

Implementation of Efficient Algorithms for Globally Optimal Trajectories*

L. C. Polymenakos[†]

D. P. Bertsekas[‡]

J. N. Tsitsiklis[‡]

Abstract

We consider a continuous-space shortest path problem in a two-dimensional plane. This is the problem of finding a trajectory that starts at a given point, ends at the boundary of a compact set of \mathbb{R}^2 , and minimizes a cost function of the form $\int_0^T r(x(t))dt + q(x(T))$. For a discretized version of this problem, a Dijkstra-like method that requires one iteration per discretization point has been developed by Tsitsiklis [Tsi93]. Here we develop some new label correcting-like methods based on the Small Label First methods of Bertsekas [Ber93], and Bertsekas et. al. [BGM94]. We prove the finite termination of these methods, and we present computational results showing that they are competitive and often superior to the Dijkstra-like method.

*Research Supported by NSF Grant No 9300494-DMI.

[†]Ph.D. Candidate, Laboratory for Information and Decision Systems, M.I.T., Cambridge, MA, 02139.

[‡]Professor, Laboratory for Information and Decision Systems, M.I.T., Cambridge, MA, 02139

1 Introduction: Problem Formulation

In this section we discuss a continuous-space shortest path problem and its discretization. Our presentation follows closely [Tsi93]. We are given a bounded connected open subset G of \mathbb{R}^2 and a point $x(0) \in G$. A *trajectory* starting at $x(0)$ is a continuous function $x : [0, T] \in \mathbb{R}^2$, where T is some positive number, such that $x(t) \in G$ for all $t \in [0, T)$ and $x(T) \in \partial G$, where ∂G is the boundary of G . A trajectory is called *admissible* if there exists a measurable function $u : [0, T] \rightarrow \mathbb{R}^2$ such that

$$x(t) = x(0) + \int_0^t u(s) ds$$

and

$$\|u(t)\| \leq 1, \quad \forall t \in [0, T],$$

where $\|\cdot\|$ stands for the Euclidean norm. The cost of an admissible trajectory is defined as

$$\int_0^T r(x(t)) dt + q(x(T)),$$

where $r : G \rightarrow (0, \infty)$ and $q : \partial G \rightarrow (0, \infty)$ are given cost functions. We want to find an admissible trajectory of least cost. Note that we have considered a two-dimensional space for simplicity. The algorithms and the analysis of this paper admit straightforward generalizations to higher dimensional spaces.

We consider a method for discretization of this problem described and analyzed by Kushner and Dupuis [KuD90], who give several earlier references. We form a discretization grid using a square centered at the origin whose corners are vectors w_1, w_2, w_3, w_4 of length h , as shown in Fig. 1.1. This grid consists of two disjoint finite sets S and B such that for each $x \in S$, the set of *neighbors* of x , defined by

$$N(x) = \{x + hw_i \mid i = 1, 2, 3, 4\}$$

is a subset of $S \cup B$. The set S is considered to be a discretization of G and the set B is considered a discretization of ∂G . We also have two functions $f : B \rightarrow (0, \infty)$ and $g : S \rightarrow (0, \infty)$ that represent discretizations of the cost functions q and r of the original problem, respectively. The function g can be usually defined by $g(x) = r(x)$ for every $x \in S$.

We now consider a finite-state optimal control problem, the states of which are the points $x \in S \cap B$, also referred to as *nodes*. This problem fits within the framework of the stochastic shortest problems discussed in [BeT89] and [BeT91], and is defined as follows: At a state $x \in S$, we must choose a quadrant spanned by the vectors w_α and $w_{\alpha+1}$, where $\alpha \in \{1, 2, 3, 4\}$, and $w_5 = w_1$, and then choose a parameter $\theta \in [0, 1]$, that specifies an element $\theta w_\alpha + (1 - \theta)w_{\alpha+1}$ of the line segment connecting w_α and $w_{\alpha+1}$. The cost of the choice (α, θ) is $hg(x)\tau(\theta)$, where

$$\tau(\theta) = \sqrt{\theta^2 + (1 - \theta)^2},$$

and

$$h\tau(\theta) = \|\theta w_\alpha + (1 - \theta)w_{\alpha+1}\|$$

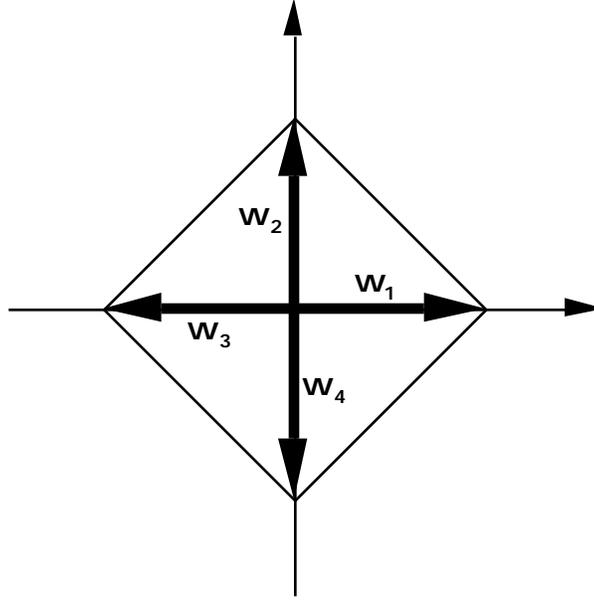


Figure 1.1: A square centered at the origin and the definition of the vectors w_1, \dots, w_4 of length h .

[figure1]

is the distance travelled from x to the point $\theta w_\alpha + (1 - \theta)w_{\alpha+1}$. The next state is $x + w_\alpha$ with probability θ_α and $x + w_{\alpha+1}$ with probability $1 - \theta$. Also, if the state $x \in B$ is reached, then a terminal cost $f(x)$ is incurred and the process terminates.

The optimal cost-to-go function $\{V^*(x) \mid x \in G \cup \partial G\}$ of the original problem, which is the infimum of the costs of all admissible trajectories that start at x , is approximated by the optimal cost function $\{V(x) \mid x \in S \cup B\}$ of the discretized problem. This latter function satisfies the following Bellman equations, which can also be viewed as a discretization of the Hamilton-Jacobi equation for the original problem:

$$V(x) = \min_{\alpha=1,2,3,4} \min_{\theta \in [0,1]} \left[\underbrace{hg(x)\tau(\theta)}_{\text{traveling cost}} + \underbrace{\theta V(x + w_\alpha) + (1 - \theta)V(x + w_{\alpha+1})}_{\text{cost-to-go}} \right], \quad x \in S \quad (1.1)$$

$$V(x) = f(x), \quad x \in B. \quad (1.2)$$

The above equations, under reasonable assumptions, have a unique solution [BeT89], [BeT91]. We assume in the sequel that indeed Eqs. (1.1), and (1.2) have a unique solution satisfying $V(x) < \infty$ for all $x \in S \cup B$. This implies in particular, that for every node $x_0 \in S$ there exists at least one sequence of nodes $x_0, x_1, \dots, x_k, x_{k+1}$ such that $x_{k+1} \in B$ and $x_{i+1} \in N(x_i)$, for all $i \in \{0, \dots, k\}$. The theory of stochastic shortest path problems [BeT91], guarantees that the successive approximation method will converge to the solution of the above equations, but does not guarantee finite termination. However, here we have special structure that is implied by

the positivity of the cost $g(x)$ and the shortest path character of the problem. A key property in this regard is given in the following proposition, first proved in [Tsi93].

Proposition 1 *Let V be the solution of the system of equations (1.1), (1.2). Let $x \in S$, and let θ and α be such that $V(x) = hg(x)\tau(\theta) + \theta V(x + w_\alpha) + (1 - \theta)V(x + w_{\alpha+1})$. If $\theta > 0$ then $V(x) > V(x + w_\alpha)$. If $1 - \theta > 0$ then $V(x) > V(x + w_{\alpha+1})$.*

Using the above proposition, it was shown in [Tsi93] that a Gauss-Seidel algorithm that cycles through the nodes terminates finitely. Furthermore, a Dijkstra-like algorithm that requires only one iteration per node was proposed. As discussed in [Tsi93], the Dijkstra-like algorithm is much faster in practice as well as in theory than the Jacobi and Gauss-Seidel methods that are typically used to solve stochastic shortest path problems.

The objective of this paper is twofold. First, we introduce a broad class of successive approximation methods, which resembles the label correcting methods used to solve deterministic shortest path problems. These methods can be viewed as Gauss-Seidel methods with the node order used for iteration being arbitrary (not necessarily cyclical). We show that these methods terminate (see Prop. 2 in the next section), although they may require more iterations than the Dijkstra-like algorithm of [Tsi93]. Second, we develop several new label correcting-like methods, which try to approximate the operation of the Dijkstra-like algorithm with smaller overhead per iteration. These methods are patterned after the Smallest Label First (SLF) method of [Ber93] and the Smallest Label First - Last Label Last (SLF-LLL) method of [BGM94], which have been shown experimentally to be very effective for deterministic shortest path problems. We provide computational results showing that these new methods are competitive and often superior to the Dijkstra-like algorithm.

2 Generic Label Correcting Algorithm

In this section we describe a general algorithm for solving the stochastic shortest path problem corresponding to the discretization discussed in section 1. The algorithm, referred to as *generic*, is patterned after a generic label correcting method for deterministic shortest path problems; see for example [GaP88], [Ber91]. It maintains a list of nodes L called *the candidate list*, and a *label* $V(x)$ for each node $x \in S \cup B$. Each label is either a real number or ∞ . Initially,

$$\begin{aligned} L &= B, \\ V(x) &= f(x), \quad \forall x \in B, \\ V(x) &= \infty, \quad \forall x \in S. \end{aligned}$$

For convenience, we also keep track of the direction $(\alpha(x), \theta(x))$ along which the current label of x was calculated. The algorithm proceeds in iterations and terminates when L is empty. The typical iteration of the algorithm (assuming that L is not empty) is as follows:

Typical Iteration of the Generic Label Correcting Algorithm

Remove a node x from the candidate list L . For each neighbor y of x that also belongs to S , if $V(y) > V(x)$, calculate

$$\tilde{V}(y) = \min_{\{\alpha=1,2,3,4|x \in \{y+w_\alpha, y+w_{\alpha+1}\}\}} \min_{\theta \in [0,1]} [hg(y)\tau(\theta) + \theta V(y+w_\alpha) + (1-\theta)V(y+w_{\alpha+1})].$$

Let $(\tilde{a}, \tilde{\theta})$ be the direction for which the minimum value $\tilde{V}(y)$ is obtained. If $V(y) > \tilde{V}(y)$, then set $V(y) = \tilde{V}(y)$, $a(x) = \tilde{a}$, $\theta(x) = \tilde{\theta}$, and add y to L if it is not already in L .

It can be seen that in the course of the algorithm the labels are monotonically nonincreasing and that $V(x) < \infty$ if and only if x has entered the candidate list at least once. The following lemma gives the main properties of the algorithm:

Proposition 2 1. *At the end of each iteration the following conditions hold:*

- (a) $V(x) = f(x)$ for all $x \in B$. Furthermore, nodes in B do not re-enter the candidate list once removed.
- (b) For all $x \in S$, if $V(x) < \infty$, then

$$V(x) \geq hg(x)\tau(\theta) + \theta V(x+w_{\alpha(x)}) + (1-\theta)V(x+w_{\alpha(x)+1}). \quad (2.1)$$

- (c) If for a node $x \in S$, there is a quadrant α such that $x+w_\alpha \notin L$ and $x+w_{\alpha+1} \notin L$, then we have:

$$V(x) \leq \min_{\theta \in [0,1]} [hg(x)\tau(\theta) + \theta V(x+w_\alpha) + (1-\theta)V(x+w_{\alpha+1})]. \quad (2.2)$$

- 2. *The algorithm terminates. The set of labels $\{V(x) \mid x \in S \cup B\}$ obtained upon termination solves Eqs. (1.1) and (1.2).*

Proof:

- 1. Condition (a) holds since by the rules of the algorithm the labels of the border nodes cannot change, and only nodes in S can re-enter the candidate list L .

We prove (b) as follows: Just after the label of x is reduced, we have

$$V(x) = hg(x)\tau(\theta) + \theta V(x+w_{\alpha(x)}) + (1-\theta)V(x+w_{\alpha(x)+1}).$$

Since the labels of the neighboring nodes of x are nonincreasing, we see that Eq. (2.1) holds in subsequent iterations until the value of x is recalculated.

To prove (c), note that initially the nodes not in L have infinite labels. Therefore, (c) holds trivially at the beginning of the algorithm. Let us now fix a node x and its two neighbors $x+w_\alpha$, $x+w_{\alpha+1}$ of some quadrant α . If neither of the nodes $x+w_\alpha$, $x+w_{\alpha+1}$ enters the candidate list L throughout the algorithm, then (c) holds since the label of x is nonincreasing. Otherwise, at least one of the nodes $x+w_\alpha$, $x+w_{\alpha+1}$ enters L at some time. Consider now an iteration k of the algorithm where node $x+w_\alpha$, or node

$x + w_{\alpha+1}$ exits L and as a result both nodes are not in L . At this iteration, the label of x is recalculated and Eq. (2.2) is satisfied. Furthermore, Eq. (2.2) is satisfied in all subsequent iterations of the algorithm until either termination is reached or one of the nodes $x + w_{\alpha}$, $x + w_{\alpha+1}$ re-enters L at some iteration $k' > k$. This is because, for all iterations performed after k until either termination or iteration k' is reached, the labels of $x + w_{\alpha}$, $x + w_{\alpha+1}$ remain unchanged while the label of x is nonincreasing. We conclude that Eq. (2.2) holds throughout the algorithm.

2. We assume that the algorithm does not terminate in order to reach a contradiction. Let I be the set of nodes that enter L a positive but finite number of times, and let \bar{I} be the set of nodes that enter L an infinite number of times. Let also J be the set of nodes that never enter L and whose labels are infinite throughout the algorithm. The sets I , \bar{I} , and J , and the labels of the nodes in I remain unchanged after some iteration denoted \bar{m} . Furthermore, the sets I and \bar{I} are nonempty since from Prop. 2.1(a), $B \subset I$, implying that I is nonempty, and the algorithm does not terminate, implying that \bar{I} is nonempty. Each time a node enters L , its label is smaller than the preceding time it entered L . Since the label of a node is bounded below by zero, we conclude that the labels of the nodes in $I \cup \bar{I}$ converge. For a node $x \in I \cup \bar{I}$, let $V^{\infty}(x)$ denote the limiting value of the label of x , and let $\alpha^{\infty}(x)$, $\theta^{\infty}(x)$ be such that

$$V^{\infty}(x) = hg(x)\tau(\theta^{\infty}(x)) + \theta^{\infty}(x)V^{\infty}(x + w_{\alpha^{\infty}(x)}) + (1 - \theta^{\infty}(x))V^{\infty}(x + w_{\alpha^{\infty}(x)+1}). \quad (2.3)$$

We also define the set $\mathcal{P}(x)$ of *desired nodes of x* as follows:

$$\mathcal{P}(x) = \begin{cases} \{x + w_{\alpha^{\infty}(x)}\} & \text{if } \theta^{\infty}(x) = 1 \\ \{x + w_{\alpha^{\infty}(x)+1}\} & \text{if } \theta^{\infty}(x) = 0 \\ \{x + w_{\alpha^{\infty}(x)}, x + w_{\alpha^{\infty}(x)+1}\} & \text{otherwise} \end{cases}$$

We observe that, in view of Eq. (2.3), we have

$$\mathcal{P}(x) \cap \bar{I} \neq \emptyset, \quad \forall x \in \bar{I},$$

since, after iteration \bar{m} , for each $x \in \bar{I}$, the labels of the nodes in $\mathcal{P}(x) \cap I$ remain constant, while the label of x decreases infinitely many times; if all nodes of $\mathcal{P}(x)$ were in I , Eq. (2.3) would be violated. Thus, each node $x \in \bar{I}$ has at least one desired node in \bar{I} . Consequently, there exists a cycle of nodes of \bar{I} , say $(x_1, x_2, \dots, x_k, x_{k+1})$ such that $x_{k+1} = x_1$ and x_{i+1} is a desired node of x_i , for $i = 1, \dots, k$. From Prop. 1 we have that $V^{\infty}(x_i) > V^{\infty}(x_{i+1})$ for all $i = 1, \dots, k$, which is a contradiction. Thus the algorithm terminates.

Next we show that all nodes $x \in S$ will exit L at least once, so that they must be finite upon termination. In particular, let us consider a node $x_1 \in S$ and a sequence of nodes $(x_1, \dots, x_k, x_{k+1})$ such that $x_{k+1} \in B$ and $x_{i+1} \in N(x_i)$, for all $i = 1, \dots, k$. As discussed in Section 1, such a node sequence exists for every $x_1 \in S$. Since the algorithm terminates, the node x_{k+1} must exit L , since it belongs to L initially. When node x_{k+1} exits L , the

label of node x_k is calculated and, if x_k has never before entered L , its label becomes finite and it enters L . Repeating this argument using x_k in place of x_{k+1} , and proceeding similarly we can prove that each of the nodes x_{k-1}, \dots, x_1 will exit L at least once.

We finally note that upon termination, the list L is empty, so that both Eqs. (2.1) and (2.2) hold for all $x \in S$, implying that the labels of the nodes satisfy the optimality conditions (1.1). **Q.E.D.**

Different algorithms can be derived from our generic label correcting algorithm by using different ways to choose the node that exits the candidate list at each iteration. In particular, if a node with the smallest label is chosen to exit the list, then we obtain an analog of Dijkstra's algorithm. The key property of this algorithm is that once a node exits the list, it never re-enters it, as shown in [Tsi93]. Other possibilities are to organize the list according to the schemes described in [Ber93] and [BGM94]. In the following section we discuss implementations of such schemes and give some computational results.

3 Implementing Dijkstra-Like and Label Correcting-Like Algorithms

In this section we compare the Dijkstra-like algorithm discussed above and two recently developed label correcting algorithms (the SLF-LLL and the SLF-LLL-Threshold methods to be presented below). The Dijkstra algorithm performs at most one iteration per node, but requires some extra overhead per iteration. In the label correcting methods, the number of iterations is larger than for the Dijkstra algorithm, but the overhead is smaller per iteration. The computational results show that there are cases where the label correcting algorithms outperform the Dijkstra-like algorithm. In particular, on computer architectures where the computation overhead is small due to the speed of the floating point operations, the label correcting algorithms are strongly favored. This is encouraging for one more reason: label correcting algorithms are parallelizable whereas the Dijkstra algorithm is not. The recent paper of Bertsekas, Guerriero, and Musmanno [BGM94] explores the parallelization of the label correcting algorithms we consider for the case of deterministic shortest path problems. Their computational results show that parallelization leads to superlinear speedup, reducing dramatically the total number of iterations the label correcting algorithms needed to converge. In any case, the algorithms that we present here are much faster than the traditional algorithms described in [KuD90], where the nodes are picked for iteration according to a certain fixed order.

3.1 The Algorithms

We describe the algorithms that we tested and we discuss several issues related to their efficient implementation.

The Dijkstra Algorithm: In our implementation, we maintain the candidate list as a binary heap. At each iteration, the node removed from the list is the top node i of the heap, which is the node with the smallest label. As discussed above, node i is permanently labeled and never enters the heap again. The labels of its temporary labeled neighbors are computed and the respective nodes are inserted in the binary heap if they are not already there. Observe that if a node i becomes permanently labeled, then at most three of its neighbors may enter the binary heap, since at least one of its neighbors, the one having smaller label than node i , is permanently labeled. The iterations proceed until the list is empty. Our binary heap code is based on the SHEAP code for deterministic shortest paths of Gallo and Pallotino [GaP88]. In the presentation of our results we will refer to the implementations of the Dijkstra-like algorithm with the prefix “DIJ-”.

The Label Correcting-like Algorithms:

1. The SLF-LLL Method: In this algorithm, the candidate list is maintained as a queue. Nodes enter the list according to the following *Smallest Label First (SLF)* criterion first proposed in [Ber93]:

SLF criterion: *Whenever a node i not already in the list enters the list, its label $V(i)$ is compared to the label $V(j)$ of the top node j of the list. If $V(i) \leq V(j)$ node i is inserted at the top of the list; otherwise node i is inserted at the bottom of the list.*

This criterion for insertion in the list is combined with another simple criterion for the extraction of a node from the list. The extraction criterion is called *Large Label Last (LLL)*, and was proposed in [BGM94]. The node at the top of the list exits the list. If its label is larger than the average of the labels of all the nodes in the list, then the node is reinserted at the bottom of the list and its successor node in the list is extracted. This procedure continues until a node with label less than or equal to the average of the labels of all the nodes in the list is found. The labels of the neighboring nodes of the node extracted from the list are recomputed and the corresponding nodes are inserted in the list according to the SLF criterion.

2. The SLF-LLL-Threshold Method: This is a combination of the SLF-LLL method with the threshold method of Glover et. al. [GGK86], (see [BGM94]). Here, the candidate list maintained by the SLF-LLL method is divided in two lists: The nodes on the first list are the ones that have labels less than or equal to some threshold value. Nodes are inserted in both lists according to the SLF criterion. Nodes are extracted only from the first list according to the LLL criterion. When the first list becomes empty, the threshold is appropriately increased and nodes with label less than or equal to the new threshold are removed from the second list and inserted to the first list according to the SLF criterion. The method we chose for increasing the threshold is the following: Initially, the threshold is set to the minimum node cost $\min_{x \in S} g(x)$ plus a user-chosen percentage of the maximum node cost $\max_{x \in S} g(x)$. Each time the first list empties, the threshold is

increased by the user-chosen percentage of the maximum node cost. If this increase was not sufficient to transfer any nodes of the second list to the first list, the threshold is set equal to the minimum node label in the second list plus the user-set percentage of the maximum node cost. In the presentation of the computational results we will refer to the SLF-LLL-Threshold algorithm with the prefix “SLF-LLL-TH-”.

In all the above implementations we have used certain tests that help avoid unnecessary recomputations of the labels of the nodes. The reduction in computation obtained by these tests is significant. Specifically, recall that the label of a node x is computed by performing minimizations along four quadrants (parameter α) and the minimization in a particular quadrant is based on the labels of two neighboring nodes. In particular, we make the following observations:

- When a node exits the list, only its neighbors with larger label need to have their label recomputed (see Prop. 1).
- The recomputation of the label of a node needs to be made only in the quadrants that involve the neighboring node that exited the list at the beginning of the current iteration.
- If the labels of the two neighboring nodes of node x in a particular quadrant have remained the same since the last time that the label of x was calculated, no minimization using that quadrant is needed. We consider implementations of the algorithms both with and without keeping track of the labels of neighboring nodes. In the presentation of the computational results we will distinguish the two implementations of an algorithm by appending the suffixes “-NEIGH” and “-NO-NEIGH”, respectively, to the name of the algorithm.
- Let two neighboring nodes of node x in a particular quadrant be i and j . If i has greater label than the current label of x , then the simpler label update $V(x) := \min\{V(x), hg(x) + V(j)\}$ can be used. Such a label evaluation will be referred to as *simplified*.

The implementations of all algorithms solve the problem of finding the optimal costs from all nodes $x \in S$ to some border point \bar{x} . In particular, we have set $f(\bar{x}) = 0$, and for all border nodes $x \neq \bar{x}$ we have set the cost $f(x)$ to be a very large number.

3.2 The Test Problems

We implemented a test problem generator called GRIDQUAD. The test problems generated are grids resulting from a square discretization of parallelograms with sides whose length is an integer multiple of the discretization step h . Thus G is the interior of the parallelogram and ∂G is the border of the parallelogram. The corresponding discretizations S and B are easy to obtain. Each node $x \in S$ has an associated cost which is the discretized version of the positive cost function $r(x)$ defined on G . The program GRIDQUAD uses a quadratic function for $r(x)$. The cost of all boundary nodes $x \in B$ is (effectively) infinity except for the two neighbors of the top righthand corner of the parallelogram border. To make the test problems more interesting,

we have introduced *obstacles*, i.e., interior points of the parallelogram with infinite cost. These are considered to be part of ∂G . In our test problems, obstacles occur in rows; all the points on a row have infinite cost except for a segment of length kh where k is some positive integer. Obstacle rows are equally spaced at a distance which is an integer multiple of the discretization step h .

3.3 Test Results

We present the results of some of the computational experiments. The test problems are described by giving the dimensions of the grid and the number of obstacles. Costs of interior points are in the range [1-1000] and are generated by a quadratic function which has its maximum at the center of the grid. For each algorithm we tested two implementations, one that keeps track of the values of neighboring nodes in order to avoid unnecessary computations (suffix -NEIGH) and one that does not maintain such information (suffix -NO-NEIGH). We report running times on a Macintosh Powerbook 170 and on an Alpha-Dec computer running a version of UNIX. We also report the number of iterations, label calculations, and simplified label calculations performed by each algorithm to facilitate the comparison between the algorithms.

Our computational experiments indicate that the label correcting algorithms are competitive with the Dijkstra algorithm. The SLF-LLL-TH algorithm outperforms consistently the Dijkstra algorithm, while the SLF-LLL algorithm outperforms the Dijkstra algorithm on the Alpha-Dec computer and is a bit slower than the Dijkstra algorithm on the Mac. On problems with obstacles, all algorithms have similar running times. This is because the addition of obstacles has an effect similar to breaking the problem to several smaller-sized problems. Therefore, the data structures maintained by the algorithms for the choice of the node to exit the list contain few nodes and thus need less overhead, leading to similar running times. The differences in running times are more evident on problems without obstacles.

One final observation is that times on the Alpha-Dec seem to favor the SLF-LLL methods a lot more than on the Mac. The extra iterations that the label correcting algorithms require are outweighed on the Alpha-Dec by the amount of time spent on maintaining the heap in the Dijkstra method. A similar effect is observed when we keep extra data to reduce calculations (-NEIGH vs. -NO-NEIGH). Although the number of time consuming minimizations as in Eq. (1.1) is reduced by keeping labels of neighbors, the running time is not improved. This may be due to a higher penalty for memory access operations than for computations in the Alpha-Dec.

3.4 Conclusions

The computational results indicate that the methods presented in this paper are very efficient and are apparently much faster than traditional Gauss-Seidel methods that cycle through the nodes in a fixed order. Furthermore, the label correcting methods developed are parallelizable as in [BGM94] and will likely lead to efficient parallel algorithms. Finally, the label correcting methods we presented may also be used in the case where the cost function is of the form $r(x, u)$. In this case the theory of [Tsi93] cannot be applied and a Dijkstra-like algorithm is not possible. However, the label correcting algorithms we proposed can be used as heuristics that

specify the order in which the label updates are performed. Their efficient implementation is an interesting subject for further research.

References

- [Ber91] Bertsekas, D. P., *Linear Network Optimization*, M.I.T. Press, 1991.
- [Ber93] Bertsekas, D. P., "A Simple and Fast Label Correcting Algorithm for Shortest Paths," *Networks*, Vol. 23, pp. 703-709, 1993.
- [BeT89] Bertsekas, D. P., and Tsitsiklis, J. N. *Parallel and Distributed Computation*, Prentice Hall, 1989.
- [BeT91] Bertsekas, D. P., and Tsitsiklis, J. N. "An Analysis of Stochastic Shortest Path Problems," *Math. of Operations Research*, Vol. 16, 1991, pp. 580-595.
- [BGM94] Bertsekas, D. P., Guerriero, F., and Musmanno, R., "Parallel Asynchronous Label Correcting Methods for Shortest Paths," *Laboratory for Information and Decision Systems Report, LIDS-P-2250, M.I.T., May 1994, to appear in JOTA.*
- [KuD90] Kushner, H. J., and Dupuis, P. G., *Numerical Methods for Stochastic Control Problems in Continuous Time*, Springer-Verlag, New York, 1992.
- [GaP88] Gallo, G., S., and Pallotino, S., "Shortest Path Algorithms," *Annals of Operations Research*, Vol. 7, pp. 3-79, 1988.
- [GGK86] Glover, F., Glover, R., and Klingman, D., "The Threshold Shortest Path Algorithm," *Networks*, Vol. 14, No. 1, 1986.
- [Tsi93] Tsitsiklis, J. N., "Efficient Algorithms for Globally Optimal Trajectories," *Laboratory for Information and Decision Systems Report, LIDS-P-2210, M.I.T., October 1993, to appear in IEEE Transactions on Automatic Control.*

Problem	DIJ-NO-NEIGH	DIJ-NEIGH	SLF-LLL-NO-NEIGH	SLF-LLL-NEIGH
150X150 No Obst.	22.02/0.38 21904/43230/43796	20.8/0.52 21904/21721/43796	25.45/0.42 32976/61605/69195	23.65/0.4 32976/32976/69195
150X150 3 Obst.	20.73/0.32 21474/41469/42977	19.60/0.45 21474/20873/42977	20.20/0.33 26131/49214/53138	18.37/0.31 26131/29234/53138
100X200 No Obst.	19.27/0.29 19404/38289/38787	18.08/0.35 19404/19193/38787	22.95/0.37 30729/54588/67360	22.18/0.35 30729/37368/67360
50X250 No Obst.	11.47/0.169 11904/23232/23794	10.73/0.177 11904/11650/23794	14.48/0.244 19505/33574/43556	14.18/0.23 19505/24553/43556
100X225 No Obst.	21.70/0.33 21854/43057/43719	20.47/0.388 21854/21609/43719	26.67/0.427 35777/63405/78777	26.15/0.4 35777/44971/78777
100X225 3 Obst.	20.42/0.32 21387/41191/42783	19.25/0.406 21387/21387/42783	19.60/0.371 25590/47856/52116	18.08/0.346 25590/29373/52116
300X300 No Obst.	*/2.05 88804/176412/177614	*/2.63 88804/88450/177614	*/1.73 132476/248243/280171	*/1.71 132476/155172/280171
300X300 3 Obst.	*/.86 87924/172819/175927	*/2.44 87924/86716/175927	*/1.5 113448/216430/232960	*/1.476 113448/131261/232960
500X500 No Obst.	*/6.99 248004/493991/496035	*/8.49 248004/247471/496035	*/5.25 394289/742420/830920	*/5.38 394289/344023/830920

Table 1: Computational results on the Mac and Alpha-Dec. On each problem line, the numbers on top are the times in seconds on the different machines, and the numbers at the bottom are the Iterations/label Calculations/Simplified label Calculations. All problems were generated by the GRIDQUAD program. The cost assigned to a grid position (x, y) which is not a boundary point or an obstacle is given by $1001 - 1000 \frac{10(x-x_0)^2 + 40(y-y_0)^2}{10(x_0+1)^2 + 40(y_0+1)^2}$, where (x_0, y_0) is the central point of the grid.