Improved Analysis of Greedy Mapping

Craig Tovey   Sven Koenig
College of Computing
Georgia Institute of Technology
Atlanta, GA 30332-0280, USA
{ctovey, skoenig}@cc.gatech.edu

Abstract

We analyze Greedy Mapping, a simple mapping method that has successfully been used on mobile robots. Greedy Mapping moves the robot from its current location on a shortest path towards a closest unvisited, unscanned or informative location, until the terrain is mapped. Previous work has resulted in upper and lower bounds on its worst-case travel distance but there was a large gap between the bounds. In this paper, we reduce the gap substantially by decreasing the upper bound from $O(|V|^{3/2})$ to $O(|V| \ln |V|)$ edge traversals, where $|V|$ is the number of vertices of the graph. This upper bound demonstrates that the travel distance of Greedy Mapping is guaranteed to be small and thus suggests that Greedy Mapping is indeed a reasonable mapping method. The guaranteed good performance of Greedy Mapping is robust in that it holds for different versions of Greedy Mapping, regardless of sensor type and sensor range.

1 Introduction

Robotics researchers have developed a large number of mapping methods [8]. In this paper, we study Greedy Mapping, a simple mapping method that assumes that the location of the robot is approximately known. Specific versions of Greedy Mapping may vary, but all employ the same basic principle: always move the robot from its current location on a shortest path towards a closest location that is of interest – that is, a location that is still unvisited, unscanned or informative. Greedy Mapping has several advantages. For example, the shortest paths can be calculated efficiently with incremental heuristic search methods, such as D* [12] or D* Lite [5], because the search is restricted to the known part of the map and can reuse information from previous searches [7]. It is therefore not surprising that several researchers have successfully used versions of Greedy Mapping on their robots. For example, Greedy Mapping has been used on a nomad-class tour-guide robot that offered tours to museum visitors [14]. It has also been used on Nomad 150 mobile robots [6] and Super Scouts [9]. Since Greedy Mapping seems to work well in practice, it is important to analyze its travel distance. In previous work, we derived upper and lower bounds on its worst-case travel distance but there was a large gap between the bounds. In this paper, we reduce the gap substantially by decreasing the upper bound from $O(|V|^{3/2})$ to $O(|V| \ln |V|)$ edge traversals, where $|V|$ is the number of vertices of the graph. Our new upper bound demonstrates, even more than the weaker previous one, that the travel distance of Greedy Mapping is guaranteed to be small and thus suggests that Greedy Mapping is indeed a reasonable mapping method. We identify four specific versions of Greedy Mapping: closest unvisited, closest unscanned, closest unscanned with replanning, and closest informative. Our bound applies to all four versions, regardless of sensor type or sensor range, although not always via the same mathematical proof. This suggests that the small travel distance of Greedy Mapping is robust with respect to implementation specifics.

2 Greedy Mapping

We can formulate the mapping task as a graph-coverage task similar to the one studied in [3], where vertices correspond to locations and edges correspond to paths between the locations. We now describe four versions of Greedy Mapping that operate on initially unknown undirected connected finite graphs $G = (V, E)$. It is possible for Greedy Mapping to have some initial information about the graph, for example, that it is a grid (graph). We say that a vertex is unvisited if Greedy Mapping has not yet been in that vertex. We say that a vertex is unscanned if Greedy Mapping has not yet learned all edges incident to it or all vertices that the edges lead to. When Greedy Mapping has scanned all vertices, it has mapped the graph. We assume that the sensors of the robot are powerful enough to enable Greedy Mapping to always scan at least its current vertex. Finally, we say that a vertex is informative if Greedy Mapping can gain information about the graph when being in that vertex, after which it becomes uninformative. Visited vertices cannot become unvisited again, scanned vertices cannot become unscanned again, and uninformative vertices cannot become informative again.

Closest Unvisited: This version of Greedy Mapping always moves on a shortest path from its current vertex to a closest unvisited vertex and repeats the process when it reaches that vertex. It terminates because it visits a previously unvisited vertex between replanning episodes and there are only a finite number of them.

Closest Unscanned: This version of Greedy Mapping is similar to the Closest Unvisited version. If the robot has powerful sensors, it may be able to scan a vertex from a distance, without visiting it. Therefore, Greedy Mapping might not have to visit every vertex, as in the Closest Unvisited version. The Closest Unscanned version repeatedly moves on a shortest path from its current vertex to a closest unscanned vertex and repeats the process it reaches that vertex. It terminates because it scans a previously unscanned vertex between replanning episodes, either when it visits the previously unscanned vertex or before, and there are only a finite number of them. This version of Greedy Mapping seems likely to be less efficient than the next one. We include it primarily to aid the exposition.

Closest Unscanned with Replanning: This version of Greedy Mapping is the same as the Closest Unscanned version, except that it immediately repeats the process instead of continuing on to its target vertex if it is able to scan the vertex before it reaches it. It terminates for the same reason as the Closest Unscanned version. The resulting behavior is equivalent to repeatedly traversing the first edge of a shortest path from the current vertex to a closest unscanned vertex. Although this version of Greedy Mapping has been used on robots [6], no upper bound besides the obvious $O(|V|^2)$ had previously been established for it.

Closest Informative: This version of Greedy Mapping repeatedly moves on a shortest path from its current vertex to a closest informative vertex and repeats the process when it reaches that vertex. It terminates because it makes a previously informative vertex uninformative between replanning episodes and there are only a finite number of them.

We clarify a somewhat subtle issue. A closest unscanned vertex in the currently known subgraph is also a closest unscanned vertex in the graph itself (since the vertices adjacent to scanned vertices are known) and the path lengths from the current vertex to the closest unvisited vertex are identical in both cases. The same is true for closest unvisited and closest informative vertices since visited and uninformative vertices are also scanned vertices. Therefore, there is no harm in defining the versions of Greedy Mapping in terms of closest vertices in the graph rather than closest vertices in the currently known subgraph.

3 A Formalization of Greedy Mapping

We now describe a framework that allows us to analyze three of the four versions of Greedy Mapping at the same time, namely the Closest Unvisited, Closest Unscanned, and Closest Informative versions. Initially, all vertices are interesting. We use x^0 to denote the start vertex of Greedy Mapping but make no assumptions about which vertex this is. When the current vertex of Greedy Mapping is x^{i-1}, it marks x^{i-1} and possibly other vertices as uninteresting. We make no assumptions about which additional vertices it marks as uninteresting. However, uninteresting vertices remain uninteresting. We use B^i to denote the set of all uninteresting vertices afterwards. Greedy Mapping then moves along a shortest path from its current vertex to a closest interesting vertex x^i, and terminates once there are no such paths any longer. We make no assumptions about how it breaks ties among interesting vertices that are equally close. Greedy Mapping must terminate since it marks at least one additional vertex as uninteresting each time it moves from its current vertex to the closest interesting vertex, and there are only a finite number of vertices. We define unvisited vertices, unscanned vertices, and informative vertices as interesting for the Closest Unvisited, Closest Unscanned and Closest Informative version of Greedy Mapping, respectively.

4 Previous Upper and Lower Bounds

To make the graph-coverage task as hard as possible, we studied the Closest Unvisited version of Greedy Mapping in previous work. We analyzed its worst-case travel distance (measured in the number of edges traversed) as a function of the number of vertices of the graph because a small worst-case travel distance provides a good performance guarantee on all graphs. We showed that a lower bound on its worst-case travel distance is $(x^{x+3} + 3x^{x+2} - 8x^{x+1} + 2x^2 - x + 3)/(x^2 - 2x + 1)$ edge traversals (steps) on graphs with $|V| = (3x^{x+2} - 5x^{x+1} - x^x + x^{x-1} + 2x^2 - 2x + 2)/(x^2 - 2x + 1)$ vertices for integers $x \geq 3$ [4]. Thus, it is $\Omega(\frac{\ln |V|}{\ln \ln |V|} |V|)$ edge traversals. This lower bound applies to all four versions described here, since the versions behave identically for robots that can scan only their current vertex, that is, robots with tactile sensors. An upper bound on the worst-case travel distance of the Closest Unvisited version of Greedy Mapping is $2|V|^{3/2} + 4|V|$ and thus $O(|V|^{3/2})$ edge traversals [6]. This upper bound has been used by robot practitioners to justify their choice of Greedy Mapping [15].

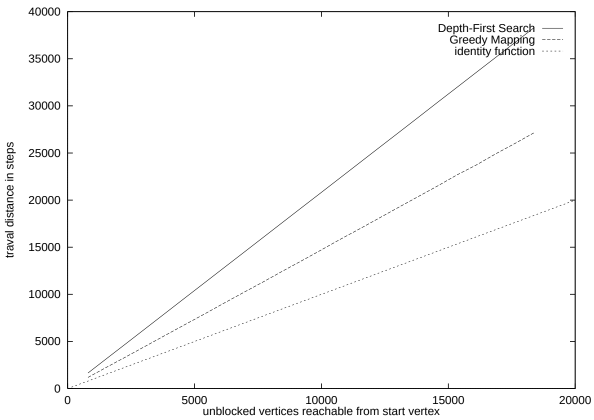

Figure 1: Experiment on Random Grids

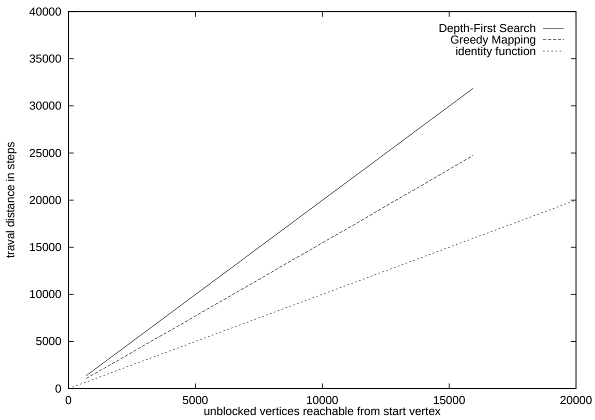

Figure 3: Experiment on Mazes

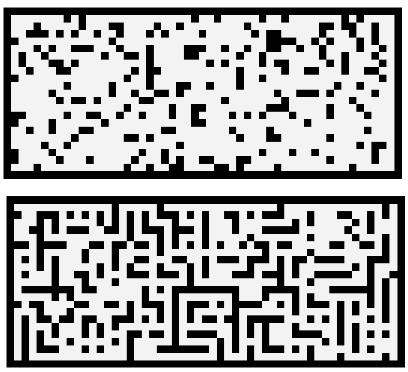

Figure 2: Example Grids

5 Greedy Mapping in Typical Terrain

The gap between the upper bound and lower bound from the previous section is rather large. It was unknown whether the true worst-case travel distance was closer to the upper bound or lower bound. We now present experimental results about the average travel distance of the Closest Unvisited version of Greedy Mapping on four-connected grids with width 51 and lengths from 21 to 481 (without the border), where the robot starts in the upper left corner. Figure 1 shows results of experiments on grids that were generated by starting with all traversable cells and then randomly making cells untraversable until the robot can reach about 75 percent of cells. Figure 2 (top) shows an example grid. Similarly, Figure 3 shows results of experiments in mazes that were generated by starting with all untraversable cells and first using depth-first search to generate an acyclic maze and then making additional cells traversable until the robot can reach about 65 percent of cells. Figure 2 (bottom) shows an example grid. Both figures show the travel distance of the Closest Unvisited version of Greedy Mapping (averaged over 500 grids of the same size each) as a function of the number of unblocked cells that are reachable from the start cell, together with the identity function. Since the Closest Unvisited version needs to visit every cell at least once, its travel distance cannot be smaller than the number of cells. Both figures show that the travel distance increases about linearly with the number of cells and is not much larger than it. These results and the good performance of Greedy Mapping in practice suggest that the known upper bound on the worst-case travel distance of Greedy Mapping can be reduced. In the following, we indeed decrease the upper bound substantially, resulting in a new upper bound that is quite close to the lower bound.

6 An Improved Upper Bound

We now prove that an upper bound on the worst-case travel distance of all four versions of Greedy Mapping is $O(|V| \ln |V|)$ edge traversals. Note that for each version, the upper bound applies to all sensor types and sensor ranges that can scan at least the current vertex. To understand why the worst-case travel distance of Greedy Mapping is small, assume that Greedy Mapping moves from its current vertex x^{i-1} to a closest interesting vertex x^i. Then, every vertex whose distance from x^{i-1} is less than the distance of x^i from x^{i-1} is uninteresting. Thus, intuitively, if the travel distance of Greedy Mapping from x^{i-1} to x^i is large, then there are many vertices whose distance from x^{i-1} is less than the distance of x^i from x^{i-1} and thus many vertices are already uninteresting. This implies that the travel distance of Greedy Mapping cannot be very large since it terminates once it has marked all vertices as uninteresting. To formalize this argument, we use $d(x, y)$ to denote the distance (measured in the number of edges traversed) from vertex x to vertex y. Instead of directly proving upper bounds on the travel distance of each version of Greedy Mapping, we first prove an upper bound on a more abstract process. To this end, we define marking sequences as follows:

Definition 1 A marking sequence on graph $G = (V, E)$ is a sequence of triples $\{v^i, r^i, M^i\}$ for $i = 1, 2, \ldots$, where $v^i \in V$, integer $r^i \geq 0$, and $M^i \subseteq V$ satisfy the following properties:

Property 1: $v^i \notin M^i$,

Property 2: $M^i \subseteq M^{i+1}$, and

Property 3: $d(v^i, v) \leq r^i$ implies $v \in M^{i+1}$.

The cost of the ith step (that is, triple) of the marking sequence is $1 + r^i$, and the cost of the marking sequence itself is $\sum_i (1 + r^i)$.

We call vertices in M^i uninteresting. Since at least v^i is marked as uninteresting in step i (and uninteresting vertices remain uninteresting) but v^{i+1} must still be interesting, marking sequences on finite graphs are finite. We now prove an upper bound on the cost of marking sequences that we can then use to derive an upper bound on the travel distance of Greedy Mapping since Greedy Mapping defines an associated marking sequence with $v^i = x^{i-1}$, $r^i = d(x^{i-1}, x^i) - 1$ and $M^i = B^{i-1}$, where B^0 is the empty set. Note that marking sequences, in general, do not require that v^i be at distance $1 + r^i$ from v^{i+1}. Instead, marking sequences consist of a sequence of choices of an interesting vertex v^i and a radius r^i. All vertices within distance r^i of v^i (and possibly some other vertices) are marked as uninteresting, and the sequence continues.

Theorem 6.1 *An upper bound on the maximum cost of any marking sequence on a given connected graph $G = (V, E)$ is $|V| + 2|V| \ln |V|$ and thus $O(|V| \ln |V|)$.*

Proof: Let $\{v^i, r^i, M^i\}$ be a maximum-cost marking sequence for the given graph. We will show that we can assume without loss of generality that $|M^{i+1}| = 1 + |M^i|$.

We first show that we can assume that $M^{i+1} = M^i \cup \{v \in V | d(v^i, v) \leq r^i\}$, that is, the marking sequence marks only those vertices as uninteresting that are within distance r^i of v^i. To prove that this assumption is valid, we show how to replace a marking sequence that does not fit the assumption by one that does but costs the same. Define $\{v^i, r^i, \hat{M}^i\}$ where $\hat{M}^1 = M^1$ and $\hat{M}^{i+1} = \hat{M}^i \cup \{v \in V | d(v^i, v) \leq r^i\}$. It is easy to show that this alternative marking sequence satisfies Properties (1), (2) and (3) and thus is indeed a valid marking sequence and has the same cost as the original one.

So far, we have shown that we can assume that $M^{i+1} = M^i \cup \{v \in V | d(v^i, v) \leq r^i\}$. We can therefore denote marking sequences by $\{v^i, r^i\}$ rather than $\{v^i, r^i, M^i\}$. We now show that we can make the more stringent assumption that $|M^{i+1}| = 1 + |M^i|$, that is, r^i is always so small that the marking sequence marks only v^i as uninteresting. To prove that this assumption is valid, we show how to replace a marking sequence that does not fit the assumption by one that does but costs the same. Suppose in step k there is a vertex $v \notin M^k$ with $0 < d(v^k, v) \leq r^k$, that is, a vertex different from v^k that gets marked as uninteresting. Replace step k with two steps: $\{v^k, d(v^k, v) - 1\}$ and $\{v, r^k - d(v^k, v)\}$. The two replacement steps together have the same cost as the replaced step, $r^k + 1$. Any vertex z marked as uninteresting during the first replacement step satisfies $d(v^k, z) \leq d(v^k, v) - 1 \leq r^k$ and thus z was marked as uninteresting in the original step. Similarly, any vertex z marked as uninteresting during the second replacement step satisfies $d(v, z) \leq r^k - d(v^k, v)$. By the triangle inequality it holds that $d(v^k, z) \leq d(v, z) + d(v^k, v) \leq r^k - d(v^k, v) + d(v^k, v) = r^k$ and thus z was marked as uninteresting in the original step. The alternative marking sequence therefore does not mark vertices as uninteresting that were not marked as uninteresting by the original marking sequence and therefore satisfies Property (1). The other properties hold by definition. The alternative marking sequence is therefore indeed a valid marking sequence and has the same cost as the original one. Finally, the replacement steps can be used repeatedly, resulting in a marking sequence with $|M^{i+1}| = 1 + |M^i|$ that has the same cost as the original one. (The replacement process terminates because Properties (1) and (3) limit the total number of steps to $|V|$.)

So far, we have shown that there is a maximum-cost marking sequence $\{v^i, r^i, M^i\}$ with $|M^{i+1}| = 1 + |M^i|$. Therefore every vertex in V appears (as the first element) in a triple of the marking sequence. Renumber the vertices in the order in which they appear. Call this kind of marking sequence an *orderly* marking sequence. We then have the following lemma.

Lemma 6.2 *Define $S^t = \{v^i \in V | r^i \geq t\}$ for an orderly marking sequence $\{v^i, r^i, M^i\}$ on a given connected graph $G = (V, E)$. Then, it holds that $|S^t| \leq 2|V|/t$.*

Proof: If v^j is at distance $d(v^i, v^j) \leq r^i$ from v^i and $i \neq j$, then $j < i$ because v^j is marked as uninteresting after step i according to Property (3) but still interesting before step j according to Property (1). Therefore, $v^i \in S^t$ and $v^j \in S^t$ with $i \neq j$ implies $d(v^i, v^j) > t$. Now, define the ball around $x \in S^t$ of radius $t/2$ as $B(x) = \{v \in V | d(x, v) \leq t/2\}$. The balls for $x \in S^t$ are pairwise disjoint, because a nonempty intersection between $B(v^i)$ and $B(v^j)$ would imply $d(v^i, v^j) \leq t$ by the triangle inequality. Each ball must contain at least $1 + \lfloor t/2 \rfloor$ vertices since G is connected. Therefore there can be at most $\lfloor |V|/(1 + \lfloor t/2 \rfloor) \rfloor$ such balls. The claim follows from $\lfloor |V|/(1 + \lfloor t/2 \rfloor) \rfloor \leq |V|/(1 + \lfloor t/2 \rfloor) \leq |V|/(t/2) \leq 2|V|/t$. ■

Applying the lemma, there exists a maximum-cost marking sequence which has $|V|$ steps and costs

$$\begin{aligned}
\sum_{j=1}^{|V|}(1+r^j) &= |V| + \sum_{j=1}^{|V|} r^j \\
&= |V| + \sum_{t=0}^{|V|} t(|S^t| - |S^{t+1}|) \\
&= |V| + \sum_{t=1}^{|V|} |S^t| \\
&\leq |V| + \sum_{t=1}^{|V|} 2|V|/t \\
&\approx |V| + 2|V| \ln |V|.
\end{aligned}$$

Note that this is a natural log. This proves the theorem. ■

Corollary 6.3 *An upper bound on the worst-case travel distance of the Closest Unvisited, Closest Unscanned, and Closest Informative versions of Greedy Mapping is* $|V| + 2|V| \ln |V|$ *and thus* $O(|V| \ln |V|)$ *edge traversals.*

Proof: All three versions of Greedy Mapping define marking sequences with $v^i = x^{i-1}$, $r^i = d(x^{i-1}, x^i) - 1$ and $M^i = B^{i-1}$, where B^0 is the empty set, on the subgraph given by the vertices that can be reached from the start vertex. (The sets B^i of uninteresting vertices were defined in section 3.) Their travel distances are the same as the costs of the associated marking sequences since $1 + r^i = d(x^{i-1}, x^i)$. Thus, the corollary follows directly from Theorem 6.1. ■

It is tempting to argue that the Closest Unscanned with Replanning version of Greedy Mapping takes no more edge traversals than the Closest Unscanned version since both versions move toward the closest unscanned vertex v^i but the latter version may save some edge traversals by stopping before it gets to v^i. However, this reasoning is fallacious because the two versions replan at different vertices and thus cannot be compared directly. The following corollary therefore derives an upper bound for the Closest Unscanned with Replanning version of Greedy Mapping using a slightly different proof idea than the preceding corollary.

Corollary 6.4 *An upper bound on the worst-case travel distance of the Closest Unscanned with Replanning version of Greedy Mapping is* $|V| + 2|V| \ln |V|$ *and thus* $O(|V| \ln |V|)$ *edge traversals.*

Proof: Let

- $w^0 = v^0$ be the start vertex of robot,
- v^1 be the closest unscanned vertex after the robot has scanned from v^0,
- w^1 be the vertex from which v^1 is first scanned,
- v^i be the closest unscanned vertex after robot has scanned from w^{i-1},
- w^i be the vertex from which v^i is first scanned, and
- O^i be the set of scanned vertices after the robot has scanned from w^i. (Define $O^{-1} = \phi$.)

Now consider the sequence of triples

$$\left\{v^i, [d(w^i, v^{i+1}) - d(w^i, v^i) - 1]^+, O^{i-1}\right\} \qquad (1)$$

for $i = 0, 1, \ldots$.

We claim that the sequence (1) is a marking sequence. By definition of the Closest Unscanned with Replanning version of Greedy Mapping, $v^i \not\in O^{i-1}$, giving Property (1). Also, scanned vertices remain scanned, giving Property (2). To verify Property (3), we claim that $d(w^i, v^{i+1}) - d(w^i, v^i) - 1 \geq -1$. By the principle of optimality, at the instant the robot is about to arrive at w^i, the closest unscanned vertex is v^i. After scanning from w^i, the closest unscanned vertex becomes v^{i+1}. Since v^{i+1} was unscanned prior to scanning from w^i, it must be the case that $d(w^i, v^{i+1}) \geq d(w^i, v^i)$ and thus $d(w^i, v^{i+1}) - d(w^i, v^i) - 1 \geq -1$. This inequality implies that there are two cases. Case 1: $d(w^i, v^{i+1}) - d(w^i, v^i) - 1 = -1$. Then $[d(w^i, v^{i+1}) - d(w^i, v^i) - 1]^+ = 0$. The only vertex at distance 0 from v^i is v^i itself, which is scanned from w^i and thus contained in O^i. Case 2: $d(w^i, v^{i+1}) - d(w^i, v^i) - 1 \geq 0$. Suppose a vertex q is within that distance to v^i, that is, $d(q, v^i) \leq d(w^i, v^{i+1}) - d(w^i, v^i) - 1$. By the triangle inequality, it holds that $d(q, w^i) \leq d(q, v^i) + d(v^i, w^i) \leq d(w^i, v^{i+1}) - 1$. By definition, v^{i+1} is a closest unscanned vertex after the robot has scanned from w^i. Since q is strictly closer to w^i than v^{i+1}, it must have been scanned already and thus is contained in O^i, giving Property (3). Thus, the sequence (1) is indeed a marking sequence.

Finally, by Theorem 6.1, it holds for the travel distance D that

$$\begin{aligned}
|V| + 2|V| \ln |V| &\geq (\sum_i 1 + [d(w^i, v^{i+1}) - d(w^i, v^i) - 1]^+) \\
&\geq \sum_i (d(w^i, v^{i+1}) - d(w^i, v^i)) \\
&\geq \sum_i (d(w^i, v^{i+1}) - d(w^{i+1}, v^{i+1})) = D.
\end{aligned}$$

■

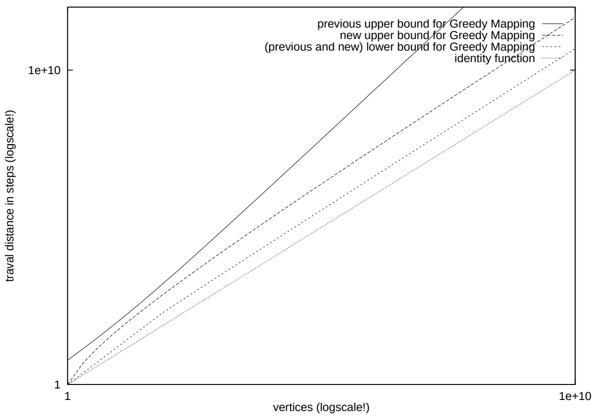

Figure 4: Upper and Lower Bounds

7 Discussion of Results

The previous upper bound on the worst-case travel distance of Greedy Mapping was $2|V|^{3/2}+4|V|$ edge traversals and our new upper bound is $|V| + 2|V|\ln|V|$ edge traversals. Figure 4 summarizes these results in graphical form, together with the (previous and new) lower bound and the identity function. Note that the figure shows a log-log plot. Consequently, it illustrates that we were able to reduce the upper bound substantially, thus reducing the gap between the upper bound and lower bound substantially. We also reduced the gap between the upper bound and the identity function substantially. This is important because depth-first search (that is, chronological backtracking) can be used to map unknown terrain with a travel distance that is linear in the size of the terrain. (In contrast, breadth-first search does not contend as an efficient mapping algorithm, because it requires a travel distance that is quadratic in the size of the terrain if, for example, the robot starts in the middle of a path.) While the lower bound on the travel distance of Greedy Mapping is super-linear in the size of the terrain, the new upper bound is sufficiently close to the identity function that Greedy Mapping is indeed a reasonable mapping method, especially since it outperforms depth-first search in practice, as shown in Figures 1 and 3, and has several advantageous properties that depth-first search does not have. In the following, we describe some of these properties:

- Greedy Mapping is reactive to changes in the location of the robot since it replans at every step, using the current location of the robot. It does not predict the current location based on its movement recommendations but rather uses the actual one. This makes the control loop of the robot simple to implement. It is also not a problem if the movement recommendations of Greedy Mapping are altered or ignored from time to time [1].
  - One consequence of this property is that Greedy Mapping does not need to have control of the robot at all times. For example, if a robot has to recharge its batteries during mapping, then it might have to preempt mapping and move to a known power outlet. Once restarted, Greedy Mapping does not have to return the robot to the location where mapping was stopped (which could be far away) to resume its operation from there. Rather, it resumes mapping from the power outlet.
  - Another consequence of this property is that Greedy Mapping can easily coexist with other modules of a robot architecture that might change its movement recommendations. For example, if Greedy Mapping suggests to pass an obstacle very closely, its movement recommendation might get changed by the obstacle-avoidance module to pass the obstacle with a larger safely margin. This is not a problem for Greedy Mapping.
- Greedy Mapping is also reactive to changes in the knowledge of the terrain since it replans at every step, using all of the information available to it. It always takes new information into account right away when it determines which unvisited, unscanned, or informative location is closest to the robot and how to get there quickly. It does not matter whether this information was learned by the robot or provided to it. Thus, Greedy Mapping can take advantage of a-priori terrain information or terrain information obtained on-line by other robots, if available [10, 2, 11].

8 Conclusions and Future Work

In this paper, we have analyzed Greedy Mapping, a simple mapping method that has been used successfully on mobile robots. Previous work resulted in upper and lower bounds on its worst-case travel distance but there was a large gap between the bounds. In this paper, we reduced the gap substantially by decreasing the upper bound from $O(|V|^{3/2})$ to $O(|V|\ln|V|)$ edge traversals, where $|V|$ is the number of vertices of the graph. We also showed that the new upper bound applies to four different versions of Greedy Mapping, regardless of sensor type and sensor range, which suggests that the good performance of Greedy Mapping is robust with respect to implementation specifics. We believe that the Closest Informative version of Greedy Mapping is conceptually attractive, because it searches in the space of terrain information, rather than the physical terrain itself. In the future we intend to explore whether this version performs better than other versions of Greedy Mapping. We also intend to analyze versions of Greedy Mapping that do not assume that the location of the robot is approximately known but maintain a probability distribution over its location instead [13]. Finally, we intend to provide a theoretical analysis that explains why the average travel distance of Greedy Mapping was smaller than the travel distance of depth-first search in the tests reported in Section 5.

9 Acknowledgment

This project is partly supported by an NSF award under contract IIS-0098807. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the sponsoring organizations and agencies or the U.S. government.

References

[1] P. Agre and D. Chapman. Pengi: An implementation of a theory of activity. In *Proceedings of the National Conference on Artificial Intelligence*, pages 268–272, 1987.

[2] W. Burgard, D. Fox, M. Moors, R. Simmons, and S. Thrun. Collaborative multi-robot exploration. In *Proceedings of the International Conference on Robotics and Automation*, pages 476–481, 2000.

[3] B. Kalyanasundaram and K. Pruhs. Constructing competitive tours from local information. *Theoretical Computer Science*, 130:125–138, 1994.

[4] S. Koenig. Exploring unknown environments with real-time search or reinforcement learning. In *Advances in Neural Information Processing Systems 12*, volume 11, pages 1003–1009, 1999.

[5] S. Koenig and M. Likhachev. Improved fast replanning for robot navigation in unknown terrain. In *Proceedings of the International Conference on Robotics and Automation*, pages 968–975, 2002.

[6] S. Koenig, C. Tovey, and W. Halliburton. Greedy mapping of terrain. In *Proceedings of the International Conference on Robotics and Automation*, pages 3594–3599, 2001.

[7] M. Likhachev and S. Koenig. Incremental replanning for mapping. In *Proceedings of the International Conference on Intelligent Robots and Systems*, pages 667–672, 2002.

[8] N. Rao, S. Hareti, W. Shi, and S. Iyengar. Robot navigation in unknown terrains: Introductory survey of non-heuristic algorithms. Technical Report ORNL/TM–12410, Oak Ridge National Laboratory, Oak Ridge (Tennessee), 1993.

[9] L. Romero, E. Morales, and E. Sucar. An exploration and navigation approach for indoor mobile robots considering sensor's perceptual limitations. In *Proceedings of the International Conference on Robotics and Automation*, pages 3092–3097, 2001.

[10] R. Simmons, D. Apfelbaum, W. Burgard, D. Fox, M. Moors, S. Thrun, and H. Younes. Coordination for multi-robot exploration and mapping. In *Proceedings of the National Conference on Artificial Intelligence*, pages 852–858, 1997.

[11] K. Singh and K. Fujimura. Map making by cooperating mobile robots. In *Proceedings of the International Conference on Robotics and Automation*, pages 254–259, 1993.

[12] A. Stentz. Optimal and efficient path planning for unknown and dynamic environments. *International Journal of Robotics and Automation*, 10(3):89–100, 1995.

[13] S. Thrun. Probabilistic algorithms in robotics. *Artificial Intelligence Magazine*, 21(4):93–109, 2000.

[14] S. Thrun, A. Bücken, W. Burgard, D. Fox, T. Fröhlinghaus, D. Hennig, T. Hofmann, M. Krell, and T. Schmidt. Map learning and high-speed navigation in RHINO. In D. Kortenkamp, R. Bonasso, and R. Murphy, editors, *Artificial Intelligence Based Mobile Robotics: Case Studies of Successful Robot Systems*, pages 21–52. MIT Press, 1998.

[15] R. Zlot, A. Stentz, M. Dias, and S. Thayer. Multi-robot exploration controlled by a market economy. In *International Conference on Robotics and Automation*, pages 3016–3023, 2002.