papadimitriou [13] and Betke [6] address the problem of learning an unknown directed graph. Bender and Slonim [5] show how two cooperating robots can

learn a directed graph. Rivest and Schapire [22] model the robot’s unknown environment by a deterministic finite automaton. They describe algorithms that efficiently infer the structure of the automaton through experimentation. Deng, Kameda, and Papadimitriou [12] consider how to learn the interior of a two-dimensional room. Blum, Raghavan, and Schieber [10] consider a robot navigating in an unknown two-dimensional geometric terrain with convex obstacles. Bar-Eli, Berman, Fiat, and Yan [3] give an efficient algorithm for reaching the center of a two-dimensional room with obstacles. Betke [7] and Kleinberg [17] address the problem of localizing a mobile robot in its environment. Blum and Chalasani [9] consider the problem of finding a “k-trip” shortest path in the environment. There are many other related papers in the literature (e.g., [16, 14]). Rao, Karet, Shi, and Iyengar [21] give a survey of work on “robot navigation in unknown terrains.”

Our techniques are inspired by the work of Awerbuch and Gallager [1, 2]. We observe that our learning model bears some similarity to the asynchronous distributed model. This similarity is surprising and has not been explored in the past.

2 Model and statement of main results

We model the robot’s environment as a finite undirected graph $G = (V, E)$ with a distinguished start vertex s . The graph is initially unknown to the robot. Each vertex in the graph represents an accessible location, and each edge represents a connection between adjacent locations. During each step of exploration, the robot moves from its current location to an adjacent location; it is not allowed to “teleport” from one vertex to another distant vertex. The robot can recognize previously visited vertices. The robot can distinguish the edges incident to its current vertex and it knows which edges it has traversed already, but it has no vision or long-range sensors. The robot incurs a cost only for traversing an edge; thinking and path planning (computation) are free.

We consider two closely related constraints on the exploration: the “piecemeal constraint” to model learning unknown environments in phases, and the “approximate BFS constraint” to model exploring an unknown graph in order to find a treasure.

PIECEMEAL LEARNING: The robot’s goal in piecemeal learning is to explore its entire (unknown) environment, while satisfying the piecemeal constraint that it must return every so often to its starting point. To assure that the learner can reach any vertex in the graph, do some exploration, and then get back to the start vertex, we assume the robot may traverse $(2 + \alpha)r$ edges in one exploration phase, where $\alpha > 0$ is some constant and r is the *radius* of the graph. The radius of the graph is the maximum of all shortest path distances between s and any vertex in G . We assume that the radius of the graph is known to the robot.

We say an exploration is *efficiently interruptible* if the robot always knows a path of explored edges of length at most R back to s .

Theorem 1 *An efficiently interruptible algorithm for exploring an unknown graph $G = (V, E)$ with n vertices and m edges that takes time $T(n, m)$ can be transformed into a piecemeal learning algorithm that takes time $O(T(n, m))$.*

The proof of this theorem is similar to one shown by Betke, Rivest, and Singh in a previous paper [8].

All the algorithms we present in this paper are efficiently interruptible, and thus give efficient piecemeal learning algorithms for undirected graphs. Our main theorem is:

Theorem 2 *Piecemeal learning of an arbitrary undirected graph $G = (V, E)$ can be done in time $O(E + V^{1+o(1)})$.*

Proof sketch: Following the RECURSIVE STRIP algorithm, given in Section 5, the robot always knows a path from its current location back to the start vertex of length at most the radius of the graph. The running time of this algorithm is $O(E + V2^{O(\sqrt{\log V \log \log V})}) = O(E + V^{1+o(1)})$. By Theorem 1, this algorithm can be interrupted efficiently to give a piecemeal learning algorithm. \square

TREASURE HUNTING: If the robot's goal is to explore an unknown environment in order to find a treasure that is believed to be near s , then the robot should explore in a breadth-first manner.

In traditional BFS, the robot may not move further away from the source than the unvisited vertex nearest to the source. At any given time in the algorithm, let Δ denote the shortest-path distance from s to the vertex the robot is visiting, and let δ denote the shortest-path distance from s to the vertex nearest to s that is as yet unvisited. With traditional breadth-first search we have $\Delta \leq \delta$ at all times. With teleport-free exploration, it is generally impossible to maintain $\Delta \leq \delta$ without a great loss of efficiency:

Lemma 3 *A robot which maintains $\Delta \leq \delta$ (such as a traditional BFS) may traverse $\Omega(E^2)$ edges.*

Proof: Consider a graph with vertices $\{-n, -n+1, \dots, -1, 0, 1, 2, \dots, n-1, n\}$, where $s = 0$ and edges connect consecutive integers. To achieve $\Delta \leq \delta$, a teleport-free BFS algorithm would run in quadratic time, traveling back and forth from 1 to -1 to -2 to 2 to 3 \dots \square

Given this lower bound, we solve the treasure hunting problem efficiently while maintaining the *approximate BFS constraint* that the robot is never more than twice as far from s as is the nearest unvisited vertex from s (i.e., $\Delta \leq 2\delta$). Our initial algorithms STRIP-EXPLORE, ITERATIVE-STRIP, and RECURSIVE-STRIP, described in

Sections 3, 4, and 5, maintain this constraint. Our final algorithm TREASURE-SEARCH, given in Section 6, satisfies the stronger condition $\Delta = \delta + o(\delta)$. Note that this algorithm is also efficiently interruptible and thus can also be used to solve the piecemeal learning problem; however, it is more complicated. Our main theorem about treasure hunting is:

Theorem 4 *Given an unknown graph with a treasure at distance δ_T from s , a robot can find the treasure while getting at most distance $\Delta = \delta_T + o(\delta_T)$ away from the start vertex, with an algorithm of running time $O(E + V^{1+o(1)})$, where E and V are the total number of distinct edges and vertices within radius Δ from the start vertex.*

Proof sketch: Algorithm TREASURE-SEARCH given in Section 6 satisfies the theorem. We discuss the properties of this algorithm in Section 6. \square

3 An exploration algorithm: STRIP-EXPLORE

This section describes an efficiently interruptible algorithm for undirected graphs with running time $O(E + V^{1.5})$. It is based on breadth-first search.

A *layer* in a BFS tree consists of vertices that have the same shortest path distance to the start vertex. A *frontier vertex* is a vertex that is incident to unexplored edges. A frontier vertex is *expanded* when the robot has traversed all the unexplored edges incident to it.

The traditional BFS algorithm expands frontier vertices layer by layer. In the teleport-free model, this algorithm runs in time $O(E + rV)$, since expanding all the vertices takes time $O(E)$, and visiting all the frontier vertices on layer i can be performed with a depth-first search of layers $1 \dots i$ in time $O(V)$, and there are at most r layers. The procedure LOCAL-BFS describes a version of the traditional BFS procedure that has been modified for our teleport-free BFS model in two respects. First, the robot does not relocate to frontier vertices that have no unexplored edges. Second, it only explores vertices within a given distance-bound L of the given start vertex s . (The first modification, while seemingly straightforward, is essential for our analysis of our more complex algorithms that use LOCAL-BFS as a subroutine at various source vertices). A procedure call of the form LOCAL-BFS(s, r), where s is the start vertex of the graph and r is its radius, would cause the robot to explore the entire graph.

Awerbuch and Gallager [1, 2] give a distributed BFS algorithm which partitions the network in *strips*, where each strip is a group of L consecutive layers. (Here L is a parameter to be chosen.) All vertices in strip $i-1$ are expanded before any vertices in strip i are expanded. Their algorithms use as a subroutine breadth-first type searches with distance L .

Our algorithm, STRIP-EXPLORE, searches in strips in a

LOCAL-BFS(s, L)

1. **for** $i = 0$ **to** $L - 1$ **do**
2. let $verts =$ all vertices at distance i from s
3. **for** each $u \in verts$ **do**
4. **if** u has any incident unexplored edges
5. **then**
6. relocate to u
7. traverse each unexplored edge
8. incident to u
9. relocate to s

new way. See Figure 1. The robot explores the graph in strips of width L . First the robot does LOCAL-BFS(s, L) to explore the first strip. It then explores the second strip as follows. Suppose there are k frontier vertices v_1, v_2, \dots, v_k in layer L ; each such vertex is a *source vertex* for exploring the second strip. A naive way for exploring the second strip is for the robot for each i , to relocate to v_i , and then find all vertices that are within distance L of v_i by doing a BFS of distance-bound L from v_i within the second strip. The robot thus traverses a forest of k BFS trees of depth L , completely exploring the second strip. The robot then has a map of the BFS tree of depth L for the first strip and a map of the BFS forest for the second strip, enabling it to create a BFS tree of depth $2L$ for the first two strips. The robot continues, strip by strip, until the entire graph is explored.




Figure 1: In the naive algorithm, the shaded areas are retraversed completely. In STRIP-EXPLORE, the shaded areas are passed through more than once only if necessary to get to frontier vertices.

The naive algorithm described above is inefficient, due to the overlap between the trees in the forest at a given level, causing portions of each strip to be repeatedly re-explored. The algorithm STRIP-EXPLORE presented below solves this problem by using the LOCAL-BFS procedure as the basic subroutine, instead of using a naive BFS. (See Figure 2.)

In STRIP-EXPLORE, the robot searches in a breadth-first manner, but ignores previously explored territory. The only time the robot traverses edges that have been previously explored is when moving to a frontier vertex it is about to expand. This results in retraversal of some edges in previously explored territory, but not as many

STRIP-EXPLORE(s, L, r)

1. $numstrips = \lceil r/L \rceil$
2. $sources = \{s\}$
3. **for** $i = 1$ **to** $numstrips$ **do**
4. **for** each $u \in sources$ **do**
5. relocate to u
6. LOCAL-BFS(u, L)
7. $sources =$ all frontier vertices

as in the naive algorithm.




Figure 2: Contrasting BFS and Local-BFS: Consider a BFS of depth 5 from s_1 , followed by a BFS of depth 5 from s_2 . (The depth of the strip is $L = 5$.) The BFS from s_2 revisits vertices a, b, c, d, e . On the other hand, if the BFS from s_1 is followed by a LOCAL-BFS from s_2 , then it only revisits d, c, e . After edge (f, d) is found, vertex e is a frontier vertex that needs to be expanded.

Theorem 5 STRIP-EXPLORE runs in $O(E + V^{1.5})$ time.

Proof: First we count edge traversals for relocating between source vertices for a given strip. For these relocations, the robot can mentally construct a tree in the known graph connecting these vertices, and then move between source vertices by doing a depth-first traversal of this tree. Thus the number of edge traversals due to relocations between source vertices for this strip is at most $2V$. Since there are $\lceil r/L \rceil$ strips, the total number of edge traversals due to relocations between source vertices is at most $2rV/L + 2V$.

Now we count edge traversals for repeatedly executing the LOCAL-BFS algorithm. First, for the robot to expand all vertices and explore all edges, it traverses $2E$ edges. Next, each time line 9 of procedure LOCAL-BFS is called, at most L edges are traversed. To account for relocations in line 6 of procedure LOCAL-BFS, we use the following scheme for “charging” edge traversals. Say the robot is within a call of the LOCAL-BFS algorithm. It has just expanded a vertex u and will now relocate to a vertex v to expand it. Vertex v is charged for the edges traversed to relocate from u to v . (We are only considering relocations within the same call of the LOCAL-BFS algorithm; relocations between calls of the LOCAL-BFS algorithm were considered above.) Source vertices are not charged anything. Moreover, the robot can always relocate from u to v by going from u to the source vertex

ITERATIVE-STRIP(s, r)

1. for $i = 1$ to \sqrt{r} do
2. for each source vertex u in strip i do
3. relocate to u
4. BFS from u to depth \sqrt{r} , but do not enter previously explored territory
5. while there are any active connected components iterate
6. for each active connected component c do
7. repeat
8. let v_1, v_2, v_3, \dots be active frontier vertices exclusively in c
9. with smallest depth among active frontier vertices in c
10. relocate to each of v_1, v_2, v_3, \dots , and expand
11. until no more active frontier vertices exclusively in c
12. determine new and active connected components

of the current local BFS, and then to v , traversing at most $2L$ edges. Thus, each vertex is charged at most $2L$ when it is expanded. LOCAL-BFS never relocates to a vertex v unless it can expand vertex v (i.e., unless v is adjacent to unexplored edges). Thus, all relocations are charged to the expansion of some vertex, and the total number of edge traversals due to relocation is at most $2LV$.

Thus the total number of edge traversals is at most $2rV/L + 2V + 3LV + 2E$, which is $O(rV/L + LV + E)$. When L is chosen to be \sqrt{r} , this gives $O(E + V^{1.5})$ edge traversals. \square

Procedure STRIP-EXPLORE, and the generalizations of it given in later sections, maintain that $\Delta \leq 2\delta$ at all times; the worst case is while exploring the second strip.

4 Iterative strip algorithm

We now describe ITERATIVE-STRIP, an algorithm similar to the STRIP-EXPLORE algorithm. It is an efficiently interruptible algorithm for undirected graphs inspired by Awerbuch and Gallager's [1] distributed iterative BFS algorithm. Although its running time of $O((V^{1.5} + E) \log V)$ is worse than the running time of STRIP-EXPLORE, its recursive version (described in Section 5) is more efficient. (It is not clear how to recursively implement STRIP-EXPLORE as efficiently, because the trees in a strip are not disjoint.)

With ITERATIVE-STRIP, the robot grows a global BFS tree with root s strip by strip, in a manner similar to STRIP-EXPLORE. Unlike STRIP-EXPLORE, here each strip is processed several times before it has correctly deepened the BFS tree by \sqrt{r} . We next explain the algorithm's behavior on a typical strip by describing how a strip is processed for the first time, and then for the remaining iterations.

In the first iteration, a strip is explored much as in STRIP-EXPLORE. The robot explores a tree of depth \sqrt{r} from each source vertex, by exploring in breadth-first manner from each source vertex, without re-exploring previous trees. Whenever the robot finds a *collision*




Figure 3: The iterative strip algorithm after the first iteration on the fourth strip. Two connected components c_1, c_2 have been explored. The collision edges e_1 and e_2 connect the first three approximate BFS trees. The dashed line shows how source vertices s_1, s_2, s_3 connect within the strip. There are three active frontier vertices with depth less than $D + \sqrt{r}$.

edge connecting the current tree to another tree in the same strip, it does not enter the other tree. Unlike STRIP-EXPLORE, the robot does not traverse explored edges to get to the active frontier vertices on other trees. Therefore, after the first iteration, the trees explored are *approximate BFS trees* that may have frontier vertices with depth less than \sqrt{r} from some source vertex. These vertices become *active frontier vertices* for the next iteration. Thus, the current strip may not yet extend the global BFS tree by depth \sqrt{r} , so more iterations are needed until all frontier vertices are inactive and the global BFS tree is extended by depth \sqrt{r} (see Figure 3).

In the second iteration (see Figure 4), the robot uses the property that two trees connected by a collision edge form a connected component within the strip. (The graph to be explored is connected, and thus forms one connected component; but we refer to connected components of the explored portion of the graph contained within the strip.) The robot need not traverse any edges outside the current strip to relocate between these active frontier vertices in the same connected component. In the second and later iterations, the robot works on one connected component at a time.

The robot explores active frontier vertices in one connected component as follows. It computes (mentally) a




Figure 4: The iterative strip algorithm after the second iteration. Now the circled vertices which were active frontier vertices at the beginning of the iteration are expanded. One of the expansions resulted in a collision edge. Now the strip consists of only one connected component (shaded area). There are six frontier vertices which become source vertices of the next strip. All frontier vertices have depth $D + \sqrt{r}$.

spanning tree of the vertices in the current strip. This spanning tree lies within the strip. Let d be the least depth of any active frontier vertex in the component from a source vertex. It visits the vertices in the strip in an order determined by a DFS of the spanning tree. As it visits active frontier vertices of depth d , it expands them. It then recomputes the spanning tree (since the component may now have new vertices) and again traverses the tree, expanding vertices of the appropriate next depth d' . Traversing a collision edge does not add the new vertex to the tree, since this vertex has been explored before. This process continues (at most \sqrt{r} times) until no active frontier vertex in the connected component has distance less than \sqrt{r} from some source vertex in the component.

The robot handles each connected component in turn, as described above. In the next iteration it combines the components now connected by collision edges, and explores the new active frontier vertices in these combined components. Lemma 6 states that at most $\log V$ iterations cause all frontier vertices to not be active any more; then the only active frontier vertices are the new sources of the next strip. The proofs of Lemma 6 and Theorem 7 are omitted here for lack of space.

Lemma 6 *At most $\log V$ iterations per strip are needed to explore a strip and extend the global BFS tree by depth \sqrt{r} .* \square

Theorem 7 *ITERATIVE-STRIP runs in time $O((E + V^{1.5}) \log V)$.* \square

5 The recursive strip algorithm

This section describes an efficiently interruptible algorithm RECURSIVE-STRIP, which gives a piecemeal learning algorithm with running time $O(E + V^{1+o(1)})$. RECURSIVE-STRIP is the recursive version of ITERATIVE-STRIP; it provides a recursive structure that coordinates

the exploration of strips, of approximate BFS trees, and of connected components in a different manner. The robot still, however, builds a (global) BFS tree from start vertex s strip by strip. The robot expands vertices at the bottom level of recursion.




Figure 5: The recursive strip algorithm processing an approximate BFS tree from source vertex s_2 to depth $d_{k-1} = L$. Recursive calls within the tree are of depth $d_{k-2} = L'$.

In RECURSIVE-STRIP, the depth of each strip depends on the level of recursion (see Figure 5). If there are k levels of recursion, then the algorithm starts at the top level by splitting the exploration of G into V/d_{k-1} strips of depth d_{k-1} . Each of these strips is split into d_{k-1}/d_{k-2} searches of strips of depth d_{k-2} , etc. We have $V = d_k > d_{k-1} > \dots > d_1 > d_0 = 1$.

Each recursive call of the algorithm is passed a set of source vertices $sources$, the *depth* to which it must explore, and a set T of all vertices in the strip already known to be less than distance *depth* from one of the sources. The robot traverses all edges and visits all vertices within distance *depth* of the sources that have not yet been processed by other recursive calls at this level. RECURSIVE-STRIP($\{s\}, r, \{s\}$) is called to explore the entire graph.

At recursion level i , the algorithm divides the exploration into strips and processes each strip in turn, as follows. Suppose the strip has l source vertices v_1, \dots, v_l . The strip is processed in at most $\log l = O(\log V)$ iterations. In each iteration, the algorithm partitions T into maximal sets T_1, T_2, \dots, T_k such that each set is known to be connected within the strip. Let S_c denote the source vertices in T_c . A DFS of the spanning tree of the vertices T gives an order for the source vertices in S_1, S_2, \dots, S_k ; this spanning tree is used for efficient relocations between these source vertices. Note that all source vertices are known to be connected through the spanning tree of the vertices in T , but they might not be connected within the substrips. Since relocations between the vertices in S_c in the next level of recursion use a spanning tree of T_c , for efficiency the vertices of T_c must be connected within the strip. After partitioning the vertices into connected components within the strip, for each connected component T_c , the robot relocates (along a spanning tree) to some arbitrary source vertex in S_c . It then calls the algorithm recursively with S_c ,

RECURSIVE-STRIP(*sources*, *depth*, *T*)

1. **if** *depth* = 1
2. **then** let v_1, v_2, \dots, v_k be the depth-first ordering of sources in spanning tree
3. **for** $i = 1$ **to** k **do**
4. relocate to v_i
5. **if** v_i has adjacent unexplored edges **then** traverse v_i 's incident edges
6. $T = T \cup \{\text{newly discovered vertices}\}$
7. **return**
8. **else** determine *next depth*
9. $\text{number-of-strips} \leftarrow \text{depth}/\text{next-depth}$
10. **for** $i = 1$ **to** *number-of-strips* **do**
11. determine set of source vertices
12. **for** $j = 1$ **to** *number-of-iterations* **do**
13. partition vertices in T into maximal sets T_1, T_2, \dots, T_k such that
14. vertices in each T_c are known to be connected within strip i
15. **for** each T_c in suitable order **do**
16. let S_c be the source vertices in T_c
17. relocate to some source $s \in S_c$
18. RECURSIVE-STRIP(S_c , *next-depth*, T_c)
19. $T = T \cup T_c$
20. relocate to some $s \in \text{sources}$
21. **return**

the depth of the strip, and the vertices T_c which are connected to the sources S_c within the strip.

The remaining iterations in the strip combine the connected components until the strip is finished. Then the robot continues with the next strip in the same level of recursion. Or, if it finished the last strip, it relocates to its starting position and returns to the next higher level of recursion. The proof of the following theorem is omitted for lack of space.

Theorem 8 RECURSIVE-STRIP runs in time $O(E + V^{1+o(1)})$. \square

6 Treasure Hunting

We now consider an application of the piecemeal learning problem where our goal is to find a treasure in an unknown, potentially infinite graph $G = (V, E)$. We give the procedure TREASURE-SEARCH, which uses the RECURSIVE-STRIP algorithm as a subroutine. If the treasure is at a location which is distance δ_T away from the source vertex, this algorithm maintains the condition that the robot is never further from the source than Δ , where $\Delta \leq \delta_T + o(\delta_T)$. Procedure TREASURE-SEARCH traverses $O(E + V^{1+o(1)})$ edges, where E and V are the total number of distinct edges and vertices within radius Δ from the source.

The robot explores the graph for the treasure in phases. In each phase i , the robot calls RECURSIVE-STRIP to explore a strip in the graph. The size of the strips changes over time. The change at phase i depends on $\epsilon_i = 1/\sqrt{i}$. Initially, the robot explores the graph out to distance

TREASURE-SEARCH(s)

1. $i = 0$
2. $r_0 = 1$
3. **Do** until treasure is found:
4. $i = i + 1$
5. $\epsilon_i = 1/\sqrt{i}$
6. $r_i = r_{i-1} \cdot (1 + \epsilon_i)$
7. let S be the set of source vertices distance
8. r_{i-1} away from s
9. RECURSIVE-STRIP($S, r_i - r_{i-1}, S$)

$r_1 = 1 + \epsilon_1$. Next, the robot extends its exploration by a factor of $1 + \epsilon_2$. That is, the size of the next strip is $(1 + \epsilon_1)(1 + \epsilon_2) - (1 + \epsilon_1)$, and the robot has learned the graph out to distance $r_2 = (1 + \epsilon_1)(1 + \epsilon_2)$. After extending the next strip, the robot has learned the graph out to distance $r_3 = (1 + \epsilon_1)(1 + \epsilon_2)(1 + \epsilon_3)$, and so on.

The proof of correctness of this procedure is omitted due to space constraints.

7 Open problems

We have presented an efficient $O(E + V^{1+o(1)})$ algorithm for piecemeal learning of arbitrary, undirected graphs. The only lower bound known for this problem is the trivial linear bound $\Omega(E + V)$. It is open whether there is a linear-time algorithm for piecemeal learning of general graphs.

References

- [1] Baruch Awerbuch and Robert G. Gallager. Distributed BFS algorithms. *The 26th Symposium on Foundations of Computer Science*, pages 250–256, October 1985.
- [2] Baruch Awerbuch and Robert G. Gallager. A new distributed algorithm to find breadth first search trees. *IEEE Transactions on Information Theory*, IT-33(3):315–322, 1987.
- [3] E. Bar-Eli, P. Berman, A. Fiat, and P. Yan. On-line navigation in a room. In *Symposium on Discrete Algorithms*, pages 237–249, 1992.
- [4] R. Behringer and S. Hötzel. Simultaneous estimation of pitch angle and lane width from the video image of a marked road. In *IEEE/RSJ/GI International Conference on Intelligent Robots and Systems*, September 1994.
- [5] Michael A. Bender and Donna K. Slonim. The power of team exploration: two robots can learn unlabeled directed graphs. In *Proceedings of the Thirty-Fifth Annual Symposium on Foundations of Computer Science*, pages 75–85, November 1994.
- [6] Margrit Betke. Algorithms for exploring an unknown graph. Master's thesis, MIT Department of Electrical Engineering and Computer Science, February 1992. (Published as MIT Laboratory for Computer Science Technical Report MIT/LCS/TR-536, March, 1992).
- [7] Margrit Betke and Leonid Gurvits. Mobile robot localization using landmarks. *IEEE/RSJ/GI International Conference on Intelligent Robots and Systems*, September 1994. To appear in *IEEE Transactions on Robotics and Automation*.
- [8] Margrit Betke, Ronald Rivest, and Mona Singh. Piecemeal learning of an unknown environment. *Machine Learning*, 18(2/3):231–254, March 1995.
- [9] Avrim Blum and P. Chalasani. An on-line algorithm for improving performance in navigation. In *Proceedings of the Thirty-Fourth Annual Symposium on Foundations of Computer Science*, pages 2–11, November 1993.
- [10] Avrim Blum, Prabhakar Raghavan, and Baruch Schieber. Navigating in unfamiliar geometric terrain. In *Proceedings of Twenty-Third ACM Symposium on Theory of Computing*, pages 494–504. ACM, 1991.
- [11] Thomas H. Cormen, Charles E. Leiserson, and Ronald L. Rivest. *Introduction to Algorithms*. MIT Press/McGraw-Hill, 1990.
- [12] Xiaotie Deng, Tiko Kameda, and Christos H. Papadimitriou. How to learn an unknown environment. In *Proceedings of the 32nd Symposium on Foundations of Computer Science*, pages 298–303. IEEE, 1991.
- [13] Xiaotie Deng and Christos H. Papadimitriou. Exploring an unknown graph. In *Proceedings of the 31st Symposium on Foundations of Computer Science*, volume I, pages 355–361, 1990.
- [14] Gregory Dudek, Michael Jenkin, Evangelos Milios, and David Wilkes. Using multiple markers in graph exploration. In *SPIE Vol. 1195 Mobile Robots IV*, pages 77–87, 1989.
- [15] John Evans. HelpMateTM: An autonomous mobile robot courier for hospitals. In *IEEE/RSJ/GI International Conference on Intelligent Robots and Systems*, September 1994.
- [16] Rolf Klein. Walking an unknown street with bounded detour. *Computational geometry: theory and applications*, 1(6):325–351, June 1992. Also published in *The 32nd Symposium on Foundations of Computer Science*, 1991.
- [17] Jon M. Kleinberg. The localization problem for mobile robots. In *Proceedings of the Thirty-Fifth Annual Symposium on Foundations of Computer Science*, pages 521–531, May 1994.
- [18] C. Y. Lee. An algorithm for path connection and its applications. *IRE Transactions on Electronic Computers*, EC-10(3):346–365, 1961.
- [19] Edward F. Moore. The shortest path through a maze. In *Proceedings of the International Symposium on the Theory of Switching*, pages 285–292. Harvard University Press, 1959.
- [20] Christos H. Papadimitriou and Mihalis Yannakakis. Shortest paths without a map. *Theoretical Computer Science*, 84:127–150, 1991.
- [21] Nagewara S. V. Rao, Srikumar Kareti, Weimin Shi, and S. Sitharama Iyengar. Robot navigation in unknown terrains: Introductory survey of non-heuristic algorithms. Technical Report ORNL/TM-12410, Oak Ridge National Laboratory, July 1993.
- [22] Ronald L. Rivest and Robert E. Schapire. Inference of finite automata using homing sequences. *Information and Computation*, 103(2):299–347, April 1993.